



ALMA
ALMATY MANAGEMENT
UNIVERSITY



GLOBAL
EDUCATION
FUTURES

UNIVERSITIES AFTER AI

Voices from the Global Frontier

Collection of essays edited by
Pavel Luksha and Sholpan Tazabek
with foreword by Sayasat Nurbek

Universities After AI

Voices from the Global Frontier

Edited by Pavel Luksha and Sholpan Tazabek

Almaty
Smart University Press
2026

UDC 378.14:004.8
LBC 74.58:32.813
L96

Recommended for publication by the University Council
EI “Almaty Management University”
(Protocol No. 11, June 1, 2026)

Reviewers:

Arman Komekbayev
Dean of the School of Digital Technologies and Economics,
Almaty Management University

Timur Imankulov
PhD, Associate Professor, Dean of the Faculty of Information
Technologies and Artificial Intelligence,
Al-Farabi Kazakh National University

Edited by Pavel Luksha and Sholpan Tazabek

L96 Universities After AI: Voices from the Global Frontier/
Edited by Pavel Luksha, Sholpan Tazabek. — Almaty: Smart
University Press, 2026.—192 p.
ISBN 978-601-08-6218-0

“Universities After AI” brings together university leaders and scholars from Kazakhstan, Kyrgyzstan, Georgia, Bangladesh, Mexico, South Africa, France, the United States, and beyond. The volume argues that AI is not merely a new educational tool but a catalyst for rethinking the industrial university model, challenging assumptions about information scarcity, stable professions, and the value of credentials. Combining theoretical analysis with practical governance approaches, it examines the implications of AI for higher education and presents institutional experiments from emerging and mid-tier universities. Contributors converge on a central idea: in an age of machine cognition, the university’s enduring role is to cultivate human judgment, agency, ethics, and the capacity to navigate complexity.

UDC 378.14:004.8
LBC 74.58:32.813

© Almaty Management University, 2026
ISBN 978-601-08-6218-0 © Smart University Press, 2026

Declaration

This volume is an original scholarly publication resulting from an international academic partnership between the International Association of University Presidents (IAUP), Almaty Management University (Kazakhstan), and Global Education Futures.

The editors and contributors hereby declare that all chapters contained in this volume are independent, original works of the respective authors and have not been previously published in whole or in part, nor are they simultaneously submitted for publication elsewhere, except where explicitly stated. All sources, references, and intellectual contributions of third parties have been appropriately acknowledged in accordance with established academic standards and publication ethics.

The authors further confirm that they bear sole responsibility for the content, accuracy, and originality of their individual contributions. The editors are responsible for the coordination and compilation of the volume and do not assume responsibility for the specific claims or interpretations presented in individual chapters.

This publication is issued within the framework of a collaborative initiative between IAUP, Almaty Management University, and Global Education Futures to advance global dialogue and knowledge exchange in higher education and related fields.



**ҚАЗАҚСТАН РЕСПУБЛИКАСЫ
ҒЫЛЫМ ЖӘНЕ ЖОҒАРЫ БІЛІМ МИНИСТРЫ**

**MINISTRY OF SCIENCE AND HIGHER EDUCATION
OF THE REPUBLIC OF KAZAKHSTAN**

Foreword

Every generation encounters a technological transformation that compels it to reconsider the purpose of education. For our generation, artificial intelligence is such a transformation.

The emergence of AI challenges many assumptions upon which modern higher education has been built. For centuries, universities have served as institutions for the creation, transmission, and application of knowledge. Today, while knowledge remains essential, the conditions under which it is created and accessed are changing at an unprecedented speed. Universities are confronted not simply with the task of adopting new tools, but with a deeper challenge: rethinking how they prepare humanity for a world in which cognitive labor is increasingly shared with machines.

Kazakhstan has chosen to engage with this challenge directly and with absolute urgency. Guided by our newly adopted national strategy, “Digital Kazakhstan” we are systematically transitioning to an AI-driven economy. We recognized early on that AI must be governed, not merely deployed. In 2025, Kazakhstan adopted one of the first dedicated AI laws in our region, establishing clear frameworks for algorithmic transparency and human oversight.

However, effective governance is not simply about regulation; it is about institutional agility. As I have emphasized in dialogues with our international partners, such as the Asian Universities Alliance, true university-industry integration in the AI era happens only when governance enables innovation rather than resists it. Universities must transform from rigid, legacy bureaucratic structures into flexible, resilient ecosystems that can pivot at the speed of technological change.

Our ambition, therefore, is fundamentally a human capital and infrastructure program designed to build this agile ecosystem. The Ministry of Science and Higher Education of Kazakhstan has completely reimaged the talent pipeline. We have introduced AI education across the higher education system, integrating it as a mandatory component in national educational standards.

Through our flagship national initiative, the AI-Sana program, we are building an AI-ready workforce at an extraordinary scale. To date, more than 673,000 university students have earned certificates in basic AI literacy. Moving beyond basic exposure, every year we are empowering 100,000 ambitious students to transition from AI consumers to AI creators, engaging them in technology entrepreneurship and industry-specific problem-solving.

The results of this proactive stance are already visible across our institutions. In 95 universities nationwide, student teams have successfully developed and implemented

308 functional AI agents —ranging from AI admissions assistants and automated library systems to AI-driven proctoring and mental health support platforms. To accelerate this momentum, we have forged strategic partnerships with global technology leaders, integrating thousands of Coursera courses into university curricula and deploying ChatGPT Edu licenses across our higher education system to empower our faculty.

Simultaneously, we are building the sovereign digital infrastructure necessary to anchor this talent. In 2025, Kazakhstan launched the national supercomputer cluster, delivering up to 6 EFLOPS of computing power. This year, we are taking the next definitive step by establishing a dedicated Artificial Intelligence Research University in Astana, directly integrated with our national exponential computing clusters and the Alem.ai platform.

We are executing this technological leap while deepening Kazakhstan's integration into global knowledge networks. By establishing 40 international branch campuses in partnership with the world's leading institutions, we are rapidly transforming Kazakhstan into a central Academic and Research Hub for Eurasia. Combined with our goal of attracting 100,000 international students by 2029, these efforts reflect our ambition to serve a broader region home to more than a billion young people.

Yet, amidst this massive investment in digital infrastructure and startup ecosystems, we remain clear-eyed about where the most decisive investment lies: in the holistic formation of the human being. The question is not only whether our classrooms have AI platforms, but whether our universities are producing graduates possessing the critical judgment and ethical grounding to govern these technologies.

For this reason, I deeply welcome this publication. The International Association of University Presidents (IAUP), Almaty Management University, and Global Education Futures have created a vital space for university leaders to engage with questions that no institution can answer alone. I particularly commend AlmaU's pioneering work with the "Tolyq Adam" framework—a holistic educational model rooted in the rich philosophical traditions of Central Asia. It powerfully demonstrates that our region is not just importing technological solutions, but actively exporting profound, culturally grounded educational innovations to the global stage.

No country yet has a complete blueprint for the university that the AI era demands. But the institutions that will shape that blueprint are those that treat transformation as their own responsibility. Artificial intelligence will undoubtedly transform universities. The deeper question is whether our societies will be the authors of that transformation or merely its subjects. This volume, much like Kazakhstan's own systemic reforms, makes a serious contribution to ensuring that the answer is the former.

Sayasat Nurbek

Minister of Science and Higher Education

Republic of Kazakhstan

June 2026



Preface

Universities at the convergence point

Universities today face not one disruption but several, arriving simultaneously. Artificial intelligence is transforming how knowledge is produced, evaluated, and economically valued. However, it arrives into a world already destabilized by geopolitical fragmentation, ecological emergencies, democratic erosion, and labor markets in structural flux. For many regions, including Central Asia, South Asia, sub-Saharan Africa, Latin America, these pressures coincide with what should have been a period of extraordinary opportunity: demographic dividends, rapid urbanization, expanding middle classes, and growing demand for higher education. The conditions that enabled three decades of university expansion are now shifting beneath institutions that have barely had time to consolidate.

At the International Association of University Presidents' (IAUP) 60th anniversary conference in Seoul in October 2025, university leaders confronted a shared realization: the pressures facing higher education are no longer isolated disruptions, but converging transformations affecting the foundations of the institution itself.

Some participants emphasized the need to defend core academic values under conditions of accelerating change. Ted Mitchell, President of the American Council on Education, was unequivocal: "Independence, intellectual inquiry and academic excellence, inclusion, the expectation that education is a good to be consumed by all of our citizens — these are things that we cannot compromise on". Others argued that defending existing structures would not be sufficient. Korea's Education Minister Choi Kyo-jin insisted that "universities must go beyond the simple transmission of knowledge to create new value and nurture global talent capable of addressing humanity's shared challenges". IAUP Vice-President Lee Eun-joo framed the stakes in institutional terms: "What we need today is not only reflection but a fundamental reimagining of the purpose and direction of higher education, even to envision something that has not yet existed".

The crisis is not only technological. Arturo Cherbowski, executive director of Santander Universities, warned the conference that the challenge has "existential dimensions", and that many institutions missed the transformation opportunity offered by the pandemic, reverting to pre-pandemic operations with little sense of urgency. The financial pressure is structural: "For the last 20 years, public money for universities has consistently gone down, and this is not a syndrome of one particular country". Meanwhile, the rising influence of politics on higher education, what can be described as "three Ps" — polarization, populism, and politicization — makes it increasingly difficult to persuade governments to increase funding, even as the demands on universities multiply.

In this context, AI-driven disruption now threatens to make institutions obsolete if they fail to reinvent their curriculum, teaching, and operating models. As Paul LeBlanc, former president of Southern New Hampshire University, observed at the Ellucian Summit in Denver in April 2026, faculty members will need to shift from being facilitators of knowledge to becoming mentors who guide the use of artificial intelligence (AI) with a

humanistic, critical, and cautious perspective that requires building stronger and more meaningful relationships with students rather than competing with machines on information delivery. Nine out of ten faculty members already believe they need training on AI and identify lack of knowledge as the main barrier to adoption. The transformation, as LeBlanc argues, requires profound human relations. AI will not substitute the transformational impact of faculty.

What emerged across these conversations was not a single solution, but a growing recognition that higher education can no longer respond through fragmented reforms or isolated technological adoption. Questions of pedagogy, governance, technology, legitimacy, and societal purpose are becoming inseparable. Under the leadership of IAUP President Shawn Chen, the IAUP — the world’s only international organization led exclusively by university presidents — has therefore made this integrated response a central priority, emphasizing that “IAUP is uniquely positioned to set standards, share best practices, and ensure that artificial intelligence serves civility, inclusivity, equity, and humanity”.

The IAUP Central Asia Regional Conference in Almaty in June 2026 continues this work under the theme “Impact-Driven Universities: Education, Leadership and Societal Change”. Bringing together university presidents, ministers, accreditation leaders, entrepreneurs, and regional innovators, the conference seeks to move beyond diagnosis toward institutional experimentation and action.

Almaty Management University (AlmaU), as host of the IAUP Central Asia Regional Conference 2026 and institutional home of this collective book project, has confronted these questions under conditions familiar to many emerging higher education systems: limited resources, multilingual student populations, and rapidly shifting labor-market demands. Rather than waiting for a stable global model to emerge, AlmaU has used this period to experiment with new educational approaches, including the development of the “Tolyq Adam” concept — a graduate profile rooted in Central Asian intellectual traditions rather than imported institutional templates. These efforts are unfolding within a broader national shift. Kazakhstan adopted one of the region’s first dedicated AI laws in 2025 and introduced a nationwide mandate for AI education across universities, placing institutions under immediate pressure to rethink curriculum, pedagogy, and graduate preparation.

In collaboration with Global Education Futures, the project explored with university leaders, faculty, and students how universities must evolve when the cognitive infrastructure of society itself is being transformed. The findings suggest that AI does not simply add new tools to an existing educational model. It destabilizes assumptions that industrial higher education was built around: that expertise is scarce, professions remain relatively stable, and degrees reliably certify capability. As knowledge production becomes increasingly automated and distributed, universities are being forced to confront a deeper question: if information delivery is no longer their defining role, what remains essential?

The contributors to this volume — from Kazakhstan, Kyrgyzstan, Georgia, Bangladesh, Mexico, South Africa, the United States, France, and beyond — respond to this question from the perspective of the vast majority of the world’s universities: institutions operating under constraint, serving highly diverse student populations, and navigating the hardest version of the transition. Across very different contexts, a common insight emerges. As AI systems become more capable, the enduring role of the university becomes more deeply human. Universities justify themselves not through information delivery alone, but through what they cultivate in people: judgment, agency, ethical orientation, and the capacity to act meaningfully under conditions of complexity.

This volume does not offer a single blueprint. It offers a structured encounter between institutions navigating a shared transition under profoundly different conditions. A companion volume provides the broader foresight framework and scenario architecture behind this work. Together, they are offered as resources for university leaders seeking not merely to react to disruption, but to shape the direction of the transition itself.

Fernando León García

President, CETYS University System (Mexico)
Board Chair, International Association of
University Presidents (IAUP)

Pavel Luksha

Founder and CEO
Global Education Futures

Assylbek Kozhakhmetov

Founder and President
Almaty Management University
(Kazakhstan)

June 2026

Acknowledgments

The editors gratefully acknowledge the International Association of University Presidents, Almaty Management University, and Global Education Futures for their collaboration and institutional support in the development of this volume.

We extend particular gratitude to Almaty Management University for its financial support, which made the preparation and publication of this volume possible. This contribution was essential in advancing the project and ensuring its timely completion.

We sincerely thank all contributing authors for their scholarly work, intellectual engagement, and commitment to this collective endeavor. Each chapter represents an independent contribution, and together they form a comprehensive exploration of transformations in higher education.

Finally, we recognize the efforts of the editorial and production teams for their professionalism, dedication, and careful attention to detail throughout the preparation of this publication.

List of contributors

Alexander M. Sidorkin is Chief AI Officer and Director of the National Institute on Artificial Intelligence in Education and formerly Dean and Professor at California State University, Sacramento, USA. A scholar in pedagogy and AI-driven instructional innovation, he is the author of *Embracing Chatbots in Higher Education* and *The Pedagogy of Relation*. He serves as President-Elect of the Philosophy of Education Society.

Aigul Sarenova is Director of “Tolyq Adam” Institute at Almaty Management University, Kazakhstan. She is co-author of Chapter 10 on the *Tolyq Adam framework*, university’s holistic graduate profile rooted in Central Asian intellectual traditions.

Assel Aryn is Vice-Rector for Academic Development at Almaty Management University, Kazakhstan. She leads reforms in curriculum design, quality assurance, and university governance. She also serves on the Board of the Academic Integrity League, a network of more than 30 universities in Kazakhstan.

Assylbek Kozhakhmetov is Founder and President of Almaty Management University, Kazakhstan. He holds a DBA from the Higher School of Economics, Russia, and completed a fellowship program under the Bolashak Presidential Scholarship at University of California, Berkeley, USA. He also serves as Vice President of CEEMAN for Central Asia and is a member of the Executive Committee of the International Association of University Presidents (IAUP)

Charles Fadel is Founder and Chairman of the Center for Curriculum Redesign in Boston, USA. Widely recognized for his work on competency-based learning and curriculum transformation, he is the co-author of *Education for the Age of AI* (2024), *Artificial Intelligence in Education* (2019), and *Four-Dimensional Education*. He has worked with education systems in more than 30 countries and is a member of the OECD AI Expert Group

Fernando León García is President of CETYS University System, Mexico, and Chair of the Board of the International Association of University Presidents (IAUP). He holds a doctorate in Administration and Educational Policy from Stanford University, USA

François Taddei is President and Founder of the Learning Planet Institute (formerly CRI), France, and Co-Director of the UNESCO Chair in Learning Sciences at Université Paris Cité. A systems biologist by training, he is one of the leading international voices on learning ecosystems, collective intelligence, and educational transformation. His work has influenced educational innovation initiatives across Europe and beyond through collaborations with UNESCO, OECD, and numerous public institutions.

Ivan Ninenko is Head of Innovation Department at International University of Central Asia, Kyrgyzstan. He is author of Chapter 9 on the “Learn with AI, test without AI” pedagogical experiment, focusing on AI integration in resource-constrained educational environments

Jack Gao is a researcher affiliated with NEOLAF Inc., New Jersey, USA, and SIAS University, China, working in AI-driven education and the AI+X paradigm, with a focus on applications of artificial intelligence in higher education systems and cross-border educational innovation

Kakha Shengelia is President of Caucasus University, Georgia. He holds an MBA from the University of Hartford and a PhD from Georgian Technical University. He is a Fellow of the World Academy of Arts and Science and a recipient of the Presidential Order of Excellence of Georgia. Formerly served as President of International Association of University Presidents (2019–2022)

Letlhokwa George Mpedi is Vice-Chancellor and Principal of the University of Johannesburg, South Africa. He holds an LLB from Vista University and an LLD in Mercantile Law from the University of Johannesburg, South Africa

Md. Sabur Khan is Chairman of Daffodil International University and Daffodil Group, Bangladesh. He is Treasurer-Elect (2027–2030) of International Association of University Presidents (IAUP) and former President of the Dhaka Chamber of Commerce and Industry and the Bangladesh Computer Samity

Natalia Cebotari — Independent Education Advisor, Former Director of Research and Strategy, Yandex Education. Built universities and national education platforms across Europe and Central Asia (Moldova).

Pavel Luksha is Founder and Director of Global Education Futures, Associate Researcher with Learning Planet Institute, and the Board member of World University Consortium. He is a foresight practitioner and systems thinker with focus on the future of education, work, and institutional transformation, authoring *The Next 100 Years* (2023) and 10 GEF key reports in 2014–25. He is chief editor of this volume and author of the editorial introduction.

Ralph A. Wolff is Founder and President of The Quality Assurance Commons and the Chairman of the Board of World University Consortium. He is a member of the Board of Trustees of Golden Gate University and a member of the World Academy of Art and Science, as well as a member of the Institutional Quality Assurance Board of Dubai University. He previously served as President of the WASC Senior College and University Commission (1996–2013). He holds a Juris Doctor from George Washington University

Raphael Costambeys-Kempczynski is Director of Education at the Learning Planet Institute, France. Previously, he held leadership roles at Sorbonne Nouvelle University, including Vice-President for Student Experience. His work focuses on learner agency, educational innovation, learning ecosystems, and the transformation of higher education in response to societal and planetary challenges. He is co-author of Chapter 5 on planetary agency and *Doughnut Learning*, focusing on sustainability and systems thinking in education

Richard Jiarui Tong is Chair of the IEEE Artificial Intelligence Standards Committee (AISC), IEEE Computer Society, and a leading contributor to IEEE initiatives in AI governance, interoperability, trust, and standards development. He is also active in IEEE AI conference leadership and cross-sector collaborations on AI deployment and standardization.

Shawn Chen is Founder and Chairman of SIAS University, China, and President of the International Association of University Presidents (2024-2027). He holds MBA from UCLA, USA and has completed executive education at Harvard University. He serves in international higher education networks including the Association of Universities of Asia and the Pacific (AUAP) and the World Federation of United Nations Associations (WFUNA) and is co-founder of Plug and Play Tech Center.

Sholpan Tazabek is Vice-Rector for Global Partnerships at Almaty Management University, Kazakhstan. She completed a doctoral fellowship at University of Pennsylvania, USA, and holds a PhD in Education from Nazarbayev University, Kazakhstan. She also serves as a national champion/consultant for the British Council and QAA's transnational education quality framework. She is co-editor of this volume and author of the editorial conclusion.

Timur Schukin is a researcher in philosophy of education and AI epistemology working with Tyumen State University AI lab and SKOLKOVO School of Management, formerly a co-founder of the NeuroNet initiative within Russia's National Technology Initiative. His work explores collective intelligence, human-AI cognitive systems, AI "xenopsychology", and the future of learning. He is author of Chapter 4 on epistemic legitimation and hybrid cognitive units in AI-augmented learning environments.

Vadim Novikov is a faculty member and education analyst at Almaty Management University, Kazakhstan. He also serves as strategic advisor to the President of Almaty Management University. His work focuses on higher education transformation, institutional strategy, and the development of analytical frameworks for understanding global AI-driven education ecosystems

Victoria Haro is Co-founder and Director of the Universidad del Medio Ambiente (UMA), Mexico, one of Latin America's first universities dedicated to sustainability and regenerative education. She holds a PhD from Columbia University and previously served as environmental advisor to the government of Valle de Bravo.

Xiangen Hu is Chair Professor of Learning Sciences and Technologies at The Hong Kong Polytechnic University and Director of the Institute for Higher Education Research and Development. He is also affiliated with NEOLAF Inc., New Jersey, USA, working on AI-driven educational systems and neuro-symbolic cognitive architectures for adaptive learning.

Contents

Declaration.....3

Foreword4

Preface.....6

Acknowledgments.....9

List of contributors.....9

Contents13

Editorial Introduction: Designing the AI-Realist University

Pavel Luksha..... 15

Analytical preface. Mapping the global discourse, manifestos, and institutional debt

Vadim Novikov..... 23

Part I. The structural and epistemic reconfigurations30

 Chapter 1. Labor economics, so-so automation, and relational university

Alexander M. Sidorkin 31

 Chapter 2. After AI, the university cannot be organized around content

Natalia Cebotari..... 40

 Chapter 3. The knowledge trapezoid: expertise, depthification, and transfer

Charles Fadel..... 51

 Chapter 4. The future of education: symbiotic intelligence and the AI+X paradigm

Shawn Chen, Richard Jiarui Tong, Xiangen Hu, and Jack Gao 59

 Chapter 5. Epistemic legitimation and hybrid cognitive units

Timur Schukin 76

Chapter 6. The architecture of Planetary Agency and Doughnut Learning <i>Raphael Costambeys-Kempczynski and François Taddei</i>	89
Chapter 7. AI meets accreditation: lessons from the U.S. for the future of quality assurance <i>Ralph A Wolff</i>	104
Bridging Chapter: Operational Mapping — The AI Ecosystem Map for Universities <i>Vadim Novikov</i>	118
Part II. Institutional Experiments From the Global Frontier	128
Chapter 8. Building AI-ready graduates in a small, fast-shifting economy <i>Kakha Shengelia</i>	129
Chapter 9. The tactical classroom: “Learn with AI, test without AI” in Kyrgyzstan <i>Ivan Ninenko</i>	140
Chapter 10. The “Tolyq Adam” framework: holistic human development in Kazakhstan <i>Assylbek Kozhakhmetov, Assel Aryn, and Aigul Sarenova</i>	148
Chapter 11. Becoming an entrepreneurial, AI-native university in Bangladesh <i>Md. Sabur Khan</i>	155
Chapter 12. Continental impact and indigenous/African legal governance <i>Letlhokwa George Mpedi</i>	161
Chapter 13. University as regenerative ecosystem: a planetary mission experiment in Mexico <i>Victoria Haro</i>	169
Editorial conclusion: Toward an agentic future <i>Sholpan Tazabek and Pavel Luksha</i>	185

Editorial Introduction: Designing the AI-Realist University

Pavel Luksha

*The ‘message’ of any medium or technology
is the change of scale or pace or pattern that
it introduces into human affairs.*

— Marshall McLuhan, *Understanding Media* (1964)

University is one of the oldest continuous social institutions in human history. Scholarly communities dedicated to collective inquiry, knowledge preservation, and elite formation existed in Ancient India, China, the Arab Caliphate, and the Mediterranean world long before the European medieval university took institutional shape. This remarkable longevity across cultures, political orders, and civilizational epochs testifies to the university’s capacity for adaptation. However, adaptation is not permanence. Each era’s university emerged as a response to particular historical conditions, and these conditions that once shaped the modern university are now destabilizing simultaneously.

The contemporary university, in its widely recognized form, is a product of Western European Modernity. It emerged as a vehicle for producing and reproducing the technological, managerial, and intellectual elites that industrial civilization required, while also serving as one of the principal institutions through which modern societies organized, validated, and transmitted knowledge. Its social contract rested on several assumptions that appeared durable: information scarcity, the relative stability of professional roles, slow-moving knowledge cycles, credential-mediated social mobility, and the expectation that economic growth would continue to expand demand for cognitive labor. Universities became gateways into professional life, legitimizing institutions for expertise, and sites for the transmission of accumulated knowledge across generations. They also carried forward, often unwillingly, the limitations of the Modernity paradigm, its colonial and extractive dimensions, reproducing them through research priorities, curriculum structures, and the worldview of their graduates.

Academic institutions have experienced crises for at least a century. From José Ortega y Gasset’s *Mission of the University* to 1960’s student riots, tensions have recurred around what universities promise — cultural leadership, advanced research, and social transformation — and what they actually deliver. Yet despite these challenges, universities flourished during the post-war era, a period sometimes called the Long Peace. Massification of higher education, coupled with growing demand for advanced knowledge, catapulted universities into a pivotal social role. Global tertiary enrolment rose from 32.5 million students in 1970 to over 260 million by 2025. No other institution could match the university’s breadth of education, credentialing, and research. For much of the world — Central Asia, South Asia, sub-Saharan Africa, Latin America — the conditions were exceptionally favorable: demographic dividends, rapid urbanization, expanding middle classes, and templates for institutional development refined through seventy years of massification. Even institutions with modest resources could grow, modernize, internationalize, and offer their students a credible path into the knowledge economy. This was an unprecedented era of higher education expansion — one that persisted despite the model’s internal shortcomings, precisely because the external conditions remained supportive.

Those conditions are now breaking down on multiple fronts. The disruption is often described through the language of artificial intelligence, but AI is better understood not as an isolated technological event but as a catalyst acting upon deeper structural transformations already underway. The internet has already weakened universities' monopoly over access to knowledge. Platform economies have already destabilized labor pathways, and credential inflation has already begun eroding the signaling power of degrees. Public trust in institutions have already fragmented: in the United States, the share of adults expressing confidence in higher education fell from 57 percent in 2015 to 36 percent by 2025. Ecological instability and geopolitical fracture have already started to challenge assumptions of linear progress upon which modern higher education systems were built.

The geopolitical order that made the modern university possible — international mobility of scholars, political protection of academic freedom, stable funding environments, a broadly shared commitment to open inquiry — is also under severe pressure. Armed conflicts are now at their highest level since the end of the Second World War. Investment flows are shifting toward defense and dual-use technologies, while programs oriented towards human development and educational cooperation are often being defunded. The post-war institutional settlement that insulated universities from direct state instrumentalization is being actively dismantled, and in some cases by the very governments that once sponsored it. The world is no longer centered on the Euro-Atlantic axis, economically or intellectually: new centers of gravity are emerging, and with them new conceptions of what universities are for and whom they serve.

Artificial intelligence arrives at this already destabilized landscape, and it is not just another technology wave. Universities have survived many technological disruptions: the printing press, radio, television, the Internet, the MOOCs. Each of these technologies changed how knowledge was distributed. However, AI is different because it reaches into the production, interpretation, synthesis, evaluation, and economic value of knowledge itself. Previous technologies transformed the distribution of knowledge, but AI begins to transform the conditions under which knowledge is created and validated. For institutions whose legitimacy has long rested on their role as custodians of rational inquiry, the implications are profound.

Increasingly sophisticated forms of symbolic, analytical, communicative, and creative labor can now be externalized into machine systems operating at a planetary scale. Language models generate fluent text, working code, competent analysis, and plausible argument. Autonomous agents are beginning to perform tasks that constituted the core demand for university graduates. The half-life of technical knowledge is contracting. Employers across major economies are dropping degree requirements, and students are arriving on campus already embedded in AI-mediated workflows that their institutions have barely begun to understand.

In a nutshell, AI transforms the cognitive infrastructure of society. Knowledge access, and increasingly many forms of knowledge production and manipulation, is ceasing to be scarce. The substrate around which universities organized their entire institutional logic is shifting beneath them. This does not eliminate the need for human intelligence. But it fundamentally alters the economic, epistemic, and institutional environment within which universities evolved. The result is not merely disruption, but a growing mismatch between institutional metabolism and civilizational tempo.

Most institutions have not yet confronted the deeper implications. The dominant response so far has been adaptation within the existing frame: adopting AI tools, updating academic integrity policies, offering faculty training, launching pilot projects. These are reasonable operational responses. However, they assume that the institutional logic

itself — lecture, course, credit hour, degree, department, examination — remains sound and merely requires updating. The question this volume asks is different: whether the industrial university model remains sufficient and viable under conditions of cognitive abundance, and if not, what comes after it.

If knowledge access is no longer scarce, if procedural cognitive labor is increasingly automated, and if societies face mounting ecological, political, and civilizational instability that require new ways of learning and collaboration, then the university can no longer justify itself primarily as a system for information transfer and credential distribution. The deeper question becomes: what forms of education, collective intelligence, and institutional legitimacy remain uniquely valuable under conditions of distributed AI-intense cognition?

This question is especially acute for the majority of the world's universities. The global discourse on AI in higher education is dominated by a small number of highly resourced institutions concentrated in North America, Europe, and East Asia. These institutions represent a fraction — less than ten percent of the world's roughly 25,000 universities. The largest share, 80-85 percent, are mid-tier. These resource-constrained, regionally embedded universities, often operating outside the English-dominant knowledge ecosystem, educate the overwhelming majority of the world's students. Many serve first-generation learners. Many operate in languages for which AI systems perform significantly worse. Many face demographic and economic pressures that leave no margin for failed experiments.

This volume begins from their position — not because elite institutions are unimportant, but because the mid-tier university is where the hardest and most consequential version of the problem lives. Elite institutions can buffer disruption through prestige, capital, and selectivity. Mid-tier universities cannot. They must renegotiate their legitimacy more directly — to students, employers, governments, communities, and regions. At the same time, this constraint may generate strategic diversity rather than collapse. Different institutions are beginning to explore different adaptive trajectories: restructured labor economics, task-based curricula, formation-centered pedagogy, regenerative campuses, sovereign AI infrastructure, culturally grounded models of the graduate, hybrid human-AI learning systems. Some may drift toward increasingly hollow credential-production functions, but others may become sites of genuine institutional reinvention.

Contributions

This collection brings together both conceptual arguments and grounded institutional experiments emerging from this transitional landscape. The volume is organized in three parts. The first part brings together conceptual arguments about the nature and depth of the transformation universities face, exploring the structural, economic, epistemic, and civilizational dimensions of the university's encounter with artificial intelligence. The second part presents an operational bridge — a conceptual matrix that can navigate universities through their AI transformations. The third part showcases working experiments from universities across different regions and contexts navigating the transition under real constraints, in real time.

The opening chapter, by **Alexander Sidorkin**, reframes the AI challenge through the lens of labor economics. His argument is that AI did not create the crisis of the university, but massification did it decades ago by making genuine relational education economically impossible at scale. Drawing on Trow's phases of higher education expansion and

Acemoglu and Restrepo’s task-based framework for automation, Sidorkin shows that AI is the first technology capable of absorbing the mechanical, procedural tasks that have consumed the majority of faculty time since universities became mass institutions. This opens the possibility of what he calls the Double Pivot: shifting curriculum from procedural to conceptual depth while simultaneously redirecting recovered faculty time toward sustained human contact: mentoring, dialogue, and formation of judgment. The relational university he proposes is not a utopian aspiration but a labor-economic argument: the tasks already exist, the need has always been there, and for the first time, the resource to fund them is available. His chapter also names the default path “so-so automation”, where AI saves costs, class sizes grow, and the relational deficit deepens — and warns that without deliberate institutional choice, the default is what will happen.

Natalia Cebotari advances the argument from economics to curriculum structure. Her chapter traces the collapse of what she calls “the content-era grammar of higher education” — the organizing assumption that the university’s primary function is to deliver disciplinary content in a structured sequence, certify its acquisition, and grant credentials accordingly. She argues that this grammar is not merely outdated but structurally locked in place by accreditation systems, funding mechanisms, and institutional inertia across very different national contexts — India, China, Uzbekistan, the United States. The alternative she proposes is a shift from content as the organizing principle of learning to the task: open-ended, consequential activities around which students discover what they need to know, rather than absorbing predetermined syllabi. Her distinction between this post-content model and competency-based education is important — competency frameworks still map onto pre-specified outcomes, while task-based learning treats the task itself as the site where the learner encounters the unknown. The chapter’s most useful diagnostic is the observation of “double accounting” at elite institutions: they already operate in post-content mode in practice, through research apprenticeships, studio work, and project-based learning, while maintaining content-era grammar for regulatory and reputational purposes.

Charles Fadel brings the curriculum argument to its most prescriptive form. Drawing on the Center for Curriculum Redesign’s framework, he proposes a specific redesign of the undergraduate degree around four layers of knowledge: declarative, procedural, conceptual, and epistemic. His core argument is that AI readily replicates the first two layers, i.e. facts and procedures, and that universities must therefore reorganize teaching around the upper two: conceptual understanding and epistemic awareness of how discipline produces, validates, and communicates knowledge. The resulting “trapezoid” of knowledge, replacing the traditional triangle that narrowed toward specialization, becomes a practical tool for any faculty member asking what to keep, what to compress, and what to deepen. Fadel also sketches a structural redesign of the undergraduate degree in the AI era: two years of broad polymathic education followed by two years of AI-assisted disciplinary depth oriented toward demonstrated professional capability.

Shawn Chen, Richard Jiarui Tong, Xiangen Hu, and Jack Gao present one of the most fully developed pedagogical redesigns in the volume through the AI+X paradigm, which treats AI not as a supplementary tool but as a structural partner in learning. Their framework integrates three dimensions of capability development: disciplinary world-modeling, where students construct transferable mental models rather than memorize content; human–AI collaboration, where learners develop the ability to work symbiotically with intelligent systems; and metacognitive self-modeling, where they learn to monitor and improve their own learning processes. These dimensions are operationalized through a dual-loop pedagogy that distinguishes conceptual understanding from procedural execution while preserving their interaction. Grounded in the ongoing

transformation of Sias University in China, the chapter argues that the central challenge of AI-enabled education is not technological adoption but institutional redesign, offering one of the volume's clearest articulations of a symbiotic intelligence pathway for higher education.

Timur Schukin's chapter operates at another level of ambition. Beginning with Aquinas and the medieval university's practice of *quaestio* and *disputatio*, he argues that the university's enduring function across its historical transformations has been epistemic legitimation: the conversion of fragile, local, or contested practices into knowledge that can be questioned, reproduced, disputed, taught, and trusted beyond the presence of its first carriers. He traces this function through the Humboldtian research university and Clark's entrepreneurial university, and then asks what practice demands legitimation in the AI era. His answer is distributed intellectual production — the increasingly common condition in which problem formulation, hypothesis generation, validation, and documentation are distributed across humans, language models, agents, tools, and external memory systems. He proposes the hybrid cognitive unit as the organizational form through which such distributed work can be governed and learned and introduces the zone of proximal degradation — a Vygotskian inversion describing the developmental risk when machine systems expand output faster than learners appropriate the operations needed for independent reconstruction. His closing argument that second-tier universities may be best positioned for this work as second-order integrators, connects the epistemological argument to the institutional realities the cases in Part Three describe.

Raphael Costambeys-Kempczynski and **François Taddei** provide the volume's most expansive vision of what the university could become. Their chapter moves from the same massification diagnosis as Sidorkin to a fundamentally different destination: the university as a site for cultivating planetary agency. Drawing on Freire, Senge, Raworth, and Campbell, they develop the concept of *Doughnut Learning* — education operating within a safe and just space between a social floor of belonging and agency and an environmental ceiling of planetary limits. Their MOOD framework (Discover, Ideate, Build) offers a pedagogical architecture for this vision with a distinctive emphasis on shadow questions: who might be harmed if a proposed solution succeeds? The chapter's central provocation is the distinction between universities that aim to be the best *in* the world and those that aim to be the best *for* the world — a distinction that, once stated, reframes every other argument in the volume.

Ralph Wolff's chapter addresses the question that follows logically from the preceding arguments: if the curriculum changes, the institutional function evolves, and the purpose expands, what happens to the systems that certify quality? Writing from decades of experience leading one of the major U.S. accrediting commissions, Wolff traces eight interconnected transformations that AI is bringing to quality assurance: the shift from general standards to metrics-based evaluation, the emergence of comparative benchmarking across peer institutions, the evolution of the self-study from periodic ritual to continuous data monitoring, AI-assisted curricular analysis, and the entry of new AI-native accrediting agencies into a field that has been remarkably stable for a century. His chapter is grounded in the U.S. regulatory landscape but illuminates a structural shift with global implications: accreditation is moving from asking whether an institution follows good processes to asking whether it produces good outcomes — a shift that AI both enables and accelerates. His closing caution is important for the entire volume: an accountability framework built entirely on earnings metrics will systematically disadvantage programs in teaching, social work, arts, and humanities, precisely the programs that develop the civic capacities a democratic society depends on.

Before the second part, we present the operational bridge chapter, in which **Vadim Novikov** argues that universities do not need more tool catalogues, but need a way to make better decisions about AI adoption. The chapter builds a practical framework around university functions: teaching, assessment, student support, research, operations, infrastructure, and data governance. For each, it asks what the tool does actually change, what conditions it requires, and who is responsible when something goes wrong. It identifies three traps that recur across institutions — buying tools before defining the problem, copying elite university strategies without their resources, and running pilots that produce announcements rather than learning. Working through seven functional domains, it separates low-risk uses from those that demand serious oversight, and names cases that should be avoided entirely. It closes with a minimum viable governance model built for universities with limited capacity: simple enough to run, but robust enough to protect students and staff.

Finally, the second part of the volume shifts from conceptual arguments to practical evidence — not evidence in the sense of proof, but in the sense of institutional experience made available for examination. Each case describes a university navigating the AI transition under particular constraints, in a particular context, with particular resources. Together they offer a geography of adaptive experimentation across regions rarely centered in the global higher education discourse.

Kakha Shengelia opens the cases with an account of Caucasus University in Georgia — a small economy where the education-to-employment pipeline is short, the brain drain pressure is intense, and the labor market is thin. His chapter is the most analytically developed of the institutional cases, articulating a concept of a “structured uncertainty” as an institutional posture: building resilience precisely through the acknowledgment that the future is unknowable. His treatment of AI assessment governance, specifying what students must declare, how AI-generated content must be cited, and what constitutes substantial use requiring methodological disclosure, offers the most operationally detailed ethical framework in the collection.

Ivan Ninenko describes what a post-content classroom looks like in practice at the International University of Central Asia in Kyrgyzstan. His “learn with AI, test without AI” model resolves the academic integrity dilemma through a clean pedagogical principle rather than a philosophical position on AI permissibility. The chapter’s deeper contribution is its argument about meta-learning — making the learning process itself visible to students as an object of reflection — and its closing inversion: peripheral institutions, precisely because they have always had to do more with less, may possess knowledge about adaptive teaching that elite institutions never needed to develop.

A chapter by **Assylbek Kozhakhmetov**, **Assel Aryn** and **Aigul Sarenova** from Almaty Management University (AlmaU) in Kazakhstan presents the case of the volume’s institutional home. Beginning not with AI but with a crisis of student wellbeing in the aftermath of the pandemic and sociopolitical turbulence, the chapter traces how AlmaU arrived at a curricular response grounded in the nineteenth-century Kazakh philosopher Abai’s vision of the Tolyq Adam — a holistic person formed through the integrated cultivation of mind, heart, and will. The resulting *LEADER* framework and its eight-course general education core demonstrate that genuine curricular transformation is possible within regulated national systems, that cultural grounding gives a program motivational depth no skills checklist can match, and that the philosophical traditions of Central Asia offer resources for human formation that belong in the front ranks of global educational discourse — not as local color but as legitimate intellectual contribution. The program’s recognition at the 2026 Times Higher Education Asia Awards, along-

side institutions from across the Asian continent, suggests that the global turn toward human-centered education is not a single tradition's export but a convergent discovery.

Sabur Khan documents the AI-native transformation of Daffodil International University (DIU) in Bangladesh, a private entrepreneurial institution that has moved from digitization through AI integration toward what Khan describes as an AI-native institutional model. The chapter traces how learning flows were redesigned from lecture-examination-grading sequences toward AI exploration, human reflection, industry application, and output creation. Operating under significant resource constraints, DIU's experience illustrates the leapfrogging argument: that constraint-driven institutions may experiment faster precisely because they have less legacy infrastructure to protect.

Letlhokwa George Mpedi writes from the University of Johannesburg, one of Africa's largest universities and the institutional home of significant AI-and-law scholarship. His chapter's distinctive contribution is the argument that AI governance frameworks are being shaped by jurisdictions with the most technical capacity, while African legal systems — customary law, Roman-Dutch law, religious personal law — remain essentially absent from both the training data and the regulatory assumptions of AI systems deployed across the continent. For universities operating within these legal traditions, the governance question is not simply how to adopt global standards but how to ensure that local epistemic and legal realities are represented in the emerging regulatory architecture.

Victoria Haro closes the cases with an account of Universidad del Medio Ambiente (UMA) in Mexico, a small regenerative campus that has organized learning around ecological relationship, community formation, and purpose-driven inquiry since 2009, long before AI made these questions urgent for everyone else. Her chapter serves as the vital empirical bookend to the volume, directly operationalizing the “planetary mission” conceptualized by Taddei and Costambeys-Kempczynski in Part One. Because UMA's model was never built on traditional information transfer, AI does not disrupt it; instead, distributed machine cognition throws UMA's core evolutionary questions into sharp relief: what does it mean to learn in relationship with a living system? The chapter asks what the rest of higher education might learn from an institution that treats the university not as a data-delivery system, but as a relational layer within a broader regenerative ecology.

Taken together, these cases do not converge on a single model. They reveal an emerging landscape of institutional diversity with different starting points, different constraints, different ambitions, and different definitions of success. What they share is a refusal to treat the AI transition as merely operational, and a willingness to ask what their institutions are for under conditions where the old answers no longer suffice.

The chapters that follow do not present a final blueprint for the university of the future. They map an emerging landscape of institutional pathways under conditions where intelligence itself is becoming infrastructural, distributed, and contested.

For centuries, universities occupied a privileged position within a civilization organized around production, preservation, and transmission of rational knowledge. If many forms of cognition can now be externalized, and if the relationship between knowledge, expertise, and human capability is being fundamentally reconfigured, the question that organizes this volume is simple but dangerous: **what becomes the enduring purpose of the university after AI?**

Analytical preface. Mapping the global discourse, manifestos, and institutional debt

Vadim Novikov

The debate about the “university after AI” looks coherent only from a distance. When the key texts are read together, they do not form a single conversation. They form an overlapping field of manifestos, practical guides, governance frameworks, competence models, maturity toolkits, product visions, and institutional anxieties. Some texts speak in the name of the university as a moral institution. Some speak in the name of the teacher. Some speak in the name of students. Others speak in the language of risk, procurement, assessment reform, data governance, or productivity.

This matters because the apparent object of the debate, “the university”, is not a natural unit. It is already an abstraction. A research-intensive university in London, a large public institution in North America, a private mid-tier university in Central Asia, a technical college, an online provider, and a mass undergraduate institution do not face the same AI problem. They differ by resources, language, disciplinary mix, student preparation, legal environment, data maturity, staff capacity, procurement power, and the degree of institutional vulnerability to political control. Yet much of the global discourse proceeds as if the same normative map could guide all of them.

The central corpus of the debate is heavily manifesto-shaped. The Manifesto for Generative AI in Higher Education (Farrell and McCarthy, 2025), Bozkurt et al.’s Manifesto for Teaching and Learning in a Time of Generative AI (2024), the LSE Student Manifesto for Assessment in the Age of AI (2024/2025), the Safe AI in Education Manifesto (2024), and related texts all attempt to name a crisis and defend educational values under pressure. Around them sit boundary texts: Mollick’s productivity-oriented account of working with AI (2024), Khan’s vision of scalable AI tutoring (2024), and a growing set of institutional frameworks from UNESCO, the UK’s Joint Information Systems Committee (JISC), European University Association (EUA), the US National Institute of Standards and Technology (NIST), and others.

The geography of this conversation is not neutral. Its most visible documents come from English-speaking, research-intensive, policy-rich, or platform-adjacent environments. That does not make them irrelevant. It does mean that their assumptions travel with them. A framework written for an institution with mature data governance, legal support, staff development capacity, and stable procurement channels cannot simply be dropped into a university where those conditions do not exist. The first methodological problem in the literature is not bias in a narrow political sense. It is the tendency to speak about “higher education” while silently presupposing a particular segment of it.

This chapter therefore treats the literature not as a neutral list of positions, but as a discourse map. Its task is to ask who is speaking, from which institutional location, with which assumptions about capacity, and for which kind of university. The argument is not that global discourse is wrong. Its diagnoses are often sharp. The argument is that it becomes misleading when it is universalized.

The review reads the corpus with a question borrowed from institutional analysis: under what conditions would this recommendation make sense? A rule that is plausible for an elite seminar may fail in a mass course. A policy that is sensible in an English-speaking institution with strong legal support may become dangerous in a multilingual university with weak appeals procedures. A tutoring model that works for mastery learning may not solve the problem of undergraduate students who have learned to generate fluent answers without stable conceptual control.

A real consensus: the collapse of trust in academic artifacts

The strongest point of consensus in the literature is that generative AI did not create the crisis of higher education from nothing. It made visible weaknesses that were already present in the architecture of assessment, authorship, and trust. The final essay, the take-home assignment, the problem set, the reflective text, and the research summary have long functioned as proxies for thinking. Institutions assumed that the submitted artifact could stand in for the student's process, understanding, judgment, and effort. Generative AI breaks that assumption not by making cheating new, but by making the old proxy unstable at scale.

The crisis is therefore deeper than plagiarism. If a student can generate an acceptable academic artifact without passing through the cognitive labor that the artifact is supposed to evidence, then the artifact no longer carries the same institutional meaning. The hard question is not only whether the student wrote the text. It is whether the university can still infer competence from the text.

Several parts of the literature converge on this point. The essay manifesto tradition questions whether the essay can remain the dominant form of both learning and assessment under conditions of abundant text generation. Practical guides and institutional frameworks increasingly recommend process evidence, disclosure, redesigned assessment, oral or practical verification, and more explicit expectations around AI use. Even vendor-facing discussions of detection now tend to acknowledge that probabilistic outputs cannot carry the whole burden of academic judgment.

This is a real and valuable consensus. The literature is right to say that universities can no longer rely on the old settlement: teachers assign work, students return artifacts, and institutions certify achievement. But the consensus becomes weaker when it moves from diagnosis to remedy. Some texts respond by defending human agency, cognitive effort, and academic integrity. Others respond by embracing AI as a partner, tutor, or productivity amplifier. Governance frameworks respond by turning the crisis into policies, risk matrices, competence models, and maturity stages. Each response captures something important. None of them, by itself, solves the operational problem of a constrained university.

The more precise question is therefore not "should AI be allowed?" and not even "how do we detect AI use?" The better question is: what must remain visible in student work for a university to trust that learning has occurred? In an AI-saturated environment, the unit of assessment cannot be the final artifact alone. It must include evidence of decision, interpretation, revision, explanation, transfer, and performance under conditions where the student's own competence can be seen. AI does not abolish academic artifacts; it forces universities to stop treating them as self-authenticating objects.

This distinction matters because it reframes the problem from cheating to competence. The cheating frame asks whether the student illegitimately used a tool. The competence frame asks whether the student can do the thing the institution is about to certi-

fy. These questions overlap, but they are not the same. A student may use AI legitimately and still learn little; another may use AI in a prohibited way and nevertheless reveal competence in a subsequent defense. The institutional design problem is to build assessment systems that can distinguish these cases without relying on suspicion as a default mode — a problem that Natalie Cebotari’s chapter in this volume examines through the lens of accreditation and curriculum structure, and that several of the institutional cases address through concrete redesign experiments.

The normative split: efficiency, agency, and the problem of segment

After the shared diagnosis, the corpus splits. But the split is not simply between optimists and pessimists. The more important division concerns the imagined user of AI.

On one side are texts that treat AI as a partner, amplifier, tutor, or co-worker. Mollick’s *Co-Intelligence* (2024) develops the most influential account of this position: AI as a strange, unreliable, but genuinely useful cognitive partner — what he calls a “co-intelligence” — that can draft, brainstorm, critique, simulate, and iterate alongside a human user. His figures of the “centaur” (human and AI taking turns on distinct subtasks) and the “cyborg” (human and AI fully interleaved in one workflow) have become reference points across the field. The account is powerful for a particular user: an adult professional or student with enough expertise to supervise the machine, judge its output, and incorporate it into a meaningful workflow. In that segment, the AI can be a collaborator because the human user remains capable of evaluating, redirecting, and correcting it. The model works when the user has judgment before the tool accelerates the task.

Khan’s *Brave New Words* (2024) makes a different but adjacent argument: AI as a scalable Socratic tutor that can provide every learner with individualized guidance, practice, and feedback that human teachers cannot deliver at scale. His platform *Khanmigo* is designed to ask guiding questions rather than give answers, scaffolding the student through mastery of structured material. The vision is strongest in its own segment: K-12, mastery learning, structured progressions, and a pedagogical environment where the tutor’s work is embedded in a larger curricular design.

However, this promise changes meaning when transferred to mass higher education, especially where students enter with fragile cognitive foundations. A tutor does not automatically make the student a principal. An individualized interface is not yet a governance structure for learning. There is always a strategic layer: someone defines the goal, the sequence, the scaffolding, the acceptable shortcut, the moments of productive difficulty, and the evidence of mastery. If that layer is not visible or is silently delegated to the platform, “personalization” may produce the appearance of agency while moving the student’s work into a system designed elsewhere. Alexander Sidorkin’s chapter in this volume develops this argument through a different lens: his “so-so automation” scenario describes precisely the default path where AI saves costs and increases efficiency without redirecting freed resources toward the relational depth that education requires.

For the same reason, the tool that can be emancipatory in one segment and can also be deskilling in another. A coding assistant may help an experienced programmer move faster and prevent a novice from understanding the code. A writing assistant may help a multilingual scholar express an argument and allow an underprepared student to bypass the formation of the argument. A tutor may scaffold productive practice or normalize dependence. The literature is at its strongest when it recognizes this ambiguity. It is at its weakest when it treats AI either as universal empowerment or universal threat.

On the other side are texts that stress human agency, cognitive effort, equity, and institutional control. Bozkurt et al.'s collective manifesto (2024), signed by over thirty scholars from multiple countries, gives the most developed version of this pole. Their argument is that generative AI is not a neutral tool but a system shaped by the data, values, and power structures embedded in its training, predominantly Western, English-language, and commercially driven. They warn that uncritical adoption risks reshaping cognition itself, marginalizing non-dominant epistemologies, reproducing the biases of the corpus, and transferring educational authority from institutions to platforms. The LSE Student Manifesto (2024/2025), written collaboratively by students and the university's education enhancement center, gives this concern a more local and experiential form: students argue that assessment should exist to support learning, not to police it, and that AI use policies should be co-designed with students rather than imposed on them. They welcome AI as a potential equalizer — for language support, accessibility, and feedback — but insist that its integration must not come at the cost of meaningful intellectual challenge or reduce assessment to a surveillance exercise.

This normative pole sees a real danger. If AI allows students to bypass difficulty, it can weaken the very capacities education is meant to build. The phrase “cognitive strain is part of the learning process” (Bozkurt et al, 2024) captures an important educational intuition. However, cognitive strain does not have the same meaning in every segment. For a student with a formed base, strain can be productive: it forces articulation, judgment, and transfer. For a student with a fragile base, the same strain may not become thinking. It may become a barrier, a performance of confusion, or an invitation to out-source the task completely. Timur Schukin's chapter in this volume formalizes this risk as the *zone of proximal degradation*, a Vygotskian inversion where machine systems expand achievable output faster than learners appropriate the operations needed for independent reconstruction.

The difficulty with both poles is universalism. Manifestos tend to write one normative map for radically different institutional territories. Practical guides tend to describe their own segment without articulating its boundaries. A claim that is pedagogically true in one segment can be institutionally false in another. Mollick does not need to be refuted; his model needs a boundary. Khan does not need to be dismissed; his tutoring vision needs a segment and a governance layer. The defense of cognitive effort does not need to be caricatured; it needs operational translation into institutions where “effort” means something different depending on the student's starting point and the institution's capacity to support it.

The same applies to the political architecture of AI in education. Calls for open, public AI infrastructure as an alternative to corporate platforms are normatively attractive but institutionally ambiguous. A centralized public system can become opaque, politically aligned, technologically stale, and difficult to leave. Under conditions of weak institutional protection, educational infrastructure acquires the force of institutional reasonableness — what is filtered or discouraged can appear not as censorship but as responsible pedagogy. The relevant comparison is not market versus virtue. It is one set of governance risks against another. For universities in the segment this book addresses, the practical strategy may be maintaining multiple channels of access — commercial, open-source, locally deployed — rather than depending on any single provider, public or private.

Frameworks, capacity, and translation problem

The procedural frameworks — competency models, risk registers, maturity toolkits, due-diligence guides — are indispensable because they move the discussion from values to institutional action. They slow down enthusiasm. They make it harder for a university to say simply: “AI is the future, therefore we need a tool”. A good framework asks for roles, risks, documentation, training, and review. It creates a language in which educators, lawyers, IT staff, librarians, administrators, and students can talk to one another.

Jisc’s maturity toolkit is useful precisely because it recognizes that organizations differ and that AI adoption moves through stages. EUA’s guidance is useful when it insists on moving from speculation to use cases. NIST-style risk vocabulary is useful because it disciplines the conversation: risks must be identified, governed, monitored, and revised. UNESCO-style competency frameworks are useful because AI literacy cannot be left to individual improvisation.

The limitation is that frameworks can become universalism in procedural form. A maturity model can imply a single ladder. A competency framework can imply a single user. A risk register can imply that the institution has the personnel to maintain it. A governance checklist can imply legal, technical, and administrative capacities that many universities do not possess. Capacity is the hidden variable in literature. It appears indirectly, through words such as governance, transparency, oversight, evidence, literacy, infrastructure, maturity, and responsibility. However, capacity is not a word in a policy document. It is a set of routines: who drafts the policy, who explains it to staff, who answers student complaints, who audits the tool, who stops a bad pilot, who negotiates with a vendor, who revises the assessment design, who decides whether a model is good enough for the work of Central Asian or Latin American academy, who bears the cost when an AI-supported process fails.

A framework must therefore be filtered twice. First by function: teaching, assessment, research, operations, support, infrastructure, governance. Second by context: resources, language, data maturity, staff capability, student preparation, and political environment. A university that lacks a reliable appeals process should not adopt a tool whose errors require appeals. A university without clean data should not build a predictive advising system as if the data were neutral. A university with multilingual students should not assume that English-centric evidence transfers automatically. These are not implementation details. They are the substance of responsible adoption. Charles Fadel’s chapter in this volume addresses the “what to teach” dimension of this challenge through a knowledge taxonomy that distinguishes what AI readily replicates from what requires deliberate human cultivation. The institutional cases in Part Three show how different universities have navigated the “what to adopt” dimension under their specific constraints.

The gap between normative ambition and operational capacity is the central weakness of global discourse. It is not enough to say that AI should be ethical, transparent, inclusive, and human-centered. The operational question is: which institution can make those values real, through which routines, with what evidence, at what cost, under which political and linguistic conditions, and with what fallback when the system fails?

What the literature sees, what it misses, and what this volume does

The book is written for a specific kind of university: not the global elite, not the platform provider, not the well-funded experimental campus, and not the individual expert

worker who can supervise AI from a position of high competence. Its central segment is the mid-tier or mass university in a non-elite jurisdiction, often with constrained resources, mixed linguistic environments, uneven staff development, incomplete data maturity, and students whose cognitive preparation is often more fragile than institutional rhetoric admits.

For this segment, the main risk of AI is not simply corporate capture and not simply missed productivity. It is the false appearance of mastery: the production of outputs that look educationally adequate before the underlying capacity has been formed. Students may learn to produce acceptable artifacts without developing the ability to question, verify, transfer, explain, or perform. Teachers may receive polished work that conceals weak understanding. Administrators may buy tools that display innovation while adding hidden operational debt. Institutions may adopt policies that express responsibility without creating the capacity to enforce or revise them.

This false appearance of mastery is not only an individual learning problem. It is an institutional problem. A university can also develop a false appearance of AI readiness: policies without enforcement, pilots without learning, dashboards without interventions, detectors without due process, tutoring platforms without pedagogy, and governance committees without operational authority. The same surface fluency that AI gives to student writing can appear in institutional strategy. The document looks competent. The underlying capacity is missing.

Language makes this problem sharper. In non-English or mixed-language environments, access to a tool does not guarantee access to the knowledge base that matters locally. A model may produce fluent English summaries and still fail to represent local scholarship, national regulations, regional labor markets, or disciplinary materials in Kazakh, Tamil, or another underrepresented language. This is not only a translation problem. It is a problem of epistemic coverage. A university in such a context cannot treat global AI discourse as if it were already its own infrastructure.

The phrase “student as principal” must therefore be used carefully. A student with an AI tutor is not automatically the principal of the learning process. The student becomes principal only if the student can define goals, judge outputs, resist shortcuts, and understand the difference between assistance and substitution. In many real settings, however, the strategy of the tutor is defined before the student enters the interaction: by the teacher, the platform, the curriculum designer, the procurement decision, or the default prompt architecture. Otherwise, agency may sit elsewhere while the interface continues to speak the language of personalization.

Such kind of contextualization also preserves the university’s own educational responsibility. If the institution simply imports a manifesto, it borrows someone else’s moral grammar. If it simply buys a tool, it borrows a vendor’s workflow. If it tries to imitate an elite university, it borrows expectations for a capacity structure it may not have. The work of strategy is to decide what can be owned responsibly. For the universities addressed in this book, ownership will often mean smaller, more selective, and more explicit choices rather than grand institutional transformation.

The same applies to universities. A university is not more sovereign because it has an AI policy, a chatbot, a detector, a procurement contract, or a maturity score. It becomes more capable only when it can connect tools to functions, functions to evidence, evidence to governance, and governance to revision. That is the bridge from this literature review to the implementation framework that follows.

Global literature gives us warnings, aspirations, categories, and frameworks. It tells us that AI changes trust, authorship, assessment, agency, and institutional responsibility. But it does not yet provide a route map for constrained universities. The chapters that

follow take up this challenge from different angles. Sidorkin asks what becomes possible when AI changes the labor economics of teaching. Cebotari asks what happens to curriculum when the content-era grammar collapses. Fadel specifies which layers of knowledge resist automation, and which do not. Schukin asks what the university's enduring function has been across its historical transformations and what new practice demands legitimation now. Taddei and Raphael ask what the university is ultimately in service of. The institutional cases in Part Three show what these questions look like when they are answered not in theory but in practice, under constraint, in real time.

The next step is the work of disciplined contextualization — deciding which global claims apply here, which require translation, which are irrelevant, and which become dangerous under local conditions. That work begins in the chapters that follow.

References

- Alier-Forment, M., García-Peñalvo, F. J., Casañ, M. J., Pereira, J. A., & Llorens-Largo, F. (2024). *Safe AI in education manifesto* (Version 0.4.0). <https://manifesto.safeaieducation.org>
- Bozkurt, A., et al. (2024). The manifesto for teaching and learning in a time of generative AI: A critical collective stance to better navigate the future. *Open Praxis*, 16(4).
- European University Association. (2023). *Artificial intelligence tools and their responsible use in higher education learning and teaching*. EUA. <https://www.eua.eu/publications/positions/artificial-intelligence-tools-and-their-responsible-use-in-higher-education-learning-and-teaching.html>
- European University Association. (2024). *Adopting AI that serves the needs and values of universities*. EUA. <https://www.eua.eu/publications/reports/adopting-ai-that-serves-the-needs-and-values-of-universities.html>
- Farrell, H., & McCarthy, K. (2025). *GenAI:N3: Manifesto for generative AI in higher education*. <https://genain3.ie/manifesto/>
- Jisc. (2024). *AI maturity toolkit for tertiary education*. <https://www.jisc.ac.uk/ai-maturity-toolkit-for-tertiary-education>
- Khan, S. (2023, May). *How AI could save (not destroy) education* [TED Talk]. <https://blog.khanacademy.org/sal-khans-2023-ted-talk-ai-in-the-classroom-can-transform-education/>
- Khan, S. (2024). *Brave new words: How AI will revolutionize education (and why that's a good thing)*. Salaam Reads.
- London School of Economics Students' Union & LSE Eden Centre. (2024). *A student manifesto for assessment in the age of AI*.
- Mollick, E. (2024). *Co-intelligence: Living and working with AI*. Portfolio/Penguin.
- Tabassi, E. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). National Institute of Standards and Technology, U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- UNESCO. (2024). *AI competency framework for students*. United Nations Educational, Scientific and Cultural Organization. <https://unesdoc.unesco.org/ark:/48223/pf0000391105>
- UNESCO. (2024). *AI competency framework for teachers*. United Nations Educational, Scientific and Cultural Organization. <https://unesdoc.unesco.org/ark:/48223/pf0000391104>

Part I.
**The structural
and epistemic
reconfigurations**

Chapter 1. Labor economics, so-so automation, and relational university

Alexander M. Sidorkin

Universities have a set of deep, structural problems that long predate artificial intelligence. Student engagement is thin. Advising is reactive and understaffed. Instruction runs on low-engagement models that compress assessment into a few high-stakes moments. Faculty who entered the profession to work with ideas and people spend large portions of their time on mechanical tasks. Students navigate systems designed for institutional convenience, not for learning. These are not new complaints. They are defining failures that decades of reform have failed to resolve.

The usual response to AI in higher education begins with the technology and asks what it can do, what it threatens, and what policies it demands. I want to start from the other end. The question that matters is not what AI can do. It is what AI is for. And the answer, I will argue, is straightforward: AI is for making universities more relational. Not more automated, not more efficient in the narrow sense, but more capable of doing what educational theory has always said matters and what institutional economics has always prevented.

To see why this matters, we need to understand how universities arrived at their current condition. The story is not primarily about bad decisions or institutional laziness — it is about a structural trap that no amount of reform could escape.

Martin Trow's framework for understanding higher education development remains the clearest account of this trap (Trow, 1973; 2006). Trow identified three stages in the growth of higher education systems: elite, mass, and universal. In the elite stage, universities served a small, socially homogeneous population. The model worked because the ratio of faculty to students was manageable, advising was informal and personal, and assessment could be substantive because the numbers allowed it. The entire system rested on an implicit assumption: that relational labor, the time faculty spend in direct intellectual contact with students, would remain roughly proportional to enrollment.

The transition to mass higher education broke that assumption. As participation rates crossed Trow's threshold of 15 percent and moved toward 50 percent and beyond, student numbers grew enormously. This was not only a story about new institutions built for mass access. The oldest and wealthiest universities grew alongside this shift, multiplying their enrollments several times over while retaining structures that had been built for a far smaller student body, so that a department which once served 200 majors now served 800 with no corresponding increase in faculty lines. Relational capacity, however, did not scale with enrollment: class sizes swelled, advising caseloads multiplied, and assessment compressed into formats that could be administered at scale, among them multiple-choice exams, standardized rubrics, and infrequent high-stakes papers. Faculty adapted by shifting time away from direct student contact and toward tasks that could be performed more efficiently — preparing lectures for large halls, managing course logistics, processing grades, so that the intimate pedagogical

norms of the elite university endured only as aspirations, rendered undeliverable by the labor economics of the mass university.

This is not simply a problem of insufficient funding, though funding matters. This is a problem that Baumol's cost disease helps explain: in labor-intensive services, productivity gains are structurally difficult to achieve (Baumol & Bowen, 1966). A professor in 2026 still needs roughly the same amount of time to mentor a student, respond to a paper with substantive feedback, or guide a difficult intellectual conversation as a professor in 1965. AI does not violate this logic. It shifts the unit of analysis from the job as a whole to the individual task, making gains possible in mechanical tasks while leaving the relational core untouched. Unlike manufacturing, where technological advances allow fewer workers to produce more output, the core service of education is the human contact itself. You cannot speed it up without degrading it.

The result is a structural bind. Universities enrolled dramatically more students while remaining unable to increase the relational labor that good education requires. At the same time, other demands on institutional resources grew steadily. Regulatory and compliance environments became more complex, requiring larger administrative and support staffs. Institutions built expensive student engagement infrastructure, from recreation centers to counseling services, that earlier generations of universities did not provide. Contemporary IT systems added another layer of cost that competes directly with instructional budgets. Each of these expenditures may be individually justified, but their cumulative effect is the same: they absorb resources that might otherwise support the relational core of education. The gap between what we know works and what institutions can deliver is not a gap of knowledge or will. It is a gap of labor. We have known for decades what good advising looks like: proactive, personalized, informed by data about student risk (Kuh et al., 2010). We have known what good assessment looks like: formative, continuous, with each assignment informing the design of the next (Black & Wiliam, 1998). We have known what good instruction requires: authentic problems, sustained intellectual engagement, and regular feedback (Barr & Tagg, 1995). None of this is controversial in the literature. All of it has been impossible at scale because the labor required to deliver it exceeds what any institution can afford.

I call this the bottleneck assumption: the belief, grounded in decades of evidence, that the quality of higher education is constrained not by what we know but by the human labor available to act on what we know. The theory has been right. The practice has lacked capacity. This is not a claim about any particular reform failing for any particular reason. It is a claim about a structural ceiling that all reforms hit. Every major reform effort of the past fifty years has run into this wall. Active learning pedagogies work, but they require more preparation and smaller groups. Formative assessment works, but it demands feedback cycles that no instructor with 120 students can sustain. Proactive advising works, but it presupposes knowledge of individual students that no advisor with 800 cases can possess. The reforms succeeded in pilot programs and failed at scale, not because the ideas were wrong but because the labor economics were impossible.

This is the context in which AI enters higher education. Not as a disruptive force that creates new problems, but as a potential response to problems that have been accumulating since the transition to mass higher education. The question is whether AI can help loosen the bottleneck, whether it can free enough human labor from mechanical tasks to make relational education viable at the scale the modern university demands.

The task-based framework: from economics to education

The pattern described above, where good theory fails at scale because labor runs out, is not unique to higher education. It is a general feature of labor-intensive services. And because it is a labor problem, it is worth looking at what labor economics has to say about what happens when technology enters the picture. But economics has recently developed a framework that clarifies what happens when technology enters such sectors, and that framework turns out to be remarkably useful for thinking about AI in the university.

Daron Acemoglu and Pascual Restrepo have built the most influential contemporary account of how automation reshapes labor markets (Acemoglu & Restrepo, 2018, 2019, 2022). Their central insight is simple but powerful: production does not consist of undifferentiated labor. It consists of discrete tasks, each of which can be allocated to either human workers or machines. Automation occurs when a technology becomes capable of performing a task previously done by humans. This creates what Acemoglu and Restrepo call the displacement effect: machines take over certain tasks, reducing labor's share of the work.

But displacement is only half the story. Automation also generates a productivity effect, because cost savings in automated tasks free resources that can increase demand for labor in the remaining non-automated tasks. More important still, technological change can create entirely new tasks where humans have comparative advantage. Acemoglu and Restrepo call this the reinstatement effect. The net outcome for workers depends on the balance between displacement and reinstatement. When automation displaces labor from routine tasks but simultaneously opens new tasks that require human judgment, creativity, or relational skill, the overall effect on employment and wages can be positive. When it merely displaces without reinstating, the result is what Acemoglu and Restrepo call “so-so automation”: modest efficiency gains with no improvement in the quality of work or the demand for human contribution (Acemoglu & Restrepo, 2019).

This framework, developed to explain wage dynamics and employment patterns in manufacturing and services, maps onto higher education with unusual precision. Consider what a university instructor actually does. The job is not one task. It is a bundle of tasks that differ sharply in their cognitive and relational character. Some are mechanical: formatting syllabi, calculating grade distributions, checking citations, providing first-pass feedback on whether a paper meets basic structural requirements, answering routine questions about deadlines and policies. Others are deeply relational: guiding a student through a conceptual difficulty in a one-on-one conversation, designing a sequence of assignments calibrated to what a particular group of students needs next, recognizing when a student's performance signals a personal crisis rather than academic weakness, modeling intellectual curiosity in a seminar discussion.

These two categories of tasks have always coexisted in the same job description, performed by the same person, consuming the same finite hours. The mechanical tasks are not trivial. They are necessary for the institution to function. But they are not the tasks that make education work. The relational tasks are where learning actually happens, where students form the intellectual dispositions and human connections that justify the existence of the university. The tragedy of the modern university is that the mechanical tasks have steadily consumed the time that the relational tasks require. Massification made this inevitable. When you multiply students without multiplying faculty, the mechanical workload grows while the relational capacity shrinks.

AI changes the economics of this allocation. For the first time, a technology exists that can absorb a significant fraction of the mechanical tasks in instruction, advising,

assessment, and administration. Not all of them, and not without human oversight. But enough to alter the balance. The tasks suited for absorption share a common shape: determinate inputs, repeating structure, attention rather than judgment. The relational tasks resist absorption for the opposite reason: unique persons, unstable contexts, judgments that depend on knowing who is in front of you. Concrete applications across advising, instruction, and administration appear later in this chapter. The conceptual move matters first, because absent that move the applications dissolve into a catalogue of efficiency tools.

Even so, the instructor, the advisor, and the staff member are not displaced. What changes is the task composition of their work. In Acemoglu and Restrepo's terms, AI displaces human labor from a set of mechanical tasks within education while simultaneously creating the conditions for reinstatement in relational tasks that were always part of the job description but never practically deliverable at scale.

This is where the education case diverges from the standard automation story. In manufacturing, the reinstatement effect is speculative. When robots displace assembly line workers, the new tasks that emerge, programming the robots, designing new products, managing automated systems, are genuinely new. They require retraining, different skills, sometimes entirely different workers. The transition is real and often painful. In education, the reinstatement effect is not speculative at all. The tasks that emerge when mechanical work is absorbed by AI are the tasks that educational theory has been calling for since at least Dewey: sustained individual feedback, mentoring relationships, intellectually demanding conversations, formative assessment cycles that respond to what students actually need. These are not new tasks. They are old tasks that the labor economics of mass higher education made impossible. AI does not invent the relational university. It removes the constraint that prevented it from existing.

This distinction matters for how we think about institutional strategy. If AI in education were primarily about displacement, the central concern would be job loss, and the policy response would be defensive: protect positions, limit adoption, negotiate safeguards. But if AI in education is primarily about reinstatement, about restoring tasks that were always part of the educational mission, then the central concern shifts. The question becomes whether institutions will use the recovered capacity for its intended purpose or waste it on cost reduction.

Acemoglu and Restrepo's concept of so-so automation names this risk precisely. So-so automation occurs when technology displaces workers from tasks without generating meaningful productivity gains or new demand for human labor. In higher education, the so-so path would look like this: AI handles grading and advising queries, class sizes increase further, adjunct positions are eliminated, and the student experience becomes more impersonal than it already is. The technology works. The institution saves money. But the bottleneck remains, and the relational deficit deepens. This is the path of least resistance for institutions under financial pressure, and it is the most likely outcome if AI adoption is driven by budget logic rather than educational purpose.

The alternative, the path toward the relational university, requires a deliberate institutional choice to redirect the freed labor toward human contact. Faculty whose grading load decreases by several hours per week gain that time not as leisure but as availability: for individual student meetings, for small-group intellectual work, for the kind of substantive feedback that formative assessment has always demanded but never received. Advisors whose routine caseload is triaged by AI gain time for the conversations that actually change student trajectories: the meeting where a first-generation student is talked through the decision to stay enrolled, the conversation where an advisor helps a student recognize that their course schedule is a path to failure and redesigns it with

them. Administrative staff whose days are no longer consumed by answering the same twelve questions gain time for the complex cases, the ones involving financial hardship, family crisis, immigration status, disability accommodation, that require human judgment and institutional knowledge.

The labor has always been there. It has been allocated to the wrong tasks. AI offers a way to reallocate it. But only if we choose to.

The double pivot and the shape of the relational university

The theoretical framework is in place. Massification created the labor gap. The task-based model explains how AI can close it. But what does the relational university actually look like in practice? And what does it ask of faculty, students, and institutional leaders? I want to answer this through a concept I have been developing for faculty audiences: the Double Pivot. The term names two deliberate trade-offs that together define what AI-enhanced education requires at the course level. Understanding them is essential because institutional transformation does not happen through enterprise strategy alone. It happens course by course, instructor by instructor, through changes in what we teach and how we spend our time.

The first pivot is in content. When students have access to AI, the traditional emphasis on procedural knowledge, on drilling skills to automaticity, loses much of its rationale. A statistics student no longer needs years of practice performing calculations when AI handles computation in seconds. A writing student no longer needs to spend a semester on grammar drills when AI provides mechanical editing on demand. This does not mean procedural knowledge disappears. Students still need enough procedural understanding to grasp why a procedure works, when it applies, and what its results mean. But the goal shifts from fluency through repetition to comprehension through engagement. I develop this argument at length elsewhere (Sidorkin, 2026), including the claim that students need procedural knowledge sufficient for conceptual understanding. What was once the bulk of a course, the procedural middle, shrinks. What expands is conceptual knowledge: the capacity to frame problems, interpret results, exercise judgment, and work with ideas at a level of complexity that procedural drill never reached. Assignments become harder, not easier. Expectations rise precisely because students have more powerful tools. The trade-off is real: we give up coverage of procedural skill in exchange for depth of conceptual understanding and authenticity of assessment.

The second pivot is in the instructor's time. Faculty work consists of a bundle of tasks, and the bundle has been badly composed for decades. A large fraction of a typical faculty member's week goes to mechanical work: preparing routine materials, formatting syllabi, writing quiz questions, providing first-pass feedback on whether papers meet basic structural requirements, calculating grades, answering repetitive questions about deadlines and policies. AI now handles most of this, or accelerates it so dramatically that the time cost drops by half or more. In my own practice, AI-assisted grading of formative assignments reduces the time by roughly three-quarters. The recovered hours do not vanish into new administrative obligations. They go toward the relational work that the bottleneck has always prevented: individual consultations with students, small-group intellectual conversations, learning who students are and what they need. Faculty move from being processors of student output to being present in students' intellectual lives.

Neither pivot works without the other. Deeper conceptual work requires the kind of sustained, personalized support that only recovered faculty time makes possible. A student attempting a genuinely complex project, the kind that authentic assessment de-

mands, needs a mentor who knows their strengths, their gaps, and their trajectory. That mentoring relationship is the product of the second pivot. And the second pivot has no purpose without the first: recovered time spent supervising the same procedural drills that AI could handle is not a gain. The two pivots form a single integrated move.

What does this look like at institutional scale? Consider advising. The current model at most universities is reactive. An advisor with a caseload of several hundred students waits for students to identify their own problems and seek help. Those who most need support are least likely to ask for it. AI changes this by scanning enrollment, academic, and financial data to detect patterns that predict difficulty: a student whose course combination historically leads to failure, a student whose engagement indicators are declining, a student whose financial aid status has changed in ways that correlate with drop-out. The AI flags these students. The advisor calls them. The conversation that follows is human: it involves listening, judgment, knowledge of the student as a person, and the institutional authority to act. AI handles surveillance and triage. The advisor handles the relationship. The pull model becomes a push model, and the advisor's expertise is finally directed where it matters.

Consider instruction more broadly. Formative assessment theory has held for decades that feedback should be continuous and that the results of one assignment should inform the design of the next (Black & Wiliam, 1998). No instructor teaching more than a hundred students has ever had time to analyze patterns across a full set of papers and then redesign the subsequent task based on what the class actually needs. AI can do this. It can identify collective strengths and weaknesses across student submissions, draft a subsequent assignment calibrated to the gaps, and present the result to the instructor for review, adjustment, and teaching. For the first time, the feedback cycle that formative assessment always promised becomes practically operational. The instructor does not disappear from this process. The instructor becomes more important, because the judgments involved in interpreting patterns and designing responses require exactly the disciplinary expertise and knowledge of particular students that no machine possesses.

Consider administration. Staff in student-facing offices currently spend large portions of their days answering questions with definitive answers: how to file a form, when a deadline falls, what a policy requires. AI handles these instantly and accurately, through retrieval-augmented systems that draw on institutional documents and provide personalized guidance based on the student's specific situation. When the question is routine, the AI resolves it. When the question is complex, involving financial hardship, family crisis, immigration status, or disability accommodation, the AI routes it to a human staff member who now has time for precisely this kind of work. The staff member's role shifts from information gatekeeper to problem-solver.

Elements of each are already operating at institutions that have moved past the paralysis stage. But they remain scattered, uncoordinated, and dependent on individual initiative rather than institutional strategy. The relational university will not emerge from organic diffusion. It requires deliberate leadership: clear decisions to allocate labor toward human contact rather than cost reduction.

The implementation path matters as much as the vision. A staged approach works best, beginning with what creates immediate, visible gains for individual faculty and staff. The first stage provides AI tools for personal productivity: grading assistance, correspondence drafting, document summarization, meeting scheduling. These applications require no institutional coordination. They meet the people where they are, and they produce results fast enough to build confidence and appetite for deeper change. The second stage moves to course-level transformation, where departments collectively

redesign curriculum and assessment around the Double Pivot. This requires coordination, faculty development, and protected time for experimentation. Departments need at least a full academic year for planning and piloting, and another for refinement. Institutions that try to compress this into a summer workshop will get compliance, not transformation. The third stage addresses enterprise-level systems: predictive analytics for advising, integrated scheduling optimization, institution-wide formative assessment infrastructure. These require substantial investment, technical expertise, and multi-year timelines. They are important but not urgent. The first two stages can proceed without them.

Two conditions determine whether any of this produces the relational university or merely produces so-so automation dressed in educational language. The first is that recovered human capital must be explicitly redirected toward relational work. If AI-assisted grading leads to larger class sizes, the bottleneck tightens rather than loosens. If automated advising triage leads to reduced advising staff, the human contact disappears precisely when the technology makes it more effective. Institutions must make a visible, accountable commitment: the time and resources freed by AI go toward more human contact, not less. The second condition is that faculty must control the pedagogical transformation. Administrative mandates to “integrate AI” without faculty ownership produce resistance, and the resistance is justified. The Double Pivot works only when instructors themselves decide what procedural knowledge to reduce, what conceptual depth to add, and how to use recovered time.

The relational university is not a utopian vision. It is a description of what becomes possible when we use AI to remove the constraint that has prevented good education from operating at scale. We already have the labor. The knowledge of what works was always there. What was missing was the capacity to act on both. AI does not replace the human work of education. It removes the mechanical burden that has kept that work from happening.

What we know, what we do not, and what leaders must decide

The argument of this chapter rests on a chain of claims, each of which I believe is well supported but none of which is beyond question. Massification created a labor gap that degraded relational education. The task-based framework from economics explains how AI can close that gap through displacement of mechanical tasks and reinstatement of relational ones. The Double Pivot translates this logic into a course-level practice that faculty can adopt. And the relational university names the institutional outcome if these moves are made deliberately rather than left to drift.

That chain holds together. But it also exposes how much we do not yet know, and how much of the necessary infrastructure remains unbuilt. We do not have mature curriculum models for the post-procedural course. The first pivot, the shift from procedural to conceptual emphasis, sounds clean in the abstract. In practice, it raises discipline-specific questions that no general framework can answer. How much procedural knowledge in statistics does a social science student need before conceptual engagement becomes productive? At what point does reducing grammar instruction in a writing course cross from liberation into negligence? Where is the line between procedural knowledge that AI has made redundant and procedural knowledge that remains foundational for understanding? These are genuine questions, not rhetorical ones. They require empirical investigation and disciplinary judgment, and the answers will differ across fields. We are at the beginning of this work, not the end.

We do not have adequate assessment frameworks for what the relational university demands. If we shift toward authentic, complex, AI-enhanced assignments, we need ways of evaluating student work that capture conceptual depth, quality of human-AI collaboration, and intellectual growth over time. Traditional rubrics designed for procedural competence will not serve. Portfolio-based approaches, competency-based progression, and formative assessment cycles all show promise, but none has been tested at scale in the context of deliberate AI integration. The assessment challenge is not merely technical. It involves rethinking what counts as evidence of learning when the tools available to students have fundamentally changed.

We do not have reliable metrics for relational quality. If we ask institutions to redirect recovered labor toward human contact, we need ways of knowing whether that contact is happening and whether it matters. How do we measure whether advising has become genuinely proactive? How do we know whether recovered faculty time is going toward mentoring or being absorbed by new administrative demands? Existing institutional metrics, graduation rates, retention numbers, course completion, are too blunt to capture these shifts. The measurement gap is real, and it creates a risk: institutions may claim the relational university while delivering automation that displaces without reinstating, and we may lack the instruments to tell the difference.

None of these gaps is a reason to wait. They are reasons to begin with intellectual honesty about the stage we are at. The vision is grounded in sound theory. The implementation is early. The distance between the two is where the work lies.

This is also why the work cannot be done institution by institution, each one improvising its own solutions in isolation. The questions that remain open, about curriculum design, assessment, learning outcomes, metrics for relational quality, are questions for the global higher education community. They require shared frameworks, comparative evidence, and the kind of collective professional judgment that no single university can generate alone. Organizations like IAUP are positioned to convene exactly this kind of work: not to impose uniform solutions, but to create the conditions under which diverse institutions can learn from each other's experiments, share what works, and build the common infrastructure that the relational university requires.

What, then, do institutional leaders need to decide now, before the open questions are resolved? Three things. First, the capacity question. When AI frees labor, where does that labor go? It is a values decision, not a technology decision. If recovered capacity flows to cost reduction, the relational university does not emerge. If it flows to larger class sizes, the bottleneck tightens. Leaders must make an explicit, visible commitment that time and resources freed by AI will be redirected toward human contact with students. This commitment must be accountable, not aspirational.

Second, the faculty ownership question. The pedagogical transformation at the heart of the Double Pivot cannot be administered from above. Faculty must control what happens inside their courses: which procedural content to reduce, what conceptual depth to add, how to use recovered time. This is not a concession to faculty governance. It is a recognition that the disciplinary judgments involved are beyond the competence of administrators and consultants. What leaders can do is create the conditions for this work: protected time, faculty development resources, departmental coordination structures, and institutional patience. Curriculum transformation takes years, not semesters. Institutions that try to force it through a single summer initiative will produce compliance at best and resistance at worst.

Third, the staging question. Not everything needs to happen at once, and attempting everything at once is a reliable path to paralysis. The staged approach outlined earlier

in this chapter gives institutions a way to begin immediately with low-risk gains while building toward deeper structural change.

I want to close with a thought about the nature of this moment. The structural problems described in this chapter, thin engagement, reactive advising, compressed assessment, misallocated labor, have been with us for so long that we have learned to treat them as features of the university rather than failures of it. We have built careers, accreditation systems, and institutional identities around accommodating these problems rather than solving them. AI does not guarantee that we will solve them now. But it removes the excuse that solving them was impossible. The labor constraint that made relational education undeliverable at scale is, for the first time, negotiable. What we do with that opening is a choice, not a technological inevitability. But for the first time, we can build it.

References

- Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488–1542.
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- Acemoglu, D., & Restrepo, P. (2022). Tasks, automation, and the rise in U.S. wage inequality. *Econometrica*, 90(5), 1973–2016.
- Barr, R. B., & Tagg, J. (1995). From teaching to learning: A new paradigm for undergraduate education. *Change: The Magazine of Higher Learning*, 27(6), 12–26.
- Baumol, W. J., & Bowen, W. G. (1966). *Performing arts: The economic dilemma*. Twentieth Century Fund.
- Black, P., & Wiliam, D. (1998). Assessment and classroom learning. *Assessment in Education: Principles, Policy & Practice*, 5(1), 7–74.
- Kuh, G. D., Kinzie, J., Schuh, J. H., & Whitt, E. J. (2010). *Student success in college: Creating conditions that matter* (2nd ed.). Jossey-Bass.
- Sidorkin, A. M. (2026). *AI-enhanced pedagogies: Rethinking learning, curriculum, and human potential in the age of intelligent machines*. Routledge.
- Trow, M. (1973). *Problems in the transition from elite to mass higher education*. Carnegie Commission on Higher Education.
- Trow, M. (2006). Reflections on the transition from elite to mass to universal access: Forms and phases of higher education in modern societies since WWII. In J. J. F. Forest & P. G. Altbach (Eds.), *International handbook of higher education* (pp. 280–243). Springer.

Chapter 2. After AI, the university cannot be organized around content

Natalia Cebotari

Most discussions of AI and higher education begin with questions about adoption: how universities should use new tools in teaching, assessment, research, or administration. These are necessary conversations, but they only scratch the surface. AI may be doing more than forcing universities to update their methods. It is putting pressure on the organizational logic of higher education itself. What is beginning to weaken is the basic unit around which universities have long been structured: content.

Content is now accessible at any time, tailored to the learner, and assembled in real time from distributed sources — from emerging research to professional networks. If a machine can explain much of the curriculum, generate literature reviews in minutes, and simulate complex cases, then the institutional question becomes: what exactly does the university now do that AI cannot?

I approach this question not as an outside observer of higher education, but from inside the work of building educational systems under real constraints. Over the past decade, I have worked across national-scale digital learning platforms and AI-driven workforce systems reaching more than three million learners and 120,000 educators, and more recently have been directly involved in the creation of both private and public universities in Europe and Central Asia. The more I build, the more I see the same problem: higher education is attempting a facelift on a structure that was never designed for today's reality. The problem goes beyond pedagogy. It is institutional. Higher education systems are still largely built to account for coverage and completion, while the world increasingly demands capability and judgment.

Most universities remain organized around a unit whose institutional power depended on scarcity. That unit is content: not only lectures, readings, and disciplinary material, but the broader grammar through which curricula are structured, credits are allocated, progression is recorded, and legitimacy is made auditable. Knowledge remains essential. However, knowledge and content are not the same thing. Universities do not govern knowledge directly; they govern it through content-based structures — courses, syllabi, modules, credits, and assessed coverage. For a long time, this made sense. Universities were designed to curate, transmit, and certify specialized knowledge. Now labor markets evolve faster than traditional program cycles, professional tasks are being reassembled, and the distance between a credential and actual capability is becoming visible. In many domains, graduates and early-career professionals struggle to demonstrate the ability to frame problems, interrogate outputs, exercise judgment, and take responsibility under conditions of uncertainty. The university is pushed to justify itself less as a distributor of knowledge and more as a place where capability, judgment, and trust are formed and tested.

The university after AI cannot be organized primarily around content. It must move toward what I call a post-content model, in which content remains necessary but becomes infrastructure rather than destination. The institutional core reorganizes around

judgment and responsibility — structuring education through tasks and problem spaces, with progressively calibrated stakes.

Legitimacy, infrastructure, and the limits of content-based authority

The crisis of the university emerges from a functional paradox: content-based structures are weakening as an intellectual core while remaining essential as a regulatory language. This dual role defines the difficulty of the current transition. Ministries, accreditation bodies, and quality assurance frameworks still require institutions to demonstrate compliance through a standardized grammar of contact hours, curricula, credits, and measurable workload. As a result, the university functions as a system of visibility. It does not manage learning directly, but through time-based records and standardized proxies — credits, full-time study equivalents, modules, and assessed coverage — that allow education to be compared, funded, and regulated. These units survive because they solve a basic administrative problem: they make learning portable, legible, and fundable, even when they no longer describe with sufficient precision how a student builds capability.

Generative AI disrupts this arrangement by making reproducible a large part of “structured retelling”: not simply the delivery of information, but the work of organizing a field of knowledge into a coherent path. If a machine can map a discipline, synthesize emerging research, and structure a complex topic into a syllabus adapted to the learner, then the university’s role as the primary architect of the content-map becomes less exclusive. The labor of organizing and explaining knowledge — one of the institutional layers that higher education systems used to monopolize — is now increasingly reproducible.

Academic work, of course, extends far beyond explanation. Faculty produce knowledge, sustain disciplinary standards, supervise inquiry, and form scholarly communities. However, a substantial part of the instructional economy of higher education has depended on structured explanation, interpretation, and sequencing. It is this layer, not scholarship itself, that AI most directly commoditizes.

This pressure appears on three core levels at the same time: accreditation, operations, and the social contract. Accreditation remains tied to time-based curricular units that only partially reflect how learning happens. Operationally, staff are often positioned as carriers of a process-centered on content delivery, even as this function becomes increasingly reproducible. Socially, universities are large employers whose stability has depended not only on research and institutional service, but also on professional explanation, interpretation, and curricular mediation. When technology can partially reproduce this labor, the issue shifts from teaching efficiency to the stability of the broader employment structure.

The point is not that the institution disappears. It is that the formal architecture through which universities have made knowledge teachable, auditable, and fundable — courses, syllabi, credits, workload, and assessed coverage — is losing its power to describe what higher education must actually produce. If the institutional spine continues to rely on managing content, the university risks becoming a system that certifies its own internal processes more convincingly than it certifies the actual judgment of its graduates.

The most significant objection to this argument typically comes from representatives of elite universities. Their response is familiar: content was never the primary objective. What mattered was socialization, ambitious faculty, project work, laboratories, informal

intellectual formation, and the development of judgment. This perspective is partially correct; at their best, universities have always provided far more than content delivery. However, that reply misses the institutional point: the formal grammar through which universities are recognized, regulated, and financed still relies heavily on content-era units of legibility, even where the most valuable forms of formation exceed them.

The best universities often operate with a kind of double accounting. Formally, they remain legible through courses, credits, modules, and assessable learning outcomes. In practice, much of their value stems from assets only partly captured by that formal structure: selective admissions, demanding peers, institutional reputation, and proximity to advanced research. Their official account records what can be specified, allocated, and audited; their substantive account contains what actually matters. This gap was present before AI, but it becomes more consequential now. The strongest institutions may continue to rely on reputation and selection effects, but this does not mean their formal architecture is well aligned with their educational value. It may only mean that their inherent advantages compensate for structural weaknesses that the rest of the sector cannot afford to sustain.

This misalignment is especially visible in the current institutional fascination with “AI skills”. Many universities are introducing prompting workshops while students, operating outside formal teaching, are already building agents, orchestrating workflows, and testing model limitations in ways that move far beyond the official curriculum. The issue is not merely institutional lag, but the attempt to integrate AI skills into an architecture still organized around content delivery. This constitutes an optimization of an obsolete system.

The gap between the two accounts would matter less if the substantive side were reliably delivering results. However, the evidence suggests otherwise. Arum and Roksa’s study (2011) of over 2,300 undergraduates found that 45 per cent showed no significant improvement in critical thinking or complex reasoning during their first two years of college. More than a decade later, an OECD assessment (Van Damme & Zahner, 2022) across seven countries found that, on average, only 45 per cent of graduating students were proficient in critical thinking — the very skills universities consistently claim to develop. AI did not create this deficiency; it merely renders the gap between formal certification and actual capability impossible to ignore.

The critical question, then, is not whether students can use AI. It is what capacities must exist before AI becomes genuinely useful rather than merely convenient. In an environment of synthetic fluency, the primary bottleneck is not access to information, but the capacity to verify knowledge, challenge its premises, and act on judgment under consequence. If the university can no longer justify itself primarily through the management of content, what should it be organized around instead?

What post-content means

If content is no longer the organizing center, what then replaces it? In a post-content university, that unit is the task.

A task is not simply an assignment, a project, or a practical exercise added to a course. It is a structured problem situation in which a learner must determine what needs to be understood, tested, built, defended, or revised in order to act. The task may be technical, professional, scientific, civic, artistic, or philosophical. What matters is not that it resembles workplace practice, but that it requires judgment under constraints.

This is why post-content education is not a niche model for applied, technical, or entrepreneurial programs. It is a shift in how universities organize learning across domains. In engineering or medicine, the task may involve designing a system, diagnosing a case, or making a decision under risk. In public administration, it may involve weighing competing interests and acting under institutional constraint. In the arts, it may involve choices of form, critique, constraint, and resolution. Across these domains, serious intellectual work is not defined by content accumulation alone. People frame, interpret, decide, justify, revise, and take responsibility for what follows. The difference between fields is not whether there are tasks, but what kind of task is at stake, what counts as a valid answer, and what forms of error are tolerable.

By post-content, I therefore do not mean education without knowledge, disciplinary depth, or intellectual rigor. I mean a structural shift in which content remains essential but is no longer the organizing principle. The core moves to problems, tasks, judgment, and the formation of capability under graduated responsibility.

This shift does not weaken disciplinary knowledge. It gives it a different institutional role. Knowledge is no longer treated as a self-justifying pathway through the institution. It is mobilized in relation to increasingly complex problem spaces. The question is not only “what must be taught?” but “what must be known, constructed, tested, or challenged in order to act responsibly in this domain?”

This is not competency-based education under a new name. In many of its dominant institutional forms, competency-based education maps learning onto pre-specified frameworks. It shifts the unit from subject matter to skill descriptors, but often preserves a prescriptive and externally anchored logic. Post-content education differs at a deeper level. The task is not merely a delivery vehicle for pre-specified competencies. It is the site where the learner discovers what must be known and done in order to act. While CBE focuses on achieving pre-determined benchmarks, post-content education focuses on the ability to determine relevant knowledge and action under shifting conditions. True structural change requires that the task serves as the curriculum’s foundation, rather than an enrichment activity attached to a traditional syllabus.

The task as the unit of progression

Once the task becomes the organizing unit, the challenge is not simply to design better assignments, but to classify the kinds of problem situations through which students learn to think, judge, and take responsibility.

A task is not a generic project, and a problem space is not a loose invitation to “do something practical”. Tasks differ by structure, uncertainty, stakes, evidence, and the forms of judgment they require. At one end of the spectrum are well-defined problems: clear conditions, known methods, correct answers, and relatively stable criteria of success. Well-defined problems establish the necessary technical precision, conceptual fluency, and procedural discipline required for foundational expertise. At the other end are open-ended problems: unclear boundaries, evolving conditions, competing interpretations, incomplete evidence, and no single correct solution. Here the learner must decide what knowledge is relevant, what evidence is sufficient, what trade-offs are acceptable, and what kind of answer can be defended. At the highest levels — doctoral research, frontier innovation, public policy, artistic creation, institutional design — the task includes defining the field of action itself, formulating the problem, and proposing what counts as a valid resolution.

A content-era university may include the full spectrum of problems, but keeps it subordinate to curriculum. Well-defined problems appear as exercises. Open-ended problems are postponed to capstones, theses, studios, laboratories, internships, or doctoral work. Progression is still imagined as movement through content first, with complex judgment arriving later as application.

The post-content model organizes progression through the growing complexity of problem spaces, treating the curriculum as a sequence of tasks rather than topics. Through this process, students learn to identify, source, and validate the knowledge required for a given challenge; in doing so, they build the capacity for decision-making under uncertainty. This is not a sudden plunge into ambiguity. The transition to independent judgment is a designed trajectory: from bounded to ambiguous situations, from low-stakes practice to consequential work, and from supervised attempts to independent responsibility.

However, this trajectory presupposes something that cannot be taken for granted. The machine has made content cheap. It did not make thinking, judgment, or responsibility cheap. AI provides answers, but does not enforce the construction of a mental model nor any cognitive tension. Two students can submit the same AI-generated answer. One has understood nothing. The other has learned to think. The output is identical. If the machine reduces the cost of output, the university must preserve the cost of thought.

This means that a post-content university must deliberately set and protect the conditions under which thinking happens. Evaluation must require justification, not reproduction. The goal is to ensure that learning remains a process of building structure in the mind, not merely consuming or generating outputs.

Post-content education also presupposes learner agency: the ability to identify gaps, search for what is needed, and recognize the limits of one's own reasoning. For many students entering universities, these capacities are not yet formed. Foundation stages and early curricula cannot simply present students with open-ended problems and expect navigation. They must first scaffold the conditions under which independent judgment becomes possible, and treat this scaffolding as a core pedagogical task.

Nor does post-content education mean learning in isolation. A significant share of professional work is irreducibly collaborative. The formation of judgment includes the capacity to think with and against others, to defend positions in dialogue, and to navigate the cognitive friction that deep collaboration produces. Deep early professional networks, then, are not a social ornament around learning. They are part of the infrastructure through which judgment is formed.

This reorientation has resource implications. The shift from content delivery to task architecture does not reduce institutional cost — it reallocates it: away from lecture-scale content production and toward the design of problem environments, the scaffolding of judgment, and the development of more sophisticated forms of assessment. Institutions that treat this transition as cost-neutral will underinvest in precisely the parts that matter most.

If post-content education is serious, assessment cannot remain an incidental by-product of content delivery. The more universities claim to certify judgment, readiness, and access to responsibility, the less plausible it becomes to treat evaluation as a casual secondary function. In time, at least in some fields, universities may need more calibrated and professionalized systems of evaluation than those inherited from the content-era course model.

From knowledge to trust, and the right to be wrong

What does all this mean for the graduate? The future graduate cannot be defined primarily by what content they have covered. Nor can they be reduced to one more generic catalogue of “future skills”. The more serious question is what kind of person a university must form if knowledge is abundant, task structures evolve quickly, and responsibility remains human.

Such a graduate must be able to work with AI without surrendering judgment to it — directing intelligence rather than being directed by it. This means knowing how to frame problems, interrogate outputs, detect weak inference, and recognize when machine fluency is useful, insufficient, or dangerous. It means tolerating uncertainty without collapsing into imitation or deference, and moving across levels of abstraction, from domain detail to systemic context and back again. Crucially, it also means being able to justify decisions to other humans. Even when a system generates options or recommendations, accountability for action remains social.

Historically, that was always part of what the university certified, even if imperfectly. A degree never only meant “I know”. It also meant “I can be trusted with a role”, “I can be trusted to exercise judgment within a defined domain”, or “I can be trusted to carry some portion of society’s responsibility”. That trust was never purely meritocratic, and never evenly distributed. Yet, it was real. Today, knowledge is easier to access than ever, but trust is not.

That shift matters more than many current debates allow. A society can tolerate wide variation in how people acquire knowledge. It is far less relaxed about who is allowed to decide, diagnose, approve, interpret, prescribe, build, govern, or be wrong in consequential settings. Universities have always been part of that allocation system. They do not merely teach. They participate in deciding who may enter domains where error has cost.

This is why I resist the opposite simplification as well: the idea that universities should stop paying attention to jobs because the future is too unstable. That would be careless. Dynamic labor-market sensing is now basic hygiene. Universities need much better ways of monitoring changing task structures, role evolution, sector signals, and employer expectations. Job descriptions are not a final blueprint for curriculum, but they are an important signal stream. Institutions that ignore them risk inertia disguised as timelessness. Institutions that obsess over them risk panic disguised as relevance.

A post-content model offers a third path. It takes labor-market sensing seriously, but refuses to translate it mechanically into static curricular lists. Instead, those signals should feed a more dynamic institutional process: redesigning tasks, updating capability pathways, revising program architectures, and distinguishing between what remains foundational and what must change rapidly.

This applies beyond the corporate sector. Research itself is increasingly AI-mediated, and even in science the structure of responsible work does not disappear — it intensifies. The same is true of the state, healthcare systems, courts, schools, museums, and civil administrations: all allocate responsibility, and in all of them graduates interpret, justify, advise, decide, and act under consequences. The post-content question is therefore broader than employability in the narrow sense. It concerns how a society licenses people into increasingly complex forms of responsible action. The harder question is: who is allowed to be wrong, in what domains, and under what conditions?

Why legibility wins

If the argument is persuasive, why do universities not change faster? The answer is not simply conservatism. It is also that higher education systems reward legibility. What can be specified, modularized, compared, and audited travels more easily through regulation than what is situational, developmental, and harder to evidence.

Part of the answer is cultural inertia. Part is organizational convenience, but the other part is more structural. Universities are not free design spaces. They are governed by accreditation frameworks, qualification regimes, credit systems, program approval procedures, and quality assurance routines that still encode an older ontology of higher education. Even where these systems now include learning outcomes, student-centered language, and more flexible design principles, legitimacy still runs heavily through program architectures made legible in terms of components, credits, workload, sequencing, and formally assessable outcomes.

The problem, then, is not that quality assurance systems have remained frozen in a crude seat-time past. Many have evolved. The problem is that even improved systems still require the university to be readable as a program composed of stable, auditable parts. Learning outcomes by themselves do not make post-content education either possible or impossible. They can survive within very traditional models and within much more experimental ones. The real issue is what kinds of evidence, progression, and institutional design they are permitted to anchor.

A useful illustration comes from leading European institutions experimenting with challenge-based learning. At Eindhoven University of Technology (2023, 2024), students increasingly work in interdisciplinary settings on open-ended, real-world challenges, often in collaboration with external stakeholders. Pedagogically, this moves much closer to a post-content model: learning organized around tasks, uncertainty, and problem resolution rather than predefined content sequences. Yet institutionally these experiences must still be translated back into ECTS-based workload units. In EUT's formal architecture, one credit corresponds to a fixed quantity of study time, and the academic year remains legible through standardized credit accumulation. The result is a revealing mismatch. Learning becomes non-linear, task-based, and outcome-driven; institutional recognition remains linear, time-based, and standardized. Innovation happens in pedagogy. Equivalence remains administrative.

This is where the conflict becomes sharper. Post-content architectures depend on fluid recombination of knowledge around tasks, non-linear progression, richer forms of situated evidence, and staged access to responsibility under conditions of uncertainty. Most accreditation systems, by design, are strongest where learning can be specified *ex ante*, modularized, compared, and audited through relatively stable program descriptions. Those systems are not irrational, but they may be better at validating formal educational descriptions than at detecting whether meaningful formation is taking place.

That constraint becomes clearer in more strongly regulated contexts. The point is not confined to any one regulatory culture. Different systems phrase the problem differently, but many converge on the same administrative preference: curriculum, credits, program design, and auditable units remain the dominant grammar of proof. In India, recent reforms under the National Education Policy have expanded flexibility through a Curriculum and Credit Framework, multiple entry and exit options, and broader learning-outcomes language, but the architecture remains explicitly framed through curriculum, credits, and program design (University Grants Commission, 2022, 2023, 2024). In China, national undergraduate teaching quality standards have been articulated through program-level standards with detailed qualitative and quantitative indicators

across hundreds of disciplines (Ministry of Education of the People's Republic of China, 2018). In Uzbekistan, state educational standards, curricula, and training program remain central to higher education regulation and accreditation, and the published quality assurance materials make the course-credit-plan logic particularly visible (Republic of Uzbekistan, 2020; The Quality Assurance Agency under the Cabinet of Ministers of the Republic of Uzbekistan, 2025). This is not a feature of less developed or transitional systems alone. In the United States, the Higher Learning Commission (n.d.-a, n.d.-b) continues to define institutional legitimacy substantially through credit-hour equivalency — the same unit that, as Laitinen (2012) documented, was originally designed not to measure learning at all, but to calculate faculty pension eligibility. The grammar of proof is global, and it is old.

This does not mean the solution is to attack accreditation from the outside or to romanticize total institutional freedom. Degrees still matter because legitimacy, mobility, status, financing, and public trust matter. The stronger question is whether all legitimate institutions must be forced into the same grammar of proof.

Some of the most interesting educational experiments already point beyond the content-first model. One example is Oneday's entrepreneurial MBA in partnership with William Jewell College (Oneday, 2025; William Jewell College, n.d.). The educational proposition is task-first: students build a business while learning what is needed in the process. What makes the case interesting is not simply its entrepreneurial theme, but its unit of learning. The learner is not first moved through a content sequence and only then invited to apply it. The business itself becomes the organizing task, and knowledge is pulled in as needed under conditions of consequence. In that sense, the educational logic is already post-content. Yet for the qualification to become formally legible, the model still has to be translated into a recognizable degree architecture (Oneday, n.d.-a, n.d.-b). The point is not that all higher education should become entrepreneurial. But it is a revealing edge case: even genuinely post-content learning models often have to wrap themselves in the old grammar in order to become institutionally legible.

A different edge case comes from models that sit largely outside traditional degree structures. The 42 Network of coding schools (The 42 Network, n.d.) organizes progression through projects and peer evaluation rather than lectures, formal instructors, or fixed curricular sequences. Here the unit is neither content coverage nor seat time, but demonstrated capability through increasingly difficult problem-solving. The fact that such models can scale and produce labor-market recognition without reproducing the classical course-credit logic suggests that post-content architecture is not a fantasy. It already exists, though often outside the zones of strongest institutional legitimacy.

Even highly redesigned institutions such as Minerva (Minerva University, n.d.) show a similar pattern: pedagogy can be radically reworked, but formal legitimacy still tends to return through recognizable credits, degrees, and approved program structures. That is the institutional lock-in problem. If the architecture of meaningful learning is changing, then the systems that authorize, compare, and validate universities will eventually need to change as well.

What leaders can do

This is where university leaders become central, not as implementers of a predetermined model, but as the people best positioned to ask different questions about what their institutions are actually for. Presidents, rectors, and provosts may not be able to redesign national quality assurance systems on their own, but they are not powerless

either. They can demand differentiated regulatory pathways rather than one universal grammar of legitimacy. They can also insist that the question is not only how institutions are reviewed, but what counts as valid evidence of formation in the first place. They can push for earlier and more varied proof points rather than relying almost exclusively on end-of-degree recognition. They can ask whether all institutions should be judged against the same architecture of evidence, or whether quality assurance should take mission, design, and intended educational form more seriously.

They can also build internal systems that are more ambitious than the external minimum. Even within today's constraints, institutions can organize progression around more demanding tasks rather than only course completion; treat assessment as demonstration, defense, and evidence-rich observation rather than grading alone; create ladders of supervised responsibility; and develop stronger external signal loops between program review and real task evolution in the world outside the university.

Some leaders may wish to go further. Where regulation permits, they can argue for experimental statuses, sandbox pathways, or pilot recognition models that allow institutions to demonstrate value earlier than the full degree cycle. They can propose that legitimate evidence of formation should include not only curricular intention and credit accumulation, but also externally validated capability milestones, situated performance, and progressively trusted responsibility. They can ask whether the degree should remain the only socially legible unit of proof.

None of this requires abandoning learning outcomes. Nor does it require pretending that responsibility can already be measured with elegant simplicity. Responsibility is harder to evidence than content mastery, but that is not a reason to exclude it from institutional design. It is a reason to develop better forms of staged judgment, supervised practice, defense-based assessment, external validation, and longitudinal observation of performance under real or near-real conditions.

This is also why the present moment may require an uncomfortable question from those who work in accreditation and quality assurance. If formal compliance so often coexists with weak educational performance and limited intellectual formation across the mass sector, then the issue may not be only institutional failure. It may also be that our dominant systems of authorization rely too heavily on proxies of quality that are too weak to govern the outcomes they claim to protect.

None of this implies a single institutional model. Some forms of deep disciplinary study remain indispensable. Some fields require long periods of structured immersion. Some institutions should remain intensely research-led. Others may become more applied, more regional, more hybrid, or more explicitly developmental in mission. Strategic differentiation will matter more than rhetorical completeness. However, what is becoming harder to defend is the assumption that content progression, course completion, and degree accumulation can remain the unquestioned center of the institution.

What a post-content university would look like

The shift to post-content is not only an educational redesign, but also a societal one. The university has never been only an instructional system. It produces and verifies knowledge — a function that post-content logic does not diminish but intensifies, particularly at the level of doctoral research and frontier inquiry, where defining the problem is itself the highest-order task. It forms citizens capable of participating in public life and reading contested evidence critically — a function that becomes more consequential, not less, as AI-generated content saturates public discourse. It enables social mobili-

ty — and in a world where professional knowledge is increasingly accessible outside institutional walls, the mobility function may gradually decouple from exclusive control over content and reattach to the capacity to certify demonstrated judgment and trusted capability. It anchors regional economies and professional communities in ways that no platform can replicate. None of these functions disappears in a post-content model, but each of them depends, ultimately, on the university's ability to certify something real. An institution that can no longer credibly validate judgment and capability gradually loses its claim to perform any of these functions with authority.

The university after AI will have to answer three questions with much greater precision than before. What capabilities are we actually here to build? What forms of responsibility are we preparing people to assume? And what essentially human and social work does this institution organize that cannot simply be outsourced to platforms, tools, and generic content providers?

In a post-content model, the central educational question is no longer simply “what content has been covered?”. It becomes “what can this learner now be trusted to do, explain, defend, and take responsibility for?”

The future of higher education will not be decided only by elite institutions. In fact, elite institutions may be able to survive strategic confusion longer than the rest. The real future may be decided by the majority: institutions under constraint, without symbolic protection, forced to justify themselves through function rather than inheritance. What remains worth organizing as a university when content is no longer scarce, but responsibility still is?

We do not yet have robust, scalable ways to assess judgment rather than coverage — and the gap between aspiration and institutional practice is real. We do not know how to build learner agency at scale in systems where students arrive with fragile foundations and institutions operate under severe resource constraints. We do not know what the transition looks like for the many educators whose professional identity and labor conditions were shaped by the content-era model. And we do not know whether accreditation systems can evolve quickly enough to authorize new forms of formation before the institutions that most need to change run out of time. These are genuine design tensions that will require institutional courage and political negotiation.

My own answer is as follows. The university should become less a place for moving students through content pathways. Instead, it should focus on sharpening judgment, certifying true capability, and vetting those ready for high-stakes responsibility. We cannot predict which job titles or domains will trend next, but we know for certain that constant change is the new baseline. In that environment, the technical details matter less than the ability to navigate social complexities and actually take ownership of the results. Knowledge is just the starting point; the university's real future lies in providing the human oversight and integration that an algorithm never will.

References

- Arum, R., & Roksa, J. (2011). *Academically adrift: Limited learning on college campuses*. University of Chicago Press.
- Eindhoven University of Technology. (2023). *Master's CS programme guide*. https://assets.w3.tue.nl/w/fileadmin/Education_Guide/Mathematics%20_%20Computer%20Science/M_CS%20-%20general%20department/Program%20Guide%202023%20-%20Masters%20CS%20-%20Version%201.1.pdf

- Eindhoven University of Technology. (2024). *Student statute 2024–2025*. https://assets.w3.tue.nl/w/fileadmin/Education_Guide/Content/Praktische%20zaken/Regelingen%2C%20gedragscodes%20en%20richtlijnen/Studentenstatuut/DEF%20Student%20statute%202024-2025%2027%20juni%202024.pdf
- Higher Learning Commission (n.d.a). *Assignment of credits, program length and tuition*. <https://www.hlcommission.org/accreditation/policies/assignment-of-credits/>
- Higher Learning Commission. (n.d.-b). *Eligibility requirements*. <https://www.hlcommission.org/accreditation/policies/eligibility-requirements/>
- Laitinen, A. (2012). *Cracking the credit hour*. New America Foundation & Education Sector. <https://www.newamerica.org/education-policy/policy-papers/cracking-the-credit-hour/>
- Minerva University. (n.d.). *Academics*. <https://www.minerva.edu/academics/>
- Ministry of Education of the People's Republic of China. (2018, January 31). *MOE publishes first national standard for teaching quality of undergraduate programs*. https://en.moe.gov.cn/News/Top_News/201801/t20180131_326118.html
- Oneday. (n.d.-a). *Accreditation*. <https://www.oneday.org/us/accreditation>
- Oneday. (n.d.-b). *Launch your business idea and earn your MBA*. <https://www.oneday.org/us>
- Oneday. (2025). *A unique collaboration with William Jewell College* [Curriculum document]. <https://website-cdn.oneday.io/docs/curriculum-us-compressed-oct25.pdf>
- Republic of Uzbekistan. (2020). *Law of the Republic of Uzbekistan “On education” (O’RQ-637)*. <https://natlex.ilo.org/dyn/natlex2/natlex2/files/download/116503/UZB-116503.pdf>
- The 42 Network. (n.d.). *About 42*. <https://www.42network.org/about/>
- The Quality Assurance Agency under the Cabinet of Ministers of the Republic of Uzbekistan. (2025). *Education quality assurance ecosystem of Uzbekistan*. https://nqaae.uz/uploads/pages/nqaae_standarts.pdf
- University Grants Commission. (2022). *Curriculum and credit framework for undergraduate programmes*. https://www.ugc.gov.in/pdfnews/7193743_FYUGP.pdf
- University Grants Commission. (2023). *National higher education qualifications framework*. https://www.ugc.gov.in/pdfnews/2990035_Final-NHEQF.pdf
- University Grants Commission. (2024). *SOPs for implementation of National Credit Framework*. <https://www.ugc.gov.in/pdfnews/SOPs%20for%20implementation%20of%20NCrF%20Final.pdf>
- Van Damme, D., & Zahner, D. (Eds.). (2022). *Does higher education teach students to think critically?* OECD Publishing. <https://doi.org/10.1787/cc9fa6aa-en>
- William Jewell College. (n.d.). *Entrepreneurial MBA*. <https://www.jewell.edu/programs/emba>

Chapter 3. The knowledge trapezoid: expertise, depthification, and transfer

Charles Fadel

*“Generative AI didn’t break Higher Ed —
it just held a mirror to its brittle foundations”
— Randy Ziegenfuss, Ed.D., Moravian University*

This conceptual paper explores the requirements for a modernized Higher Education Bachelor’s degree capable of responding to the profound transformations introduced by Artificial Intelligence (AI) and its rapid acceleration.

Emergence of generative AI has exposed structural weaknesses in existing higher education models that had long been accumulating beneath the surface. Universities increasingly face a world in which information access is abundant, procedural cognitive work is automated, and many traditional forms of academic assessment are losing their meaning. The question is therefore no longer simply how universities can “integrate AI tools” but rather what universities should cultivate in human beings in the first place, when AI systems can already perform large portions of the cognitive work once associated with expertise. In other words, rethinking the “what” not just the “how”.

This paper explicitly does not cover a number of relevant issues, in order to maintain focus: for instance, the tension between the economic and psychosocial needs of education — that is a tired and unnecessary dichotomy, as the needs are similar if not identical. Nor does it address specific employability challenges across a 20-year horizon, the reluctance of some employers to offer internships, the resistance to change within Higher Education, or the implementation challenges related to faculty expertise and budget allocation. It does not attempt to forecast the precise trajectory of AI’s technological development — one must assume substantial progress over coming decades while recognizing the irreducible unpredictability of AI’s “jagged profile”. And it sets aside the global diversity of university systems in order to maintain focus on the U.S. 4-year Bachelor’s degree as a tractable case.

The central question therefore remains: *What should universities teach in the Age of AI?*

Behind this practical question, however, lies a deeper one. AI is not simply another technological wave comparable to earlier educational technologies. For the first time ever, humanity has at its disposal not just a communication technology (from scripts to the Internet) nor just a calculating technology (computers from the 1950’s until today), it has now access to a cognitive technology. This cognitive technology functions as an externalized reasoning capacity embedded within everyday work, communication, research, and decision-making — an “exo-brain” of sorts. As a result, AI changes not only labor markets but the structure of knowledge itself: what humans need to know, how expertise functions, how judgment is formed, and what kinds of human capability remain distinctively valuable. Employability is one manifestation of this shift, but it is not the whole story.

The uber-goal: wisdom

In the book *Education for the Age of AI* (Fadel et al., 2024), the Center for Curriculum Redesign (CCR) advocated for Wisdom as the essential goal of education — a concept existing in a feedback loop with the increasingly influential notion of human flourishing, itself rooted in the ancient conception of *eudaimonia* (Fadel, 2024).

“Wisdom” is often imagined as a contemplative or solitary activity. CCR’s synthesis research shows that, instead, Wisdom is fundamentally active: knowing “what?” must be followed by understanding “so what?” and then acting through “now what?”. This distinction becomes especially important under conditions of AI abundance. When information retrieval and procedural generation are increasingly automated, the educational challenge shifts toward discernment, judgment, contextual understanding, ethical orientation, and meaningful action under uncertainty. The challenge is no longer merely acquiring information. It is developing the capability to orient oneself wisely within rapidly changing systems where cognitive authority itself becomes unstable.

This active conception of Wisdom was recently echoed by the Cosmos Institute’s notion of the “Philosopher-Builder” — a person capable not only of reflection but of constructive action in the world, with Benjamin Franklin offered as an archetypal example (Cosmos Institute, 2025). The future university cannot merely prepare students to consume knowledge or execute procedures. It must cultivate individuals capable of understanding complex systems, collaborating with AI, exercising sound judgment, and acting constructively in environments characterized by ambiguity and accelerating change.

Structure: $Y_{1-2} + Y_{3-4}$

Given the uncertainties posed by future labor markets and AI capability growth, what should universities teach — and what should students learn? This can be turned into a deeply personal question: what would *you* study today, entering university in an age where AI increasingly automates cognitive work?

One could easily become trapped in philosophical debates regarding the future value of human intelligence itself. Yet educational systems cannot wait for certainty. This paper therefore proposes a practical “insurance policy” approach: in a world of rapid disruption, students should graduate not just as narrow specialists optimized for fragile labor-market niches, but also as multidimensional individuals with broad capabilities, transferable understanding, and the ability to continually re-specialize throughout life. This is not merely labor-market adaptation. It is preparation for operating within unstable and rapidly evolving knowledge ecologies — and the responsibility for this preparation belongs primarily to universities, not to students alone. A simple metaphor is that of a Swiss army knife: a polymathic student, good at a number of tasks, and capable of “sharpening the blade” (deepening expertise) as circumstances require.

A common reaction to AI disruption is that university duration should be shortened, given rising costs and mounting employability pressures. This paper argues the opposite. The challenge posed by AI is not an efficiency problem. It is a human capability formation problem under conditions of accelerating complexity. Students require time for maturation, for conceptual and epistemic understanding to develop, for transfer capabilities to be built through repeated application, and for real-world interdisciplinary integration to occur. The U.S. cost crisis should be addressed through political, institutional and financial redesign — not by collapsing educational depth.

Years 1–2: Polymathy

The first two years of a Bachelor’s degree should function as a broad polymathic foundation — resembling the traditional strengths of Liberal Arts education, but deliberately redesigned for the Age of AI, and not as a refuge for the STEM-phobic. Students should encounter multiple disciplines not as collections of facts but as distinct ways of understanding reality, producing knowledge, framing evidence, and solving problems. The purpose is not shallow exposure. It is the development of cognitive breadth, interdisciplinary orientation, and robust foundations for later specialization.

General Education (GE) should therefore be redesigned around broad, synthetic understanding rather than fragmented specialization. Too often, existing GE reflects the over-specialization of faculty rather than the developmental needs of students. The result is courses taught at a depth calibrated for future specialists rather than future polymaths — mathematics taught as if producing professional mathematicians, history as if producing professional historians. This reliably produces alienation rather than engagement.

The corrective is not simply less depth but *right-sized* depth. CCR proposes three levels of engagement adaptable to student trajectory: Appreciate, Interpret, and Produce. Students should understand what a discipline contributes, how it structures knowledge, how it interprets reality, and where its methods become useful — without necessarily mastering its full technical apparatus.

Within this framework, three domains should also become mandatory for all students regardless of eventual specialization. Technology and Engineering, because economic systems increasingly revolve around technological infrastructures — AI, biotechnology, advanced manufacturing, digital platforms, and cyber-physical environments — and students need not become engineers but must understand these systems deeply enough to participate intelligently within them. Social Sciences, because as AI transforms social organization, labor, communication, and identity, the ability to understand institutions, incentives, human behavior, and governance becomes increasingly important. And Entrepreneurship — not merely as a business specialty but as a mode of operating under uncertainty, requiring initiative, self-direction, and the capacity to create rather than merely fill roles.

During these first two years, students should also explore their own capabilities, motivations, values, and possible trajectories, described in the CCR framework as “Identity and Belonging; Agency and growth mindset; Purpose and passion”. The capstone of this phase should be an explicit examination of AI’s implications for different professions, human comparative advantage, and emerging societal needs. Students then select majors not merely according to market trends but according to the intersection of capability, purpose, social contribution, and realistic employability — the spirit captured by the concept of *Ikigai*: what one can be paid for, what one can become good at, what one loves, and what the world needs.

Years 3–4: jump above the missing rungs

Once specialization begins, the challenge shifts fundamentally. AI increasingly threatens to eliminate the lower rungs of professional careers: junior analysts, entry-level coders, routine legal research, basic content production, and many forms of predictable cognitive labor. This changes the meaning of “job-ready”. Students are no

longer competing only with other graduates — they are competing with AI systems capable of performing significant portions of entry-level work.

The logical response is not to outrun AI but to ride it. This means an AI-assisted “depthification” of education: using AI systems not as substitutes for thinking but as accelerators for experimentation, simulation, iteration, and real-world problem-solving. A recent example from Harvard mathematics is instructive (Boles, 2025). In the 2025 version of Applied Mathematics 201, the instructor eliminated traditional homework problems. Instead, students had to create their own problems, have classmates verify them, and attempt to outsmart an AI. By semester’s end, students had generated nearly 700 problems of increasing difficulty. As the instructor observed: *“It’s just amazing, because I’m teaching math, right? And we’ve got a situation where the students are inventing math problems that are harder and harder and harder, and trying to solve them. That is the dream”*.

Educational challenge in the Age of AI is therefore no longer simply solving known problems. It is framing problems, questioning assumptions, designing systems, evaluating outputs, transferring understanding across contexts, and operating effectively alongside intelligent systems. The final two years should include repeated internships, multidisciplinary projects, AI-augmented workflows, and increasing exposure to real-world complexity — ensuring students graduate having already participated in environments requiring collaboration, adaptation, and judgment under uncertainty. This is how universities help students jump above the rungs that AI is eliminating.

Curriculum: expertise AND transfer

The role of education is twofold: the development of *expertise*, and the ability to apply that expertise across contexts — also known as *transfer* (described for instance by David Perkins (2014) work). In the Age of AI, this distinction becomes critically important. AI systems increasingly make inert expertise insufficient. What matters is not merely possessing knowledge, but understanding its conceptual structure, detecting its patterns, recognizing its limitations, applying it flexibly, and using it wisely in novel situations.



Figure 1: Expertise AND Transfer (CCR, adapted from (Perkins, 2014))

Deeper knowledge: from triangle to trapezoid

Traditional education focused heavily on *declarative knowledge* — facts, definitions, information — and *procedural knowledge* — methods, algorithms, execution. AI systems increasingly perform both with considerable facility. Educational value therefore migrates upward toward Conceptual knowledge and Epistemic understanding.

Conceptual knowledge means understanding the deeper organizing principles behind a discipline — not merely manipulating the formula for exponential growth, but grasping its behavior as “deceiving then explosive”. Epistemic knowledge means understanding how a discipline generates, validates, and revises knowledge: what counts as evidence, how evidence is generated, how claims are justified or refuted, and how knowledge is communicated and updated. This epistemic layer becomes increasingly important precisely because generative AI can produce fluent but unreliable outputs. Students require not merely access to knowledge but the ability to evaluate claims, interrogate assumptions, challenge AI-generated outputs, and reason about uncertainty itself.




Figure 2: Deeper Knowledge (Fadel et al., 2024)

CCR therefore proposes moving from the traditional knowledge triangle — wide at the declarative base, narrow at the top — toward a *knowledge trapezoid* in which the upper layers of conceptual and epistemic understanding are substantially expanded (Fadel et al., 2024). Deep expertise requires all four layers; the question is augmenting the ones education has historically neglected, and which ones AI now makes urgent, while simultaneously and carefully pairing back the lower layer, which are still necessary scaffolding. The analogy here is of the game “Jenga”, where one deconstructs a tower of blocks by removing the non-load-bearing beams.

This transformation thus forces a reconsideration of what counts as foundational knowledge. Not all declarative and procedural content remains equally important when AI systems increasingly externalize memory and automate routine cognition. The question is no longer “what content has traditionally been taught?” but rather “what knowledge remains essential for conceptual understanding, transfer, and epistemic robustness?” This is fiendishly difficult to determine, and personal biases abound — faculty specialists have a hard time disentangling what is truly foundational from what simply reflects their own intellectual formation. CCR is developing criteria for these judgments, organized around both expertise (disciplinary coherence, scaffolding for learning, insti-

tutional priorities) and transfer (conceptual substrates, epistemic grounding, interdisciplinary applicability).

Exercising transfer

Knowledge remains inert unless students repeatedly apply it across multiple settings. Transfer develops through multidisciplinary projects, internships, simulations, and AI-assisted problem-solving in increasingly complex contexts. But transfer also requires the explicit cultivation of competencies alongside disciplinary expertise. CCR's Four-Dimensional Education framework (Fadel et al., 2015) identifies three competency domains that become especially important under AI conditions.




Figure 3: Four-Dimensional Education (CCR, 2015)

Skills — Creativity, Critical Thinking, Communication, Collaboration — because these capacities increasingly define the distinctively human contribution within AI-augmented workflows. Character — Curiosity, Courage, Resilience, Ethics — because rapid change requires not only cognitive capability but dispositional stability and ethical orientation. And Meta-Learning — the capacity to reflect, adapt, and learn how to learn — because information changes rapidly, work reorganizes continuously, and human value increasingly depends on flexibility, discernment, and self-direction.

Research has shown that the majority of students graduate after four years being no better critical thinkers than when admitted (Van Damme, Zahner, 2022). This is an indictment not of student capability but of educational design. Teaching competencies *explicitly* — deliberately, systematically, comprehensively, demonstrably — requires significant professional development by faculty and sophisticated (AI-assisted) lesson-plan design. It cannot be assumed to happen organically.

The special case of meta-learning

Meta-Learning occupies a uniquely important position in the Age of AI, and deserves extended treatment. CCR defines Meta-Learning as the capacity to reflect, adapt, and learn how to learn (and do so continuously). Each of these facets has growing importance as AI systems evolve faster than static educational models can follow.

Reflection — thinking about thinking — becomes essential because humans must critically evaluate AI-generated outputs rather than accepting them uncritically. This requires metacognition, conceptual clarity, mental models, and the ability to reason about one's own reasoning. The most cognitively agile formulation comes from Moldoveanu (2010), who describes a “diamind” that *thinks about thinking while thinking* — a kind of second-order cognition that becomes increasingly valuable when the first-order cognitive work is partially delegated to machines.

Adaptability is both a cognitive and a dispositional challenge. One useful definition of intelligence simply is adaptability — and adaptability is at a premium during turbulent times. Yet universities tend to be controlled, low-risk environments, and some have become places where even dissenting ideas face suppression. Preparing students for an unpredictable world requires deliberate exposure to uncertainty, experimentation, failure, and real-world unpredictability. This can be achieved through improvisation, project-based learning, crisis simulations, physical activities requiring adaptation (such as scouting and sailing), entrepreneurial experiences, and inquiry practices (such as “Five Whys” of Total Quality Management, followed by the Five Hows of CCR) — not as extra-curricular enrichment but as mandatory components of the educational architecture.

Learning how to learn (continuously) is perhaps the most durable of these capacities. In a world of constant and rapid change, the ability to self-direct one's learning becomes a foundational survival skill. Projects and internships should increasingly place students into environments requiring rapid knowledge acquisition, just-in-time learning, and continual upskilling. The future graduate must be capable not merely of applying existing expertise, but of continuously rebuilding it throughout life.

Traditional universities will face significant challenges adapting rapidly enough to AI-driven transformation. This will open the door for new educational challengers willing to design curricula explicitly oriented toward robust human capability formation — institutions that take seriously what AI disruption actually demands, rather than applying incremental reforms to structures built for a different era.

Yet even if the recommendations proposed in this paper are fully implemented, two large questions remain unanswered — and they matter enormously.

1. Are all human beings genuinely capable of adapting to the level of cognitive and societal transformation that AI introduces — both emotionally and in terms of general cognitive capacity?

2. And will economic structures continue to provide meaningful, sustainable roles for large portions of humanity as AI capability expands — or are we facing a fundamental restructuring of the relationship between human labor and economic value?

No curriculum can fully answer these questions. Nevertheless, universities still possess a uniquely important responsibility. As AI increasingly automates information retrieval and procedural cognition, the central educational challenge becomes the cultivation of judgment, wisdom, transfer, adaptability, epistemic discernment, and meaningful human agency. These may prove to be the most important “future-proof” capabilities of all — and developing them is, and should remain, distinctively human work.

References

- Cosmos Institute. (2025, July 16). *The philosopher-builder: The AI age demands a new kind of technologist*.
- Fadel C., Black A., Taylor R., Slesinski J., Dunn K. (2024). *Education for the age of AI*. Center for Curriculum Redesign.
- Fadel C. (2024). *The relationship between wisdom and flourishing*. CCR Opinion Paper. <https://curriculumredesign.org/wp-content/uploads/The-Relationship-Between-Wisdom-and-Flourishing-CCR.pdf>
- Fadel, C., Bialik, M., & Trilling, B. (2015). *Four-dimensional education: The competencies learners need to succeed*. Center for Curriculum Redesign.
- Boles S. (2025). *AI leaps from math dunce to whiz*. Harvard Gazette. 23 July 2025 <https://news.harvard.edu/gazette/story/2025/07/ai-leaps-from-math-dunce-to-whiz/>
- Moldoveanu, M. (2010). *The design of insight: How to solve any business problem*. Stanford University Press.
- Perkins, D. (2014). *Future wise: Educating our children for a changing world*. Jossey-Bass.
- Van Damme, D., & Zahner, D. (Eds.). (2022). *Does higher education teach students to think critically?* OECD Publishing.

Chapter 4. The future of education: symbiotic intelligence and the AI+X paradigm

Shawn Chen, Richard Jiarui Tong, Xiangen Hu, and Jack Gao

The prevailing industrial-age model of education is not merely outdated; it is structurally incapable of preparing learners for the complexities of the emerging era of symbiotic intelligence (SI). An alarming alignment crisis has unfolded a widening chasm between what traditional schooling produces and the skills demanded by a world increasingly defined by human–AI cognitive partnerships. This alignment gap is more than a pedagogical concern — it poses a socio-economic and demographic threat to the stability of advanced economies. Educational outcomes focused on static knowledge and test scores are misaligned with a labor market that values adaptability, creative problem-solving, and the ability to leverage AI as a collaborator. To bridge this gap, we must fundamentally reorient education from its 20th-century focus on knowledge transmission to a 21st-century model centered on the orchestration of intelligence.

Exponential change and the human–AI partnership

Recent advances in AI, especially generative AI, have transformed it from a niche tool into a pervasive cognitive partner in nearly every field. AI systems can now produce human-like text, code, designs, and decisions, reshaping the nature of intellectual labor and professional work. Unlike previous technologies that mainly automated physical or routine tasks, modern AI encroaches on complex cognitive tasks in white-collar domains, disrupting job roles that were once considered the exclusive province of human experts. Traditional curricula, however, still prioritize content memorization and standardized testing, metrics that are increasingly irrelevant in an AI-augmented landscape.

As AI automates foundational procedures and information retrieval, organizations are flattening their structures and evolving from rigid hierarchies into agile, networked teams. In this new environment, every professional is envisioned as part of a human–AI cell: a symbiotic working unit where human judgment and creativity are amplified by AI’s computational power and speed. Companies adopting such “obelisk organizations” (tall and flat in structure, in contrast to broad pyramids) can achieve unprecedented efficiency and innovation. The onus is now on education systems to produce graduates who can thrive in these human–AI partnerships rather than being displaced by them. Early evidence already shows that collaborating with AI can significantly boost human productivity and change how we approach problem-solving. This context of exponential technological change demands that education move beyond imparting facts and routine skills and instead focus on higher-order capabilities: critical thinking, ethical reasoning, creativity, and the meta-skill of learning itself, all in tandem with intelligent machines.

The education–demography trap

Misalignment between education and workforce needs is further exacerbated by demographic trends in many advanced economies. In countries like South Korea, Japan, and China, the legacy education system has inadvertently contributed to a vicious cycle we term the education–demography trap. This trap is driven by three interlocking forces that depress fertility rates and delay the onset of productive adulthood:

1) Educational inflation: As the total domain of human knowledge expands exponentially, individuals are funneled into ever-longer periods of schooling to attain competitive credentials. Entry into the workforce is delayed into the mid-20s or later, as each generation pursues higher degrees considered necessary for employment.

2) Economic pressure: The time and financial investment required for extensive education—tuition, tutoring, examinations, and degrees—places a heavy burden on families. Parents often feel compelled to pour substantial resources into their children’s education, which in turn leads many to postpone or reduce family formation due to economic constraints.

3) Delayed autonomy: Prolonged educational pathways mean young adults remain economically dependent on parents or institutions well into their third decade of life. This postpones key life milestones—entering stable employment, marriage, and parenthood—resulting in an aging population with fewer children.

This cycle is unsustainable. In South Korea, a university degree has become the minimum qualification for white-collar jobs, and nearly 80% of schoolchildren receive private after-school tutoring (consuming about 10% of household income) in hopes of attaining elite university admission. This intense “education fever” yields diminishing returns: many graduates are overqualified or mismatched for the jobs they take. At a societal level, the consequence is a demographic implosion—fewer births and a rapidly aging population—as young people defer marriage and child-rearing in order to first clear the high bar set by educational inflation. To reverse these trends, a new educational paradigm must dramatically compress time-to-productivity, enabling individuals to become productive and self-sufficient at a younger age while reducing the financial and time burdens of education. In essence, education should no longer be a long detour delaying adulthood, but rather a springboard launching students into economic and social participation much earlier.

Symbiotic education as a demographic stabilizer

We propose that the AI+X program offers both a strategic solution to the education–demography trap and a profound philosophical shift in the purpose of education. At its core lies the concept of *symbiotic education* — a paradigm premised on the mutual enhancement and co-evolution of human and artificial intelligences. This stands in stark contrast to the outdated notion of AI as a human replacement; instead, AI is envisioned as a force multiplier for human learning and creativity. By reorienting educational objectives from the possession of static knowledge to the cultivation of dynamic capabilities, AI+X aims to produce graduates who can orchestrate diverse intelligences — their own and AI’s — to innovate and solve complex problems.

Crucially, the AI+X model can act as a “demographic stabilizer”. By accelerating students’ readiness for professional work, it enables meaningful workforce entry by age 16–18, years earlier than the conventional model. A student in an AI+X curriculum spends less time on rote learning and more time co-creating with AI, achieving in high

school what might otherwise be expected only of college graduates. This early professional competence, coupled with dramatically lower costs of advanced training, thanks to AI tutors and automated resources, means young adults can attain economic independence sooner. In turn, they can make life choices — starting a family, for instance — on a more natural timeline, without being locked into prolonged dependency. Fundamentally, the AI+X approach prioritizes capability generation and innovation cultivation over memorization, aiming to graduate individuals who are agile, creative, and symbiotically intelligent.

Program overview: from knowledge accumulation to symbiotic productivity

The AI+X program redefines “readiness” by making students the productive contributors years earlier than in the traditional model. Six core shifts define the transformation. First, where traditional schooling has students spend years acquiring abstract knowledge before they can do real work, AI+X has students co-learn and co-produce with AI from the outset: AI replaces rote learning and procedural drudgery, while students focus on conceptual reasoning, design intent, ethical decision-making, and problem framing — all known as uniquely human cognitive strengths. Just as tractors once liberated farmers from back-breaking manual labor and GPS freed drivers from memorizing routes, AI+X liberates learners from repetitive procedural cognition. Cognitive energy is redirected from rote tasks to creativity and critical thinking, enabling adolescents to contribute in ways that used to wait until mid-career.

Second, the conventional education-work progression (learn → graduate → work) is replaced by a blended model where learning happens through authentic work with AI from early on. An AI+X high school student might co-design an app, write a research brief, or prototype an engineering design — achievements typically expected only of college graduates — building industry-ready portfolios by graduation.

Third, since AI handles repetitive practice and information retrieval, students can progress sooner to analysis, evaluation, and creation. Even middle schoolers learn to direct and critique AI reasoning, exercising metacognitive skills typically seen at the graduate level.

Fourth, AI acts as a democratizing force: every student has access to personalized tutoring and sophisticated tools once limited to elite institutions, and human-AI co-modeling can amplify individual productivity by an order of magnitude (10×).

Fifth, the entire educational system is transformed: learning objectives shift from mastery of static knowledge to development of adaptive, co-creative intelligence; curriculum becomes project-based and interdisciplinary; assessment shifts from standardized exams to portfolios evidencing process and real-world results; and teachers evolve from lecturers to mentors and co-designers of learning experiences.

Sixth, cumulatively, these changes point toward a society that learns, works, and innovates in fundamentally different ways: a generation of lifelong learners who grow up already producing, innovating, and learning in symbiosis with intelligent systems.

Theoretical framework of AI+X: defining symbiotic intelligence

To successfully re-architect education for the AI era, we must establish a clear theoretical foundation that ensures pedagogical changes are not ad-hoc additions of technology but part of a coherent, systemic transformation of learning. The AI+X theoretical

framework revolves around two core concepts: symbiotic intelligence (SI) and the AI+X paradigm for education. Together, these define a new ecosystem of learning in which human and AI cognition are deeply interwoven.

Symbiotic intelligence: human–AI cognitive synergy. Symbiotic Intelligence can be defined as a collaborative cognitive ecosystem in which human and artificial intelligences mutually enhance each other to achieve goals neither could attain alone. It is a paradigm of co-creation, not competition. In an SI system, each participant contributes complementary strengths: the human offers creativity, contextual understanding, moral judgment, and the ability to set meaningful goals; the AI offers pattern recognition at scale, vast memory and knowledge retrieval, speed in computation, and consistency in executing procedures. By allocating roles based on these strengths, the human–AI pair functions as a unified cognitive unit greater than the sum of its parts.

In the context of education, symbiotic intelligence implies that students are not mere consumers of information delivered by teachers or computers; rather, they become active co-constructors of knowledge alongside AI agents. They learn how to learn with AI, developing a form of “dual intelligence” literacy. As Tinholt (2025) argues, human intelligence must evolve into a symbiotic relationship with AI, wherein humans retain agency and leverage AI to extend their cognitive reach. In education, this translates to students and AI tutoring systems iteratively engaging in dialogue, problem-solving, and creation, with each learning from the other’s feedback: the student learns domain knowledge while the AI system improves its modeling of the student.

Symbiotic intelligence reframes the debate on AI in education: the goal is no longer to see if AI can replace teachers or if automation can teach skills faster, but to design environments where AI and human intellects co-evolve. This means shared understanding (students and AI agents build models of knowledge together), adaptive role allocation (the division of labor in learning tasks is fluid, shifting based on whether a task requires creativity or computation), and a continuous feedback loop (ongoing, dynamic interaction where both parties contribute to advancing the knowledge state). Symbiotic intelligence in the classroom therefore means: a student working on a science experiment might propose a hypothesis, and an AI partner might suggest relevant data or simulate the experiment. Through iteration, the AI comes to anticipate the student’s conceptual gaps, and the student learns the limits and capabilities of the AI. If a task requires creativity or ethical consideration (e.g., writing an essay on a controversial topic), the human takes the lead while the AI offers supportive information. If a task requires brute-force computation or remembering many facts (e.g., analyzing a large dataset), the AI takes charge under human direction. The SI paradigm prepares students for a future in which distributed, symbiotic cognition is the norm. Learning to learn with AI becomes as fundamental as learning to read or write. Recent pedagogical theory supports this shift, suggesting that “distributed cognition” — intelligence distributed across people and tools — is a more realistic model of how advanced problem-solving occurs in the real world.

The AI+X paradigm: tripartite re-architecture of education

We formalize the integration of AI into all disciplines and learning activities with the expression AI+X, where “X” can be any subject or domain of study. This is not merely about using AI as a tool in existing subjects; it is about fundamentally re-architecting the components of education — the who, what, and how — in a tripartite model to enable symbiotic intelligence:

Who (Participants): Students evolve from passive absorbers of content into “Super Individuals” (shuang nao ren cai, literally “dual-brain talent”) who leverage both a biological brain and an AI “cognitive brain”. Each student learns to be the director of an ensemble of AI agents, effectively extending their cognitive abilities. Teachers shift correspondingly from transmitters of information to designers of learning experiences and orchestrators of intelligent tools—a meta-mentor who ensures that each student’s human–AI team is productive and on track.

What (Objectives): The goal is the development of dynamic cognitive capacities structured into three dimensions: (L1) building an internal world model of disciplinary knowledge, (L2) mastering AI collaboration skills, and (L3) cultivating metacognitive self-awareness to continuously improve one’s own learning process. Content knowledge is still present but reframed as part of L1; it is learned in service of solving problems and understanding systems, not as an end in itself.

How (Pedagogy): Learning is driven by Interactive Learning Tasks (ILTs) — open-ended, authentic challenges that students tackle with continual input from AI mentors, advisors, and collaborators. Instead of passively reading a chapter, a student engages in a Socratic dialogue with an AI tutor; instead of following a lab recipe by rote, the student designs an experiment and asks an AI to help simulate or analyze it; instead of writing an essay alone, the student brainstorms with an AI writing partner and then iteratively refines the result. Because AI can automate routine assessments, students receive instant feedback and can iterate rapidly, making learning more agile.

This tripartite re-architecture provides a blueprint for systemic change. It is not enough to sprinkle AI on top of existing curricula; the entire system’s purpose and operation must be rethought. In AI+X, learning itself is redefined as cognitive orchestration: the ability to dynamically marshal both human and machine intelligences to address complex tasks. This becomes the new literacy of the 21st century.

Philosophical underpinnings: knowledge as dynamic functorism

The shift in “what” and “how” of learning is underwritten by a shift in our concept of knowledge itself. Traditional education treats knowledge as a static stockpile of facts and formulas to be transferred. In the AI+X paradigm, influenced by ideas from category theory in mathematics, particularly the notion of functors that map structures from one context to another, we embrace a view of knowledge as a dynamic, transferable capability rather than a fixed inventory. We call this perspective *dynamic functorism* in knowledge. Knowing something is defined by the ability to use it as a mapping to generate results in new situations: a student truly “knows” Newton’s laws not when they can recite $F=ma$, but when they can apply that principle to predict or explain diverse physical phenomena or transfer the abstract principle to a completely different domain. Education is no longer about chasing “the right answers” (which an AI can retrieve in milliseconds). Instead, the purpose of education becomes cultivating the ability to construct, adjust, and transfer understanding to deal with uncertainty and novel situations.

What to learn: the three dimensions of learning objectives

Transitioning to a framework of capability construction requires a redefinition of educational goals. The AI+X paradigm organizes learning objectives into three interconnected dimensions that correspond to the major domains of cognitive development

necessary for symbiotic intelligence. These dimensions are not separate subjects; they are intertwined facets of an integrated learning architecture that ensure a student can simultaneously understand the world, harness AI tools, and direct their own growth.

L1: Disciplinary knowledge as world modeling. In the first dimension, traditional subject matter knowledge is reframed as the process of world modeling. Instead of treating a discipline as a static set of facts to memorize, AI+X treats it as a lens for interpreting reality and a toolkit for predicting outcomes. The learner's task is to construct an internal, structured model that mirrors how an expert in that discipline understands the world, i.e. the key entities, relationships, and rules that explain phenomena in that field.

The core methodology for L1 is first-principles thinking: training students to break complex systems down to their fundamental elements and governing principles, then reconstruct their understanding from the ground up. By doing so, students transition from being users of formulas or facts to being builders of knowledge structures. They learn to ask: "What are the basic assumptions here? If I rebuild this concept from scratch, what essential truths would I start with?" Through world modeling, knowledge becomes not just declarative (knowing that something is the case) but generative (knowing how to use knowledge to produce predictions and solutions).

Concretely, implementing L1 in the classroom involves several interrelated practices. Students are explicitly taught to construct models: for a given topic, they identify core entities and variables, map out relationships, and articulate the rules that connect those entities. They learn to distinguish different layers of abstraction — the concrete observable layer vs. an abstract theoretical layer — and see how the discipline connects these layers. Teachers and AI tutors continuously prompt students to justify their models from basic principles, encouraging them to derive formulas or results themselves whenever feasible. Depth is prioritized over breadth: by deeply modeling a few representative systems, students learn the process of modeling itself, which they can transfer to new topics. Assessment measures how well their models perform: can they predict outcomes accurately, explain a real-world phenomenon using their model, debug their model when it fails?

Consider a student learning classical mechanics. Rather than memorizing equations, they are tasked to "build a world" where gravity affects motion: they identify core entities (mass, position, velocity, acceleration, force, time), map relationships (Newton's second law, definitions of acceleration and velocity), and establish hierarchical layers connecting physical intuition to mathematical representation. They use first-principles thinking to ask: What are the minimal assumptions here? What invariants exist? Through this, they derive that gravitational motion is independent of mass (absent air resistance). Their assessment asks them to extend their model to a new, more complex scenario to demonstrate that they can adapt their base model rather than merely applying memorized formulas.

AI supports world modeling by providing dynamic simulations, visualizations, and by handling low-level computations, while the student remains the driver of model construction and revision. AI can generate multiple scenarios for a student to test their model against, or highlight inconsistencies in a student's model by comparing it with empirical data. To conclude L1: to learn is to model, and to master is to simulate and shape the world through the language of a discipline.

L2: AI Collaboration skills as human-AI co-modeling. The second dimension focuses on developing students' proficiency in working symbiotically with AI systems — not merely learning specific software or programming, but mastering the higher-order skill of cognitive teaming with machines. We often talk about collaboration as a human-to-human skill; here we extend it to human-AI collaboration. The student must

learn to effectively communicate with AI, leverage AI's strengths, compensate for AI's weaknesses, and integrate AI outputs into their own reasoning.

We introduce the concept of a Human–AI co-modeling loop that formalizes how a human and an AI can iteratively solve a problem together across five phases: (1) Context Framing—defining the situation, objectives, and constraints, while the AI interprets the human's input and retrieves relevant knowledge; (2) Role Allocation—deciding which subtasks to delegate to AI (e.g., data processing, brute-force search) and which to handle oneself (e.g., ethical decisions, final choices); (3) Interaction & Prompting—communicating intent clearly to the AI and executing joint work; (4) Validation & Reflection—critically evaluating AI outputs against expectations, theory, or external criteria, and identifying any errors or biases; and (5) Integration & Refinement—synthesizing AI-generated results with human insight, making final decisions, and iterating.

Key sub-skills that students develop include: context sharing and prompt engineering (learning how to “talk” to AI—clearly setting context, posing good questions, giving structured instructions); understanding AI's capabilities and limits (knowing that a language model predicts likely text and may sound confident but could be wrong, or that AI vision systems might be confused by unusual perspectives); strategic delegation (deciding what to offload to AI—a student solving a design problem might delegate generating 100 possible design variations to an AI, but retain the task of evaluating which designs are ethical, user-friendly, or aligned with the initial vision); and critical evaluation of AI output (cultivating healthy skepticism and checking AI answers against their L1 world model, thus creating a direct interplay between L1 and L2).

Learning these collaboration skills is staged through progressive levels of AI fluency: Awareness (knowing what AI tools exist and what they broadly can do), Understanding (conceptual grasp of how AI systems function, including their biases and limitations), Application (using AI tools for concrete tasks in a relatively linear way), and full Collaboration (integrating AI into creative, multi-turn dialogues for complex tasks, treating the AI almost like a colleague whose suggestions are valued but also critically appraised). In AI+X, the curriculum is designed to move students up this ladder systematically.

L2 is also where ethics comes into play heavily. Students learn that AI can produce misinformation or can be used maliciously. As part of their training, they discuss cases and set personal and community norms for responsible AI usage: practicing giving credit when AI helps create content, learning not to over-rely on AI in a way that would dull their own skills, and understanding the broader societal implications of AI deployment. Achieving L2 mastery turns what might be perceived as a threat — AI automating tasks — into an opportunity: students become not the ones replaced by AI, but the ones amplified by it.

L3: Metacognitive intelligence as self-modeling. The third dimension addresses the highest level of cognitive development in the AI+X paradigm: learning about one's own learning. Metacognitive intelligence involves a student's ability to understand, monitor, and adapt their own cognitive processes. In an age where fixed knowledge quickly becomes obsolete and AI can answer many routine questions, the capacity for self-directed, lifelong learning becomes paramount. L3 ensures that students become self-improving systems — they continuously reflect on what and how they learn, identify their weaknesses, and adjust strategies to become more effective over time.

We refer to L3 as self-modeling, in parallel to L1's world modeling and L2's co-modeling: the learner constructs and refines a model of their own mind, including their habits, biases, strengths, and areas for growth. To operationalize this, we incorporate a structured reflective-adaptive loop consisting of five phases: (1) Awareness—observing how one approaches a task, including emotions, tool usage, and strategies; (2) Re-

flection—analyzing the effectiveness of the approach, comparing expected outcomes to actual outcomes, and performing causal attribution about one’s own performance; (3) Adjustment—devising specific changes to methods or strategies; (4) Habit Formation—practicing the adjusted approach until it becomes more automatic; and (5) Growth Evaluation—periodically reviewing outcomes over a longer term and identifying new areas to work on.

In AI+X, this loop is actively taught and scaffolded. Students keep learning journals; AI-based reflective companions prompt them with questions like “How confident are you about this answer? Why?” or “What was your strategy for this problem?” After significant tasks, reflection sessions guide students through error analysis: did they misremember a fact (memory issue), misapply a concept (understanding issue), or misread the question (carelessness issue)? Such distinctions are critical because they lead to different remedies. Personal KPI tracking allows students to see improvement or stagnation in their own learning metrics over time, with AI aggregating data to provide insights.

Metacognitive development can be categorized into levels of proficiency: Recognition (the student can acknowledge their own cognitive and emotional states and articulate some of their learning habits); Analysis (the student can not only notice but analyze and explain their approach, identifying root causes for outcomes and seeing patterns); Regulation (the student can intentionally change their approach—if they know they make arithmetic mistakes, they budget extra time to double-check); and Optimization and Personalization (at the highest level, students design new methods for themselves and can even mentor peers in metacognitive strategies). In AI+X, we want every student to reach at least the Regulation stage, and ideally the level where they are self-transforming learners. This positions them for lifelong learning, because once you leave formal education, no one else is structuring your learning — you have to do it yourself.

How to learn: the agentic dual-loop methodology

Translating the ambitious learning objectives of AI+X into classroom practice requires an innovative pedagogical design. The AI+X framework is implemented through a process-driven methodology that integrates the three learning dimensions into a cohesive, iterative learning experience. Two core design principles anchor this methodology: (1) the use of specialized AI agents to scaffold different stages of learning, and (2) the separation and integration of conceptual learning vs. procedural execution—the “dual-loop” approach.

Three-stage agentic learning cycle. In AI+X, learning is envisioned as a continuous cycle of preparation, performance, and reflection, with each stage facilitated by AI agents playing distinct roles:

Stage 1: Socratic Playground of Learning (Pre-task, Conceptual Focus). An AI Socratic Learning Companion transforms what would traditionally be passive study into interactive, inquiry-driven exploration. The AI companion plays a Socratic role — asking thought-provoking questions, encouraging students to explain concepts in their own words, providing instant clarifications or analogies when stuck, and generating visualizations or simulations to demonstrate concepts dynamically. For example, if a student is learning about cellular biology, the AI can simulate cell processes or quiz the student on “what-if” scenarios. The outcome is validated conceptual understanding: concept maps the student and AI built together, or a transcript of a dialogue where the student’s answers demonstrate comprehension.

Stage 2: Interactive Learning Task (Execution, All Dimensions Integrated). The student tackles authentic, open-ended tasks that require applying disciplinary knowledge (L1), collaborating with AI (L2), and using metacognitive strategies (L3) in tandem. An AI Co-Creative Task Agent is available throughout as a brainstorming partner, simulator, and coach—suggesting design ideas, fetching relevant data, modeling energy consumption of a student-designed system, reminding the student of project requirements if they are stuck. The student leads in creativity and decision-making while the AI handles heavy computations and routine subtasks. The outcome is a concrete artifact or solution—a written report, a prototype, a presentation—co-created by human and AI with transparent acknowledgment of each contribution.

Stage 3: After-Action Review (Post-task, Metacognitive Focus). An AI Reflective Companion guides structured reflection to consolidate learning and improve future performance. The AI prompts the student to review what happened—“What were your initial goals? Did you achieve them? Where did you encounter difficulties?”—and might show a timeline of the student’s actions. Using such data, the AI can prompt deeper reflection: “I notice you tried three different approaches in coding. What made you change course each time? What did you learn from those attempts?” The emphasis is on distilling lessons: both about the subject matter (“I realize I still don’t understand X fully”) and about the process (“Next time I will start coding with simpler test cases”). The outcome is actionable insight that feeds forward into the next cycle.

These stages form a repeating, continuous cycle. The insights from Stage 3 feed into the next Stage 1, ensuring students approach new topics with improved strategies and identified knowledge gaps addressed. Over multiple iterations, students become less reliant on AI scaffolding—they internalize the Socratic questioning, become fluent in working with AI during tasks, and habituate reflective practice. Unlike the periodic exam system, the portfolio of artifacts and reflections becomes the living record of learning.

Conceptual vs. procedural dual-loop design. A cornerstone of the AI+X pedagogy is recognizing and leveraging the distinction between conceptual knowledge and procedural skill. Cognitive science tells us that understanding “why” (declarative knowledge) and knowing “how” (procedural knowledge) are related but distinct processes in the brain. In the AI+X paradigm, we intentionally separate these into two interlocking learning loops:

Conceptual Loop (Human-centered): Focuses on understanding, reasoning, and transferable knowledge. The human is in the driver’s seat, using mental effort to build schemas, question assumptions, and make sense of ideas. AI supports this loop by acting as a Socratic questioner, generating simulations, spotting logical inconsistencies in the student’s reasoning, offering analogies to known ideas, and generating thought experiments. These ensure the student does not just get the right answers, but builds a robust schema.

Procedural Loop (AI-augmented): Focuses on executing tasks, following steps, and practicing routines until they’re efficient. AI takes a large role here as executor or guide. For tasks where the human must still learn the procedure (like writing a geometric proof), the AI can guide the student step-by-step: “Try considering triangle ABC. What do you know about angle A? ... Good, now given that, what can you say about angle B?” The AI acts as a Personal Trainer: keeping track of which sub-skills the student has mastered and which need more practice, adjusting difficulty dynamically, providing immediate error correction, and noting when a student is solving something in a needlessly complicated way.

The availability of AI changes the balance between these loops dramatically. Since AI can now handle many procedural tasks (from calculations to drawing graphs to writ-

ing code given instructions), the role of the human learner can shift more towards the conceptual loop without losing the ability to get things done. However, it is crucial that students still understand procedural thinking, even if they aren't doing all steps manually, so they can effectively supervise and guide AI. The dual-loop design addresses the common challenge that students may be conceptually strong but make procedural errors, or may be procedurally competent but conceptually lost. A benefit of AI here is mass customization: if a student is conceptually strong but making careless errors, the AI funnels them into targeted drills; if a student is doing procedures blindly, the AI focuses them back on concept questions.

An important philosophical point underlies the dual-loop design: we treat knowledge as the context for skill, and skill as the demonstration of knowledge. They feed each other. Without concept, skill is rote and brittle; without skill, concept is inert and untested. AI enables us to maintain the interplay while reducing cognitive overload. The real power comes when students learn to consciously switch between loops themselves: thinking conceptually first ("What principles might apply here?"), then shifting to procedural mode ("Now I need to crunch the numbers"), then returning to concept ("Does the answer make sense in context?").

From exams to evidence: transforming assessment. A natural question arises: how do we assess learning in this model, especially when AI is deeply involved? Traditional exams are individual, timed performances that isolate the student from resources, but isolating from AI would be counterproductive in AI+X. Moreover, a one-off exam cannot easily measure metacognitive growth or collaboration skill. Therefore, AI+X shifts to a dual-dimensional assessment system valuing both outcomes (what can the student produce or solve?) and process (how did the student approach it?).

We replace episodic high-stakes exams with an AI+X Digital Portfolio, which is a living repository of a student's work, reflections, and achievements. The portfolio includes variety: concept maps drawn, code written (with comments on which parts were AI-generated), lab reports, essays with evidence of drafts and revisions, engineering designs with version history and commentary. Crucially, it also includes process evidence: transcripts of Socratic dialogues, prompts and AI responses used (showing skill in AI collaboration), and self-reflection logs. To ensure fairness and academic integrity, we include process verification: for example, a student might present a polished report (which AI helped write), but can also submit a video in which they verbally explain the key points or answer unscripted questions about it.

In AI+X, failure is part of the evidence of learning. Unlike traditional grading where mistakes are directly penalized, an interesting failure followed by a reflection and adjustment is viewed positively as a demonstration of metacognition and resilience. The portfolio becomes a more comprehensive evidence of capability for external audiences—universities or employers see actual projects completed, the evolution of skills, and how the student handled challenges—and a better predictor of future success than grades or test scores. It also incentivizes students to focus on genuine learning and growth, as nearly everything they do contributes to their portfolio.

The role of AI in assessment is also significant. AI systems can assist in evaluating certain aspects of the portfolio: by checking the coherence of a student's reflection with their work (does the student accurately understand their mistakes?), by tracking growth metrics ("This student's code correctness on first try improved from 50% to 90% over 6 months"), or by generating personalized feedback continuously ("next challenge for you: try a project involving uncertainty, since all your projects so far had deterministic outcomes, to stretch your skills").

The mapping of learning dimensions to portfolio evidence is straightforward. Conceptual understanding (L1) is demonstrated through mind maps and conceptual diagrams, excerpts of Socratic dialogue showing explanations in the student's own words, and challenging conceptual questions answered. The human articulated reasoning and crafted connections between concepts; the AI generated visualizations and asked prompting questions. Procedural fluency (as developed through L2 contexts) is demonstrated through completed tasks—working code, engineered devices, solved problem sets—along with a log of iterative attempts. The AI executed routine sub-tasks and provided hints; the human directed the approach, verified and refined AI outputs, and provided creative inputs. Metacognitive growth (L3) is demonstrated through reflection journals, self-assessment essays, after-action review summaries, and portfolio annotations (“I struggled here and here’s why”). The AI summarized performance data and prompted self-analysis questions; the human reflected on experience, identified lessons learned, and adjusted habits. Ultimately, the success of this assessment approach is measured by the fact that the portfolio becomes a better predictor of future success than grades or test scores. Early pilots have indicated that when employers or evaluators review such portfolios, they can more reliably select candidates who are prepared for real work.

Macroeconomic and societal impact

The AI+X educational paradigm is not just an academic exercise; it is conceived as a strategic response to pressing macroeconomic and societal challenges. By fundamentally realigning the relationship between education and productivity, AI+X has the potential to accelerate economic growth, alleviate demographic pressures, and promote social equity.

Accelerating the education-to-work pipeline. One of the most immediate macroeconomic benefits of AI+X is the shortening of the lag between education and productive employment. In many advanced economies, individuals spend their entire early adulthood attaining educational credentials before contributing substantially to the economy. By enabling students to be productive in real-world projects during their teens, AI+X effectively yields a “time dividend” of perhaps 5–7 years. Under an AI+X model, it is plausible that by age 18 or 19, many young adults will already have portfolios demonstrating competence equivalent to early-career professionals in fields like software development, design, research assistance, or technical writing.

Moreover, the productivity amplification from human–AI symbiosis means that even when young entrants join the workforce, they are not starting at the low end of productivity. With AI tools at their disposal and the orchestrative skills honed through L2 learning, a team of 18-year-olds might accomplish what would traditionally require older experts or larger teams. AI+X flips the chronic skills gap by merging training with education: industries receive an influx of talent that is not only knowledgeable but has practical project experience and is adept at using the latest AI tools for high productivity. There are also implications for innovation: younger minds in the workforce bring fresh perspectives, and equipped with AI, they might contribute to R&D and creative industries earlier, potentially accelerating the rate of innovation.

AI+X as a demographic stabilizer. Revisiting the education–demography trap: AI+X can act as a demographic stabilizer in advanced economies by breaking the negative feedback loop where long education depresses birth rates and strains aging societies. Several mechanisms are at work. First, reduced educational cost burden on families:

if AI+X produces pathways that are shorter and more efficient, families do not need to invest as much money for as long a period. AI tutors substitute for many private tutoring hours—a massive industry in places like Korea or Japan—and by reducing the “education arms race” of credential inflation, AI+X could relieve pressure on families and on the fertility decision. Second, earlier economic independence: achieving financial self-sufficiency at 18–20 instead of the late 20s means young adults can consider forming families earlier. The opportunity cost of taking time off for children is lower when one has already gained a career foothold by that age. Third, alignment of work-life balance with biology: AI+X compresses time-to-competence, allowing the phases of life to overlap more naturally—one could be a productive worker and a young parent simultaneously in one’s 20s, rather than facing the conflict between continuing education/career and starting a family. To quantify with a scenario: if a country’s average age of first childbirth moved from 32 to 28 on average, and average children per family rose modestly from 1.1 to 1.5, that difference in population trajectory over decades is dramatic.

Equity and global competitiveness. AI+X inherently carries a promise of democratizing quality education. Since AI tutors and resources can be replicated at near-zero marginal cost and personalized for each learner, the model could dramatically increase access to top-tier educational support in underserved communities—rural areas, poorer regions, or developing countries. If implemented globally, AI+X might allow countries to leapfrog traditional industrial-age education, similar to how some regions skipped landlines and went straight to mobile phones. Within a country, every student with an AI tutor means pacing and depth can adapt to them, potentially mitigating the disadvantages of large class sizes or variable teacher quality that hit lower-income school districts hardest.

AI+X is also an answer to the risk of technological unemployment. Rather than training people to compete with AI at tasks AI can automate, it trains people from a young age to be the ones that AI augments. If done broadly, you create a generation not competing with AI but using AI to create value in ways AI alone cannot—securing their employability and positioning the economy to create new roles and industries around human–AI collaboration. The nation that first fully implements AI+X at scale could see a surge of well-prepared, AI-savvy, innovative youth entering their industries, giving them an edge in the global knowledge economy. Perhaps most importantly, AI+X could help foster what some call a “learning society”: instead of viewing education as a one-shot early-life event, it becomes a lifelong continuum seamlessly blending with work and personal development.

Enabling infrastructure, standardization, and SIAS case study

Implementing the AI+X paradigm on a broad scale requires more than curriculum changes and pedagogy; it demands a robust technological infrastructure and a framework of standards to ensure interoperability, safety, and quality.

Cognitive architecture and knowledge representation. At the heart of the AI+X infrastructure is a cognitive architecture for the AI agents that facilitate learning. The architecture we adopt is inspired by neuro-symbolic AI: combining the pattern recognition prowess of neural networks with the logical reasoning and transparency of symbolic AI. One instantiation of this is the Never-Ending Open Learning Agent Framework (NEOLAF), designed to allow AI agents to continuously learn and improve in their educational strategies over time. A NEOLAF-based tutor could notice that a particular anal-

ogy worked well for many students struggling with a concept and thus incorporate that analogy into its repertoire for future students.

An essential part of the architecture is the KSTAR model (Knowledge, Situation, Task, Action, Result), which provides a structured format for representing the knowledge needed, the context, the task at hand, the action taken, and the outcome. By logging learning encounters in this format, the AI builds up a memory that it can use to adapt instruction. For example, an entry might record: Knowledge: Pythagorean theorem; Situation: student solving a ramp problem in physics; Task: find resultant force; Action: student applied theorem to velocity components; Result: incorrect because friction was ignored. The AI uses such memories to adapt—perhaps next time it sees a similar situation, it will remind the student about friction. Because of the symbolic component, these KSTAR logs are human-auditable: teachers or developers can inspect what the AI “remembers” and how it is making decisions, which is crucial for trust and debugging. Scalability is a design goal: NEOLAF agents can be deployed across millions of students with federated learning approaches that gather improvements without centralizing sensitive data.

Interoperability and governance through international standards. To prevent the educational technology landscape from fragmenting, establishing standards is critical. We align with and contribute to emerging standards efforts: IEEE P3428 (Standard for Large Language Model Agents in Education), which defines how AI educational agents should behave, what components they should have, and how they interface with learning platforms; IEEE P3394 (Standard for LLM Agent Interfaces), which ensures that interfaces between AI agents and external systems are standardized, so a student’s learning companion can follow them across institutions; and IEEE P2807.6 (Standard for Educational Knowledge Graphs), which outlines how to design and use ontologies and knowledge graphs for education, ensuring a concept like “Newton’s Second Law” has a standard identifier with its prerequisites and applications so any compliant AI tutor “knows” what it is and how it connects to other concepts.

Beyond technical interoperability, ethical standards ensure that AI systems explain their uncertainties, protect student data rigorously, and are auditable for bias. The effect of standards is trust and adoption: governments and institutions are more likely to adopt AI+X approaches when they see a clear standards-based pathway rather than a fragmented landscape of unproven applications. Standards also allow an ecosystem of contributors—a small startup could create a brilliant AI simulation tool that slots into the AI+X framework because it speaks the same “language” as the AI agents. It is also worth noting that infrastructure includes teacher training and organizational structure. The success of AI+X depends on teachers being trained to work with AI tools and orchestrate them, and on schools having the IT infrastructure—reliable internet, devices for students, AI software—to implement the model at scale.

Practical implementation: SIAS University case study. To ground the framework in a live, real-world deployment, we turn to SIAS University (Henan, China), where the AI+X paradigm has been under continuous institutional implementation since the beginning of 2025. This is not a thought-experiment or a composite sketch: it is an ongoing program led jointly by Dr. Shawn Chen, Chair of the Board of SIAS, and the NEOLAF Inc. Team — co-authors of this chapter — who together have carried the roll-out from initial policy design through to student-facing operation.

The transformation has proceeded in three deliberately sequenced phases, each designed to de-risk the next. Phase 1 (Institutional Experiment, early 2025) focused on the university as a system rather than on any single classroom. Governance, academic-integrity policy, data infrastructure, and assessment rules were rewritten to admit AI as a

legitimate co-learner and co-producer rather than an outside tool to be banned. A small number of cross-disciplinary pilot cohorts ran end-to-end through the three-stage agentic learning cycle, giving the administration concrete evidence of what the paradigm looks like when policy and technology are aligned.

Phase 2 (Faculty Training at Scale, mid-2025 onward) launched once the institutional scaffolding was in place: a large-scale faculty development effort to move the teaching corps from lecture-centric delivery toward the mentor-and-orchestrator role that AI+X requires. Training combined hands-on practice with the agent stack, co-design of project-based Interactive Learning Tasks, and discipline-specific workshops. The goal was not to retrain teachers to operate software but to re-skill them as designers of human–AI learning experiences.

Phase 3 (The Elite 20 Program, late 2025 onward) opened with the full AI+X curriculum end-to-end. Students in Elite 20 engage simultaneously in L1 disciplinary world-modeling, L2 human–AI co-modeling, and L3 metacognitive self-modeling, with progress evaluated through continuously-updated portfolios. Elite 20 serves both as a showcase and as an in-house laboratory for calibrating the paradigm before wider diffusion across the university.

The student benefits observable after little more than a year of operation have been substantial. Elite 20 learners routinely produce professional-grade artifacts—code, analyses, research briefs, design prototypes—at a level of sophistication well beyond what their age and traditional course history would predict. External industry panels reviewing the resulting portfolios have consistently rated them ahead of conventional same-year peers, and in several cases ahead of graduates two cohorts senior. Students report markedly higher engagement and a clearer sense of authorship over their learning, while faculty report a shift of their own time away from rote grading and exposition toward higher-value mentoring and research collaboration with the students.

Crucially, implementation at SIAS underscores a point that a purely technological rollout would have missed: for AI+X to work, the institution itself has to change. SIAS revised its academic-integrity code to permit—and in fact require—transparent, reflective use of AI in student work; its admissions and progression criteria now accommodate portfolio evidence alongside traditional transcripts; and faculty incentive structures have been adjusted to reward mentorship and course design rather than lecture volume. Building on this, SIAS and NEOLAF are now establishing a Symbiotic Intelligence Education Center to package the experience into transferable materials, train educators at partner institutions, and seed similar transformations elsewhere. The broader lesson is that institutional transformation, not tooling, is the rate-limiting step. Phase 1 created the conditions under which agents could legitimately teach; Phase 2 built the human capacity to work with them; only then did Phase 3 make the full student benefit visible. Institutions that attempt to leapfrog to a student-facing product without first doing this institutional work tend to plateau at surface-level adoption. SIAS’s sequenced approach offers an empirically grounded template for avoiding that trap.

Conclusion: accelerating the transition to a symbiotic civilization

The transformation envisioned by the AI+X paradigm is far-reaching. It is not a mere enhancement or digitization of the current educational system, but a fundamental revaluation of the very purpose and process of education in the 21st century. In conclusion, we synthesize the core outcomes of this transformation and outline a call to action

for stakeholders to collaboratively accelerate progress toward a symbiotic civilization—one in which collective human and AI intelligence work in concert to advance society.

From knowledge transmission to intelligence orchestration. Traditional education has been centered on the transfer of knowledge from teacher to student, with success measured by how much of that knowledge the student can retain and recite in predictable settings. The AI+X framework signals the definitive end of that era. In a world where AI can store and retrieve an effectively infinite library of information and even perform routine problem-solving, the premium shifts to what humans can do that machines cannot do alone.

AI+X reframes expertise as the capacity for cognitive orchestration: the ability to direct, align, and integrate multiple forms of intelligence—human intuition, algorithmic analysis, domain knowledge, ethical reasoning—to address novel problems. It is the skill of a symphony conductor rather than a soloist. Education becomes about nurturing versatile thinkers, adaptive collaborators, and lifelong learners. The three dimensions—world modeling (L1), AI collaboration (L2), and self-modeling (L3)—ensure that graduates are deeply knowledgeable yet not rigid (because they learn knowledge as models to be refined, not edicts to be memorized), technologically empowered but not over-reliant (because they practice the dance of delegation and supervision with AI), and self-aware and resilient (because they continually reflect and adjust). Such individuals essentially future-proof themselves: prepared not just for one career, but for the many careers or roles they may need to take on as technology and society evolve.

Symbiotic intelligence as a societal engine. If we populate our institutions—companies, research labs, governments—with such symbiotic individuals, we edge closer to a *symbiotic civilization* (gong zhi wen ming): a society where the norm is humans and AIs working together at every level to enhance productivity, creativity, and well-being. In such a civilization, the fears of AI causing mass unemployment or loss of human value are allayed, because human roles have shifted to higher-level orchestration and uniquely human contributions that AI by itself cannot replicate authentically. Work is redefined: many mundane tasks are handled by machines, while humans focus on creative, strategic, and interpersonal aspects. Innovation accelerates, as more people can innovate with AI assistance and more time is freed for exploration through efficiency gains. Lifelong learning becomes reality, as the line between education and work blurs and AI mentors make continuous up-skilling accessible anytime. Social inclusion can improve if we manage the transition well, with AI democratizing access to quality education.

Call to action. Realizing the AI+X vision requires bold action from multiple stakeholders. Policymakers and governments must re-evaluate outdated regulations that govern education and work. Could a 17-year-old with an AI+X portfolio be allowed to bypass certain diploma requirements to start an apprenticeship or job? Policy can facilitate earlier transitions and lifelong learning accounts that people can use at any time to learn new skills via AI platforms. Governments should invest in the necessary digital infrastructure so that every school can implement AI+X, and must update assessment standards at the national level to move beyond standardized tests to more holistic evaluations. Importantly, governments will need to collaborate on international standards through bodies like UNESCO or OECD to ensure a globally competitive and compatible system.

Educational institutions and educators should pilot AI+X methods now, iterating to find best practices. Teacher training programs need revamps to prepare new teachers as AI orchestra leaders. Teachers should embrace the chance to offload routine work to AI and focus on mentorship and high-level guidance—the chronic education–work

gap diminishes when teachers can concentrate on what only humans do best. Universities may shift from time-based curricula to competency-based frameworks, validating learning through portfolios and performance rather than credit hours. We can envision consortia of schools and employers collaborating on project ecosystems where students solve real problems for real stakeholders as part of their education.

Industry and employers should start valuing the outputs AI+X generates: instead of requiring a bachelor's degree, articulate what a well-crafted portfolio should demonstrate. Companies could partner with AI+X programs to provide real project opportunities and scout talent early. Technology developers must focus on education as a domain with pedagogical soundness, adhere to ethical standards, be transparent about limitations, and contribute open-source frameworks for educational AI so resource-constrained schools can implement them. Students and families must cultivate intrinsic motivation to use AI tools constructively, not as a crutch for cheating but as a genuine amplifier of their own capabilities. All these stakeholders must come together to dismantle outdated structures—the factory-model classroom, the obsession with standardized tests, rigid job credentialism—and embrace new roles and systems.

The cost of inaction could be severe: a widening alignment gap where education falls further behind technological reality, leading to greater unemployment, social unrest, and wasted human potential. Conversely, by proactively realigning now, we set up our societies not only to weather the AI revolution but to harness it for a new era of prosperity and human flourishing.

We stand at a crossroads reminiscent of past great transformations. The AI+X paradigm offers a clear blueprint for navigating it—ambitious yet attainable, disruptive yet deeply constructive. It promises to elevate human potential rather than diminish it. If we succeed in this transition, future historians may mark it as the moment we truly entered a symbiotic age: when we stopped fearing intelligent machines as rivals and started embracing them as extensions of our own collective mind. In that age, learning never stops, creativity knows fewer bounds, and the partnership of minds and machines tackles challenges that once seemed insurmountable. Building that future begins now, in our classrooms and curricula. It is time to orchestrate the symphony of human and artificial intelligence for the betterment of all.

References

- Hu, X., Xu, S., Tong, R., & Graesser, A. (2024). *Generative AI in education: From foundational insights to the Socratic playground for learning*. <https://doi.org/10.48550/arXiv.2412.00315>
- IEEE P2807.6 Working Group. (2024). *EduKG — Educational knowledge graph: Architectural framework and application*. IEEE.
- IEEE Standards Association. (2024a). *IEEE 2807.6: Guide for K–12 educational knowledge graphs*. IEEE.
- IEEE Standards Association. (2024b). *IEEE P3428: Standard for large language model agents for AI-powered education*. IEEE.
- IEEE Standards Association. (2025a). *IEEE P3394: Standard for large language model agent interfaces*. IEEE.
- IEEE Standards Association. (2025b). *IEEE P3428: Standard for large language model agents for AI-powered education* (Working draft). IEEE.

- Kinder, M., Briggs, X. de S., Muro, M., & Liu, S. (2024). *Generative AI, the American worker, and the future of work*. Brookings Institution.
- Miao, F., & Cukurova, M. (2024). *AI competency framework for teachers*. UNESCO.
- Miao, F., & Shiohira, K. (2024). *AI competency framework for students*. UNESCO.
- Naglee, K. (2023). *Symbiotic intelligence: A roadmap for AI adoption in K–12 education*. Self-published.
- OECD. (2025). *Korea's unborn future: Understanding low-fertility trends*. OECD Publishing.
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4
- Tinholt, D. (2025, January). *Symbiotic intelligence theory (SIT): Rethinking human potential in the age of AI*. Medium. <https://medium.com/>
- Tong, R. J., Li, H., Liang, J., Xu, B., et al. (2025). An architectural framework for educational knowledge graphs (IEEE P2807.6): Ontology design, LLM integration, and adaptive learning applications. In *Proceedings of the IEEE Conference on Artificial Intelligence (CAI 2025)*. IEEE.
- Tong, R., Li, H., Liang, J., & Wen, Q. (2024). *Developing and deploying industry standards for artificial intelligence in education (AIED): Challenges, strategies, and future directions*. <https://doi.org/10.48550/arXiv.2403.14689>
- Tong, R. J., Cao, C. C., Lee, T. X., Zhao, G., Wan, R., Wang, F., Hu, X., Schmucker, R., Pan, J., Quevedo, J., & Lu, Y. (2023). *NEOLAF: An LLM-powered neural-symbolic cognitive architecture*. <https://doi.org/10.48550/arXiv.2307.12057>
- Tong, R. J., & Hu, X. (2024). Future of education with neuro-symbolic AI agents in self-improving adaptive instructional systems. *Frontiers of Digital Education*, 1(2), 198. <https://doi.org/10.23977/fode.2024.010201>
- Wang, Y., Pan, Y., Su, Z., Deng, Y., Zhao, Q., Du, L., Luan, T. H., Kang, J., & Niyato, D. (2024). Large model based agents: State-of-the-art, cooperation paradigms, security and privacy, and future trends. <https://doi.org/10.48550/arXiv.2409.14457>

Chapter 5. Epistemic legitimation and hybrid cognitive units

Timur Schukin

Artificial intelligence first becomes a university problem when academic work no longer reliably demonstrates the capabilities universities are supposed to certify. Students can now produce sophisticated essays, code, analyses, and research outputs with AI assistance even when they do not fully possess the underlying practices of reasoning, interpretation, judgment, or technical competence that those outputs were previously assumed to reflect. Assessment becomes unstable because the visible result no longer clearly reveals how the work was produced or what capabilities were actually formed. The same shift affects faculty work and research. Explanation, drafting, coding, literature review, and even parts of scientific inquiry increasingly move into shared human–AI workflows in which intellectual work is distributed across people, models, tools, memory systems, and institutional procedures. As a result, the historical relationship between academic output, formed capability, assessment, and degree certification begins to weaken.

Most universities initially respond to this disruption at the administrative level. They introduce rules for AI use, redesign assignments, develop detection policies, and train faculty to work with new tools. These measures may reduce immediate disorder, but they do not address the deeper institutional problem. The university historically derived its legitimacy from its ability to connect intellectual work to publicly recognizable procedures of judgment, competence, and transmission. AI systems increasingly reorganize those procedures from outside the university itself. Platform and vendor infrastructures now shape how questions are formulated, how evidence is evaluated, how authorship is assigned, and what kinds of outputs appear credible or competent. If universities respond only through compliance measures or workflow adaptation, they risk allowing external technological infrastructures to define the conditions under which intellectual work becomes legitimate (Bommasani et al. 2021; Bender et al. 2021; Suchman 2007).

The deeper problem is therefore not simply technological adaptation. It is institutional legitimacy. If AI systems can increasingly generate convincing intellectual outputs outside traditional academic procedures, then the university can no longer rely on inherited assumptions about what makes its judgments credible. The question becomes more fundamental: *what function does the university perform that society still needs?*

Foundation of university: epistemic legitimation

The university becomes institutionally necessary when a fragile, local, contested, or emerging practice has to acquire public form. A practice becomes university knowledge when it can be taught to new participants, exposed to criticism, tested through recognized procedures, and transmitted beyond the first people who performed it. This function can be called epistemic legitimation.

In the age of AI, this function is no longer secured by tradition, accreditation, or institutional prestige. The university retains legitimacy only if it can legitimate the practice now emerging at scale: distributed intellectual production, where intellectual work is carried across humans, models, memory systems, evaluators, and institutions (Haskins 1923; de Ridder-Symoens 1992; Kerr 2001).

The university does not create the importance of a practice simply by naming it. Practices usually emerge outside the university or at its edge: scholastic interpretation, experimental inquiry, clinical judgment, engineering design, and entrepreneurial project work all entered universities from wider fields of action. What the university changes is the mode of existence of the practice. A local technique, solution, or insight becomes durable only when it can be separated from the authority of its original carrier, reconstructed by others, challenged through public procedures, and transferred into new situations without losing its object (Rashdall 1936; Clark 1998; Etzkowitz and Leydesdorff 1997).

The university performs this transformation by placing practices under organized procedures of questioning, comparison, revision, and transmission. A technique held by a master becomes a teachable method. A successful solution becomes a case that can be compared with alternatives and failures. A tradition becomes capable of absorbing criticism without collapsing into repetition or arbitrariness. These procedures do not simply preserve knowledge; they change how knowledge can circulate, accumulate, and develop.

Epistemic legitimation is not the same as certification. Certification attaches a mark to a person, program, or product. Legitimation builds the conditions under which a practice can be inspected, taught, challenged, revised, and inherited. This is why universities historically rely on forms such as disputation, dissertation, laboratory report, studio critique, and clinical supervision. Across different disciplines, these practices perform the same institutional function: they make intellectual work visible enough for others to enter it, test it, and carry it further.

University procedures often appear slow because they force practices to acquire the conditions of transmission. They stabilize language, expose errors, preserve traces of reasoning, and situate new work in relation to prior knowledge. A practice may already be useful, profitable, or impressive before passing through this process. It becomes university knowledge only when it can survive organized criticism and be carried beyond its original context.

The university's historical role therefore extends beyond preserving knowledge inside institutional walls. Once a practice acquires shared standards, recognized carriers, procedures of validation, and routes of training, it becomes transferable across professions, firms, public agencies, and civic institutions. The university transforms local or fragile practices into durable institutional resources: disciplined forms of action that others can use, test, modify, and govern.

The historical genealogy of this function begins before the modern research university. The medieval university transformed doctrinal interpretation into a teachable and disputable intellectual practice. The Humboldtian university later did the same for experimental inquiry. The entrepreneurial university extended this logic to technological innovation and applied research. The AI-era university now faces the same challenge in a new form: distributed intellectual production must itself become a practice that can be taught, tested, criticized, revised, and inherited (Aquinas 1947; Chenu 2002; Torrell 1996).

Generations of the university: one function, different practices

The university's function has remained surprisingly stable across very different historical forms. What changes is the kind of practice that requires legitimation, the procedures through which it is tested, and the institutions that must trust the result. The work of Thomas Aquinas gives the first historical form of the university function: external intellectual practices enter scholastic procedures and become teachable, disputable, and inheritable knowledge. That function did not remain confined to medieval theology. Later university formations changed the practices that required legitimation, the procedures through which those practices were tested, and the institutions that had to trust the result. The university function persisted across these changes: it gave a practice public form by naming its object, exposing it to dispute, forming carriers, and securing procedures through which the practice could be taught, challenged, revised, and inherited (Aquinas 1947; Leff 1968; Torrell 1996).

The medieval university made interpretation and doctrinal dispute institutionally usable. Commentary, quaestio, disputation, and degree procedures formed carriers capable of working with authority in church, law, administration, and education. The university did not simply preserve inherited texts; it placed them inside public procedures of argument and judgment. Rival authorities could be compared, doctrinal limits could generate distinctions rather than terminate inquiry, and intellectual conflict could pass through recognized forms of dispute. The result was not only theological knowledge, but a disciplined practice of judgment that could travel beyond the university into wider institutional life (Haskins 1923; Rashdall 1936; de Ridder-Symoens 1992).

The Humboldtian research university shifted the central practice from interpretation to experimental inquiry. Laboratory work, journal publication, dissertation, seminar, and disciplinary career made inquiry into nature reproducible, teachable, and publicly credible. The laboratory turned an observed event into a repeatable result; the dissertation trained carriers capable of reconstructing the path from question to evidence. Through these procedures, the university transformed experiment into an inheritable institutional practice rather than a collection of isolated discoveries (Humboldt 1970; Rüegg 2004).

Burton Clark's entrepreneurial university model extended this logic to technological innovation and economic translation. Prototype development, partnership, incubators, patent procedures, and start-up pathways gave technological risk a recognizable institutional form. Innovation became not only commercially valuable, but also teachable, governable, and publicly legitimate within university procedures of judgment and recognition (Clark 1998; Etzkowitz and Leydesdorff 1997).

The university function remains stable across these formations while the practice requiring legitimation changes. The medieval university gave interpretation a public form after authority could no longer operate as an undivided ground. The research university gave experiment a public form after nature had to be approached through reproducible inquiry. The entrepreneurial university gave innovation a public form after research became bound to technological and economic risk. Each formation emerges when an older certainty weakens and a new practice can survive only if it acquires procedures of dispute, testing, formation, and inheritance.

The unstable center is no longer the individual carrier, but the academic or professional role-position as a formed bundle of judgment, method, authorship, competence, and responsibility. AI systems fracture that bundle. They redistribute how problems are framed, how alternatives are generated, how results are validated, and how criteria are revised across human participants, language models, external memory, and institution-

al norms. The emerging practice is distributed intellectual production: intellectual work carried by shifting configurations whose responsibility can no longer be read from the visible author of the final output. The university function returns in a harder form. This practice must be named, tested, taught, bounded, and inherited before platform infrastructures define its norms by default.

Collective intellectual work: the Moscow methodological circle and Engelbart

Thought in the university was always collective: disputation, commentary, seminar, laboratory, and school already worked between people, texts, instruments, and inherited procedures. The twentieth-century problem is different. Collective intellectual work becomes an explicit organizational problem when multiple participants have to work on the same object across different disciplines, vocabularies, institutional interests, and standards of proof. Without shared procedures, collective intellectual work either fragments into disconnected expert monologues or collapses into administrative compromise. It therefore requires forms that stabilize the object outside private opinion and preserve disagreement long enough for it to become method rather than conflict. Only then can a result be transmitted as an inheritable practice rather than as the memory of a meeting.

Douglas Engelbart, known for the computer mouse, windowed interaction, hypertext, and the idea of networked computing as an environment for collective work, approached intellectual production through the augmentation of human intellect. These were not isolated interface inventions. They belonged to a project of reorganizing the working system through which people address complex problems. The group's memory could be externalized. Navigation through practice could become part of the work itself. Revision could be shared. Tools could even be used to redesign the procedures through which intellectual work itself was organized. Engelbart's strongest claim lies there: intellectual work can become an object of its own redesign (Engelbart 1962; Engelbart and English 1968).

The Moscow Methodological Circle addressed the same problem from a different institutional and political position. It emerged in the Soviet Union in the 1950s around university seminars and gradually became a school of thinking and practice. Its object was not the computer interface but the organization of collective thought: the group, the scheme, the position, the seminar, and later the organizational-activity game. The Circle worked where no single participant, discipline, or institution could own the object in advance. It developed a quasi-university function outside the classical university by forming participants capable of collective work, developing a shared methodological language, and transmitting procedures for entering complex practices. Distributed intellectual production requires this line because tools alone cannot form positions, public memory, or teachable procedures of participation (Shchedrovitsky 1995; Shchedrovitsky 2004).

Scholastic and Soviet constraints worked through different institutions, but both re-directed novelty into procedure. Christian doctrine did not make the work of Thomas Aquinas free; it forced Aristotelian and Arabic materials to enter Latin Christian thought through doctrinal limits and scholastic university procedures of distinction, question, and reply. Soviet institutional constraints did not make the Moscow Methodological Circle free; they closed many direct institutional and techno-engineering trajectories and forced methodological work to develop as an organization of collective intellectual ac-

tivity. In both cases, constraint blocked some paths and forced a new practice to become durable through articulated differences, public procedures, and disciplined forms of clarification.

Engelbart and the Moscow Methodological Circle prepare different sides of the same AI-era problem. Engelbart showed how digital infrastructure can augment intellectual work by externalizing memory, navigation, representation, collaboration, and the improvement of tools. The Moscow Methodological Circle showed how collective intellectual work can form participants as positions inside a shared practice through scheme, seminar, contradiction, and reflective reconstruction. One line strengthens the technical environment of work; the other forms the carrier of participation and the norms through which a group can think together.

AI turns that convergence into an institutional problem for the university. Intellectual work can now be produced through configurations of people, language models, external memory, evaluators, organizations, and norms. Engelbart's augmentation project and the Moscow Methodological Circle's organization of collective thought cease to be marginal experiments; they become conditions of ordinary intellectual production. The university has to decide whether these configurations will remain vendor-governed service environments, or whether they can become university practices with visible procedures, trained carriers, preserved errors, public criteria, and inheritable responsibility. Careers, degrees, statuses, schools, and forms of recognition also have to be rebuilt, because intellectual work can no longer be reliably attributed to a single stable carrier.

AI as distributed intellectual production

AI brings together two earlier trajectories: Engelbart's vision of technologically augmented intellectual work and the Moscow Methodological Circle's vision of organized collective thinking. Engelbart showed how intellectual work changes when memory, navigation, representation, and revision move into computational environments. The Moscow Methodological Circle showed how collective thinking requires procedures, positions, and shared forms through which groups can hold an object of work together. AI turns this convergence into an ordinary condition: intellectual production now passes through technical and organizational arrangements that exceed the individual mind. The university must therefore treat distributed intellectual production as the new practice requiring legitimation (Engelbart 1962; Engelbart and English 1968; Shchedrovitsky 2004).

The relevant transformation is not simply the presence of AI tools inside existing workflows. It is that framing, generation, validation, memory, and criterion revision can now move across different human and nonhuman carriers. A language model may propose a framing, a simulation tool may test a constraint, and an external memory may preserve failed alternatives. Intellectual work increasingly emerges from configurations of learners, models, tools, traces, and institutional norms rather than from a single isolated actor (Hutchins 1995; Suchman 2007; Nonaka and Takeuchi 1995). The service image of AI therefore misleads the university. A service assists an existing practice by saving time, lowering cost, or automating routine operations. An infrastructure reorganizes the practice itself. It shapes which procedures become easy, which traces remain visible, which criteria become embedded in workflows, and which forms of judgment can operate without a stable human carrier. When AI is treated as a service, universities ask how teachers and students should use it. When AI functions as infrastructure, the

real question becomes who controls the conditions under which intellectual work appears credible (Bommasani et al. 2021; Bender et al. 2021; Suchman 2007).

This shift is already visible in ordinary academic work. Students can generate sophisticated outputs through distributed workflows that combine model-generated framing, automated analysis, simulation, revision, and external memory. The final artifact may appear to be an individual submission, while the intellectual practice that produced it has been distributed across heterogeneous carriers. The visible author is no longer the whole carrier of the intellectual work.

Research intensifies the same condition. Literature review, hypothesis generation, coding, statistical checking, simulation, visualization, and drafting can now be partially automated and linked into shared workflows. The researcher may still set direction and assume responsibility for conclusions, but those roles can no longer be inferred automatically from authorship alone. If systems increasingly frame questions, select materials, generate analyses, and produce conclusions, the human signature may mark only initiation and approval rather than substantive intellectual organization of the work (Bommasani et al. 2021; Brown et al. 2020).

The same pressure enters professional education because the university's legitimating function historically depended on distinguishing output from formed capability. A future engineer, manager, doctor, analyst, or designer can now produce advanced outputs before acquiring the practices historically implied by those outputs. A report may appear coherent while the student cannot reconstruct the reasoning behind it. A plan may be persuasive while the team cannot explain why one criterion was chosen over another. The university can no longer treat the final artifact as transparent evidence of competence; it must examine how the distributed practice produced it (Zawacki-Richter et al. 2019; Kasneci et al. 2023; Holmes, Bialik, and Fadel 2019).

AI becomes a competitor to the university's historical function when external technological environments begin organizing intellectual practice faster than universities can legitimate it. Universities historically formed intellectual practice through commentary, disputation, seminar, laboratory, and degree structures that stabilized criteria, preserved errors, and trained carriers of judgment. AI infrastructures can now perform parts of this organizing function externally through interface design, automated evaluation, platform defaults, and workflow optimization. As a result, external systems increasingly shape what counts as an acceptable answer before universities reconstruct their own procedures of legitimation (Bommasani et al. 2021; Bender et al. 2021).

This transformation also changes the economic organization of intellectual labor. In the research university, the socially valuable product was reproducible knowledge. In the entrepreneurial university, it was technological innovation and investable risk. In the AI era, markets increasingly sell continuous access to the production of coherent intellectual outputs. The platform becomes valuable because it controls the layer where distributed intellectual functions are assembled, normalized, and priced (Clark 1998; Etzkowitz and Leydesdorff 1997; Mazzucato 2021).

The division of intellectual labor changes accordingly. Some participants formulate requests, others provide context, others validate plausibility, and others formally assume responsibility. A smaller number retain authority to redefine the problem, revise the criterion, expose hidden error, or redesign the practice itself. Universities therefore cannot speak about "AI use" in general terms. The central question becomes which role-position preserved the criterion, exposed the error, maintained the trace, and carried responsibility for the result.

Human centrality cannot simply be assumed under these conditions. A human remains central when he defines the object, recognizes the error, holds responsibility for

consequences, and can reconstruct the practice being performed. The same person becomes a weak operator when he merely triggers workflows and accepts plausible outputs generated by processes he cannot explain. Biological presence alone no longer guarantees meaningful participation in intellectual work. Universities that want to defend human development must specify which operations cannot be delegated and which forms of reconstruction, validation, transfer, and responsibility remain necessary for genuine participation (Suchman 2007; Bender et al. 2021).

For this reason, the university's task cannot be reduced to AI literacy. A person may understand prompting techniques and know the risks of hallucination while still outsourcing judgment and becoming dependent on external completion. The institutional question is not simply whether AI is used, but which parts of the practice are delegated, reconstructed, validated, and retained under human responsibility. Universities must therefore move from literacy toward regimes of participation (Luckin et al. 2016; Holmes, Bialik, and Fadel 2019; Kasneci et al. 2023).

Three regimes organize participation. Formulation determines how the problem is defined and what ontology makes the situation actionable. Validation subjects generated outputs to appropriate forms of resistance: experiment, data, simulation, source criticism, stakeholder response, legal constraint, or transfer to another case. Responsibility and transmission preserve the trace of how the practice moved across participants, tools, criteria, and institutional procedures. Without such a trace, responsibility collapses into formal signature; with it, universities can assess not only outputs but the organization of the practice that produced them.

The next analytical object therefore cannot be the isolated student, the model, the prompt, or the final artifact alone. Intellectual work is increasingly distributed across learners, model configurations, memory systems, validation tools, and institutional norms. The university needs a smaller and more precise unit through which it can observe how framing, validation, trace, and responsibility are organized together. The next section introduces this unit: the hybrid cognitive unit, where "cognitive" refers not to a theory of mind but to distributed intellectual practice.

Hybrid cognitive units and the zone of proximal degradation

Once intellectual production becomes distributed, the final output can no longer serve as reliable evidence of learning, competence, or responsibility. A report, design proposal, codebase, policy brief, or research hypothesis may appear coherent while concealing how the work was actually produced. The artifact no longer reveals how the problem was framed, how alternatives were tested, how decisions were preserved, or where responsibility was carried. In the older assessment regime, universities could often treat the submitted artifact as a compressed sign of the student's capability. In the AI era, that compression becomes unreliable. The artifact may be stronger than the carrier who submits it.

A university that wants to legitimate distributed intellectual production therefore needs a more precise unit of analysis than the course, the prompt, the model, or the finished assignment alone. The student remains a legal and educational carrier, but the practice no longer resides entirely within the student. Model configurations generate options, datasets constrain them, validators return errors, and institutional norms determine what counts as sufficient. The relevant unit is the distributed arrangement in which framing, validation, trace, and responsibility are actually organized together.

A simple educational case makes this visible. Imagine a capstone project in urban planning at a regional university. A student team redesigns pedestrian access around several schools in partnership with a municipality. The team works with maps, interviews, traffic observations, regulatory constraints, and generative AI systems. A polished final presentation is still insufficient to show whether the students have actually learned to think and act within the practice.

The central educational problem lies not in generating outputs but in governing the process through which they emerge. Models can generate plausible framings such as school safety, congestion reduction, walkability, or public-space equity. Each framing opens some possibilities while obscuring others. The same applies to generated routes, crossings, schedules, or visual layouts. The educational task is not simply selecting plausible outputs, but exposing trade-offs, testing constraints, preserving rejected alternatives, and revising criteria when earlier assumptions fail.

A strong project therefore preserves the reasoning process itself: which alternatives were rejected, why criteria changed, which constraints proved decisive, and how validation altered the direction of the work. This trace allows teachers to assess whether students learned to frame, test, and revise a problem rather than merely produce a polished artifact. It also allows other participants to inherit the practice instead of receiving only the final result.

A hybrid cognitive unit can now be defined. “Cognitive” here refers not to cognitive science, but to distributed intellectual work. A hybrid cognitive unit is a cross-stack unit of practice in which a participant, a model configuration, an operation trace, and a validation norm jointly carry a task that none of them can perform responsibly alone. The participant brings contextual judgment, stake, and accountability. The model configuration generates, reformulates, compares, or checks possible moves. The trace preserves decisions, refusals, and errors. The validation norm determines what counts as sufficient within the practice rather than within generic fluency (Hutchins 1995; Suchman 2007).

This definition avoids two false reductions. The first attributes the result entirely to the student and treats AI as invisible assistance. The second attributes the result to the machine and reduces the student to a superficial operator. Both misunderstand the practice itself. The relevant educational object is the organized relation through which problems are framed, alternatives generated, results validated, and responsibility assigned.

Student development therefore depends on learning to assemble, constrain, validate, and govern such units. Assembly means coordinating human judgment, model configurations, evidence, and institutional norms without confusing one for another. Constraint means defining what may be delegated and what may not. Validation means exposing outputs to data, consequences, competing criteria, or external resistance. Governance means preserving the trace, locating responsibility, and revising procedures after failure.

These operations are far more demanding than tool use. A student may become highly skilled at prompting while remaining weak at framing. A team may generate many alternatives while failing to constrain the space of action. A report may include detailed process logs while recording activity rather than judgment. Universities must therefore assess whether learners gained access to the practice itself rather than merely receiving outputs produced by the configuration.

The same hybrid unit that enables development can also generate dependency. Models may provide framings before students learn to formulate problems, generate alternatives before learners understand the field of action, or produce polished explanations

before participants acquire the language of the practice. Output quality improves while access to the underlying practice weakens.

This is the beginning of the *zone of proximal degradation*. The phrase deliberately echoes Vygotsky's zone of proximal development. In development, external support allows the learner to perform beyond the current level while still making the underlying procedure available for appropriation. In degradation, the external environment completes the work while the learner remains outside the practice that produced the result (Vygotsky 1978; Wertsch 1985).

The machine-mediated "other" is especially powerful because it can simultaneously provide formulation, evaluation, completion, and authority. In developmental configurations, this can help learners explore alternatives and revisit unfinished work. In degrading configurations, it closes the process too early. Smoothness begins to substitute for understanding before the learner has appropriated the practice itself (Vygotsky 1978; Vygotsky 1986; Cole 1996).

The university cannot solve this problem simply by prohibiting external support. Human development has always relied on external signs, teachers, instruments, diagrams, and institutions. Externalization becomes harmful only when learners lose access to the relation between problem, criterion, procedure, and result. AI raises this problem not only for calculation, but for writing, reasoning, design, research, and judgment.

Several safeguards therefore become essential. The first is the operation trace: learners must preserve how problems were framed, how models were used, which alternatives were rejected, and why validation decisions mattered. The second is reverse reconstruction: learners should be able to explain how the result could be produced with weaker systems or on smaller cases. The third is transfer: a capability has formed only when the same discipline of framing, validation, trace, and responsibility can move to a different context.

Failure archives complete this loop. Hybrid cognitive units fail in recurring ways: weak framings stabilize too early, validation checks only what is measurable, traces record prompts instead of decisions, and responsibility collapses into formal presentation. When these failures are preserved and classified, they become teaching material for the practice itself. Assessment must therefore judge not only the final output, but also the learner's ability to reconstruct, test, and transfer the practice that produced it (Lave and Wenger 1991; Wenger 1998).

The zone of proximal degradation is not unique to AI. It appears whenever institutions reward outputs while blocking appropriation of the underlying practice. AI intensifies this older problem because platform infrastructures, institutional incentives, and workflow automation scale it globally. The more fluent the system becomes, the harder the degradation is to detect: weak outputs reveal incapacity immediately, while strong outputs may conceal it. The university's authority in the AI era will depend on making that concealment visible and distinguishing accelerated production from formed capability.

Second-tier universities as second-order integrators

The zone of proximal degradation clarifies the institutional question after AI. The decisive capacity is not ownership of the most advanced models or infrastructures, but the ability to turn human-AI work into practices that can be taught, tested, corrected, and inherited. Frontier universities may remain closer to major platforms, elite talent, and expensive research systems, but proximity to technological discovery does not au-

tomatically solve the problem of legitimacy. A university can possess powerful systems and still fail to form students capable of reconstructing, validating, and transferring the practices those systems support (Bommasani et al. 2021; Pinheiro, Benneworth, and Jones 2012).

Second-tier universities begin from constraints that elite institutions often do not share. They usually lack control over foundation models, global platforms, and large research infrastructures. They operate with thinner budgets, uneven student preparation, and local labor markets that cannot wait for ideal institutional design. Yet these same constraints place them closer to a different density of practice: regional employers, municipal institutions, professional communities, and concrete social problems where AI systems encounter real consequences. A second-tier university tests human–AI practices where a municipal partner needs a usable decision, where a small firm cannot afford a wrong diagnosis, and where students must work under conditions of budget, regulation, conflict, and accountability rather than technological abstraction (Pinheiro, Benneworth, and Jones 2012; Goddard and Vallance 2013).

A frontier university can invent advanced systems. A second-tier university can observe how those systems reorganize ordinary work. Institutions encountering AI after its initial deployment face three possibilities: dependence, imitation, or strategic selectivity. Dependence imports external tools and criteria without reconstructing local practice. Imitation copies elite strategies without the ecosystem that made them viable. Strategic selectivity identifies which trajectories matured, which failed, and which can be translated into local educational and institutional conditions (Gerschenkron 1962; Perez 2002).

This is the logic of second-order integration. First-order integration introduces AI tools into courses, projects, or administrative systems. Second-order integration organizes the conditions under which human–AI practices become trustworthy, teachable, and inheritable. The central question is no longer whether AI is adopted, but which practices it transforms, which capabilities it develops or weakens, and where responsibility remains located.

A *second-order integrator* therefore translates external AI capacity into local regimes of validation. It determines when machine-generated outputs count as drafts, evidence, decision support, or unacceptable results. It trains students to distinguish fluent output from validated practice, preserves traces of reasoning and failure, and connects local consequences to standards of competence. The university becomes an integrator when it keeps students, faculty, employers, public institutions, and professional communities inside the same cycle of legitimation.

This also changes governance and assessment. Universities can no longer rely on generic policies that merely allow or prohibit AI use. They must specify which human–AI practices are being formed and how those practices are validated. One course may train AI-supported problem framing; another may train validation of model-generated analysis against data or stakeholder constraints; another may train reconstruction of AI-assisted research processes. The educational question becomes primary: what practice is being formed, how its trace is preserved, and what counts as evidence that the learner can carry the practice beyond the immediate case.

For second-tier institutions, investment should therefore follow validation and inheritance rather than adoption spectacle. Licenses, dashboards, and symbolic platform partnerships may have operational value, but they do not by themselves create legitimacy. More important are project seminars, local validation laboratories, failure archives, and regional case corpora that allow practices to become teachable and trustworthy (Mazzucato 2021).

The main institutional risk is platform absorption. Universities may gradually become downstream extensions of vendor-defined workflows, competence models, and evaluation systems. In that case, universities preserve credentials and administrative structures while losing authority over the formation of intellectual practice itself (Bommasani et al. 2021; Bender et al. 2021).

A second risk is competence extraction. Firms and platforms may absorb highly skilled faculty and students while weakening the university procedures that originally produced those competencies. Advanced capabilities do not emerge spontaneously; they depend on seminars, disputes, assessment traditions, schools, and long institutional processes of formation. If external systems extract the skill while abandoning the university machinery that formed it, they undermine the reproduction of future competence.

Epistemic justice also becomes central under AI conditions. Global AI infrastructures tend to normalize dominant languages, professional formats, and platform-visible data. Second-tier universities often remain closer to forms of knowledge that exist below that layer: tacit professional heuristics, multilingual practice, institutional memory, and locally specific constraints. Their task is not to preserve such knowledge as folklore, but to translate it into validation procedures that can be tested, revised, and transmitted (Fricker 2007; Santos 2014).

Distributed intellectual production also changes the university's relation to its region. Student projects become local research devices, municipal problems become educational environments, and professional workflows become sites for testing competence. Under AI conditions, this porosity becomes a structural advantage rather than a peripheral outreach function. Knowledge is increasingly produced through regional consequence rather than delivered to regions after the fact (Kerr 2001; Goddard and Vallance 2013).

The final institutional risk is delay. If universities wait until AI practices are stabilized externally by platforms and markets, they will inherit norms they did not create. Assessment, competence, authorship, and local knowledge will increasingly be reorganized through vendor-defined systems and global datasets. Universities may continue granting degrees while losing authority over the formation and validation of intellectual practice itself (Bommasani et al. 2021; Perez 2002).

The university's historical function therefore returns in a new form. The medieval university transformed doctrinal interpretation into teachable intellectual practice. The research university transformed experiment into reproducible inquiry. The entrepreneurial university transformed technological innovation into governable risk. The AI-era university must transform distributed intellectual production into practices with validation, conflict, responsibility, and inheritance. Institutions capable of this transition will not necessarily be those closest to frontier technology, but those most capable of translating external technological capacity into durable local practices.

Second-tier universities stand where this translation becomes unavoidable. Their possible strength lies in second-order integration: forming the conditions under which people, models, tools, traces, and institutions can produce knowledge that others may trust. In the AI era, they may become laboratories of legitimacy for distributed intelligence rather than weaker copies of frontier universities.

References

- Altshuller, G. (1984). *Creativity as an exact science: The theory of the solution of inventive problems* (A. Williams, Trans.). Gordon and Breach.

- Aquinas, T. (1947). *Summa theologiae* (Fathers of the English Dominican Province, Trans.). Benziger Brothers.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). Association for Computing Machinery.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., et al. (2021). *On the opportunities and risks of foundation models*. Center for Research on Foundation Models.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- Chenu, M.-D. (2002). *Aquinas and his role in theology* (P. Philibert, Trans.). Liturgical Press.
- Clark, B. R. (1998). *Creating entrepreneurial universities: Organizational pathways of transformation*. Pergamon.
- Cole, M. (1996). *Cultural psychology: A once and future discipline*. Harvard University Press.
- Engelbart, D. C. (1962). *Augmenting human intellect: A conceptual framework*. Stanford Research Institute.
- Engelbart, D. C., & English, W. K. (1968). A research center for augmenting human intellect. In *Proceedings of the December 9–11, 1968, Fall Joint Computer Conference* (pp. 395–410). Association for Computing Machinery.
- Etzkowitz, H., & Leydesdorff, L. (Eds.). (1997). *Universities and the global knowledge economy: A triple helix of university-industry-government relations*. Pinter.
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
- Gerschenkron, A. (1962). *Economic backwardness in historical perspective: A book of essays*. Belknap Press of Harvard University Press.
- Goddard, J., & Vallance, P. (2013). *The university and the city*. Routledge.
- Haskins, C. H. (1923). *The rise of universities*. Henry Holt and Company.
- Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial intelligence in education: Promises and implications for teaching and learning*. Center for Curriculum Redesign.
- Humboldt, W. von. (1970). On the spirit and the organisational framework of intellectual institutions in Berlin. *Minerva*, 8(2), 242–250.
- Hutchins, E. (1995). *Cognition in the wild*. MIT Press.
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, Article 102274.
- Kerr, C. (2001). *The uses of the university* (5th ed.). Harvard University Press.
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.
- Leff, G. (1968). *Paris and Oxford universities in the thirteenth and fourteenth centuries: An institutional and intellectual history*. John Wiley and Sons.

- Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2016). *Intelligence unleashed: An argument for AI in education*. Pearson.
- Mazzucato, M. (2021). *Mission economy: A moonshot guide to changing capitalism*. Allen Lane.
- Newman, J. H. (1996). *The idea of a university* (F. M. Turner, Ed.). Yale University Press.
- Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. Oxford University Press.
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.
- Perez, C. (2002). *Technological revolutions and financial capital: The dynamics of bubbles and golden ages*. Edward Elgar.
- Pinheiro, R., Benneworth, P., & Jones, G. A. (Eds.). (2012). *Universities and regional development: A critical assessment of tensions and contradictions*. Routledge.
- Rashdall, H. (1936). *The universities of Europe in the Middle Ages* (F. M. Powicke & A. B. Emden, Eds.; 3 vols.). Oxford University Press.
- de Ridder-Symoens, H. (Ed.). (1992). *A history of the university in Europe: Vol. 1. Universities in the Middle Ages*. Cambridge University Press.
- Rüegg, W. (Ed.). (2004). *A history of the university in Europe: Vol. 3. Universities in the nineteenth and early twentieth centuries (1800–1945)*. Cambridge University Press.
- Santos, B. de S. (2014). *Epistemologies of the South: Justice against epistemicide*. Paradigm Publishers.
- Shchedrovitsky, G. P. (1995). *Izbrannye trudy* [Selected works]. Shkola Kulturnoi Politiki.
- Shchedrovitsky, G. P. (2004). *Organizatsionno-deiatelnostnaia igra: Sbornik tekstov* [Organizational-activity game: Collected texts] (Vol. 1). Nasledie MMK.
- Suchman, L. (2007). *Human-machine reconfigurations: Plans and situated actions* (2nd ed.). Cambridge University Press.
- Torrell, J.-P. (1996). *Saint Thomas Aquinas: Vol. 1. The person and his work* (R. Royal, Trans.). Catholic University of America Press.
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman, Eds.). Harvard University Press.
- Vygotsky, L. S. (1986). *Thought and language* (Rev. ed., A. Kozulin, Ed.). MIT Press.
- Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press.
- Wertsch, J. V. (1985). *Vygotsky and the social formation of mind*. Harvard University Press.
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? *International Journal of Educational Technology in Higher Education*, 16, Article 39.

Chapter 6. The Architecture of Planetary Agency and Doughnut Learning

Raphael Costambeys-Kempczynski and François Taddei

Introduction: the eloquence of desertion

AgroParisTech, one of France's most selective Grandes Écoles, occupies a symbolic position within a higher education system historically structured by competition, credential prestige and elite professional pathways. The implicit promise is familiar: pass the entrance examinations, enter the right institution and you are set for life. In May 2022, however, the graduation ceremony was disrupted by an act of profound rhetorical and symbolic weight. A group of graduates took to the stage not to celebrate their entry into the excellence of the agro-industrial complex, but to call for a collective *désertion* from an academic program which they claim “generally encourages participation in the ongoing social and ecological devastation”, urging their peers to abandon the very career paths their rigorous education had prepared them for within a productivist agro-industrial complex (Passelac 2026). Their speech was more than a protest; it was a performance of dissent, a refusal to be the high-performing cogs in a machine they identified as fundamentally broken. They were not rejecting the knowledge they had acquired; they were rejecting the academic predestination that went with it. By early 2026, this movement had extended into forms of on-campus occupation and environmental mobilization.

This is not merely an example of a localized student protest; it can be read as a front-line manifestation of a deep-seated fracture in the social contract of the modern university. For much of the twentieth century, at least in many industrialized higher education systems, the implicit agreement between university and student appeared comparatively linear: the institution provided the disciplinary training, credentials of excellence and status (becoming the best *in* the world), and the labor market translated them into relative stability. However, as the 21st century unfolds, this contract has been disrupted by the twin forces of ecological collapse and the rise of Artificial Intelligence (AI). Towards the end of the 20th century, as the Cold War order began to collapse from the mid-1980s onwards, things had already begun to shift towards a multi-lateral VUCA world or, to be more explicit, a volatile, uncertain, complex and ambiguous reality. As we approached the quarter-way mark of the twenty-first century, some commentators argued that the language of VUCA no longer captures the affective and systemic turbulence of the present. Jamais Cascio's BANI heuristic — Brittle, Anxious, Non-linear and Incomprehensible — offers one vocabulary for naming this shift (Cascio 2020). In this landscape, the 19th-century organization of knowledge into relatively bounded disciplines appears increasingly misaligned with problems whose causes and consequences are systemic: climate change, democratic fragility, technological disruption, public health and biodiversity loss.

Not long ago, we might have argued that education was *failing the future*. We would have pointed to a competitive system focused on reproducing 19th-century silos for a

carbon-fuelled industrial age, while the world urgently needed collaborative co-construction of new knowledge. Today, the picture is more contradictory: access to education has expanded dramatically in many regions, yet the conditions that make meaningful education possible are under renewed threat. Access is being restricted through violence, political ideology, and economic fragility. From the literal exclusion of women and girls from classrooms in Afghanistan to the more insidious erosion of academic freedom from Russia to the United States of America (cf. the Scholars at Risk *Free to Think Report 2025*). Simultaneously, we face a global teacher shortage of staggering proportions. UNESCO estimates that 44 million additional primary and secondary teachers will be needed globally by 2030, including 15 million in Sub-Saharan Africa alone. These figures reframe teacher supply not as a sectoral inconvenience but as an education-security issue. Teacher shortages are now a global crisis, extending across Europe and North America, where even well-resourced systems struggle to attract and retain qualified educators. Attrition rates among primary school teachers have nearly doubled in just seven years and up to 40% leave the profession within five years of starting (UNESCO 2024). This is not sustainable.

Alongside these pressures, the public release and rapid diffusion of generative AI systems from November 2022 onwards brought automation into the symbolic heartland of higher education: language, reasoning, writing, coding and other forms of cognitive labor. Up until then, universities felt largely unchallenged by the mechanization that was replacing blue-collar jobs. Unlike earlier waves of mechanization that were often understood through the displacement of routine manual or administrative work, generative AI has made the automation and augmentation of white-collar jobs newly visible.

Whether it is ecological collapse or omnipresence of AI, the transitions that are fundamentally transforming our known world must act as a wake-up call for higher education institutions: we need urgently to regenerate the idea of the university. It can no longer just act as a machine for credentials, nor, indeed, innovation. The university must reinvent the Socratic education ideal of human flourishing and become a space for nurturing the capacity to care, imagine, relate and act for the planet. Universities are called to rethink what knowledge is, how it is generated and for whom. And we, as leaders in higher education, are called to protect the very conditions that make meaningful education possible.

This chapter aims to offer a vocabulary, a diagnosis and a pedagogical architecture for rethinking the university's planetary mission in the age of AI. We argue that within this mission, one of the defining tasks of higher education is to regenerate the university as an architecture of agency. Following Albert Bandura, agency can be understood as the capacity of individuals and groups to influence their own functioning and life circumstances through intentional action, self-reflection and collective effort (2018, 2020). Following Amartya Sen's capability approach, agency also implies the freedom to pursue goals one has reason to value, rather than merely to perform efficiently within externally imposed systems (Sen 1999, 18-19). In this sense, student agency is at once cognitive, ethical, social and political: it concerns what learners know, what they can do, what they value, and how they act with others. It is also the capacity to become both the author and the actor of one's own future.

Here, the university must then become an institutional space in which students learn not only to acquire knowledge and credentials, but to exercise judgement, care, imagination and collective responsibility in conditions of uncertainty, thereby training students not just to be the best *in* the world, but the best *for* the world.

The weight of the crowd: massification and its discontents

An important aspect of the development of the university over the course of the 20th century is that of massification. Sociologist Martin Trow articulated this transition not merely as a demographic expansion, but as a series of three shifting archetypes: democratization, massification and universal access (Trow 1962, 1970, 1973). Each of these represents a distinct social contract. According to Trow, the elite system initially functioned as a rarefied institution designed to cultivate a small circle of social and professional leaders. Intended to strengthen the “mind and character” (Trow 1973, 7) of a future ruling class, it operated within a narrow social enclave, encompassing no more than 15% of an age cohort. Here, the university is a space of pedagogical ritual designed to reproduce the guardians of the status quo.

As the twentieth century progressed, however, the stage broadened into the mass system. While still anchored in a social elite, the definition of that elite expanded to include the technocratic and managerial strata of a rapidly industrializing society. This model aims to serve up to 50% of a generation with the focus of universities shifting from the intangible cultivation of character to the quantifiable acquisition of skills (Trow 1973, 8). This marks the era of professionalization, in which the cultivation of scholarly character is increasingly supplemented, perhaps even at times displaced, by an emphasis on employability, technical competence and managerial formation.

Finally, we arrive at the Universal model, where access exceeds the 50% threshold. In this landscape, the university is no longer a destination for the few but a requirement for the many. It is important to note that these thresholds are not absolute or mutually exclusive categories. Indeed, drawing a loose parallel from quantum mechanics, where a system can exist in a superposition of multiple states simultaneously, Trow’s thresholds function as heuristic markers. They represent concurrent shifts in institutional purpose, curriculum, student identity, and governance, rather than rigid, chronological boundaries.

The quantitative scale of this transition is nothing short of dizzying. At the dawn of the 20th century, global higher education was a boutique affair, catering to a mere 500,000 students worldwide (Schofer 2005, 898). As of 2024, that figure has increased to 269 million, up from 100 million in 2000 (UNESCO & UNESCO IESALC, 2026). Yet, this narrative of so-called universality is a story of profound and uneven distribution, reflecting more a fragmented geography of exclusion. As the UNESCO and UNESCO IESALC (2026) *Higher Education Global Trends Report* states:

Technological disruption, demographic change and shifting labour market expectations are reshaping higher education globally. Although global enrolment more than doubled over the last 20 years to reach nearly 269 million students in 2024, major inequalities persist in access, completion, funding and mobility. Fewer than 3% of students participate in international mobility, only 9% of refugees access higher education, and many institutions face mounting financial pressure, growing threats to academic freedom, and the challenge of adapting to digital transformation and artificial intelligence (UNESCO 2026, 3).

Calderon’s analysis usefully reminds us that the movement towards universal participation remains uneven: some 34% of countries seem to have crossed the threshold of universal access, while roughly 27%, predominantly in Sub-Saharan Africa, remain structurally constrained in an elite model by infrastructure, financing, demography and historical inequality (Calderon 2018).

The stakes of this expansion are not merely economic; they are existential. OECD data suggests that a tertiary degree is associated with a range of social benefits, including employment, health, civic participation and self-reported well-being. Indeed, a university graduate can expect to live longer, vote more frequently, and report higher levels of civic engagement than those excluded from the system (OECD 2013). Education, in this sense, is the ultimate tool for social resilience.

However, since the 1960s, this march towards universalism has exposed deep-seated structural fractures. The sheer cost of widening participation while preserving research intensity, student support, academic freedom and pedagogical quality has placed many systems under strain. The economic pressure this represents has led to a frantic diversification of institutional identities: the rise of vocationally-orientated colleges and the proliferation of for-profit private providers. In this frantic acceleration towards universality, have we swapped a transformational educational experience for increased recruitment? Have we traded the soul of the institution as a space for critical inquiry and creative thinking for the perceived efficiency of the credential?

As the university attempts to reconcile its 19th-century legacies with the ferocious demands of technological integration, it finds itself in a precarious position. The challenge is no longer just about opening the doors, but about ensuring that what lies behind them remains a space for meaningful agency.

AI: the new social infrastructure of uncertainty

To understand the impact of Artificial Intelligence, we must reject both the techno-optimist's utopianism and the Luddite's nostalgic denial. AI should no longer be understood only as a tool. It is increasingly becoming a layer of social, cognitive and institutional infrastructure (see Miao & Holmes 2023). This invisible architecture reshapes how we relate, how we think, and how we imagine the future.

For the university to navigate this shift systemically, we must engage with a fundamental triptych of educational response: the distinction between learning about AI, learning with AI, and learning in the age of AI. This extends naturally across the academic and professional landscape, encompassing how we teach and, ultimately, how we work.

In many institutional strategies, attention has so far gravitated towards the first two dimensions:

- 1) Learning about AI, teaching AI and working on AI: developing technical, civic and critical AI literacy, from basic understanding of algorithmic systems to questions of data, bias, governance and accountability
- 2) Learning, teaching, working with AI: using AI systems to support feedback, exploration and productivity, while maintaining human judgement, transparency and pedagogical purpose.

While these are necessary, they are insufficient. We, so far, have perhaps given far less focus to the third, most vital dimension: learning, teaching, and working in the age of AI. Yet it is precisely here, in the broader sociotechnical context, that the university's systemic mission resides. To learn in the age of AI is to examine how AI reshapes knowledge, work, identity and power; it is also to move from the passive consumption of algorithmic outputs towards a more discerning and ethically situated form of agency. In a BANI world, agency is often paralyzed by the 'I' of incomprehensibility: the sense that reality has become too opaque, too accelerated or too mediated to act upon with confidence. If technical power has grown faster than human wisdom, then the univer-

sity's task is not simply to accelerate adaptation, but to cultivate judgement. *The International Forum on Education for the Post-2030 Agenda: Outcome Report 2026* similarly identifies “deep knowledge of facts, perspectives, and values” (3) as a crucial defense against the erosion of truth. Agency, in this light, is no longer only the capacity to act; it is the prior capacity to discern. It is the capacity to distinguish evidence from hallucination, knowledge from noise, and dreams of shared flourishing from dreams built on the quiet displacement of others. Without this epistemological and ethical anchor, AI risks becoming a weaponized engine of collective nightmares; with it, it may become a loom for weaving compatible, ethically grounded visions of a shared future.

For mid-tier, regional and resource-constrained institutions that educate large numbers of students worldwide, AI is a ferocious challenger to the white-collar jobs higher education prepares for and that were previously perceived to be immune to mechanization. It risks commodifying the very hard skills that were once the gateway to the middle class, while accelerating the reconfiguration of skills demanded by employers (World Economic Forum 2025). This forces us to rethink what knowledge is, how it is generated, and for whom. Rather than competing with AI on cognitive tasks where these silicon systems can now, under certain conditions, outperform carbon-based human graduates, the university must ensure that learning environments remain spaces for authentic, context-sensitive resilience.

The question for the post-AI university is not only ‘what can the student do that the machine cannot?’ but also ‘who can the student become, with others, in a world increasingly mediated by intelligent systems?’ We must pivot from the excellence of the expert alone, increasingly exposed to forms of AI competition, and turn to the agency of the planetary citizen or, rather, what we call the “planetizen”: individuals who collectively treat the entire planet as their homeland and who consider stewardship of the commons (natural, cultural, and digital) as their primary responsibility (Taddei 2022).

The eloquence of being human: redefining intelligence in the age of AI

If AI systems can now perform a growing range of standardized cognitive tasks, universities are forced to revisit what they mean by intelligence, expertise and human educational value. This question has been approached through various frameworks that seek to identify the cognitive, social, creative and self-regulatory capacities learners need in complex societies and that should require our attention in designing successful 21st-century education systems, such as the 4Cs identified by Partnership for 21st Century Skills (P21) or the more exhaustive 7Cs identified by its senior fellows Bernie Trilling and Charles Fadel: Critical thinking and problem solving, Creativity and innovation, Collaboration, teamwork, and leadership, Cross-cultural understanding, Communications, information, and media literacy, Computing and ICT literacy, and Career and learning self-reliance (Trilling & Fadel 2009, 176).

Rose Luckin raises two issues here. Firstly, there exist multiple frameworks offering different skill sets. Secondly, there is a recurring divide between educators who believe that our primary focus should be on skills acquisition and those who believe that it should be on knowledge. For Luckin, this is a false opposition. We cannot separate knowledge from skills, since education requires both “knowledgeable skills and skillful knowledge” (Luckin 2018, 97).

Luckin, therefore, introduces her own framework, broadening intelligence beyond academic performance alone. She identifies seven elements of intelligence: academic, social, meta-knowing, metacognitive, metasubjective and metacontextual intelligence,

alongside perceived self-efficacy. These can be read as three overlapping movements. First, learners need knowledge about themselves as learners: the capacity to monitor understanding, regulate attention, recognize emotion, judge confidence and revise assumptions. Second, they need knowledge about their relation to others and to context: the ability to understand social interaction, perspective, embodiment, and the conditions in which knowledge is produced and used. Third, they need knowledge about the world: facts, concepts, disciplines, histories, evidence and our impact upon the environment and planetary resources (Luckin 2018, 65–66). Put simply, and in our own terms, education in the age of AI means learning to take care of ourselves, each other and the planet.

This is a helpful starting point, but in the age of AI it needs to be extended. Human intelligence cannot be reduced to individual cognition, however sophisticated. Nor can it be understood only as the capacity to solve problems efficiently. The challenge is increasingly collective: to deliberate across difference, to build shared understanding, to test knowledge in action, and to solve problems with others under conditions of uncertainty. This is where the university's mission must connect Luckin's account of human intelligence with the OECD's emphasis on student agency, co-agency and transformative competencies, as well as with the kinds of collaborative problem-solving cultivated in civic education and youth movements. Intelligence, in this sense, is not merely something a learner possesses; it is something learners practice together.

What such an expanded framework also reminds us is that human intelligence cannot be reduced to computational performance. Human beings do not arrive in the world as finished organisms equipped for immediate autonomy. We are born radically dependent: vulnerable, unfinished, and reliant on others for survival, language, meaning and orientation. This prolonged dependency is not a weakness to be overcome, but one of the conditions of human intelligence itself. Because we must be cared for, taught and initiated into shared worlds, intelligence becomes relational, embodied, cultural and ethical. This is what distinguishes it from the operational efficiency of machines and from any narrow account of intelligence as problem-solving capacity alone.

This, in turn, helps us resist a sterile opposition between knowledge and skills. In the age of AI, the question is not whether students need knowledge or competencies, but how knowledge becomes usable, situated and ethically directed.

We believe that if we engage in this debate holistically, it allows us to understand that the goal of the university cannot be reduced to guessing which jobs of the future will exist, as this is an exercise that usually ends up simply identifying the needs of today. Instead, the goal is to foster capacity-building and the durable, transferable capacities that allow learners to interpret situations, act with others, revise assumptions and pursue purposes they have reason to value. This is the process of empowering a student to move from being a spectator of global crises to becoming an agent of change. It is an education of the *self* in relation to the *other*, fostering a collective intelligence that is synopated in its agility, polyphonic in its perspectives, and resilient in its purpose.

The hero's journey: a narrative of becoming

One way to understand higher education as a process of agency formation is through narrative. The university journey can be read, cautiously, through the structure of the 'hero's journey' (Campbell 1949; Brown & Moffett 1999). Joseph Campbell summarizes his take on this monomyth as: "A hero ventures forth from the world of common day into a region of supernatural wonder: fabulous forces are there encountered and a deci-

sive victory is won: the hero comes back from this mysterious adventure with the power to bestow boons on his fellow man” (Campbell 1949, 30). For our purposes, the student is not a heroic individual in the romantic sense, but a learner who enters into uncertainty, encounters systemic limits, and returns with an expanded capacity for collective action. They receive a ‘Call to Adventure’ (manifested through a research question or social challenge) that urges them to cross a ‘Threshold’ into an unknown world of new experiences. This journey inevitably leads to an ‘Abyss’: a reflexive moment of uncertainty, ambiguity, and disorientation. It is here that the experiences must be scaffolded so as to be transformed into educational value, ultimately allowing the student to claim the ‘Gift’, not of personal glory, but of profound systemic understanding (Campbell 1949, *passim*).

If we are to bring learners to engage in a BANI world with resilience, this adventure cannot be a martyr’s journey. Indeed, we should not romanticize struggle, nor confuse transformative learning with exhaustion, alienation or burnout. Rather, we should seek transformative learning through a fractal and recursive process. In the age of AI, this requires universities to move beyond single-loop learning, i.e. correcting an action in order to receive a grade (‘I must revise better!’), towards deeper transformative cycles. Drawing initially on organizational learning theory, particularly Argyris and Schön’s distinction between single-loop and double-loop learning (Argyris & Schön, 1978), we can move from corrective learning to reflective learning about governing assumptions, and then extend this logic towards identity-level and planetary learning. In this sense, the hero’s journey becomes a passage from a ‘Seed Dream’, the learner’s first image of success, towards a regenerative reality in which that dream is made answerable to the lives of others. The question is no longer simply ‘What do I want to become?’ but ‘What does my becoming make possible, and for whom?’.

This opens onto three deeper learning loops:

1) Double-loop: questioning the underlying values of the system and identifying the shadow of one’s ambition: ‘Who might be harmed if I pursue success without questioning what success means?’

2) Triple-loop: transforming the learner’s identity by learning to see the self from other positions, especially those of the vulnerable, the excluded and future generations: ‘Who am I in this world? What ancestor am I becoming?’

3) Quadruple-loop: operating at the scale of the planetary system, moving from individual excellence towards a collective tapestry of *planetizens*: ‘How do I act for the planet?’

This is precisely the question raised, in another register, by the AgroParisTech graduates: not whether students should succeed, but what kind of success they are being trained to desire.

While triple-loop and quadruple-loop learning are less stabilized concepts, they are useful here as heuristic extensions of this recursive logic. Reframed in this way, the hero’s journey is less a myth of individual triumph than a narrative grammar for transformative learning: a way of describing how learners encounter complexity, rework assumptions, and return with a stronger sense of responsibility. This is how the initial seed dream can become an ethical dream: not a private fantasy of success, but a vision of flourishing that refuses to become someone else’s nightmare, today or tomorrow.

The MOOD framework: from information to sense-making

The question remains: how do we operationalize this agency? In *Pedagogy of the Oppressed* (1968), Paulo Freire delivers a searing critique of the traditional educational archetype, which he calls the “banking concept of education” where “knowledge is a gift bestowed by those who consider themselves knowledgeable upon those whom they consider to know nothing” (Freire 2005, 72). Within this transactional logic, the learner is relegated to the status of a hollow vessel and, as a passive repository divested even of their own prior knowledge through the teacher’s projection of “absolute ignorance” (Freire 2005, 72), they are denied the fundamental power of agency.

For Freire, the classroom must be transformed into a space of dialogue: a collaborative site where educator and learner co-construct a body of knowledge that contributes to a capacity to act upon the world. It is through this critical pedagogy or “pedagogy of hope” (Freire, 2014) that we pave the itineraries for true engagement, enabling us to imagine and engage in a future that is at once more ethical, more compassionate, and more deeply human.

Building on this lineage of critical, experiential and project-based pedagogies (see also Dewey 1916, 1938; Kilpatrick 1918; Kolb 1984), the Learning Planet Institute has developed the MOOD framework: Meaningful Open Opportunities for Discovery. This replaces the static, banking model of education with an experiential odyssey that aims to rebalance educational cultures that over-privilege competition, disciplinary enclosure and performance anxiety by foregrounding collaboration, interdisciplinary inquiry and hope-orientated action. MOOD should be understood not as a rejection of existing pedagogical traditions but as a synthesis and operationalization of several of them: project-based learning, challenge-based learning, research-led learning, design thinking, social-emotional learning (SEL) and reflective practice. It also seeks to articulate the vital and porous intersections between formal, non-formal, and informal learning scapes. Its specific contribution lies in combining these approaches into a structured journey that links personal meaning, interdisciplinary inquiry, AI-supported exploration and civic or planetary action.

MOOD programmes are designed primarily as a three-act structure:

1) Discover: Anchored in structured reflexive work through an Ikigai-inspired analysis, the student aligns their personal reason for being (SEL) with the world’s needs, grounded in research. In a BANI world, this is the internal compass that prevents drifting.

2) Ideate: Students engage in collective intelligence practices and use AI systems based on vast curated data sets to explore evidence, generate alternatives, test assumptions and map possible futures. But ideation is not mere brainstorming. It creates “safe spaces for complex dialogue” (International Forum on Education for the Post-2030 Agenda 2026, 3), where the initial seed dream is exposed to systemic and critical thinking. Here, learners begin to identify the shadow of their projects: who might be harmed, displaced or silenced if this vision succeeds?

3) Build: Students prototype a response, treating knowledge, skills and abilities not as decontextualized learning outcomes but as resources mobilized in situated action. In this phase, AI must remain a guardian rather than a master: a means of testing assumptions, surfacing unintended consequences and ensuring that emerging solutions remain compatible with human dignity, social justice and ecological limits.

In this framework, student engagement is understood as meaningful participation in knowledge-making, self-reflection and collective action. By combining civic engagement and deep academic immersion, the student engages meaningfully in their learn-

ing. Through this type of pedagogy, learners can develop what we might call ARC capacities: agency, resilience and co-construction (Costambeys-Kempczynski 2022). These should not be treated as automatic outcomes, but as capacities that require intentional design, scaffolding, feedback and institutional support. This transformation offers one possible antidote to the anxiety of the BANI world.

The university as a learning organization

The university after AI cannot be a static hierarchy. The institution itself must mirror the transformation it demands of its students. Drawing on Peter Senge's work on the learning organization (Senge 1990), the university must become a learning organization governed by five disciplines that can be reinterpreted for the university's planetary mission.

Mental Models: Universities must examine the assumptions through which they define excellence, success, employability, knowledge and impact. As the student must begin with an exercise in reflexivity, so must the university begin by "turning the mirror inward; learning to unearth our internal pictures of the world, to bring them to the surface and hold them rigorously to scrutiny" (Senge 1990, 8).

Personal Mastery: A commitment to lifelong learning and the objective assessment of the gap between planetary realities and an ethical vision, allowing to clarify "the things that really matter to us, of living our lives in the service of our highest aspirations" (Senge 1990, 8).

Shared Vision: A collective image of an ethical future motivated by educational purpose, planetary responsibility and institutional justice, in which no institution's progress is built on the destruction, exclusion or silencing of others. Such a vision fosters "genuine commitment [...] rather than compliance" (Senge 1990, 9).

Team Learning: Recognition that many of the most important educational capacities are relational rather than purely individual: dialogue, distributed expertise, mutual recognition and the capacity to learn with those who know what we do not. It is here that Senge's emphasis on dialogue becomes central: the "discipline of team learning starts with 'dialogue', the capacity of members of a team to suspend assumptions and enter into a genuine 'thinking together'" (Senge 1990, 10).

Systems Thinking: The cornerstone discipline that integrates the others. The university must understand itself not as an ivory tower outside society, but as an actor embedded in ecological, technological, political and economic systems. Indeed, it too is "bound by invisible fabrics of interrelated actions" (Senge 1990, 7). A planetary view is now vital in allowing the institution to see the big picture of interdependencies.

The learning organization, however, cannot be invoked naïvely. Universities are not simply communities of shared learning; they are also sites of hierarchy, resource constraint, epistemic conflict and political contestation. A planetary university must therefore combine organizational learning with democratic governance, academic freedom and institutional accountability. By treating the university itself as a learning organization, we move towards a Learning Society. Here, the logic of the *Réseaux d'Échanges Réciproques de Savoirs* becomes newly urgent: no one knows everything, everyone knows something, and we all gain by learning together (Héber-Suffrin & Héber-Suffrin 1994). This is especially true when education enters what Richard Elmore describes as the distributed and collective quadrant of learning, where learners are connected by shared interests, mutual resources and a common domain of inquiry, and where the task is no longer to transmit known answers but to face uncertainty together (Elmore

2019). In such a society, AI should facilitate networking, reciprocity and decision-making rather than replace human judgement.

Doughnut learning: the ethical compass

Finally, drawing on Kate Raworth's doughnut model, we can imagine '*doughnut learning*' as an ethical compass for higher education to navigate the transitions of the 21st century. Doughnut learning gives ethical form to a concern already present in MOOD: that no project should solve one problem by creating or deepening another, whether social or ecological. Raworth's model defines a "safe and just space" between a "social foundation" and an "environmental ceiling" (Raworth 2012, *passim*). Translated into educational terms, this suggests that universities must attend simultaneously to learner flourishing, global connectedness and planetary limits. It also prepares a different question of value: not how to be ranked as best in the world, but how to become best for the world.

The university must ensure we do not fall below the social floor, that is to say, the indispensable foundation which encompasses not only access and academic support, but the more profound textures of belonging, well-being, psychosocial competencies and the capacity for "collaborative and responsible action" (International Forum on Education for the Post-2030 Agenda 2026, 3). It is a floor sustained by the active exercise of student voice, global connectedness and meaningful participation across borders. A short-fall here can contribute to forms of educational alienation, including disengagement, exclusion and transitions into NEET status (Not in Education, Employment, or Training).

Simultaneously, the university must ensure we do not overshoot the environmental ceiling. In educational terms, this environmental ceiling asks whether the pursuit of excellence, growth, mobility, infrastructure and technological innovation is being organized in ways compatible with planetary boundaries, social justice and "the needs and interests of future generations" (United Nations 2024, *passim*). The university must not act at the cost of the environment, reinforce social hierarchies that stifle agency, or educate leaders whose success is purchased at the expense of those who come next. In this sense, it must become a space where those too often absent from educational design, namely the environment, the excluded and future generations, are given a seat at the table.

The goal is to produce graduates who are not just experts, but citizens for the planet who understand that their skills must serve the regeneration of the global commons. They are not merely trained to succeed in the world as it is, but to become responsible ancestors for the generations still to come.

Evaluating the best for the world: toward a grammar of value

If 'best in the world' is measured through rankings, credentials, selectivity, growth and visibility, 'best for the world' requires a different grammar of value. It asks not only whether students succeed, but what their success makes possible for others. It asks whether knowledge reduces harm or merely displaces it; whether technologies expand agency or increase dependency; whether institutions strengthen commons or enclose them; whether educational achievement remains compatible with human dignity, social justice and environmental limits.

This is not a call to replace one dashboard with another. The danger would be to translate the planetary mission of the university into yet another regime of indicators, benchmarks and competitive performance. Rather, evaluation must become an exercise in practical wisdom: a way of asking what forms of evidence help us judge whether education is enlarging the conditions for human and planetary flourishing.

Here, universities can learn, cautiously, from adjacent movements that have already challenged narrow accounts of value. B Corp-style approaches remind us that organizations can be assessed not only by growth, revenue or reputation, but by their social and environmental accountability (B Lab, n.d.). OECD WISE similarly reflects a broader shift away from GDP-style thinking towards well-being, inclusion, sustainability and equal opportunity (OECD, n.d.). For higher education, the equivalent question is not simply whether graduates are employable, visible or high-performing, but whether their learning enables them to generate positive handprints that exceed the footprints they leave behind.

This means that evaluation must include both quantitative and qualitative forms of judgement. We can ask whether projects create measurable benefits for communities, ecosystems or institutions; but we must also ask whether learners become more capable of listening, collaborating, revising assumptions and acting responsibly under uncertainty. We can ask whether a project produces outputs; but we must also ask whether it increases dignity, reduces harm, strengthens commons and leaves future generations with more room to act.

The ultimate evaluative question of an architecture of agency is therefore simple but demanding: if this learning disappeared tomorrow, what would the world lose? If the answer is only another credential, another ranking, or another high-performing cog, then the university has reproduced the world as it is. If the answer is a strengthened commons, a repaired relation, a new capacity to act, or a project that allows others to flourish, then education has begun to form learners who are best for the world.

Conclusion: reclaiming the common good

There is, however, a necessary tension to acknowledge in our critique of credentialism and the fractured logic of disciplinary silos. Indeed, we must be wary of falling into a simplistic dismissal of expertise. The university's planetary mission, with its daunting, intergenerational odyssey, cannot be navigated through civic enthusiasm or the mere performance of project-based learning alone.

The wicked problems of our age, from the intricacies of climate science and public health to the emerging grammars of artificial intelligence, demand more than superficial literacy; they require the rigors of deep disciplinary knowledge and a sustained apprenticeship into the specific methodologies of the field. The crisis we face is not a surplus of expertise, but rather an expertise that has become untethered from planetary realities and divorced from an ethical orientation, blind to systems thinking, and shielded from the weight of collective responsibility.

The task, therefore, is not a simple substitution of knowledge for agency but the design of learning ecologies in which knowledge itself becomes agentic. Aristotle's account of the intellectual virtues is useful here, not as a framework to be imported wholesale, but as a reminder that knowing has several modes. Alongside *nous*, the grasp of first principles, and *sophia*, theoretical wisdom, Aristotle distinguishes *epistēmē*, scientific or theoretical knowledge, *technē*, the capacity to make and produce, and *phronēsis*, the practical wisdom required to judge well in contingent situations. For the university after

AI, this distinction matters because the crisis is not a lack of knowledge or technique alone, but a deficit of judgement about what knowledge and technique are for. The university must therefore move beyond the technical and the cognitive alone, towards the ethical and the practical: the wisdom to decide how, when and why knowledge should be put to work. We must move towards an expertise that is situated, reflexive and profoundly accountable. We must move towards a knowledge that is no longer a private asset but a resource directed towards the reclamation of the common good.

The *désertion* of the AgroParisTech graduates was not only a refusal; it was a demand for a university capable of helping students make sense of the worlds into which it sends them. It was a demand for a system that recognises that in a world of AI and BANI, the degree cannot be the only measure of value. The deeper question is whether higher education enables learners to become knowledgeable, capable and responsible agents in the world.

The university must be regenerated: not by abandoning knowledge, expertise or disciplinary depth, but by reconnecting them to care, imagination, relation and action. By prioritizing psychosocial competencies, by engaging in capacity building such as agency, resilience and co-construction, and by embracing the learning organization as a lever of transformation towards a learning society, the university can finally bridge the gap between ‘forming the best *in* the world’ and ‘forming the best *for* the world.’ In this sense, the architecture of agency is also an architecture of hope: one in which deep factual knowledge, compassion and ethical discernment ensure that the dreams of our graduates do not become someone else’s nightmare.

The university is not a right to be defended; it is a common good to be reinvented. The architecture of this reinvention must ensure that universities do not merely teach students how to use intelligent machines, but how to bring human intelligence and artificial intelligence into responsible relation so that learners may think, judge, care and flourish in worlds those machines are already helping to remake.

References

- Argyris, C., & Schön, D. (1978). *Organizational learning: A theory of action perspective*. Reading, MA: Addison-Wesley.
- Aristotle. (2009). *The Nicomachean Ethics*. Translated by David Ross, revised by Lesley Brown. Oxford University Press.
- Bandura, A. (2018). Toward a psychology of human agency: Pathways and reflections. *Perspectives on Psychological Science*, 13(2), 130–136. <https://doi.org/10.1177/1745691617699280>
- Bandura, A. (2020). Social Cognitive Theory: An Agentic Perspective. *Psychology: The Journal of the Hellenic Psychological Society*, 12(3), 313–333. https://doi.org/10.12681/psy_hps.23964
- B Lab. (n.d.). *What is a B Corp?* <https://www.bcorporation.net/en-us/certification/>
- Brown, J. L., & Moffett, C. A. (1999). *The Hero’s Journey: How Educators Can Transform Schools and Improve Learning*. ASCD.
- Calderon, A. (2018). Massification of higher education revisited. https://www.researchgate.net/publication/331521091_Massification_of_higher_education_revisited
- Campbell, J. (1949). *The Hero With A Thousand Faces*. Pantheon Books.

- Cascio, J. (2020). *Facing the Age of Chaos*. Medium.
- Code, J. (2020). Agency for Learning: Intention, Motivation, Self-Efficacy and Self-Regulation. *Frontiers in Education*. 5:19. doi: 10.3389/educ.2020.00019
- Cooperman, L. (2014). From elite to mass to universal higher education: From distance to open education. *RIED. Revista Iberoamericana de Educación a Distancia*, 17(1).
- Costambeys-Kempczynski, R. (2022, March 10). Expérience étudiante: une vision globale de l'université [Conference session]. JIPES 2022 : L'expérience étudiante, Canal-U.
- Dewey, J. (1916). *Democracy and education: An introduction to the philosophy of education*. New York: Macmillan.
- Dewey, J. (1938). *Experience and Education*. Kappa Delta Pi.
- Elmore, R. F. (2019). The Challenges of Learning and Design. In P. Nair, R. Zimmer Doctori, with R. F. Elmore, *Learning By Design: Live, Play, Engage, Create* (pp. 160–195). Education Design Architects. <https://www.ecosisteam.cl/wp-content/uploads/2019/10/Paper-principal-Elmore-Learning-and-Design.pdf>
- Fink, L. D. (2013). *Creating Significant Learning Experiences: An Integrated Approach to Designing College Courses*. Jossey-Bass.
- Freire, P. (2005). *Pedagogy of the oppressed*. New York: Continuum.
- Freire, P. (2014). *Pedagogy of Hope: Reliving Pedagogy of the Oppressed*. Great Britain: Bloomsbury.
- Helker, K., Bruns, M., Reymen, I. M., & Vermunt, J. D. (2025). A framework for capturing student learning in challenge-based learning. *Active Learning in Higher Education*, 26(1), 213-229.
- Héber-Suffrin, C., & Héber-Suffrin, M. (1994). *L'école éclatée*. Epi-Desclée de Brouwer.
- International Forum on Education for the Post-2030 Agenda. (2026). *Outcome Report 2026*. Ban Ki-moon Foundation.
- Kilpatrick, W.H. (1918) The Project Method. *Teachers College Record: The Voice of Scholarship in Education*, 19, 319-335. <https://doi.org/10.1177/016146811801900404>
- Kolb, D. A. (1984). *Experiential Learning: Experience as the Source of Learning and Development*. Englewood Cliffs, NJ: Prentice Hall. <http://academic.regis.edu/ed205/Kolb.pdf>
- Luckin, R. (2018). *Machine Learning and Human Intelligence: The Future of Education for the 21st Century*. UCL Press.
- Miao, F., & Holmes, W. (2023). *Guidance for generative AI in education and research*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>
- OECD (2013). 'What Are the Social Benefits of Education?', Education Indicators in Focus, No. 10, OECD Publishing, Paris, <https://doi.org/10.1787/5k4ddxnl39vk-en>.
- OECD. (2019). Student agency for 2030 (OECD Education 2030 Project Concept Note). https://www.oecd.org/content/dam/oecd/en/about/projects/edu/education-2040/concept-notes/Student_Agency_for_2030_concept_note.pdf
- OECD (2021). *Education at a Glance 2021: OECD Indicators*. OECD Publishing.
- OECD. (n.d.). *Centre on Well-being, Inclusion, Sustainability and Equal Opportunity (WISE)*. OECD. <https://www.oecd.org/en/about/directorates/centre-on-well-being-inclusion-sustainability-and-equal-opportunity.html>

- Partnership for 21st Century Learning (2019). Framework for 21st-Century Learning Definitions. Batelle for Kids.
- Passelac A. (2026, February 26). “Mettre fin à l’inaction agroécologique” : chez les étudiants agronomes la colère s’enracine derrière la révolte. *L’Étudiant*. <https://www.letudiant.fr/etudes/ecole-ingenieur/des-futurs-ingenieurs-agronomes-veulent-repenser-l-agriculture.html>
- Raworth, K. (2012). *A safe and just space for humanity: can we live within the doughnut?* Oxfam International. <https://policy-practice.oxfam.org/resources/a-safe-and-just-space-for-humanity-can-we-live-within-the-doughnut-210490/>
- Raworth, K. (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. Chelsea Green Publishing.
- Santos, A. C., Simões, C., Melo, M. H. S., Santos, M. F., Freitas, I., Branquinho, C., Cefai, C., & Arriaga, P. (2023). A systematic review of the association between social and emotional competencies and student engagement in youth. *Educational Research Review*, 39, Article 100535. <https://doi.org/10.1016/j.edurev.2023.100535>
- Schofer, E., & Meyer, J. W. (2005). The Worldwide Expansion of Higher Education in the Twentieth Century. *American Sociological Review*, 70(6), 898–920.
- Scholars at Risk (2025). *Free to Think Report*. <https://www.scholarsatrisk.org/resources/free-to-think-2025/>
- Sen, A. (1999). *Development as Freedom*. Oxford University Press.
- Senge, P. M. (1990). *The Fifth Discipline: The Art & Practice of The Learning Organization*. Doubleday.
- Sinha, D., & Sinha, S. (2020). Managing in a VUCA World: Possibilities and Pitfalls. *Journal of Technology Management for Growing Economies*, 11, 17-21. <https://doi.org/10.15415/jtmge.2020.111003>
- Taddei, F. (2022, September 13). *Débat : Repenser la citoyenneté pour mieux entendre la jeunesse*. The Conversation. <https://doi.org/10.64628/AAK.cvy9fxduc>
- Trilling B. & Fadel C. (2009). *21st Century Skills: Learning for Life in Our Times*. Jossey-Bass.
- Trow, M. (1962). The Democratization of Higher Education in America. *European Journal of Sociology*, 3(2), 231–262. doi:10.1017/S0003975600000618
- Trow, M. (1970). Reflections on the transition from mass to universal higher education. *Daedalus* 99(1): 1–42.
- Trow, M. (1973). Problems in the transition from elite to mass higher education. Berkeley: Carnegie Commission on Higher Education.
- UNESCO (2023). *Higher education: figures at a glance*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000394112.locale=en>
- UNESCO (2024). *Global Report on Teachers: Addressing teacher shortages and transforming the profession*. UNESCO.
- UNESCO & UNESCO International Institute for Higher Education in Latin America and the Caribbean. (2026). *Higher education global trends report: Towards inclusive, equitable and quality higher education in an internationally mobile landscape*. <https://doi.org/10.54674/IN.2026.94>

United Nations. (2024). *Pact for the Future, Global Digital Compact and Declaration on Future Generations*. United Nations.

World Economic Forum. (2025). *The Future of Jobs Report 2025*. World Economic Forum.
https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf

Chapter 7. AI meets accreditation: lessons from the U.S. for the future of quality assurance

Ralph A Wolff

We are at the early stage of what is already proving to be a fundamental transformation in how accreditation functions, driven by the rapid disruption of AI, alongside a parallel transformation in how higher education is delivered. In these early days of AI, evolutionary changes are already occurring, but the significant disruptions AI is bringing to many fields, such as law, accounting, coding, medical research, and other fields, suggest that higher education and accreditation will not be immune and will, in the foreseeable future, be disrupted in significant ways.

This chapter will explore likely changes that AI will bring to US accreditation by suggesting eight interconnected areas of transformation over the next decade. Some are already in the early stages of implementation among innovative providers and early adopters. They are likely to be just the beginning of the significant changes ahead in the ways quality is defined and accreditation is practiced.

These changes include: the shift in quality definitions from general standards to metric-based data analyses; the ability to make hyper-localized as well as comparative benchmarking judgements of institutional quality; the evolution of the self-study from a labor-intensive and costly ritual to a data-driven, continuous process; the emergence of real-time automated monitoring in place of periodic peer review; the ability of AI to evaluate curriculum content, quality and rigor against peer and national norms; the reconceptualization of faculty roles in this new space; and a redefinition of ways of knowing.

While the focus of this chapter is on institutional accreditation in the US, many of these changes are likely to occur in programmatic accreditation and in accrediting bodies outside of the US.

From general standards to a metrics-based system of quality assurance

Accreditation has become a major form of quality assurance globally. In the United States, it is over 100 years old and a creature of the higher education community itself to ensure the quality and integrity of institutional performance. Accrediting agencies in the US are nongovernmental and governed by their members. Since 1952, when the federal government first formally relied on accreditation to ensure institutional quality and integrity for the award of financial aid, there has been in place a detailed “recognition” process by which accreditors are formally reviewed to determine their compliance with an extensive set of regulations.

As a result of this recognition, approximately \$114 billion in federal student aid flowed to students in fiscal year 2023 alone. Given the huge influence of this allocation of financial aid, new federal and state rules have pushed accrediting agencies toward greater public accountability and transparency, even as accreditation seeks to serve and

protect member institutions from regulatory overreach and maintain respect for the diversity of institutions they serve. As will be shown, machine learning and algorithms used by AI providers will bring new opportunities for accreditation to serve public policy interests and make multiple practices more efficient and effective.

A common framework for accreditation is utilized across nearly all agencies. It is an intensely human process relying primarily on peer review at each level. The typical process requires multiple steps: a comprehensive self-study organized around the agency's accrediting standards, a peer review team "visiting" the institution (either in person or online, or a combination), a team report submitted to the institution and accrediting decision-making body, a response from the institution and, in some cases, an appearance of key institutional leaders before the decision-making body. At the end of this process is an accrediting decision specifying whether accreditation is granted, reaffirmed, or continued—and if so, under what conditions, if any, and what type of follow-up is expected before the next comprehensive review, when the same cycle is repeated.

As such, this process is time-consuming, subject to variability and subjective judgments, and costly. It is not well understood by many constituencies within institutions or by the public, and is seen by most well-established institutions as a compliance exercise of limited educational value to gain access to federal financial aid. Common criticisms include such concerns as: recommendations for improvement in self-studies or team reports are forgotten or ignored, team reports are considered highly subjective and not evidence-based, and decisions are inconsistent and subjective. As has been said of democracy itself — paraphrasing Churchill — accreditation may be flawed but remains the best available alternative. That said, accreditation has been effective in holding institutions accountable to general standards of organization and functioning, but the balance between self-regulation and federal oversight has increasingly shifted to more federal oversight, especially with respect to outcomes.

More recently, accreditation is increasingly criticized for failing to address key issues of quality and effectiveness with respect to uneven and unsatisfactory completion rates, uneven learning results, lack of alignment with changing employability skills and workforce needs, and the failure of graduates in many programs to achieve earnings beyond those of high school graduates. To many, the accrediting system is no longer fit for purpose and in urgent need of structural reform. This has been the view of the Trump Administration, which issued an Executive Order titled "Reforming accreditation to strengthen higher education" on April 23, 2025. In response, the Department of Education launched an Accreditation, Innovation, and Modernization (AIM) negotiated rulemaking process, which reached consensus in May 2026 on a proposed regulatory framework to expand federal oversight of accrediting agencies.

Over the years, and in the face of episodic calls for reform, accreditation has survived by adapting, periodically revising its standards and processes. For example, it shifted during COVID to allow virtual reviews, and with increasing calls for accountability around completion rates, earnings following graduation, and debt load, accreditation is beginning to shift again toward a more data-centric process of evaluation. As described below, AI is likely to accelerate this shift and offer ways in which judgments about quality move from "standards" to "metrics".

Accreditation is organized around standards, which serve as the reference point for self-studies, team reviews, and all decisions. These standards reflect a traditional view of how institutions of higher education should be structured, governed, and operate. They are modeled on a mid to large-scale university with an independent governing board, a substantial core of full-time faculty leading curricular design and instruction, and are typically based on what used to be a residential institution. While many ad-

aptations have been made to accommodate online and hybrid institutions, specialized single-purpose institutions (such as art colleges, schools of psychology, etc.), the core mindset for defining institutional structure and quality is embedded in accrediting standards and federal law.

In the U.S., federal law in CFR Section 602.16 identifies areas in which standards are required. These include:

- Success with respect to student achievement in relation to the institution's mission;
- Curricula;
- Faculty;
- Facilities, equipment and supplies;
- Fiscal and administrative capacity;
- Student support services;
- Recruiting and admissions practices;
- Measures of program length and objectives of the degree or credential offered;
- Record of student complaints;
- Record of compliance with responsibilities under Title IV of federal law regarding financial aid.

As accrediting agencies formulate standards in each of these areas, they need to be drafted in such a way as to serve multiple purposes, depending on the type of institution to which they are being applied. Are they minimums to be met, based on clear quantitative or qualitative measures or results? Or are they normative expectations of how an institution should be organized in terms of structures and processes that traditionally accredited institutions typically have? Or are they aspirational, setting forth new areas to lead and promote change?

Part of the confusion around standards is that, depending on the situation, they can be any or all of these types. As a result, standards are typically written in general language to be applied to a wide range of institutional types: in terms of mission, size, organization, curricular emphasis, or other characteristics.

For example, a new institution seeking accreditation must demonstrate that it meets the agency's standards at a minimum level. Innovative or nontraditional institutions must demonstrate they meet these standards, either by "meeting their intent" or by an alternative approach. This leads to the concern that these standards limit innovation and creativity. For well-established institutions, the focus of the accrediting review process typically shifts from "compliance" to "improvement", based on the assumption that minimum standards are met and exceeded.

This adaptability in interpretation has served institutions and accreditation well, given the focus on each institution individually and the evaluation of each institution under its own terms. Many innovative and nontraditional institutions have been accredited by such adaptability. For example, if an institution with a student body of 1000 has only three full-time faculty and relies extensively on adjunct faculty for course delivery, is that sufficient to meet a standard requiring a "core of full-time faculty"? The answer, of course, is "it depends", and peers at each level of the accrediting process must review the circumstances in that or similar areas under the standards where no single quantitative measure fits the multiple configurations of each institution. Typically, the iterative (though time-consuming) nature of the accrediting process allows for issues of meeting the standards to be explored and resolved, and accreditors are under increasing public pressure to allow, if not promote, innovation and change in the institutions they accredit.

Lacking clear quantitative or qualitative benchmarks, accreditation standards have historically treated inputs and processes as proxies for quality — on the assumption that sound practices produce sound outcomes. Some common expectations have included requiring faculty with advanced degrees, support for research, processes for curriculum development and review, or programs for assessing student learning, to name but a few.

Without quantitative measures or clear “bright lines”, accreditation has relied heavily on the judgment of peers drawn from similar institutions that can understand the institutional type and typical issues of that type (e.g., size, public or private, faith-based or specialized, etc.). The fact that accreditation is a peer-driven process has led to trust and willing participation by institutions in the process, even with all its concerns and faults.

The recent availability of a substantial body of outcomes data, however, shows wide variability in outcomes even when good practices are in place, calling into question the substantial resources invested in these processes.

As more federal and state data on outcomes have become available, highlighting institutions and programs with low completion rates or uneven earnings after graduation, among other areas, the conversation on what accreditation should be focusing on is shifting to demonstrating positive metrics.

With this data becoming more public, there is the risk that conclusions will be drawn regarding quality and effectiveness, and judgments made, without adequate analysis and contextualization. Yet, it cannot be denied that many accredited institutions and programs have outcome results that are questionable or even unacceptable if confirmed upon further investigation.

What is clear is that the field is shifting from determining quality primarily through professional and peer judgment toward defining it through measurable, data-driven outcomes. As described below, AI can be used to transform the ritual of self-study and other aspects of the accrediting process into a more efficient and focused process, and provide new ways to address the focus on metrics.

From individual institutional judgments to comparative benchmarking

At the core of accreditation practice is the evaluation of each institution individually, based on its mission and defining characteristics. The rationale was to respect institutional diversity and shield institutions from unfair comparisons across different institutional types. Thus, a small regional public university would be reviewed differently from a major elite private research university. As such, formal comparisons are not made to other institutions despite the fact that there are so many similarities across institutions. Instead, institutions are evaluated on their own terms, and through the peer review process, reviewers bring their varied experiences to the evaluation and may make recommendations for improvement based on those experiences, but historically, no public comparisons were made between the performance of one institution against another, at least until now.

Accrediting agencies are the repositories of tremendous amounts of quantitative and qualitative data about nearly all aspects of institutional functioning amongst their members through the collection of annual reports, self-studies, and other reports. This model was designed for another era when each institution worked to create distinctive, even unique, characteristics to fulfill its mission and define itself internally and to the public. With AI, the possibility of benchmarking performance in key areas is already a

reality, providing the foundation for evaluating institutions not only on their own terms, but in relation to national and peer standards of performance.

Institutional and program-level data are now available from the Integrated Postsecondary Education Data System (IPEDS), College Scorecard, and state longitudinal data systems tracking labor market outcomes and workforce readiness and success. The application of AI to these data sets has made the focus on metrics more visible and relevant to the public than compliance with traditional standards. Third-party aggregators now make it possible to track students, individually and by cohort, from enrollment through their university experience and post-graduation (or upon leaving without a degree) with enough reliability to make judgments about quality and rigor far more accurately and continuously than episodic traditional accrediting reviews can do. Data analyses can be conducted at a scale and timeliness not possible by human reviewers and can identify patterns not readily visible through traditional accrediting reviews.

AI algorithms can define peer groups, create benchmarks for each institution in key areas, and compare them with like institutions and with national averages. This can be done rapidly and scaled to apply not only across accrediting membership but across regions or the entire country. Given the diversity of institutional types, it may not always be possible to set a single quantitative benchmark applicable to all institutions on a given metric, such as completion rates, though some policymakers call for exactly this, but AI makes it feasible to compare performance against institutions with similar characteristics.

Thus, an open-enrollment institution might have a six-year completion rate of 28%. Is this good or bad? Using the traditional peer review model, an accrediting team might accept this outcome given the nature of the student body, economic circumstances, and other contextual factors. Using benchmarking data, however, it becomes possible to compare this institution against others with similar characteristics; if the average among those peer institutions is well above 28%, there is a principled basis for identifying this as a low figure and requiring the institution to investigate what better-performing peers are doing differently and take steps to improve. AI can also assist in identifying what specific practices peer institutions use to achieve stronger outcomes.

Several accreditors have already moved in this direction. The Accrediting Commission for Community and Junior Colleges (ACCJC) publishes data and return-on-investment metrics for each institution, including how long it takes graduates to recover tuition and fees through post-completion earnings as well as student achievement data. ACCJC does not yet systematically benchmark institutions against each other or publish peer group comparisons, though institutions can make their own cross-institutional comparisons using the published data.

The Higher Learning Commission (HLC) has integrated benchmarking into its accrediting criteria and effective with its 2025 Criteria for Accreditation, now formally requires institutions to benchmark student success outcomes with reference to peer institutions. In its directory of institutions, HLC publishes links to data from the College Scorecard, College Navigator, and institutional outcomes pages.

The most robust public benchmarking system has been developed by the WASC Senior College and University Commission (WSCUC), which has been building a dashboard of outcomes and comparative data since 2021. Called the Key Indicators Dashboard (KID), it is interactive and regularly updated, allowing users to explore metrics including enrollment, retention, graduation and degree attainment, student debt, earnings, and institutional finances.

WSCUC recently added a Price-to-Earnings Premium (PEP) metric into KID at the program level. The PEP metric reveals how long it takes graduates of different programs

to recover their education costs, allowing students, institutions, policymakers, and other stakeholders to compare recovery times both within institutions and against national averages. In addition, KID now organizes program-level return-on-investment data by field of study, enabling users to compare discipline-specific graduate outcomes across institutions against national benchmarks. WSCUC also publishes a Student Outcomes Overview (SOO) focused on key metrics relating to student progress, completion, and post-graduation results. While the SOO was initially optional, WSCUC now requires all qualifying institutions to incorporate KID and SOO data into their self-reviews, assess their results against peer and national benchmarks, and develop improvement plans where needed.

What is still developing is how the institutional self-evaluation and team review of this data will translate into changes in accrediting actions. Using AI to evaluate and analyze institutional data in this way gives accreditation a more substantive basis for answering the question “what is good enough?” grounded in both institutional context and comparative data. More importantly, it is hoped that focusing on actual, real-time results in key areas will drive meaningful improvement. Not all students are captured by federal and state data systems, and institutions are still learning how to work with the volumes of data now made public by governments and third parties. Additionally, most accrediting websites remain oriented toward institutional members rather than the general public, emphasizing standards, processes, and procedural information over accessible performance data. Over time, however, college guides, rating systems, and ranking services will integrate this data into their public communications, leading institutions to present– and sometimes defend — their performance metrics.

A 2026 Ithaka S+R study examining how accreditors are using new outcomes data concluded that the adoption of data-driven practices is associated with only modest, short-lived improvements in student outcomes. Specifically, the study found:

Accreditors emphasize student outcomes in their standards and public materials, but their performance expectations for institutions remain uneven and at times underspecified... [and] there is limited evidence that they use the data to make decisions. Transparency alone does not appear to produce sustained improvements in student outcomes (Ithaka S+R, 2026, para. 3).

It takes time and experience to ensure that data are accurate, appropriately contextualized, and effectively deployed in improvement efforts. Lessons from the last two decades show that improving outcomes requires extended time, resources, and sustained commitment — progress cannot be made in just a semester or two — though models of success do exist and need to be more broadly emulated.

That time may be running out. Frustration over what is seen as the failure of both higher education and accreditation to address these outcome metrics has driven recent legislative and regulatory action. The One Big Beautiful Bill Act (OBBBA), signed into law on July 4, 2025, requires all higher education programs receiving federal financial aid to meet minimum earnings benchmarks measured four years after graduation. Undergraduate programs lose loan eligibility if graduates’ median earnings, measured four years after completion, fall below the median earnings of same-state high-school-educated workers ages 25–34; graduate and professional programs lose eligibility if graduates earn less than comparable bachelor’s-degree holders. Programs formally fail if they miss the benchmark in two out of three years.

This earnings-based accountability framework raises important questions about the sustainability of programs in the arts, early childhood education, social work, and many other fields where graduates frequently choose to enter lower-paying service careers

or work independently. The implications for accreditation are direct: accreditors will need to engage with program-level earnings data as part of their oversight function, and AI-enabled benchmarking will be essential to doing so at scale.

In addition, in the recently concluded negotiated rulemaking process, consensus was achieved on modifications to many regulations affecting higher education institutions and accreditation. Changes to Section 602.17 of the Education Code were made as follows (changed in italics):

§ 602.17 Application of standards in reaching accreditation decisions. The agency must have effective mechanisms for evaluating an institution's or program's compliance with the agency's standards before reaching a decision to accredit or preaccredit the institution or program. The agency meets this requirement if the agency demonstrates that it —

(a) Evaluates whether an institution or program—

(1) Maintains clearly specified educational objectives, *which may include credit for prior learning*, that are consistent with its mission and appropriate in light of the degrees or certificates awarded *that are developed and regularly reviewed and updated, using reliable data (which may include Federal or state data)* (U.S. Department of Education, 2025, § 602.17).

The Department of Education initially attempted to insert regulatory language that would have required accreditors to define “minimum expectations of student performance” for each institution’s programs, which would have been a dramatic mandate for accreditors and usurp the role of faculty in defining minimum standards. The adopted consensus language still needs to be made public for comment and possible changes over the coming months, but these negotiations show the government’s desire for more clearly defined performance standards. AI will certainly play an ever-increasing role in using data to regularly update performance against each institution’s programs.

From self-study to continuous data monitoring

Self-study is the cornerstone of the accrediting process, a vehicle for periodic institutional reflection on mission, purpose, and quality, organized around an accreditor’s standards. Typically, committees are formed for each standard, sometimes with subgroups for subareas, involving dozens of faculty and staff, the institutional research office, and many others, along with an editor to synthesize the reports into a coherent whole.

This model was developed when institutions were smaller and less complex. It works well at small to medium-sized institutions, where reaching consensus on how the institution addresses each standard is feasible and the process serves as a useful exercise in community building and reflection, though costly in time and resources. At larger institutions, with multiple schools, more than a hundred degree programs, and thousands of students in residence or online, creating a coherent institutional portrait is far more challenging. The focus consequently shifts to institution-wide processes such as strategic planning, faculty governance, academic planning, program review, and student learning outcomes assessment. Where compliance is not in question, many accreditors permit a focus on themes or improvement priorities rather than comprehensive standards review.

Increasingly, self-studies incorporate institutional data, but many remain largely textual documents without deep analysis of available internal or benchmarking data. The process is still primarily human-driven, episodic, and resource-intensive.

AI has the potential to transform multiple elements of accrediting review. At a minimum, AI can review institutional and program policies to verify conformance with applicable federal and state laws and regulations, a compliance function required by all accreditors that can be automated, checked, and kept current as laws change. This is already underway: the National Accreditation Commission (NAC), a startup accreditor for workforce and short-term certificate programs, has built this capability into its AI Hub platform. Similar tools can be extended across all program types and institutional sectors.

Beyond compliance, AI platforms can maintain continuous monitoring of key performance indicators — academic outcomes, student progression, financial health, assessment of student learning, and other trigger points — enabling early intervention and surfacing trends relative to institutional goals and peer benchmarks. Platforms such as Civitas Learning and EAB's Navigate already offer AI-powered predictive analytics, generating real-time risk scores for individual students and cohorts and forecasting multi-year outcome trajectories. These are powerful tools that are establishing a new norm: accrediting reviews will increasingly be part reflection, part data analysis, and part identification of improvement priorities, all supported by AI analytics.

The need for AI-enabled monitoring will only intensify as accrediting standards and federal rules place greater weight on performance metrics and, eventually, minimum standards. That said, the human element of periodic reflection remains indispensable: accreditation at its best creates space for institutions to examine their place in higher education and society, using performance data to inform rather than displace that conversation. The Middle States Commission on Higher Education (MSCHE) has addressed this directly in its July 2025 Use of Artificial Intelligence Policy, which prohibits self-studies from being prepared entirely by AI and requires that any AI-generated content be reviewed and verified by institutional personnel. This is a reasonable and important guardrail as institutions, accreditors, and peer reviewers collectively learn to work alongside AI tools. The Council of Regional Accrediting Commissions (C-RAC) reinforced this spirit in its October 2025 statement affirming that AI use in learning evaluation is fully compatible with accreditation standards, while underscoring that accreditation judgment remains a human responsibility.

From faculty curriculum control to AI-assisted curricular analysis

At the core of any institution is its faculty and their responsibility for designing and overseeing curricula. Accreditors have long required institutions to define program outcomes and assess their achievement through direct and indirect measures of student learning. Each institution is free to design curricula in alignment with its mission and distinctive characteristics, using program review to maintain currency and effectiveness. In practice, however, faculty exercise this responsibility with considerable variability — course offerings sometimes reflect departmental preferences or individual faculty interests more than curricular coherence, and enrollment patterns can cause delays in student progression or sustain low-enrollment courses at significant cost. Faculty oversight of curricula, moreover, too often focuses on inputs and resources rather than outcomes. The public conversation about quality has moved decisively toward out-

comes — student learning, skill development, employability, and earnings — and new federal accountability requirements will sharpen this focus further.

AI can significantly expand the scope and depth of curricular review. By integrating data from within the institution with comparisons to peer and aspirational programs, AI can highlight gaps and strengths that would not be visible through traditional committee-based review. Specifically, AI can synthesize stated program and course outcomes, pedagogical approaches, assessments at the course and program levels, enrollment and progression data, completion rates, grading patterns, placement outcomes, and student and employer satisfaction, across an entire institution, continuously, and at a scale no faculty committee could match.

Grade inflation is one area where AI-enabled analysis offers particular promise. Machine learning systems can review grading results within and across courses and programs, identify patterns, and evaluate student work against stated course and assignment outcomes to validate grading consistency and rigor. Grading moderation using second opinions or external review is common in the United Kingdom and Australia and is already being supported by AI. Woolf University, for example, uses AI to review all student work against course outcomes and grading rubrics to verify consistency and academic integrity across its distributed faculty network. If adopted more broadly and required by accreditors, such approaches could bring greater validity to grading and provide a principled response to grade inflation concerns.

There are real risks to navigate. AI-driven curricular analysis could promote homogenization, incentivize institutions to optimize for metrics rather than genuine rigor, or reduce the space for pedagogical experimentation and innovation. This is precisely where accreditors have an important role: holding institutions accountable for using AI-generated insights to improve rigor, efficiency, and effectiveness in contextually appropriate ways, rather than simply conforming to averages.

The use of AI in curricular analysis will almost certainly shift the role of faculty from serving as the primary, or sole, stewards of curriculum design and oversight to working alongside and with AI-generated analyses to make judgments about needed changes. Standard accrediting requirements regarding faculty governance and curricular oversight will need to adapt accordingly, recognizing that faculty authority now operates in relation to AI-based analyses of progression, retention, completion, grading, earnings, and curricular coherence. Just as institutional research offices became essential partners in generating data for faculty deliberation — with mixed success in actual data use — AI offices and third-party analytics providers will play increasingly powerful roles in shaping what data is collected and how it is interpreted.

The public availability of much of this data, already analyzed and sometimes converted into ratings or critical profiles of institutions, creates new communication imperatives. Institutions will need to contextualize their results, explain comparisons to peers, and manage how their performance data is presented publicly. The stakes are high: under the One Big Beautiful Bill Act's earnings accountability framework, loss of federal financial aid eligibility — likely meaning program closure — could result from sustained poor earnings outcomes. A purely economic measure of quality, however, poses a serious threat to programs in the liberal and performing arts, education, social work, and other fields that develop graduates who choose careers of social contribution over maximum earnings.

A further dimension of the federalization of higher education quality deserves mention. Both the White House Executive Order on accreditation and the new regulatory framework require accreditors to address viewpoint diversity and freedom of expression in their standards and review processes. AI is likely to be deployed to assess in-

stitutional performance in these areas: analyzing policies, faculty composition, course content, and campus climate data. These are genuinely new and contested territories for accreditation, fraught with potential conflict between accreditors, institutions, and the federal government, and their implementation will require considerable care.

From periodic site visits to AI-mediated continuous monitoring

Accreditation rests on the fundamental premise that peers are best positioned to judge institutional quality, combining document review with on-site inquiry through structured interviews and observation. Review teams are carefully constructed to include members from similar institutions who understand the mission and context of the visited institution, pairing generalists with subject matter experts. A typical site visit involves meetings with senior administration, faculty and staff leaders, key committees, trustees, and students. Working from self-study, teams develop and test impressions across the visit and render judgments on how well the institution's processes, i.e. governance, planning, program review, student learning assessment, budgeting, are functioning.

Honest appraisal reveals the limits of this model. A site visit is not a comprehensive institutional review: it cannot cover all programs, all committees, or all constituent groups. It is a well-honed sampling exercise, and undertaken every five to ten years, it cannot track the many changes that occur in between. During COVID, accreditors pivoted to virtual visits; today there are numerous variations like hybrid visits, preliminary offsite reviews, shortened compliance tracks, thematic reviews, but all retain peer judgment at the center of the process.

The growing availability of public and institution-generated data is changing this picture. An agency like WSCUC, through its Key Indicators Dashboard, can now monitor key dimensions of institutional performance in real time, with algorithms flagging gaps or trends that warrant deeper review. This kind of continuous, multi-indicator monitoring goes far beyond the traditional annual institutional report and fundamentally re-frames what a comprehensive visit is for. Rather than attempting a broad survey of institutional functioning, a visit can focus on the issues surfaced by continuous monitoring, on deeper reflection about the institution's future in a rapidly changing environment, and on evaluating how well the institution is responding to and using its data.

The role of peer reviewers is evolving rather than diminishing. New specializations are emerging: evaluating how institutions deploy AI; assessing the validity of algorithmically generated analyses; examining privacy and data governance frameworks; and scrutinizing how institutions contextualize and present their performance data to avoid misleading conclusions. Human judgment remains indispensable — and arguably more important than ever — precisely because institutions of higher education are multidimensional organizations serving diverse constituencies in the pursuit of knowledge generation, student development, and civic participation. Performance metrics capture important dimensions of quality, but they cannot substitute for the kind of contextual, relational judgment that an experienced peer brings to an evaluation. The human center of accreditation must be preserved even as its tools and methods evolve.

From legacy accreditors to an expanding accreditation marketplace

The United States accreditation landscape is both extensive and, until recently, remarkably stable. According to the Council on Higher Education Accreditation (CHEA), the country is served by 19 recognized institutional accreditors and approximately 70 programmatic accreditors. The largest and most established institutional accreditors are the seven historically regional bodies, organized as the country developed to provide quality assurance within defined geographic boundaries. Over recent decades, institutions have operated well beyond those boundaries, physically and online, and in 2019, federal regulations removed the geographic restriction recharacterizing the agencies as national in scope. Institutions may now seek accreditation from any recognized agency, and the regulatory framework facilitates accreditor switching under defined conditions. One stated purpose of this change was to promote competition and allow institutions to align with accreditors more closely matched to their missions and approaches.

Despite this opening, concerns persist that legacy accreditors constrain innovation and that the field needs new models that actively foster experimentation. This view was explicitly stated in the April 2025 White House Executive Order on accreditation. In response, the Department of Education revised its recognition regulations to create faster, less onerous pathways for new agencies, and the Fund for the Improvement of Postsecondary Education (FIPSE) allocated \$7 million in November 2025 specifically to support the creation of new accrediting agencies. In January 2026, FIPSE awarded grants to the Postsecondary Commission (PSC) and the Commission for Public Higher Education, each receiving \$1 million toward building outcomes-focused accreditation models.

The Postsecondary Commission is partnering with Mathematica to develop a rigorous, metrics-based accreditation model centered on economic outcomes, piloting its approach with Texas institutions and using state longitudinal earnings data to validate its methods.

The Competency-Based Education Network (C-BEN) has also received FIPSE support to develop a skills-based accrediting model with the stated goal of establishing a new agency focused on employer-verified competencies, performance-based assessment, mastery documentation, labor-market alignment, and employment and mobility outcomes.

A third emerging entity, the National Council for Workforce AI Program Accreditation (NCWAIPA), is developing the first programmatic accreditor specifically for short-term and certificate programs in artificial intelligence.

Operating outside the FIPSE grant framework but perhaps the most immediately consequential new entrant is the National Accreditation Commission (NAC), which has built its AIHub platform—powered by AristAI and supported by the GitLab Foundation, Ballmer Group, and OpenAI—to evaluate the quality of short-term and certificate workforce programs. The AIHub automates policy and performance compliance review, verifies Workforce Pell eligibility, validates alignment with employer demand, and coordinates with state agencies to reduce duplicative regulatory burden. In February 2026, NAC announced a landmark partnership with three established institutional accreditors — MSCHE, WSCUC, and NWCCU, collectively serving more than 900 institutions and 6.3 million students — under which those agencies will accept NAC’s AIHub-based quality reviews for workforce programs at their member institutions. This represents a structural innovation: using AI-enabled accreditation infrastructure to bridge programmatic and institutional quality assurance at scale.

It will take several years for new agencies to develop, pilot, and seek formal federal recognition under the revised regulatory framework. However, their emergence signals

a meaningful diversification of the accreditation marketplace — one oriented toward outcomes, efficiency, and AI-enabled review processes. The NAC partnership model may prove the most influential early template, demonstrating how new AI-native quality assurance infrastructure can work alongside legacy accreditors rather than simply competing with them.

Conclusion: from adaptation to transformation

Accreditation has outlasted many predictions of its demise. It has survived the introduction of the computer, the internet, the rise of online learning, and the disruptions of COVID-19. Each time, it adapted by revising its standards, adapting its processes, and finding a way to remain relevant. AI presents a challenge of a different character, not because it is more threatening, but because it is faster and more pervasive than anything that has come before. The pace of change may well outrun the capacity of both institutions and accreditors to respond thoughtfully.

That said, the direction of change is becoming clearer. Standards built around inputs and processes are giving way to expectations centered on outcomes. Comparative benchmarking is replacing the purely individualized institutional review. The periodic self-study is being supplemented, and in some respects supplanted by continuous, data-driven monitoring. Site visits are being refocused on the issues that data surfaces rather than attempting a broad sampling of institutional functioning. Faculty governance of curricula is being challenged to engage with AI-generated analyses. And a new generation of accreditors is entering the field with outcomes-first, AI-enabled models that may reshape the competitive landscape.

Managing the changes ahead will require clear-eyed and collaborative engagement of institutions, state and federal government agencies, and accreditors. AI is genuinely powerful. It can identify patterns in institutional data that no peer review team could detect, benchmark performance at a scale and speed no human process can match and predict trends that give institutions time to be proactive. Used well, it can make accreditation more rigorous, more continuous, and more credible as a public accountability mechanism.

The risks deserve equal candor. The algorithms that drive AI systems are often opaque — even to their designers — and their outputs can reflect biases embedded in the underlying data. An accountability framework built entirely on earnings metrics, as the One Big Beautiful Bill Act's earnings benchmarks impose, will systematically disadvantage programs in teaching, social work, the arts, and the humanities — precisely the programs that develop the civic capacities a democratic society depends on. And the use of AI to monitor viewpoint diversity and faculty expression on campus, as called for in new regulations, introduces forms of oversight that could easily curb the academic freedom that makes universities worth accrediting in the first place.

What I find myself returning to, after decades of work in and around accreditation, is a conviction that the peer review model endures not because it is efficient, but because it brings something to the evaluation of quality that data alone cannot: judgment, context, and a shared professional commitment to the integrity of higher education. The task ahead is not to choose between AI and peer review, but to integrate them in ways that make each stronger. AI can free peer reviewers from laborious data gathering and allow them to focus on the harder questions: Is this institution doing what it says it is doing? Is it serving its students and its community well? Is it prepared for what comes next?

Will accreditation survive the upheavals ahead? I believe it will — not unchanged, but recognizable. Transformed but still functional and necessary. Institutions need the legitimacy that accreditation confers, and the public needs an independent check on institutional quality that neither rankings nor government regulation alone can provide. The new agencies emerging from the FIPSE grants and the NAC partnership model suggest that the field is capable of innovation. The challenge is to ensure that innovation serves students and the public interest, not merely efficiency or the political priorities of any given moment. That has always been accreditation's central tension, and it remains so now.

References

- Accrediting Commission for Community and Junior Colleges. (2025, August 7). *ACCJC launches student achievement dashboards and ROI metric*. <https://accjc.org/announcement/student-achievement-dashboards-roi-metric/>
- American Institute for Research. (2025, September). *Reimagining higher education accountability*. AIR.
- Childress, S., Pisacreta, E., & Ward, J. D. (2026, February 12). *Understanding the impact of data-driven accreditor practices on student outcomes*. Ithaka S+R. <https://doi.org/10.18665/sr.324688>
- Civitas Learning. (2025). *AI-powered student success: How Civitas Learning uses predictive analytics*. <https://www.civitaslearning.com/blog/ai-powered-student-success/>
- Council for Higher Education Accreditation. (2025). *Institutional accrediting organizations*. <https://www.chea.org/institutional-accrediting-organizations>
- Council of Regional Accrediting Commissions. (2025, October 6). *C-RAC statement on the use of artificial intelligence (AI) to advance learning evaluation and recognition*. <https://www.msche.org/2025/10/06/c-rac-releases-statement-on-the-use-of-artificial-intelligence-ai/>
- Fruition Group. (2026, February 7). *NCAIWPA: Revolutionizing AI workforce training accreditation*. <https://www.ncaiwpa.org/post/ncaiwpa-revolutionizing-ai-workforce-training-accreditation>
- Hegji, A. (2020, October). *An overview of accreditation of higher education in the United States* (CRS Report No. R43826). Congressional Research Service.
- Higher Learning Commission. (2025). *Student success outcomes initiative*. <https://www.hlcommission.org/about-hlc/initiatives/student-success-outcomes/>
- Inside Higher Ed. (2026, January 6). *New accreditors, civic discourse programs win FIPSE grants*. <https://www.insidehighered.com/news/government/2026/01/06/new-accreditors-civic-discourse-programs-win-fipse-grants>
- Ithaka S+R. (2026, March 10). *Accreditation, outcomes, and accountability*. <https://sr.ithaka.org/blog/accreditation-outcomes-and-accountability/>
- Middle States Commission on Higher Education. (2025, July 2). *New use of artificial intelligence accreditation policy and procedures*. <https://www.msche.org/2025/07/02/new-use-of-artificial-intelligence-accreditation-policy-and-procedures/>
- Middle States Commission on Higher Education. (2026, February 27). *Four national accreditors unite to accelerate support for workforce education and state economic development*.

<https://www.msche.org/2026/02/27/four-national-accreditors-unite-to-accelerate-support-for-workforce-education-and-state-economic-development/>

Moultrie, T. (2023, October). *The accreditation system isn't working*. The Century Foundation.

National Accreditation Commission. (2026). *AIHub: AI-enabled workforce education quality assurance*. <https://nationalaccreditation.org/>

National Center for Education Statistics. (2024). *IPEDS data center*. U.S. Department of Education. <https://nces.ed.gov/ipeds/>

One Big Beautiful Bill Act, Pub. L. No. 119-21, 139 Stat. (2025).

Trump, D. J. (2025, April 23). *Executive order on reforming accreditation to strengthen higher education*. The White House. <https://www.whitehouse.gov/presidential-actions/>

U.S. Department of Education. (2023). *College Scorecard*. <https://collegescorecard.ed.gov/>

U.S. Department of Education. (2024). *Federal Student Aid FY2024 annual report*. <https://studentaid.gov/sites/default/files/fy2024-fsa-annual-report.pdf>

U.S. Department of Education. (2025, November 12). Fund for the improvement of post-secondary education special projects (FIPSE-SP): Applications for new awards. *Federal Register*, 90(218). <https://www.federalregister.gov/documents/2025/11/12/2025-19843/>

U.S. Department of Education. (2026, May 21). *U.S. Department of Education reaches consensus to reform and strengthen America's higher education accreditation system* [Press release].

U.S. Department of Education. (2026). Accreditation and student outcomes [Amended regulation]. 34 C.F.R. § 602.17.

WASC Senior College and University Commission. (2021–2026). *Key indicators dashboard (KID)*. <https://www.wscuc.org/resources/kid/>

WASC Senior College and University Commission. (2025, August 6). *WSCUC expands program-level ROI metrics in public dashboard*. <https://www.wscuc.org/>

WASC Senior College and University Commission. (2026, February 10). *WSCUC enhances public dashboard with field of study ROI view*. <https://www.wscuc.org/>

WASC Senior College and University Commission. (2026, March 4). *WSCUC partnership to support workforce program approval and funding*. <https://www.wscuc.org/post/wscuc-partnership-to-support-workforce-program-approval-and-funding/>

Woolf University. (2024). *Woolf University*. <https://woolf.university>

Bridging Chapter: Operational Mapping — The AI Ecosystem Map for Universities

Vadim Novikov

Madeleine was a 22-year-old final-year nursing student at Australian Catholic University when she received an academic integrity notice during an emergency ward placement. The trigger was an AI-writing detection signal. What had been marketed as an institutional tool for protecting academic honesty became, in practice, a high-stakes judgment about a student's competence, integrity, and future. The problem was not only whether the detector was accurate. The deeper problem was that a probabilistic tool had been placed inside an institutional process that needed evidence, due process, and an understanding of what the student could actually do.

This case is a useful entry point because it shows why universities do not need another catalogue of AI tools. They need an *ecosystem map*. A list says what tools exist. An ecosystem map asks which university function is being changed, what evidence the tool provides, what hidden capacity it requires, what failure mode it creates, and who is accountable when the tool is wrong.

Three traps appear repeatedly in university AI adoption. The first is tool-first thinking: the institution begins with a product rather than a task. The second is elite mimicry: a mid-tier or resource-constrained university tries to imitate the stack, governance structure, or innovation language of MIT, ASU, Toronto, or another institution with very different capacities. The third is hype capture: a sequence of pilots creates the appearance of transformation without producing scalable institutional learning.

These traps are not only diagnoses after failure. They are questions to ask before adoption: What function are we trying to improve? Which elite practice are we imitating, and do we have its supporting capacity? What will we learn from this pilot that can survive beyond the pilot?

The chapter starts from a deliberately plain rule: function-first, not technology-first. The unit of analysis is not the chatbot, detector, CRM system, tutoring platform, or discovery tool. It is the university function: teaching, assessment, student support, research, operations, infrastructure, data governance, and the institutional routines that connect them. AI matters because it changes how these functions are performed. It does not remove the conditions under which those functions succeed.

This distinction is particularly important for mid-tier universities. Elite institutions can often absorb failed experiments: they have staff, legal support, data teams, teaching centers, procurement offices, and reputational buffers. A constrained university has less room for theatrical innovation. A bad tool can consume scarce staff time, create a fragile dependency, expose student data, distort assessment, or redirect money away from more basic needs. For such institutions, a list of promising tools is not enough. The decision has to begin with the university's own bottlenecks.

The ecosystem metaphor is useful because tools interact. A detector changes assessment policy. A tutoring platform changes teacher workload and student expectations. A CRM system changes advising and data governance. A discovery tool changes research

supervision and citation practice. A local model deployment changes infrastructure and procurement. No major AI tool remains only by itself once it enters the university. It becomes part of a system of incentives, evidence, responsibility, and failure.

What existing taxonomies already know

The field has not ignored classification. Several taxonomies already help universities think about AI in higher education. Zawacki-Richter et al. (2019) provide one of the most influential higher-education-specific maps. Their systematic review organizes AI applications into four broad areas: profiling and prediction, intelligent tutoring systems, assessment and evaluation, and adaptive systems and personalization, with seventeen subcategories. Crompton and Burke (2023) update this functional mapping through five usage codes: assessment/evaluation, predicting, AI assistant, intelligent tutoring systems, and managing student learning.

Baker, Smith, and Anissa (2019) for Nesta provide a role-based distinction between learner-facing, teacher-facing, and system-facing tools. Governance frameworks from Jisc, EDUCAUSE, UNESCO, UCL, NIST, and others add a different question: under what conditions should a use be permitted, restricted, monitored, or prohibited?

Existing taxonomies make the field more legible by classifying AI applications according to function, user group, or governance category. Their limitation is not conceptual but operational: classification alone does not tell an institution whether a tool should be adopted, restricted, redesigned around, or avoided. The gap appears between category and routine, where questions of capacity, autonomy, accountability, and implementation determine outcomes.

Four blind spots recur across functional, role-based, and governance classifications. First, they often lack an operational layer: who will run the tool, maintain it, interpret its output, train users, handle appeals, and revise the process? This layer is not administrative detail. It is where pedagogical intention becomes institutional routine. Second, they rarely distinguish levels of AI autonomy. A grammar checker, a formative feedback bot, a detector, a predictive risk score, and an autonomous advising agent are not the same kind of institutional actor. Third, they often lack a contextual filter: the same tool has different implications depending on resources, language, data maturity, staff capacity, legal environment, and student preparation. Fourth, generative AI is cross-functional. It does not stay neatly inside “teaching”, “research”, or “administration” — it cuts across all of them.

These blind spots matter because implementation failure usually occurs between the category and the routine. A tool is bought as “student support”, but no one has defined what counts as a successful intervention. A detector is adopted as “assessment integrity”, but no one has designed the appeals process. A research assistant is promoted as “multilingual”, but no one has checked whether it can see local scholarship. A chatbot is launched as “student service”, but no one has determined when a conversation must be escalated to a human. The missing layer is not theoretical sophistication. It is operating discipline.

This chapter therefore builds on existing taxonomies but changes the decision logic. Instead of asking only, “*which category does this tool belong to?*”, it asks: which university function is being altered, under which constraints, at what level of autonomy, and with what governance capacity?

The Proposed Matrix

The ecosystem map has two main axes and two additional filters.

The first axis is functional, dividing university work into seven domains: teaching and learning; assessment and verification of competence; student support; research and knowledge work; operations and administration; core infrastructure; and data governance and ethical infrastructure. These domains differ in risk, cost, maturity, and speed of change. A university may be advanced in AI-supported research while immature in assessment governance or use AI informally everywhere while lacking capacity to audit any of it.

The second axis is contextual: what kind of institution is deciding. Relevant parameters include resources and procurement capacity, language and local corpus, staff capability, data maturity, student cognitive base, and governance vulnerability. This is what universal taxonomies miss: a low-risk productivity tool in one institution may be a high-risk dependency in another. Resources means maintainability, not just budget. Language means local literature visibility and model performance quality, not just interface translation. Staff capability means the concrete ability to redesign tasks and respond to failure, not abstract AI literacy. Data maturity means whether data are clean, governed, and fair enough to support decisions.

The first filter is autonomy. AI uses range from assistive (helping humans write, search, draft) to advisory (recommendations a human may accept) to semi-automated (triggering workflows or nudges) to decisional (effectively determining outcomes). Higher autonomy demands stronger evidence, oversight, appeal, and audit. Critically, autonomy should be assessed by institutional effect, not vendor description. An “advisory” system becomes practically decisive when staff are overloaded or outputs cannot be contested. A detector called a “signal” functions as a gatekeeper if it reliably triggers misconduct procedures.

The second filter is cross-functionality. Generative AI rarely stays in one box. A writing assistant affects learning, assessment, equity, and integrity. A student success platform affects advising, data governance, and resource allocation. A local model deployment affects infrastructure, privacy, and staff capability simultaneously.

The matrix prevents both premature procurement and premature prohibition. A tool too risky as misconduct evidence may be valuable as formative feedback. One that fails as an autonomous tutor may excel as a practice generator. The same tool — ChatGPT, for instance — may be appropriate for brainstorming, restricted in assessment, useful for staff drafting, unsafe for confidential advising, and valuable for known-corpus analysis. Decisions attach to function, context, autonomy, and evidence, not to the brand.

The matrix also imposes a sequence. Before asking “which AI product is best?” a university should ask: which function is failing or newly possible? What autonomy would the tool have in practice? What evidence would show genuine improvement rather than merely accelerated output? Who owns the failure mode? What must be in place before scaling? These questions are less glamorous than adoption announcements, but they are the difference between an AI strategy and an AI shopping list.




Figure 1: AI Ecosystem Map for Universities (OpenAI, 2026)

Seven functional domains

1) Domain 1. Teaching and learning

The near-term opportunity is improving feedback loops, practice, explanation, and revision within existing structures, and not replacing teachers. Quick wins include rubric-based formative feedback, lesson design support, practice-question generation, and AI-assisted revision cycles under human supervision.

Adaptive tutoring demands caution. A 2025 study comparing Khanmigo with Google search in undergraduate physics found no significant difference in learning outcomes — personalization is not magic (Slijepcevic & Yaylali, 2025). Without learning design, teacher integration, and student agency, an adaptive tutor becomes a conversational search engine.

The strategic question is not “Should we buy an AI tutor?” but “Which learning bottleneck are we addressing, and do we have the conditions to change it?” Investment is justified only when the institution can specify the bottleneck, identify success metrics, and keep teachers in control of the pedagogical sequence.

2) Domain 2. Assessment and verification of competence

Generative AI destabilizes the link between artifact and competence. Detector-first discipline is the wrong response. Turnitin itself states it does not determine misconduct; HEA Ireland warns against detectors as sole proof; OIA UK case CS072504 illustrates the procedural risks once detector evidence enters misconduct processes.

The right question shifts from “Did the student write this alone?” to “What can the student do?” This requires a two-lane structure: an open lane where declared AI use is permitted and tasks evaluate questioning, revision, and application; and a secure lane —

oral defenses, in-class problems, code walk-throughs, practical demonstrations — that verifies actual competence.

The practical alternative to detection is evidence architecture: drafts, logs, reflections, checkpoints, and visible student decisions. Selective verification is more realistic and more humane than universal surveillance. A programming course requires code walk-throughs; nursing requires scenario explanation; writing requires revision memos and oral defenses. The goal is not pre-AI purity — it is ensuring AI assistance does not erase the student's own capability.

3) Domain 3. Student support

AI is attractive here because universities have too many students and too little human attention, but predictive tools can quietly turn administrative data into consequential judgments.

Predictive success platforms work only when the institution has clean data, trained advisors, clear protocols, and equity checks. Without these, prediction automates existing blind spots. A model can appear accurate overall while under-identifying risk for specific groups; proxy variables reproduce inequality under a neutral interface.

The operational test is intervention capacity: a risk score has value only if the university can act responsibly on it. Overloaded advisors, generic interventions, biased data, or no appeals process each undermine the tool's value.

Mental-health AI requires even more caution. Tools like Woebot or Wysa may have legitimate use in clinically supervised contexts, but no chatbot response is a care pathway. Clinical oversight, escalation routes, data protection, and liability arrangements must come first.

Low-stakes navigation, appointment routing, FAQ support, and reminders are appropriate when transparent and easy to escalate. The boundary is high-stakes interpretation of vulnerability — once a system claims to assess risk or prioritize intervention, it enters the domain of institutional responsibility.

4) Domain 4. Research and knowledge work

AI can support summarization, coding, data analysis, literature organization, and research administration. Yet this domain is vulnerable to confusing language processing with corpus coverage.

A tool may summarize a local-language PDF well once provided — that does not mean it can discover local literature adequately. Interface multilingualism is not epistemic coverage; databases often underrepresent local journals, institutional repositories, and non-English scholarship. The KazMMLU work illustrates that local context requires its own benchmarks and sources.

The decisive distinction is between known-corpus work and open discovery. Given a bounded document set, AI can summarize and extract themes with manageable risk. Asked to discover literature independently, it may overrepresent English sources, miss local work, or fabricate citations. Strategic investment should therefore focus on institutional knowledge management — curated corpora, local repositories, citation discipline — not just another discovery subscription.

5) Domain 5. Operations and administration

AI can reduce routine administrative work, but only where workflows are already clear. An AI add-on to an ERP is not an operational strategy without answering: What process is slow or error-prone? What data does it need? Who owns the workflow? What

happens when the output is wrong? Unanswered, the university accumulates institutional debt rather than capability.

Quick wins include internal knowledgebase search, drafting routine communications, procurement support, meeting summarization, and controlled analytics. Strategic investments require data governance, process redesign, and staff training.

Institutional memory also matters. Many universities run on tacit knowledge about processes, exceptions, and informal workarounds. AI can help document and retrieve that knowledge, but only if the institution first determines what is authoritative.

6) Domain 6. Core infrastructure

Infrastructure choices should be plural, not ideological. A single institutional “AI solution” is usually a warning sign. The realistic question is which models suit which functions: public tools for low-risk tasks, managed accounts for staff productivity, private environments for sensitive data, and open or local models where language, privacy, or cost makes them necessary.

Building a foundation model is rarely rational for a mid-tier university. What matters is flexibility, because the model market changes faster than university procurement cycles.

7) Domain 7. Data Governance and ethical infrastructure

Data governance is the condition that makes responsible adoption possible, not a decorative layer added afterward. A university needs to know what data enters tools, where it is stored, who can access it, how outputs are logged, and how errors are appealed.

The realistic target is minimum viable governance: defined ownership, a use-case register, risk tiers, procurement review, human oversight for high-stakes uses, staff and student guidance, an appeal mechanism, and a review cycle. It should be small enough to operate and strong enough to matter.

The test is whether governance changes decisions. A policy that does not affect procurement, assessment design, or staff behavior is symbolic. A committee with no authority to stop a tool is theater. A constrained university should prefer a small number of enforceable rules over a comprehensive policy no one can operate.

Prioritization: what to adopt first

The ecosystem map should lead to sequencing. Not all AI use cases deserve the same institutional response. Quick wins are low-cost, low-autonomy, human-supervised tools that improve existing workflows without making high-stakes decisions. Examples include formative feedback support, staff-facing drafting, known-corpus summarization, internal knowledgebase search, and structured AI literacy exercises.

Strategic investments are using cases that can change institutional capacity but require infrastructure and governance. These include assessment redesign, advising systems with equity monitoring, institutional research repositories, data governance, and selective private deployment for sensitive functions.

Restricted or high-risk uses include tools that influence misconduct findings, student risk labeling, mental-health intervention, admissions, grading, or disciplinary outcomes. These require strong evidence, oversight, appeals, and audit. Avoid or wait cases include detector-first discipline, ERP AI add-ons without a use case, mental-health AI without

clinical oversight, discovery tools marketed as a solution for non-English literature visibility, and any institutional subscription bought mainly to signal innovation.

Before adopting any AI system, a university should ask four questions. What function is being changed? What evidence shows the tool works in this context? What hidden capacity does it require? What failure mode becomes unacceptable once the institution adopts it?

These questions should be asked before the pilot, not after it. Pilots often become a way to postpone judgment: the university tries a tool because it is available, then struggles to define success after the fact. A disciplined pilot begins with an exit condition. What would count as success? What would count as failure? Who will decide? What data will be collected? What harm would stop the pilot? Without these answers, the pilot is not experimentation. It is unmanaged exposure.

A practical adoption sequence can therefore be modest but strict. First, identify the function and the pain point. Second, choose the lowest-autonomy tool capable of improving that function. Third, test it in a bounded setting with clear success and stop criteria. Fourth, review evidence with the people who carry the work: teachers, advisors, IT staff, students, librarians, or administrators. Fifth, decide whether to stop, revise, scale, or replace. This sequence is slower than hype, but faster than recovering from a bad institutional dependency.

Anatomy of hype

Personalized learning shows the hype pattern clearly. Adaptive platforms promise to scale individual tutoring and close learning gaps. The more sober version is conditional: they work when task structure, data, feedback loop, and learner agency are aligned. The Khanmigo study is not a verdict against tutoring; it is a warning against assuming that a tutoring interface automatically produces learning gains. For students with fragile cognitive foundations, personalization without agency may produce fluency without understanding.

Administrative replacement is the second case. Here the promise is that AI will dramatically reduce staff needs. But universities are not merely text-processing machines. Administrative work includes authority, exception handling, judgment, compliance, coordination, and institutional memory. AI can reduce routine drafting and retrieval. It cannot safely replace a process that has not been mapped. The decision rule is simple: automate only after the workflow is understood.

AI-content detection is the clearest case because the evidence and policy direction already point away from detector-first governance. A detector can be one signal in a broader process, but it cannot substitute for assessment design. If the institution cannot explain what competence the student has failed to demonstrate, it has not solved the academic integrity problem. It has only moved suspicion into software.

Across the three cases, the promise has a kernel of truth. Tutoring can help. Automation can reduce routine work. Detection can identify suspicious patterns. Yet between the tool and the institutional effect stand conditions like task design, data quality, staff capacity, due process, language context, oversight, and the ability to revise. Hype begins when those conditions disappear from the story.

Much of higher education AI rhetoric has this structure. “AI will personalize learning” leaves out who defines the learning path. “AI will improve retention” leaves out who can intervene. “AI will detect misconduct” leaves out evidence and appeal. “AI will support research” leaves out corpus coverage and citation reliability. “AI will automate

administration” leaves out authority and exception handling. The ecosystem map puts those missing conditions back into the sentence. It turns hype claims into implementation questions.

Horizon 2: Platform logic and minimum viable governance

In the medium term, AI becomes a platform layer entering LMS, CRM, library systems, ERP, research workflows, assessment design, and staff productivity environments. At that point, isolated pilots are no longer enough. The university needs an architecture.

For mid-tier institutions, architecture means deciding which parts of the AI stack must be institutionally controlled, which can remain external, which require private or local deployment, and which should not be adopted at all. The institution must distinguish experimentation from infrastructure.

Minimum viable governance is the practical answer: an owner for each high-stakes use case, a register of AI systems, risk tiers, procurement review, data rules, human oversight, appeal mechanisms, and periodic review. It must also include a language and equity check. A tool validated on English academic writing or North American student data should not be assumed to work in multilingual Central Asian, Russian-language, or Kazakh-language environments. If no evidence exists, that absence is itself a governance fact.

This governance layer is not a barrier to innovation, but what prevents innovation from becoming institutional theater. Without it, AI adoption will be driven by vendors, enthusiasts, and crises. With it, the university can learn from pilots, retire weak tools, scale useful ones, and protect students from avoidable harm.

The key is sequencing: governance must match the level of risk. Low-stakes drafting requires guidance. Student-facing feedback requires pedagogical review. Predictive advising requires equity monitoring. Misconduct evidence requires due process. Mental-health support requires clinical oversight. The higher the stakes, the less acceptable it is to rely on enthusiasm, vendor assurance, or informal practice.

This approach also protects experimentation. When low-risk uses are clearly separated from high-stakes ones, staff and students can explore AI freely — while the institution holds hard lines around decisions affecting grades, discipline, support access, mental health, or student records. A good ecosystem map creates room for learning while preventing the most consequential shortcuts. That balance is the point.

Closing Synthesis

The ecosystem map is a route-planning device. It does not ask whether AI is good or bad in general. It asks what the university is trying to do, under which constraints, with which tools, at what level of autonomy, and with what governance capacity. The most important discipline is to keep the order straight: functions before tools, context before procurement, governance before scale, and evidence before institutional trust. For the universities this book addresses, strategy cannot be copied from elite institutions, vendor roadmaps, or universal frameworks. It must be built from local functions and real constraints.

This does not make the strategy smaller in ambition. It makes it more serious. A university that chooses a few AI uses well, governs them honestly, and learns from their limits is doing more institutional work than a university that announces a broad AI transformation it cannot operate. The map is meant to protect that seriousness. It gives administrators a way to resist vendor pressure, gives teachers a way to connect tools to

learning, gives students a clearer account of what counts as competence, and gives the institution a language for saying no when no is the responsible answer.

AI will change the university. Yet it will not remove the need for institutional judgment. It will make that judgment more visible, more consequential, and harder to outsource.

References

- Baker, R., Smith, L., & Anissa, N. (2019). *Educ-AI-tion rebooted? Exploring the future of artificial intelligence in schools and colleges*. Nesta.
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2025). Generative AI can harm learning. *Management Science*. Advance online publication. <https://doi.org/10.1287/mnsc.2023.03679>
- Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education*, 20(1), 22. <https://doi.org/10.1186/s41239-023-00392-8>
- EDUCAUSE. (2023). 2023 EDUCAUSE horizon report: Teaching and learning edition. EDUCAUSE. <https://www.educause.edu/horizon-report>
- HEA Ireland. (2023). Guidance on AI and academic integrity. Higher Education Authority. <https://hea.ie>
- Jisc. (2024). Generative AI in education: A guide for staff. Jisc. <https://www.jisc.ac.uk>
- Myrzakhmet, M., Abildinova, G., Moldagali, B., Shamoï, P., & Pak, A. (2024). KazMMLU: A benchmark dataset for evaluating large language models in Kazakh. arXiv preprint arXiv:2406.19368.
- National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- Office of the Independent Adjudicator. (2023). Case summary CS072504. OIA. <https://www.oiahe.org.uk>
- Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist*, 57(10), 1510–1529. <https://doi.org/10.1177/0002764213479366>
- Slijepcevic, N., & Yaylali, A. (2025). Leveraging “Khanmigo” generative AI-powered tool for personalized tutoring to learn scientific concepts. *Journal of Teaching and Learning*, 19(4), 155–178.
- Sweeney, C., Lester, J., & Rangwala, H. (2017). Next-term student performance prediction: A recommender systems approach. *Journal of Educational Data Mining*, 9(1), 19–45.
- Turnitin. (2023). AI writing detection: Accuracy, fairness, and how it works. Turnitin. <https://www.turnitin.com>
- UCL. (2023). UCL artificial intelligence (AI) guidance for staff. University College London. <https://www.ucl.ac.uk>
- UNESCO. (2023). Guidance for generative AI in education and research. UNESCO. <https://doi.org/10.54675/PCSP7350>

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? *International Journal of Educational Technology in Higher Education*, 16(1), 39. <https://doi.org/10.1186/s41239-019-0171-0>

**Part II.
Institutional
Experiments
From the Global
Frontier**

Chapter 8. Building AI-ready graduates in a small, fast-shifting economy

Kakha Shengelia

In the autumn of 2022, as ChatGPT captured global attention, we at Caucasus University (CU) faced a familiar dilemma with unfamiliar urgency. Georgia, a nation of 3.7 million people nestled at the crossroads of Europe and Asia, had weathered successive economic shocks, regional instability, and the full force of the COVID-19 pandemic. Our university, like many in small economies, had learned to adapt quickly. But the arrival of what we might call “heavy” AI felt different. It was not merely another tool to incorporate into our teaching; it represented a fundamental necessity in the conditions of national lockdown. We faced a stark choice: compromise standards like some of our competitors, or persevere and innovate. The decision we made would shape what we were preparing our graduates to do, and whether the skills we had been cultivating would remain relevant in a world where machines could write code, draft essays, and analyze data with startling competence.

The questions we faced were not unique to Georgia, but they carried particular weight in our context. Small economies experience disruption differently from their larger counterparts. Labor markets are thinner, meaning job displacement can have more immediate and visible consequences. The distance between education and employment is shorter, in the sense that graduates often move directly into roles where their capabilities are tested without the buffering layers of extensive graduate training programs. International competition for talent is fierce, with brain drain a persistent concern: research from comparable small economy contexts suggests that graduate out-migration can strip billions from national GDP over time, with one study estimating that a single cohort of departing graduates can cost a state economy over \$500 million in foregone wages within just four years (Common Sense Institute, 2026). And yet, small economies also possess advantages: agility, the ability to experiment without navigating vast bureaucratic structures, and proximity between educational institutions and the employers who will hire their graduates.

This chapter reflects on how one university in such an environment has approached the challenge of building AI-ready graduates. It is not a prescriptive manual. The situation on the ground is shifting too rapidly for definitive answers, and the specific context of each institution, each national system, and each regional economy will shape what is possible and what is appropriate. This is an exploration of themes and tentative insights drawn from our experience at Caucasus University, particularly through our involvement in EU funded projects and our efforts to develop ethical frameworks for AI use in academic work. These initiatives connect us to a broader movement, which aims to design ethical standards for AI use in university education across European and partner countries, emphasizing responsible practices and sustainable policies (European Commission, 2022).

It is organized around three interconnected themes: curriculum design, employability pathways, and internationalization. These are not separate concerns but overlapping

territories where universities must make strategic choices about how to prepare students for a world in which AI is a present and pressing actor. In each domain, I consider both the practical challenges we have encountered and the broader questions they raise about the purpose of higher education in an age of intelligent machines.

ackground: Georgia's economic and education landscape

To understand the choices facing universities in small economies, we must appreciate the economic terrain they inhabit. Georgia's economy has undergone remarkable transformation since the early 2000s, transitioning from post-Soviet collapse to a market-oriented system with growing integration into neighboring and global markets. The country has pursued aggressive reforms to improve its business environment, and Tbilisi has emerged as a regional hub for technology, finance, and services. Yet the economy remains vulnerable to external shocks, regional instability, and the persistent challenge of providing meaningful employment for increasingly educated young professionals.

Our institutional context is shaped by the broader legacy of post-Soviet higher education transformation. Scholarship on de-Sovietization emphasizes that higher education systems in our region have had to re-envision values, practices, and institutions in pursuit of aspirations to overturn authoritarian and repressive educational legacies (Chankseliani, 2022; Oleksiyenko et al., 2018). This process involves decentering and deconstructing inherited structures—moving from Soviet-centric or Eurocentric models toward approaches that reclaim agency and bring about contextually appropriate change. For Caucasus University, established over 25 years ago as an international collaborative project to bring Western-style education to the freshly independent Republic of Georgia, this has meant navigating between global standards and local relevance.

The technology sector in Georgia demonstrates both current opportunities and constraints. An active startup ecosystem has emerged, supported by government initiatives, international investment, and a generation of young Georgians who have grown up with digital technologies. Tech companies, both local and international, have established operations in Tbilisi, drawn by competitive costs, a favorable regulatory environment, and a pool of technically skilled graduates. Yet the sector remains small in absolute terms, narrowly specialized, and the demand for AI-related skills, while growing, is not yet at the scale seen in larger economies. Research on Georgia's AI landscape confirms that private universities are playing an important role in thinking strategically about AI readiness and providing practical initiatives to improve capacity, even as a comprehensive national AI strategy remains under development (Oxford Insights, 2025).

This creates a particular challenge for universities. We must prepare graduates for jobs that may not yet exist in large numbers within our own economy, while also ensuring they possess skills that are immediately relevant to local employers. The balance between international competitiveness and local relevance is delicate. Graduates who are too narrowly focused on the current needs of the local market may find their skills obsolete as AI advances and reduces these jobs. Graduates who are prepared only for theoretical future roles may struggle to find employment in the present, while transition is still ongoing. As the OECD (2026) has emphasized, making AI work requires sustained investment in skills that enable people to use AI effectively, sensibly, and safely—a challenge particularly acute in emerging economies where infrastructure and access limitations persist.

The COVID-19 pandemic accelerated many of these tensions. At Caucasus University, we were forced to make a rapid shift to online learning. The crisis became an unexpect-

ed catalyst for innovation. We developed and deployed our own AI-assisted systems for managing the transition, which we still use. This experience taught us something important: necessity can drive innovation faster than planning, but the innovations that emerge under pressure require careful evaluation and integration into long-term strategy. This mirrors broader findings from post-Soviet contexts, where research indicates that technological conditions—particularly internet stability and digital readiness—are primary drivers of positive digital learning experiences, while demographic characteristics contribute minimally to student satisfaction models (Rane et al., 2023).

Curriculum: beyond tool training

The most immediate response to AI in higher education has been a rush to add AI-related courses and programs. New degrees in artificial intelligence, machine learning, and data science have proliferated globally, and Georgia has been no exception. At Caucasus University, we have expanded our offerings in these areas, recognizing that technical skills in AI development and deployment will be increasingly valuable. Yet we have also been cautious about equating AI readiness with technical proficiency alone.

Contemporary scholarship on AI-driven curriculum development emphasizes that while AI can enable dynamic, responsive, and adaptable curricula through personalized learning pathways, there are significant risks of overuse potentially diminishing students' critical thinking and reflective capabilities (Strielkowski et al., 2024). The deeper curriculum challenge is not teaching students to use AI tools, or even to build them, but to think critically about their appropriate use, their limitations, and their broader implications. This is more like a form of judgment that integrates technical understanding with ethical reasoning, contextual awareness, and the ability to navigate changing environments. Developing this capacity requires a different approach to curriculum design—one with more integrated, problem-based learning that challenges traditional models concerned with uniformity at the expense of considering alignment with economic indicators, learning styles, and contextual factors (Yadav, 2025).

Our experience with the ongoing projects has been instructive in this regard. The project brought together universities from Europe and partner countries to develop ethical frameworks and standards for research, which we used as our chance to also develop standards for AI use in higher education. For us, this was more than an academic exercise in producing guidelines and documents. It was a rare opportunity to engage our academic community in a sustained conversation about AI usage in practice, supported by European expertise. The ethical norms we developed, which now govern how students and faculty may use generative AI tools in their academic work, emerged from this collaborative process.

The framework we established reflects several key principles that align with emerging international standards. First, transparency is non-negotiable. Students must declare when they have used AI tools, and they must do so with specificity, identifying which tools were used, for what purposes, and how the outputs were modified and verified. This principle resonates with the European Commission's (2022) ethical guidelines on AI and data in teaching and learning, which emphasize transparency, accountability, and human agency. Second, original intellectual work remains the student's responsibility. AI may assist with research, with structuring arguments, with checking grammar, but the thinking, the analysis, the conclusions must be the student's own. Third, AI outputs must be treated with appropriate skepticism. The tools can generate plausible-sounding but incorrect information — a phenomenon well-documented in the literature as “hal-

lucinations” (Ellison et al., 2025) — and students must learn to verify and validate rather than accept.

These principles are embodied in specific requirements. Students must attach a declaration to any academic work in which AI tools were used. They must cite AI-generated content appropriately, using a format we developed that identifies the tool, version, and date of use. For substantial use of AI in research processes, they must describe in their methodology section what prompts were given, what outputs were received, and how those outputs were transformed into the final work. In some cases, they must attach the actual prompts and outputs as appendices.

Implementing these requirements has not been straightforward. Faculty needed training to understand the tools themselves and to recognize the boundary between legitimate assistance and academic misconduct. Students needed guidance on how to use the tools effectively and ethically. Assessment practices had to be reconsidered, with more emphasis on process, on oral defense of work, and on tasks that require students to demonstrate understanding in ways that are difficult to fake with AI assistance. This aligns with broader international trends: the University of Sydney (2025), for instance, has embedded a “two-lane” Assessment Framework distinguishing between secure assessments (in-person, supervised tasks) and open assessments (tasks that reflect authentic challenges and scaffold responsible AI use), with Assessment Plans required for every degree component.

The curriculum implications extend beyond individual courses to ensure that AI-related competencies are developed progressively across the duration of degree programs, not concentrated in a single course or left to students to figure out on their own. This means identifying where in the curriculum students should first encounter AI tools, where they should develop critical evaluation skills, and where they should demonstrate mastery of ethical, responsible use. Recent research emphasizes that institutions are prioritizing critical thinking, communication, and adaptability skills over traditional knowledge retention metrics, with innovative approaches including project-based assessments, real-time problem-solving scenarios, and collaborative evaluations that are inherently resistant to AI manipulation (Marshall, 2024).

One area where we have invested significant effort is in the development of AI-assisted test banks. This might seem like a minor administrative matter, but it reflects a larger principle. By using AI tools to help generate and refine assessment questions, we free faculty time for higher-value activities while building institutional capacity. The process of developing these systems has required us to think carefully about what makes a good assessment question, how to ensure quality and consistency, and how to maintain academic integrity in an environment where students have access to powerful AI tools. This approach mirrors findings that AI can strengthen integrity by promoting transparency, fairness, and learning support when implemented responsibly, with nearly 80% of institutions now using AI tools in some capacity (Cranfield et al., 2024).

However, we remain cognizant of the limitations of authenticity as a safeguard. Recent experimental research from UK universities indicates that markers are generally unable to distinguish assessments with generative AI input from those without, and that the level of authenticity in an assessment has no impact on the ability to safeguard against or detect AI usage (Ellison et al., 2025). This suggests that authentic assessments alone are not the panacea some have suggested; the focus must remain on comprehensive assessment design and pedagogical approaches that emphasize process over product (Budde & Tobor, 2025).

Employability pathways: class to desk pipeline

If curriculum design is about what students learn, employability pathways are about how that learning connects to the world beyond the university. In a small economy, this connection is both more immediate and more complex. Employers are often known personally to faculty, or sometimes are faculty or alumni. The feedback loop between graduate outcomes and program design can be short. But the limited scale of the local labor market also means that graduates may need to seek opportunities internationally, or that their jobs may be transformed by global forces, including AI adoption by multinational companies operating in Georgia or by international competitors.

Our approach to employability in the AI era rests on three pillars. The first is maintaining strong, ongoing relationships with employers across sectors. This involves sustained engagement through joint projects, guest lectures, curriculum consultation, and planned networking events. These relationships keep us attuned to how AI is actually being used in workplaces, not how we imagine it might be used. This co-creation model aligns with emerging best practice: recent research emphasizes that the age of AI requires parallel learning, continuous employer input, and new models of experiential education, with leading employers and institutions embracing shared data, shared uncertainty, and shared accountability to build AI-ready talent pipelines (Practera, 2025).

What we have learned from these engagements is that the impact of AI on employment is more nuanced than headlines suggest. Some routine tasks are indeed being automated, but this often changes jobs rather than eliminating them. The workers who thrive are those who can work effectively with AI tools, who understand their capabilities and limitations, and who can apply human judgment to situations where AI alone is insufficient. The skills that employers most frequently mention as becoming more important are not primarily technical; they are critical thinking, communication, adaptability, and ethical reasoning. AI may automate routine tasks, but it still needs human oversight and monitoring. This finding is consistent with global surveys indicating that 96% of human resources leaders believe schools need to take more responsibility for workforce training, with critical thinking, creativity, resilience, and problem-solving abilities topping the list of in-demand skills (World Economic Forum, 2025).

The second pillar is experiential learning that bridges the gap between academic study and professional practice. Internships, project-based courses with industry partners, training programs, and startup acceleration programs all play a role here. We have been particularly focused on creating opportunities for students to work with AI tools in realistic professional contexts, under supervision that helps them develop good judgment about when and how to use these tools appropriately. Research on experiential learning emphasizes that it connects classroom theory to real-world practice, building capability, confidence, and connection—the three currencies of employability in a changing world (Practera, 2025). The World Economic Forum (2025) has reported that 40% of current workers' core skills are expected to change in the next five years, with 50% of all employees needing reskilling by 2025, making experiential approaches vital for workforce readiness.

The third pillar is career guidance that helps students navigate uncertainty. The current job landscape will be completely overhauled in ten years. The pathways from education to employment are becoming less linear. Students need support in thinking strategically about their careers, in building portfolios of skills and experiences that provide flexibility, and in developing the capacity for lifelong learning that will be essential as technology continues to evolve. Recent scholarship emphasizes that students now need to demonstrate experience, confidence, and transferable skills to secure employment

post-graduation, with project-based professional experience providing scalable pathways for inclusive work-integrated learning (Practera, 2025).

A particular challenge in the Georgian context is the balance between preparing graduates for local employment and for international mobility. Many of our most talented graduates will seek opportunities abroad, particularly in European markets where Georgian qualifications are increasingly recognized. We have a responsibility to prepare them for this mobility, which means ensuring that their skills and credentials are internationally competitive. At the same time, we are committed to contributing to Georgia's development, which requires graduates who are prepared to build careers and businesses here. Research on university-industry collaboration in post-Soviet states reveals five persistent challenges that inform our approach: state dominance, weak industry demand, underdeveloped commercialization infrastructure, human capital constraints, and cultural barriers—alongside emerging opportunities for curriculum reform, work-based learning, and employability-focused teaching models (Kuzhabekova & Ruby, 2018).

AI complicates this balance in interesting ways. On one hand, AI-related skills are globally transferable, potentially enhancing the international mobility of graduates who possess them. On the other hand, the specific applications of AI that are most relevant may differ between contexts. A graduate who understands how AI is being used in European financial services may need to adapt considerably to apply that knowledge in Georgia's emerging tech sector. The foundational skills—the capacity to learn and adapt—become more important than specific technical competencies. This aligns with findings that programs embedding AI across disciplines ensure that students, regardless of their field, can understand and leverage this transformative technology (Holmes et al., 2023).

Internationalization: industries blue and red

Internationalization has long been a strategic priority for universities in small economies. For us in Georgia, it is partly a matter of necessity. Our domestic student market is limited, and attracting international students provides both financial sustainability and the diversity of perspectives that enriches the educational experience for everyone. It is also a matter of quality. Engagement with international partners, participation in global academic networks, and alignment with international standards are essential for maintaining and improving the quality of what we offer. Thus, internationalization has been a cornerstone of CU strategy from the very days of its establishment, more than 25 years ago, as an international collaborative project to bring first Western-style education to the freshly independent Republic of Georgia (Knight, 2004).

AI is transforming the landscape of international higher education in ways that create both opportunities and challenges for institutions like ours. On the opportunity side, AI-powered tools can reduce barriers to international engagement. Language translation, once a significant obstacle, is increasingly facilitated by AI (though less so with Georgian language). Virtual collaboration tools, enhanced by AI features, make it easier to maintain relationships with partners across distances. AI-assisted administrative systems can manage the complexity of international partnerships more efficiently.

EU funded projects exemplify how international collaboration can address shared challenges. Working with partners from several European countries, we have been able to develop ethical frameworks that draw on diverse perspectives while being adapted to our specific contexts. The project has also provided valuable professional development opportunities for our faculty, exposing them to different approaches and building

relationships that will outlast the formal project period. Similar initiatives, such as the Erasmus+ Trustworthy AI project, aim to expand the teaching of AI in higher education institutions from the perspective of ethics and values, reaching broad audiences across STEM and other disciplines (European Commission, 2022). The AI-ERIT/AI ETHICS project focuses specifically on developing comprehensive frameworks for AI implementation in higher education, with emphasis on ethical considerations, capacity building, and adaptation of institutional structures (European Commission, 2022).

Yet AI also intensifies competition in ways that can disadvantage smaller institutions. The major players in online education, many of them large universities with substantial resources, are investing heavily in AI-powered adaptive learning systems, automated assessment, and personalized student support. These investments could widen the gap between well-resourced institutions and those with more limited means. Students may increasingly turn to global providers for at least parts of their education, potentially disrupting traditional models of international student mobility. The European Commission's 2025 funding call for ethical AI integration in education, with a total budget of €13 million and maximum grants of €1 million per project, signals the scale of investment flowing toward large-scale AI educational infrastructure (European Commission, 2025).

Our response to this competitive pressure has been to focus on what we can offer that large-scale online providers cannot. This includes the close personal relationships between students and faculty that are possible in a smaller institution. It includes the contextual knowledge and local connections that are valuable for students who intend to build careers in Georgia or the broader region. And it includes the experiential learning opportunities that require physical presence and real-world engagement. Research on international education trends emphasizes that institutions that weave AI, experiential learning, and collaboration into their DNA will lead transformation—not by resisting change, but by learning through it (Practera, 2025).

At the same time, we have been strategic in adopting AI tools that can enhance our international competitiveness. Our AI-assisted test bank development, mentioned earlier, is one example. We have also invested in AI-powered systems for international student training and support, recognizing that providing excellent service to international students requires efficient processes as well as personal attention.

Internationalization in the AI era also raises important questions about the content of what we teach. As AI systems are developed and deployed globally, there is a risk that the perspectives and values embedded in these systems reflect primarily the contexts of their developers, who are concentrated in a few major technology hubs. Universities outside these hubs have a responsibility to cultivate critical perspectives on AI, to surface questions and concerns that might be overlooked in dominant narratives, and to ensure that the global conversation about AI's future includes diverse voices. This aligns with broader critiques that AI training data often contains under-represented sources, disproportionately populated by English-language and Western-centric content, creating gaps in models' ability to understand cultural and linguistic nuances of diverse communities (Ellison et al., 2025).

It is also the main challenge of universities everywhere to distinguish between AI-assisted learning and AI substituting students, with academic integrity coming under increasing threats from AI. In Georgia, this means engaging our students with questions about how AI might be developed and deployed differently in different cultural and economic contexts. It means examining the implications of AI for small economies, for post-Soviet societies navigating complex transitions, for countries at the geopolitical crossroads. These are not merely academic questions; they are matters of practical importance for graduates who will shape how AI is used in our society.

Recent research on post-Soviet higher education emphasizes that de-Sovietization requires re-claiming agency to bring about change, re-envisioning values and institutions to overturn authoritarian legacies (Chankseliani, 2022). In this context, our internationalization strategy is not about mimicking Western models but about contributing distinctive perspectives from our transitional context to global conversations about AI and education.

Reflections on institutional strategy

The experience of navigating AI's impact on curriculum, employability, and internationalization has taught us several lessons about institutional strategy in a time of rapid technological change. The first is the importance of agility without recklessness. Small institutions can move quickly, experimenting with new approaches and adapting as circumstances change. But speed without direction is merely chaos. We have found value in establishing clear principles, such as those embedded in our ethical framework, that provide guidance even as specific practices evolve. At the same time, we have to take care and closely examine new approaches as a still growing private institution.

Contemporary scholarship on AI in higher education emphasizes that leading institutional responses require comprehensive strategic planning, with universities continuing to refine assessment frameworks, validate new methodologies, and explore applications across diverse educational contexts (Marshall, 2024). The “Educate, Enable, Expect” model proposed by Löfström et al. (2015)—where students are explicitly educated in ethical principles, enabled to practice and refine skills within safe formative spaces, and then held to high expectations—offers a useful framework for thinking about institutional development in the AI era.

Let us not forget the value of collaborative engagement. No single institution can master all the dimensions of AI's impact on higher education. If one tries, other universities would have to develop their own standards and practices leading to chaos across the industry. Partnerships, whether through formal projects or through informal networks of colleagues facing similar challenges, multiply our capacity to learn and adapt. These partnerships must be genuine two-way exchanges, not merely arrangements in which larger or better-resourced institutions dictate terms to smaller partners. The EU framework explicitly supports this through capacity-building initiatives that enhance lecturer capabilities in handling and condensing ethical practices in AI use (European Commission, 2022).

One of the most important lessons is the continuing centrality of human relationships in education. AI can augment many aspects of what universities do, from administrative processes to certain forms of learning support. But the core educational relationship, between a student who is struggling to understand and a teacher who can provide real expertise, remains stubbornly resistant to automation. Our strategy has been to use AI to free up resources for these irreducibly human dimensions of education, not to replace them. This cannot be achieved without maintaining a critical perspective on technology itself.

Universities have a distinctive role in society as places where new developments can be examined thoughtfully, where enthusiasm can be tempered with analysis, and where the questions that technology companies might prefer to avoid can be raised. It has become a sort of private joke how quickly they flip the scripts when talking to investors versus regulators. This critical role is not opposition to technology; it is the responsible exercise of our institutional mission to pursue truth and serve society. Recent

surveys indicate that while 67% of students think using AI is “essential” in today’s world, only 36% have received training in AI skills from their institution, and students express confusion about mixed messages regarding AI use (Freeman, 2025). This gap between adoption and education represents both a challenge and an opportunity for institutions willing to take leadership.

Finally, we have learned that the challenges posed by AI are not separate from the broader challenges facing higher education in small economies. Questions of funding, of access, of quality, of relevance, all intersect with the AI question. A strategy for AI that does not address these broader contexts will fail. Conversely, addressing these broader challenges well positions an institution to navigate AI’s impact effectively. Research on the digital transformation of higher education in AI-driven contexts emphasizes that educational systems of the future will operate as decentralized digital ecosystems in which decision-making processes are shared among multiple stakeholders, including administrators, employers, and students (Rane et al., 2023).

Conclusion

As I write this, in early 2026, the pace of AI development shows no signs of slowing. New models, new capabilities, new applications continue to emerge. The specific tools our students are using today will likely be obsolete before they graduate. The jobs we are preparing them for will continue to evolve in ways we cannot fully predict. This uncertainty is not a temporary condition to be endured until things stabilize; it is the permanent condition of education in an age of rapid technological change. The AI revolution is here and it is not going back.

What we can do, what we must do, is prepare graduates who are capable of navigating this uncertainty. This means developing their technical skills, certainly, but it also means cultivating their capacity for critical thinking, for ethical reasoning, for learning throughout their careers. It means connecting their education to the world of work, while also helping them understand that their careers will likely span multiple roles and sectors. It means engaging internationally, while also rooting their education in the specific context of our society and its needs.

For universities in small economies, these challenges are intensified but also clarified. We cannot compete with major global institutions on every dimension. But we can leverage our agility, our proximity to local needs, and our capacity for personal engagement to offer something distinctive and valuable. The graduates we produce will shape how AI is developed and used in our society. Their education is therefore not merely a matter of institutional interest but of national and social importance.

The work we have begun at Caucasus University, through the ethical frameworks developed within framework of our international collaboration, through our curriculum innovations, through our engagement with employers and international partners, is ongoing. We do not have final answers, and we would be suspicious of anyone who claimed to. What we have is a commitment to thoughtful, principled engagement with the questions AI raises for higher education, and a willingness to learn and adapt as the terrain continues to shift. This chapter is offered in that spirit, as a contribution to a conversation that must involve university leaders, faculty, students, employers, and policymakers across all the diverse contexts that make up the global higher education landscape.

The path forward requires us to embrace what we might call “structured uncertainty”: building institutions and pedagogies that are resilient precisely because they

acknowledge that the future is unknowable. In this, small universities in fast-shifting economies may paradoxically hold advantages: we are already accustomed to navigating instability, to building relationships that sustain us through change, and to valuing adaptability as a core competency. These qualities, cultivated through necessity, may prove essential assets as the entire global higher education sector confronts the transformative implications of artificial intelligence.

References

- Bregvadze, T. (2012). *Shadow education in Georgia: Prevalence, perceptions, and implications*. Tbilisi: Georgian Ministry of Education.
- Budde, J., & Tobor, J. (2025). (Gen)AI and its implementation in German higher education examination practices towards a new (sustainable) assessment culture? *Higher Education: Linking Research, Policy and Practice*.
- Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20(1), 38.
- Chankseliani, M. (2022). De-Sovietization in higher education. *International Journal of Educational Development*.
- Common Sense Institute. (2026). *Brain drain in Iowa's higher education system*. CSI Report.
- Cranfield, D. J., et al. (2024). Conversational artificial intelligence: A catalyst for rethinking assessment in higher education. *Proceedings of the European Conference on Knowledge Management, ECKM*.
- Ellison, C., et al. (2025). "What's really going on out there? Exploration of staff and student experiences of generative artificial intelligence at the University of Liverpool." *Perspectives: Policy and Practice in Higher Education*.
- European Commission. (2022). *Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators*. Publications Office of the European Union.
- European Commission. (2025). *ERASMUS-EDU-2025-PI-FORWARD-DIGITAL-AI: Ethical AI integration in education*. Funding call documentation.
- Freeman, J. (2025). *Student generative AI survey 2025*. HEPI Policy Note 61. Higher Education Policy Institute.
- Gogberashvili, S., Latsabidze, N., & Gogiashvili, S. (2026). AI in English language teaching at higher education institutions in Georgia: Usage, training gaps, and institutional support. *The JALT CALL Journal*, 22, 102849.
- Holmes, W., Bialik, M., & Fadel, C. (2023). *Artificial intelligence in education: Promises and implications for teaching and learning*. Boston, MA: Center for Curriculum Redesign.
- Khlaif, Z. N., et al. (2025). Redesigning assessments for AI-enhanced learning: A framework for educators in the generative AI era. *Education Sciences*.
- Knight, J. (2004). Internationalization remodeled: Definition, approaches, and rationales. *Journal of Studies in International Education*, 8(1), 5-31.
- Kuzhabekova, A., & Ruby, A. (2018). *Higher education in post-Soviet countries*. Comparative Education Review.

- Löfström, E., Trotman, T., Furnari, M., & Shephard, K. (2015). Who teaches academic integrity and how do they teach it? *Higher Education*, 69(3), 435-448.
- Marshall, S. J. (2024). Leading an institutional response to artificial intelligence and its impact on academic work. In *Artificial intelligence applications in higher education: Theories, ethics, and case studies for universities*.
- Matiashvili, A., & Kutateladze, T. (2006). Private tutoring in Georgia. *CERC Working Papers*.
- OECD. (2026). *Making AI work: Why investing in skills matters*. OECD Publishing.
- Oleksiyyenko, A., et al. (2018). *Higher education in post-Soviet societies*. Palgrave Macmillan.
- Oxford Insights. (2025). *Education at the centre stage: Georgia's bottom-up approach to building national AI capacity*. Oxford Insights Spotlight Series.
- Perkins, M., & Roe, J. (2024). Decoding academic integrity policies: A corpus linguistics investigation of AI and other technological threats. *Higher Education Policy*.
- Practera. (2025). *AI and experiential learning: Navigating change in international education*. AIEC 2025 Conference Materials.
- Rane, N., et al. (2023). AI in personalized learning. *International Journal of Educational Technology*.
- Silova, I. (2011). *Globalization on the margins: Education and post-socialist transformations in Central Asia*. Information Age Publishing.
- Strielkowski, W., et al. (2024). Sustainable curriculum development with AI. *Springer Education Innovation Series*.
- Thite, S. D., et al. (2024). A three-step approach to shift from assessment integrity to assurance of learning in the age of AI. *2024 World Engineering Education Forum — Global Engineering Deans Council*.
- UNESCO. (2023). *Guidance for generative AI in education and research*. Paris: UNESCO Publishing.
- University of Sydney. (2025). *Navigating AI in higher education: Tasks ahead for 2025 and 2026*. Teaching@Sydney Policy Brief.
- World Economic Forum. (2025). *The future of jobs report 2025*. Geneva: World Economic Forum.
- Yadav, A. (2025). Data-driven approaches to curriculum optimization. *Journal of Higher Education Innovation*.

Chapter 9. The tactical classroom: “Learn with AI, test without AI” in Kyrgyzstan

Ivan Ninenko

Institutional context: IUCA as a case of constrained transformation

The International University of Central Asia (IUCA) was founded in 2008 in Tokmok, Kyrgyzstan - a city of 70,000 people, 60 kilometers from Bishkek, in the foothills of the Tian Shan mountains. The institution operates on the American model of higher education with an early and explicit commitment to academic integrity - a commitment that is far from routine in the regional context. Approximately 500 students are currently enrolled across programs in business, international relations, information technology, law, and linguistics. Instruction runs in three languages: Kyrgyz, Russian, and English. The national education system within which IUCA operates is still working through its post-Soviet inheritance: lecture-heavy, with limited traditions of student-directed study.

Three institutional priorities: questioning, flipped classroom, AI literacy

Problem with the existing model

At the opening of the 2025 academic year, IUCA Provost Kamila Sharshkeeva named three priorities for the institution: building a culture of questioning in the classroom, rolling out the flipped classroom model across faculties, and developing real AI literacy among faculty and students. All three matter, and traces of all three run through what follows. The first -- the culture of questioning -- is substantial enough to deserve its own treatment; the constraints of this volume and the nature of the author's evidence pull toward the second and third. They are addressed together here because, in practice, separating them is artificial: a flipped model without AI literacy collapses, and AI literacy without a restructured classroom has nowhere to go.

The pressure behind these priorities is not difficult to identify. For decades, the default pedagogical model in post-Soviet universities has followed a consistent logic: the instructor holds knowledge, the student receives it, and the examination measures how much was retained. The internet already strained this model - once information became freely accessible, the lecture's core justification weakened considerably. Generative AI has undermined it further. AI can now produce adequate answers to most examination questions on demand. Three institutional responses have become common, none sufficient: treating AI as a compliance problem and tightening enforcement; accepting AI-generated work without adjusting assessment design; or appending an “AI literacy” module to an otherwise unchanged curriculum. The three priorities announced by Provost Sharshkeeva represent a different kind of response -- one that requires reconsidering not only how students are assessed, but the structure of the learning encounter itself.

Flipped Classroom as Structural Response

The flipped classroom – students engage with content before class, class time goes to application and discussion – predates AI by decades (Bergmann & Sams, 2012; Bishop & Verleger, 2013; Lage, Platt, & Treglia, 2000). What AI has changed is the stakes. When a student can get a competent explanation of any concept from an AI tutor at midnight, the live lecture loses its core justification. What it cannot replace is the structured human encounter: a teacher watching a student reason through a problem, pushing back on a half-formed answer, catching a misunderstanding that no AI would have flagged — or the curated peer interactions, games, and case analyses that require a room full of people to be meaningful. Recent work confirms that generative AI strengthens both ends of the flipped model – richer pre-class engagement, more productive in-class time – while cutting instructor preparation load (Monzon & Hays, 2025; Triantafyllidis, 2024; Ouyang et al., 2025).

The significance of this shift is amplified in post-Soviet educational contexts, where the instructor-as-authority model has deep institutional roots. Flipping the classroom is not a scheduling adjustment. It requires the instructor to relinquish the role of sole knowledge authority and requires the student to take active responsibility for preparation before arriving in class. Neither transition is straightforward, and neither occurs without deliberate institutional support.

At IUCA, the flipped classroom now organizes the entire approach to AI-augmented instruction. Preparation materials – readings, short videos, AI-generated podcasts, purpose-built AI teaching assistants – are distributed before class. Contact time is reserved for three activities that AI cannot adequately substitute: questioning, collaborative problem-solving, and formative assessment with real-time feedback. Assessment logistics – question generation, variant creation, grading, and personalized feedback – are fully automated through QuantaQuiz (described in Section 3.2). The automation handles the procedural work. The instructor’s time goes to the work that requires human judgment.

AI literacy as institutional infrastructure

A flipped classroom that depends on students engaging productively with AI before class is unsustainable if neither students nor faculty possess the relevant competencies. IUCA’s response was to treat AI literacy not as an optional professional development offering but as foundational infrastructure – embedded at the institutional entry points for both faculty and students.

Faculty receive AI literacy training during orientation week at the start of each academic year, supplemented by seminars throughout the semester. These sessions are not technology demonstrations. The organizing question is not “what can this tool do?” but “what does this tool change about what I should be doing?” The aspirational frame in these sessions is that every educator should eventually function as a manager of AI agents — delegating routine cognitive work, reviewing outputs, and concentrating attention on what only a human educator can do. The current reality is mixed. Faculty proficiency varies widely; some instructors have integrated AI into their workflow extensively, others remain skeptical or untrained, and meaningful institution-wide adoption is a multi-year trajectory rather than a one-semester initiative. The relevant skill is not prompt engineering but workflow design — knowing what to delegate, what to review, and what must remain a human judgment — and that skill, like any other professional competence, develops unevenly.

For students, AI literacy is embedded in University 101, the required introductory course for all first-year students. Rather than treating AI as a discrete topic, the course integrates AI fluency into the broader question of how to learn in a university context.

Students practice verifying outputs, identifying hallucinations, and distinguishing between AI-assisted inquiry and AI-substituted thinking. They are introduced to the institution's norms governing AI use -- and, critically, to the reasoning behind those norms. A student who understands why AI cannot substitute for their own understanding is better positioned to resist that substitution than one who has only been told it is prohibited.

The “Learn with AI, test without AI” model

AI-assisted home preparation

The core idea of the flipped classroom is that instruction happens at home, not in the classroom. Before each session, students are assigned preparation materials - existing readings, videos, open courseware, or purpose-built resources - and encouraged to use AI tools actively alongside them. The instructor's task is rarely to create content from scratch: the abundance of high-quality educational material already available online shifts the primary challenge to curation and adaptation - calibrating existing texts and lectures to the specific cohort, course angle, or session focus. AI makes that adaptation tractable. A reformatted reading, a set of guiding questions, a podcast-style summary, a configured teaching assistant - work that previously required three to four hours can now be completed in thirty to sixty minutes. A purpose-built AI teaching assistant for a specific topic can be configured in approximately twenty minutes. For instructors carrying full teaching loads, that reduction is what makes the preparation-intensive flipped model sustainable.

Students are not merely instructed to use AI for preparation - they are taught specific strategies for doing so productively: how to use AI as an interlocutor rather than an answer source, how to generate practice problems, how to recognize when an AI explanation warrants skepticism. These strategies are modeled in class by the instructor and refined through peer exchange. Students who develop effective approaches share them with classmates, making preparation a subject of collective development rather than purely private effort.

Two examples illustrate the format. For a session on software development methodologies in an English for IT course, students work through a pre-class interactive page built around narrative scenarios: a casino heist to illustrate the sequential logic of Waterfall, a pirate crew for the iterative flexibility of Agile (Ninenko, 2025; ninenko.github.io/apps/lectures/waterfall-vs-agile.html). The page includes embedded video, interactive comparisons, and a glossary, alongside a direct link to an AI teaching assistant for students who engage more readily through dialogue than through text.

The second example comes from the Introduction to Machine Learning module of the same course. Students watch a short video about Albert - a simulated humanoid agent learning to walk through reinforcement learning in a 3D environment - then work with an AI teaching assistant (IUCA-Alma) built specifically around that video (chat.almalearning.ai/create-chat?uuid=35c53453-5c2f-4cbf-b604-37fc2adddb68). The assistant does not check whether students watched the video. It asks them to explain what Albert looked like, what his goal was, why he initially learned to do the worm instead of walk, and what that reveals about reward design in reinforcement learning. It enforces English, adjusts to the student's level, refuses to hand over answers, and sends transcripts back to the instructor. The scenario took under twenty minutes to configure - no code, just a structured natural-language description of the learning flow. Students arrive in

class having already wrestled with the material; the session opens with a short test, then moves into discussion.

AI-enhanced in-class assessment

Each session opens with a short test - conducted without AI access, on paper or via Google Forms, typically under ten minutes. Questions target the lower levels of Bloom’s taxonomy: recall, recognition, and basic comprehension of the assigned preparation material. The tests are deliberately designed so that a student who has actually engaged with the assigned material can score 100%. This is not an accident of calibration but a structural choice: a student who prepares should experience an immediate positive feedback loop, not the chronic uncertainty produced by tests where some questions are designed to trip up even prepared students. The combination of hard preparation requirements and reachable scores is what makes the model psychologically sustainable. The tests serve two concurrent functions: holding students accountable for preparation and providing the instructor with immediate diagnostic information about the class’s actual understanding before the session’s main activities begin.

Before AI, preparing test variants - essential to prevent copying -- and grading responses consumed approximately five minutes per student. Both tasks are now fully automated. The instructor submits source material; the system generates questions, produces multiple variants, grades responses, and delivers each student a targeted explanation of their specific errors. The time cost of grading has dropped to zero.

At IUCA, this infrastructure is provided by QuantaQuiz, an AI-powered assessment platform co-developed by the author (QuantaQuiz, n.d.). The system uses large language models to generate questions from course materials, produces multiple test variants for simultaneous distribution, delivers assessments as printable PDFs or via Google Forms, and generates personalized tutor prompts for each student based on their error patterns. QuantaQuiz is one of several products in this category; comparable tools exist commercially, and the underlying capabilities are no longer the bottleneck. With the current state of AI development, any institution with modest technical capacity is in a position to build its own equivalent. The governing principle is consistent with the model as a whole: AI manages the procedural infrastructure surrounding assessment; the assessment itself — the determination of what a student has actually learned — remains a human encounter.

AI-enhanced in-class activities

With preparation, accountability established through the opening test, the remainder of the session is available for activities that require human presence: case analysis, structured debate, applied problem-solving, exercises that demand active performance rather than passive recall. AI substantially reduces the time required to design these activities - from two to four hours to fifteen to forty minutes. The reduction does not produce thinner activities; it shifts the instructor’s attention from material assembly to pedagogical design - from building the activity to deciding what kind of thinking it should provoke.

Two examples from the Neuroscience and Emotional Intelligence course illustrate the range. The first is a memory card matching game in which students pair neuron types with brain structures - a two-player competitive format that operates under mild time pressure and generates immediate peer discussion when a match is contested (Ninenko, 2025; ninenko.github.io/apps/apps/memory_neurons.html). Built with AI assistance, the game plays in approximately five minutes; the discussion it produces typically runs longer. The second example, from a computer science foundations course, is an interactive

punch card emulator: students type text and observe in real time how each character is encoded as a pattern of punched holes on an IBM-style card (Ninenko, 2025; ninenko.github.io/apps/apps/punchcard-emulator.html). Binary encoding is straightforward to describe and difficult to internalize abstractly; the emulator makes it manipulable. Both tools are published via static hosting – no server infrastructure, no login system, no institutional IT involvement required.

Meta-learning: making the process visible

Beyond deep-dive discussion and applied work, the classroom serves a third function under this model: it is the place where learning itself is examined. The deliberate practice of making learning itself an object of discussion reflects a broader insight about the skills that matter in complex, rapidly changing environments — among them the capacity to learn, unlearn, and relearn (Loshkareva et al., 2017). Students discuss their preparation strategies openly - what worked, what failed, how they used AI, what they did when AI explanations were inconsistent or insufficient. This peer modeling accelerates the development of effective preparation habits across the cohort and builds a shared vocabulary for talking about learning itself.

This kind of meta-learning is especially important in the initial weeks of a flipped course, and especially with students who have not yet developed self-directed learning habits — a common starting condition in post-Soviet educational contexts. Encouraging students to talk openly about their preparation strategies, what they tried with AI, what worked and what didn't, accelerates the cohort's adjustment more than any individual instructor intervention can. Even with this support, it typically takes three to five weeks for a class to settle into the flipped model — to begin studying materials seriously before class, engaging actively in discussion, and following up on points that remain unclear. The first weeks of a flipped course tend to feel less controlled than lecture delivery: preparation quality is uneven, and the gaps are immediately visible in test results.

The adaptation period is real, predictable, and should be communicated explicitly to both faculty and students before implementation begins, not retrospectively explained as a phase to endure. In the author's first implementation of the model, one of the three groups was particularly slow to adjust — preparation had deteriorated by week three, scores were low, and the case for reverting was difficult to dismiss. The intervention was simple but consequential: a substantial portion of the next session was given over to open meta-learning discussion — students sharing their preparation approaches, what was failing, what they had not yet tried. From week four onward the group's trajectory shifted markedly, and final results were comparable to the other two groups. The episode is a useful reminder that the meta-learning component is not a stylistic preference. When the model is in difficulty, more discussion of how learning is happening — not less — is what tends to bring it back on track.

Case study: neuroscience and emotional intelligence at IUCA

The case examined here - Neuroscience and Emotional Intelligence, taught to three groups of first-year students during Fall 2025 and the first half of Spring 2026 was selected because it makes the effects of the model measurable and because it captures the particular challenges that arise when working with students who have not yet developed habits of independent study.

IUCA first-year students arrive from a school system that shares, with most post-Soviet systems, a strong orientation toward teacher-directed learning. Independent study

- sustained engagement with material outside class, self-testing, self-directed inquiry
- is not a well-established habit. The Neuroscience and Emotional Intelligence course posed an additional challenge: the subject is terminology-dense. Amygdala, hypothalamus, neural plasticity - students cannot reason about the content at any meaningful level until the vocabulary is in place. AI tutors are well-suited to precisely this kind of preparatory work; vocabulary building and concept clarification are tasks they handle reliably, and they are available at the hours when students actually study.

Each session followed the same structure: AI-assisted home preparation, a short paper test at the start of class, QuantaQuiz-automated grading, personalized tutor prompts distributed to students afterward, and class time given over to case discussion, applied exercises, and instructor-led questioning that built on whatever the opening test had revealed. The tests were not primarily evaluative instruments. They functioned as shared diagnostics - telling both students and instructor what had and had not been retained from the preparation phase.

Test scores were tracked across all three groups, comparing each group’s baseline (early September) with their score after approximately twelve weeks of implementation (November). All three groups showed substantial improvement. Group 1 moved from 68.95% to 91.35% — a gain of 22.4 percentage points and near-mastery by November. Group 2, which began with the lowest baseline (53.80%), recorded the largest absolute gain — 34.2 percentage points — finishing at 88.00%, above Group 1’s starting score. Group 3 moved from 58.70% to 75.71%, a gain of 17.0 percentage points. Standard interpretive cautions apply: this is a practitioner case study without a control group, and the model’s contribution cannot be cleanly isolated from instructor attention, cohort characteristics, or subject-matter effects. The pattern that emerges, however — large and consistent gains across groups with substantially different starting points, in a resource-constrained Central Asian context — is not one that can be attributed to measurement artifact.

A properly designed empirical study would significantly strengthen these findings. The ideal design would involve parallel sections of the same course taught under different pedagogical conditions — flipped with the “learn with AI, test without AI” model on one side, conventional delivery on the other — with instructors blinded to the research framing and assigned approaches without being told they are being studied, in order to minimize Hawthorne effects on either side. End-of-semester outcomes would be the primary measure, with a follow-up assessment at six or twelve months to evaluate long-term retention — arguably the more important metric, given that surface-level memorization is one of the model’s named risks. Such a study is beyond the resource capacity of a single innovation department at a small university, but it is exactly the kind of work the field needs, and the author would welcome collaboration with research groups in a position to conduct it.

Conclusion

The flipped classroom is a structural commitment, not a pedagogical technique. Adopting the label without simultaneously redesigning contact hours, assessment, and the pre-class learning experience does not produce the model’s benefits. AI literacy must be embedded at institutional entry points - orientation week for faculty, the required first-year course for students - not appended as optional enrichment. Assessment redesign is the most consequential and most demanding element of the transition. Institutions that invest in making rigorous assessment tractable at scale -- through platforms

such as QuantaQuiz, through deliberate design effort -- are building the structural foundation on which the rest of the model depends.

The “learn with AI, test without AI” principle offers a practical resolution to what would otherwise be an unresolvable institutional dilemma. It does not require agreement on when AI use becomes academically impermissible. It separates the two contexts in which that question arises: the learning environment, where AI functions as a productive tool, and the assessment environment, where what matters is what the student has independently internalized. In the first context, AI is a tutor available at any hour. In the second, that tutor is absent by design.

Peripheral universities have historically looked to elite institutions for models to emulate. The experience documented here suggests that the direction of learning may, in at least some respects, run the other way. Institutions that have been required to determine what actually matters under conditions of genuine constraint -- without recourse to prestige, infrastructure, or scale as substitutes for pedagogical clarity - may have developed knowledge that well-resourced institutions operating under different pressures, rarely have occasion to acquire.

References

- Bergmann, J., & Sams, A. (2012). *Flip your classroom: Reach every student in every class every day*. International Society for Technology in Education.
- Bishop, J. L., & Verleger, M. A. (2013). The flipped classroom: A survey of the research. In *Proceedings of the 120th ASEE Annual Conference & Exposition*. ASEE.
- Lage, M. J., Platt, G. J., & Treglia, M. (2000). Inverting the classroom: A gateway to creating an inclusive learning environment. *The Journal of Economic Education*, 31(1), 30–43. <https://doi.org/10.1080/00220480009596759>
- Loshareva, E., Luksha, P., Ninenko, I., Smagin, I., & Sudakov, D. (2017). Navyki budushchego. Chto nuzhno znat' i umet' v novom slozhnom mire [Skills of the future: What you need to know and be able to do in a new complex world]. Global Education Futures. <https://globaledufutures.org/futureskills>
- Monzon, N., & Hays, F. A. (2025). Leveraging generative artificial intelligence to improve motivation and retrieval in higher education learners. *JMIR Medical Education*. <https://doi.org/10.2196/59210>
- Ouyang, F., Zheng, L., & Xu, W. (2025). Reimagining flipped learning via Bloom's taxonomy and student-teacher-GenAI interactions. *Education Sciences*, 15(4), 465. <https://doi.org/10.3390/educsci15040465>
- QuantaQuiz. (n.d.). *QuantaQuiz: AI-powered assessment platform* (Version unknown) [Computer software]. <https://quantaquiz.com>
- Triantafyllidis, A. (2024). Introducing AI in flipped classroom: A preliminary study. In *Proceedings of the 16th International Conference on Education and New Learning Technologies (EDULEARN24)*. IATED. <https://doi.org/10.21125/edulearn.2024.2129>
- Vasileva-Stojanovska, T., Malinovski, T., Vasileva, M., Jovevski, D., & Trajkovik, V. (2015). Impact of satisfaction, personality and learning style on educational outcomes in a blended learning environment. *Learning and Individual Differences*, 38, 127–135. <https://doi.org/10.1016/j.lindif.2015.01.006>

World Bank. (2024). *AI in schools: Opportunities, challenges and realities*. Vienna Development Knowledge Center.

Yan, L., Sha, L., Zhao, L., Li, Y., Martinez-Maldonado, R., Chen, G., Li, X., Jin, Y., & Gašević, D. (2024). Practical and ethical challenges of large language models in education: A systematic scoping review. *British Journal of Educational Technology*, 55(1), 90–112. <https://doi.org/10.1111/bjet.13370>

Chapter 10. The “Tolyq Adam” framework: holistic human development in Kazakhstan

Assylbek Kozhakhmetov, Assel Aryn, and Aigul Sarenova

As artificial intelligence automates an expanding range of cognitive tasks — from data analysis and legal research to coding and medical diagnosis — universities face two deceptively simple but deeply consequential questions: *what*, exactly, should we be teaching, and *how*? If AI can perform many of the technical skills that have historically defined professional training, the graduate profile of the future must shift toward something more durable, more distinctly human. And if the content of education must change, so must its form: a curriculum built around human formation cannot be delivered through the same lecture-and-assessment logic that was designed to transmit information. The *what* and the *how* are inseparable.

Researchers and employers across sectors increasingly converge on the same answer: the skills that matter most in an AI-saturated world are precisely those that machines struggle to replicate. Critical and creative thinking. Ethical judgment. Empathy and the capacity for genuine dialogue. Resilience in the face of ambiguity. The ability to find meaning and act with purpose. These are not supplements to a “real” education — they *are* the education that the coming decades demand.

Yet translating this insight into actual curriculum is far harder than articulating it. Most universities are constrained by national standards, departmental silos, and the gravitational pull of existing structures. Against this backdrop, the story of Almaty Management University (AlmaU) and its *Tolyq Adam* program offers a compelling case study — not only of what such a curriculum can look like, but of how it can be built within the real constraints of a regulated higher education system, grounded in local cultural foundations rather than imported frameworks.

Origins: a crisis, a journey, and a name already waiting

The Tolyq Adam program did not emerge from a strategic planning cycle or an institutional response to AI disruption. It began with a more immediate human problem — and a question that a university president put to himself.

In the aftermath of the pandemic and a period of acute sociopolitical turbulence in Kazakhstan, AlmaU’s leadership observed a marked deterioration in students’ inner lives: heightened anxiety, declining attention, difficulty engaging in sustained dialogue. A psychological support center was a necessary first response, but it quickly became clear that the challenge was deeper. Students were navigating questions of identity, meaning, and purpose — questions that no counseling service could address. What was needed was an educational response.

To find one, university leadership traveled to observe how the world’s best institutions approached the formation of whole human beings. In the United States, they visited Bard College in the US, one of the most distinguished universities in the liberal arts

tradition, and later the Mind Institute in Seoul, South Korea, and the Hong Kong Polytechnic University (PolyU). The observation was consistent: the strongest educational approaches were not built around the transmission of disciplinary content, but around the formation of *thinking* — the capacity to ask questions, hold uncertainty, engage in intellectual dialogue, and see the world whole.

It was also clear that no foreign model could simply be transplanted. The anchor had to be found at home.

AlmaU did not seek to reproduce a ready-made liberal arts model imported from another educational context, however successful or prestigious that model might be. The question was not how to build a Kazakhstani version of Bard College, PolyU, or any other global institution. The deeper question was different: what kind of educational philosophy could speak authentically to Kazakhstan’s own cultural, intellectual, and civic context, while remaining relevant to the global challenges of the twenty-first century?

It arrived from an unexpected direction — and from the highest level of a national discourse. In a speech to youth leaders, Kazakhstan’s President Kassym-Jomart Tokayev invoked Abai directly: “At one time, Abai introduced the concept of a ‘person of integrity.’ Abai was ahead of his time when he spoke of an integrated person in knowledge, comprehension, and societal demands”. In his address *Abai and Kazakhstan in the Twenty-First Century*, Tokayev went further, calling for the concept of *Tolyq Adam* to become foundational across all spheres of national life, including education — and explicitly linking it to his own civic vision of the *Adal Azamat*, the responsible citizen. The name that AlmaU had been searching for turned out to have been waiting all along: not borrowed from abroad, not invented in a strategic planning meeting, but embedded in the country’s own philosophical heritage and affirmed at the level of state policy. For a university that had just returned from Bard College, PolyU and Seoul asking what genuine human formation looks like, this convergence was clarifying. The answer had been articulated, with remarkable precision, a century and a half earlier.

In this sense, *Tolyq Adam* became an example of an educational idea born not through imitation, but through the search for identity and meaning. The program emerged from the conviction that genuine curricular innovation cannot be reduced to borrowing external frameworks. A university cannot become distinctive by copying even the most advanced models. It becomes distinctive when it can interpret global trends through its own intellectual tradition, institutional mission, and cultural memory.

The turn to Abai’s philosophy gave AlmaU exactly such a foundation. The concept of *Tolyq Adam* offered not merely a name, but a philosophical center of gravity. It connected contemporary discussions on human-centered education, liberal arts, social responsibility, and the future of work with a deeply rooted Kazakh understanding of human formation.

The philosophical grounding: Abai and an insight that predates AI

The nineteenth-century Kazakh philosopher and poet Abay Kunanbayuly articulated a vision of human wholeness that speaks with startling directness to the educational challenges of the present moment. In his conception of *Tolyq Adam*, a person becomes complete through the integrated cultivation of three capacities: *aqyl* (mind — the faculty of reason, discernment, and the pursuit of knowledge), *zhürek* (heart — the seat of moral sensibility, empathy, and conscience), and *qairat* (will — steadiness, discipline, and the endurance to act on one’s convictions).

His words, rendered in English, read: “*Hold wisdom, will, and heart in equal measure — then you will be complete, set apart from the rest.*”

In the seventeenth of his *Words of Edification*, Abai dramatizes this triad: mind, will, and heart each argue before Knowledge for their own primacy, and Knowledge delivers its judgment — that none can flourish without the others, and that it is the properly formed heart that must ultimately govern. The formulation is not merely poetic. It is a pedagogical claim: intellectual attainment without moral grounding produces something incomplete, and potentially dangerous.

What Abai identified in the nineteenth century, contemporary research on human-AI collaboration is rediscovering in the twenty-first. The capacities most resistant to automation — ethical judgment, empathic understanding, resilience of character, the ability to hold complexity without collapse — are precisely those that Abai placed at the center of human formation. The *aqyl-zhürek-qairat* triad maps, with striking precision, onto what the most sophisticated accounts of AI-era graduate profiles identify as essential: cognitive flexibility (*aqyl*), emotional and ethical depth (*zhürek*), and adaptive agency (*qairat*).

This parallel is not coincidental. It reflects something that Kazakhstani educators are increasingly reclaiming: that the tradition of thinking about human formation in this part of the world runs deep, and that it belongs in the front ranks of global educational discourse. The educator and poet Magzhan Zhumabayev wrote his treatise *Pedagogy* in 1922 — describing four dimensions of human development: intellectual, emotional, spiritual, and physical. More than eight decades later, Stephen Covey’s *The 8th Habit* would describe the same four dimensions and receive worldwide acclaim. The frameworks are, as AlmaU’s president has noted, essentially identical. Zhumabayev arrived first, writing in Kazakh, in Arabic script, in Tashkent — without international journals or institutional prestige. The insight, however, was entirely his.

Tolyq Adam is, among other things, an act of recovering this tradition and placing it where it belongs: at the foundation of a modern university curriculum.

The LEADER framework: philosophy operationalized

At the operational heart of Tolyq Adam is the LEADER framework — six competency domains that translate Abai’s philosophical triad into concrete educational aims.

L — Lifelong Learning addresses the disposition toward continuous growth, intellectual curiosity, and the metacognitive capacity to manage one’s own learning. In an era when the half-life of specific technical knowledge is shortening rapidly, the ability to learn, unlearn, and relearn is not a supplement to professional training — it is its foundation. This domain corresponds to *aqyl* at its most generative.

E — Empathy goes beyond interpersonal warmth to encompass emotional intelligence, social awareness, and what the program terms global consciousness. Empathy is not assumed — it is taught, through structured encounter with communities, perspectives, and texts that expand students’ moral imagination. This domain, and the Ethics domain that follows, express the *zhürek* dimension of Abai’s vision.

A — Analytical, Critical, and Creative Thinking forms the intellectual core. Crucially, this competency is framed not merely as logical reasoning but as intellectual independence — the ability to formulate one’s own questions, work with uncertainty, evaluate evidence, and generate original approaches. This is precisely the domain where human cognition remains most distinctively valuable alongside AI tools.

D — Dialogue and Discourse address the communicative dimension: academic writing, argumentation, public speaking, and the culture of reasoned disagreement. As AI-generated text becomes ubiquitous, the human capacity for authentic, situated, accountable communication becomes more — not less — valuable. One cannot outsource one’s voice.

E — Ethics and Ecology build the foundations of moral reasoning, civic responsibility, and environmental consciousness. Ethical judgment is irreducibly about what human beings care about and choose to prioritize — a domain where no algorithm can substitute for human deliberation.

R — Resilience encompasses psychological flexibility, self-discipline, and adaptability. The graduates of today will face not only technological disruption but social, ecological, and geopolitical volatility that demands a stable inner foundation. This domain expresses *qairat* — the disciplined will to remain whole under pressure.

Together, the six LEADER domains also correspond to four developmental axes used throughout the program: IQ (intellectual capacity), EQ (emotional intelligence), SQ (ethical and spiritual grounding), and PQ (physical and psychological resilience). These axes, incidentally, reproduce almost exactly the framework that Magzhan Zhumabayev proposed in 1922.

Curriculum in practice: eight courses, one arc

The LEADER framework is not an abstract statement of intent. It is embedded in eight courses that constitute the general education core at AlmaU, reaching more than one thousand students annually since the program launched in 2023. In just three years, the program has served more than 3,000 students across all undergraduate programs.

Before those courses could launch, AlmaU trained approximately fifty faculty members — in methodology, text selection, and the culture of academic dialogue. The preparation was as important as the curriculum itself.

Critical Thinking and Creative Writing (5 ECTS credits) builds intellectual independence and written voice through a reader spanning more than twenty authors from twenty countries — including Seneca, Kafka, Benjamin, Toffler, and Abai Kunanbayuly alongside contemporary thinkers. Students develop fourteen distinct writing practices, from free writing and annotation to loop writing and philosophical essay. The course is the program’s foundational intellectual experience.

Research Design (5 ECTS credits) — deliberately relocated from the senior year to the first year — introduces students early to the methods of inquiry: formulating questions, evaluating sources, working with data, and drawing warranted conclusions. The repositioning reflects a principled conviction: research culture is not a capstone skill but a foundational one.

Service Learning (3 ECTS credits) places students in direct partnership with communities and organizations — orphanages, ecological initiatives, homes for the elderly, social enterprises — to address real social challenges.

For AlmaU, social responsibility has never been treated as an auxiliary activity, separate from the academic mission of the university. It has been understood as part of education itself — a way of forming not only professionally competent graduates, but also citizens capable of empathy, responsibility, and meaningful action. In this sense, AlmaU has become one of Kazakhstan’s leading universities in embedding service learning and civic engagement into the culture of higher education.

This work began in a systematic form in 2015, when AlmaU introduced Service Learning as a mandatory course within the undergraduate curriculum, drawing on the experience of The Hong Kong Polytechnic University (PolyU). The decision was significant not only because it added a new discipline to the curriculum, but because it re-framed the relationship between the university and society. Students were no longer expected merely to study social problems from a distance. They were invited to encounter them directly, to work with communities, and to translate academic knowledge into practical forms of support.

Service Learning at AlmaU is built around a simple but powerful educational premise: students learn by serving. Through the course, they engage with nursing homes, organizations supporting people with disabilities, NGOs, and other social institutions. In these settings, knowledge becomes relational and practice-based. Students begin to see that education is not only a path to individual success, but also a responsibility toward others.

The introduction of the course was closely aligned with AlmaU's institutional mission: to educate qualified business professionals who understand their civic responsibility and are prepared to contribute to society. In a business university, this choice carried meaning. It signaled that leadership and entrepreneurship should not be reduced to efficiency, competitiveness, or market success alone. They must also include ethical judgment, social sensitivity, and the ability to create value for communities.

Over eleven years as a mandatory course at AlmaU, more than 9,914 students have participated, generating over 540 social projects and more than 252,000 hours of volunteer engagement. This is not a course about empathy. It is a practice of it.

Design Your Life (3 ECTS credits) adapts the life-design methodology developed by Bill Burnett and Dave Evans at Stanford University — best known through their 2016 book *Designing Your Life* — to the Tolyq Adam philosophical framework, helping students develop their own life strategies, values maps, and adaptive approaches to uncertainty. The aim is not career planning but self-authorship.

Great Books (4 ECTS credits) transforms Kazakhstan's mandatory cultural studies requirement into seminar-based engagement with primary texts — from Ortega y Gasset's *The Dehumanization of Art* — written in 1925, yet eerily prescient for an age of generative AI and Harari's *Sapiens* to Zira Naurzbayeva's *The Eternal Sky of the Kazakhs* — as vehicles for ethical reflection and cultural formation.

Philosophical Dialogue (5 ECTS credits) reimagines the mandatory philosophy requirement not as a survey course, but as what the program calls a “studio-lab”: students do not study philosophy — they practice it, constructing and defending their own philosophical positions, culminating in a personal manifesto grounded in the *aqyl-zhürek-qairat* triad. The course is not about Abai. It is in the spirit of Abai.

I, World (4 ECTS credits) offers an interdisciplinary inquiry into global complexity, political structures, social inequality, and civic identity — developing the analytical and communicative capacities required for responsible citizenship.

Biohacking (4 ECTS credits, based on the physical education requirement) reconceives physical culture as a science-informed practice of self-development — building the sustained habits of body and mind that *qairat* demands.

The pedagogical approach across all eight courses is deliberately consistent: small groups of no more than twenty-five students, instructors as facilitators rather than transmitters, formative assessment and peer feedback embedded throughout, and an explicit “culture of productive error” in which intellectual risk-taking is expected and mistakes are treated as data for growth.

Evidence of impact and external validation

The program’s outcomes are documented through student surveys covering 659 respondents in the 2024-2025 cycle. Across all six LEADER dimensions, strong majorities reported significant personal and academic development. Growth in critical, creative, and ethical thinking was reported by 92% of students. Enhanced capacity for reasoned dialogue and respectful disagreement by 89%. Gains in empathy and perspective-taking by 87%. Resilience and adaptability improved by 82%, while 85% reported stronger ethical awareness and 85% a deepened orientation toward lifelong learning. More than 100 community projects have been implemented by students in partnership with NGOs, schools, and social enterprises.

The external recognition of this work is significant not merely as institutional prestige but as signal. In 2026, at the Times Higher Education Asia Awards in Hong Kong — a competition that received over five hundred submissions from across seventeen countries — AlmaU was named winner in the Teaching and Learning Strategy of the Year category. The jury included representatives from Harvard Graduate School of Education, MIT, the Chinese University of Hong Kong, Tohoku University, and the World Bank’s higher education division. Among the finalists in the same category were Seoul National University, Singapore Management University, Hong Kong Polytechnic University, Hong Kong Baptist University, and Hiroshima University — a measure of the company AlmaU was keeping.

As the THE editor noted, the 2026 winners represented ten different countries and territories — a sign, as AlmaU’s president observed in his award address, that “what used to speak in one voice is now speaking in ten.” The global turn toward human-centered education is not a single tradition’s export. It is a convergent discovery, and the Kazakh philosophical school is among its legitimate contributors.

What this case teaches

The Tolyq Adam program matters for the broader conversation about curricular transformation in higher education for three reasons that go beyond the details of any one program.

First, it demonstrates that genuine curricular innovation is possible within regulated national systems. AlmaU operates under Kazakhstan’s State Compulsory Educational Standards yet found within those structures sufficient room to transform mandatory courses, introduce new disciplines, and reframe the entire general education curriculum around a single philosophical vision. Regulatory constraint, while real, is rarely an absolute barrier — it is often a design problem.

Second, it illustrates the pedagogical power of grounding. The program’s rootedness in Abai’s philosophy gives it coherence and motivational depth that a skills checklist would lack. Students are not being trained in competencies for labor market purposes. They are being invited into a living tradition of thought about what it means to become a whole person. That invitation, as the survey data suggests, generates a quality of engagement that is difficult to manufacture by other means. One student’s response — submitted in the award application and read by the jury in Hong Kong — captures it simply: *“The Seventeenth Word of Abai made me rethink my life.”*

Third, the program offers a direct answer to the question this chapter poses. What should universities teach in the age of AI? AlmaU’s answer is neither technophobic nor naive: it is a considered institutional claim that the capabilities most essential for nav-

igating an AI-transformed world are the ones that Abay Kunanbayuly identified in the nineteenth century, that Magzhan Zhumabayev systematized in 1922, and that AlmaU has now translated into a curriculum that a global jury of educators has recognized as among the best in Asia.

In this convergence, something important is visible for the field of higher education more broadly: the turn toward human-centered, value-driven curriculum does not require the erasure of cultural identity. It can be its fullest expression.

The graduates that Tolyq Adam is working to form are not in competition with AI systems. They are the people who will need to use, govern, interpret, and take responsibility for those systems — and for the communities those systems affect. For that role, no technical curriculum alone will suffice.

Chapter 11. Becoming an entrepreneurial, AI-native university in Bangladesh

Md. Sabur Khan

The next decade represents not another phase of gradual reform, but a period of structural discontinuity in higher education. AI systems increasingly challenge long-standing assumptions regarding expertise, learning, assessment, and employability. Simultaneously, universities are operating within broader transitions involving labor-market restructuring, demographic shifts, geopolitical fragmentation, and rapidly evolving skill ecosystems.

For many universities, especially in the Global South, the issue is not whether AI should be integrated, but whether institutions can redesign themselves fast enough to remain socially relevant and educationally legitimate. The prevailing global discourse on “AI-ready universities” remains structurally distorted because it is largely shaped by highly resourced institutions operating within exceptional ecosystems. Yet most universities globally operate under financial, infrastructural, demographic, and institutional constraints.

The transformation of DIU evolved through three interconnected phases. The first, digitalization, focused on initial institutional digitization: the introduction of learning management systems, online student services, and digital academic processes. The second phase, AI integration, saw AI systems begin to support specific institutional functions — among them an AI Professor, AI Proctor, Students Hub, and PerfectHR. The current third phase, the AI-native transition, involves embedding AI into the university’s broader institutional architecture across learning flows, assessment systems, administrative intelligence, career pathways, and entrepreneurship ecosystems. This trajectory reflects a broader shift from “using AI tools” toward redesigning institutional functions around environments where cognitive labor is increasingly abundant.

In this context, Daffodil International University (DIU) in Bangladesh represents a more representative institutional case: entrepreneurial rather than elite, adaptive rather than legacy-bound, and constraint-driven rather than capital-intensive. This paper therefore does not present DIU as a universal template, but as an exploratory case of institutional transformation under real-world conditions.

AI-native pedagogy: organizing teaching, learning, and assessment

One of the central questions raised in the book framework concerns how pedagogy changes when AI becomes a regular co-teacher and knowledge is no longer scarce. DIU’s experience suggests that AI-native pedagogy cannot simply involve adding AI tools to existing lecture models. Instead, learning itself must be reorganized. The AI era forces universities to reconsider not only what students should know, but what kinds of humans universities are cultivating. This aligns with the broader discussion in the book re-

garding the transition from purely technical education toward more integrative human development models.

As AI increasingly externalizes procedural and cognitive tasks, uniquely human capacities become more valuable: sense-making, ethical discernment, empathy, creativity, and cultural intelligence. This suggests that the future university may become less focused on information transfer and more focused on cultivating human flourishing within AI-mediated societies.

Where the traditional model proceeded from lecture to examination to grading, the emerging DIU model moves instead through AI exploration, human reflection, industry application, and output creation. This emerging learning flow repositions AI as exploratory and adaptive infrastructure, faculty as facilitators and sense-makers, and industry as a validation environment. The role of the university shifts from transmitting information toward orchestrating capability development.

Importantly, this transformation occurred under real constraints — limited faculty expansion, diverse student readiness, and resource limitations. Rather than building expensive new infrastructures, DIU experimented with hybrid approaches combining capstone projects, industry-integrated work, and entrepreneurship activities with AI, alongside studio-style collaborative learning environments. This aligns closely with the broader question posed in the book regarding what meaningfully replaces lecture-centric education in resource-constrained institutions.

Next, the emergence of generative AI challenges the legitimacy of many traditional assessment systems. Standardized examinations increasingly measure tasks that AI can already perform effectively. The DIU response has therefore shifted from memory-oriented evaluation toward capability-oriented assessment through portfolio-based evaluation, AI-supported project work, real-world problem solving, and industry-integrated capstones.

The critical issue is no longer whether students use AI, but whether institutions can evaluate judgment, discernment, creativity, ethical reasoning, and human–AI collaborative capability. This shift reflects one of the central tensions identified in the book framework: *how can universities assess meaningful human contribution when AI is embedded in almost every professional activity?*

Finally, AI-native universities require substantial redesigning of the faculty roles. Traditionally, faculty functioned primarily as content deliverers, subject authorities, and evaluators of student output. However, when content generation and certain forms of assessment become increasingly automated, these roles evolve toward learning architecture, curation of inquiry, ethical guidance, and the facilitation of reflection.

At DIU, faculty adaptation remains an ongoing process rather than a finished transition. Key challenges persist around readiness, technological anxiety, workload pressures, and rapidly changing expectations. At the same time, AI opens genuine possibilities: more personalized feedback, reduced repetitive administrative tasks, and enhanced interdisciplinary collaboration. Experience suggests that faculty transformation is not primarily technical — it is cultural and organizational.

AI-first administration and the integrated institutional ecosystem

The book framework emphasizes that AI-native universities must rethink operations, not merely teaching. At DIU, this meant moving beyond isolated administrative improvements toward a deliberate effort to build institutional intelligence at scale. Early experimentation with AI-supported systems — PerfectHR for predictive workforce

and performance management, StudentsHub for student engagement and behavioral insights, and AI Proctor for scalable assessment monitoring — demonstrated real gains in operational efficiency, student support, and institutional response times. Yet it also revealed a fundamental limitation: fragmented AI adoption, however well-intentioned, produces fragmented transformation.

This recognition led DIU toward developing a more integrated institutional ecosystem organized across four interrelated layers. Learning intelligence, anchored in the AI Professor and adaptive learning systems, connects directly to assessment intelligence through AI Proctor and AI-supported evaluation. These academic layers are in turn supported by institutional intelligence drawing on HR analytics, student data systems, and administrative decision support, while career intelligence — anchored in the Skill.jobs ecosystem and its integration with entrepreneurship and startup development — extends the university's reach beyond graduation. The central insight emerging from this experience is that AI transformation succeeds when ecosystems, not isolated tools, are integrated.

Building such an ecosystem, however, raises governance questions that remain unresolved not only at DIU but across higher education globally. How much decision-making should be delegated to AI systems? How should accountability be maintained in AI-mediated environments? And what forms of human oversight remain essential? These questions do not diminish the value of institutional AI ecosystems — but they make clear that constructing them is as much a matter of governance design as technical implementation.

Ethics, inclusion, and the paradox of constraint

AI introduces substantial risks for higher education systems, particularly in developing contexts. Institutional inequality, language barriers, data governance concerns, algorithmic bias, and AI dependency are not abstract threats — they are live pressures already shaping how universities in the Global South engage with AI transformation. DIU's experience underscores the importance of human oversight, inclusive access to AI tools, transparent governance, and local adaptation as practical responses to these pressures. The deeper structural challenge, however, remains unresolved: how can universities in the Global South avoid becoming dependent consumers of externally produced AI systems while lacking the capacity to shape them? This is one of the book's central concerns — how to prevent AI-driven stratification, both between elite and mid-tier institutions, and between Global North and Global South ecosystems more broadly.


Yet the DIU case simultaneously points toward a less expected possibility: that constraint itself may carry a form of institutional advantage. Operating under budget limitations, faculty preparedness gaps, infrastructure pressures, and regulatory uncertainty, DIU was pushed toward faster experimentation, lower-cost innovation, and entrepreneurial problem-solving precisely because the conventional pathways were unavailable. Where well-resourced institutions can afford deliberate, incremental reform, constrained institutions are often compelled to improvise — and in doing so, may move faster and more adaptively than their better-funded counterparts. The experience suggests that AI may enable certain universities to leapfrog traditional institutional development pathways rather than simply imitate elite models. The risk of dependency and the possibility of leapfrogging are not mutually exclusive — they are the twin realities that Global South universities must navigate simultaneously, and doing so thoughtfully may itself become a distinctive form of institutional expertise.

Blueprints in practice: designing the AI-native university

Based on DIU's experience, five interconnected stages of AI-native transformation emerge. The first is awareness — a leadership mindset transformation that preconditions everything that follows. The second is integration, involving foundational AI systems and institutional experimentation. The third is curriculum transformation, in which AI collaboration is embedded across disciplines. The fourth is governance, addressing ethics, transparency, accountability, and human oversight. The fifth and most advanced stage is ecosystem development — the integration of institutional intelligence across learning, operations, and employability into a coherent whole. This framework is not presented as prescriptive, but as an evolving institutional pathway shaped by experimentation, adaptation, and contextual realities.

While the institutional AI architecture described earlier operates at the level of the university as a whole, one of the most instructive examples of AI-native education in practice emerges at the departmental level. The Department of Innovation and Entrepreneurship at DIU has developed an integrated AI-based project tracking ecosystem that illustrates, in concrete terms, what it means to redesign a learning environment around AI rather than simply add AI tools to an existing one.

The system addresses a challenge familiar to any university running multi-year project-based curricula: how to supervise, support, and document student development at scale without reducing that process to a series of disconnected submissions and grades. The solution DIU developed is organized around a web-based dashboard that serves as a central hub connecting students, faculty supervisors, industry advisors, and program coordinators within a single environment. Each student enters the system through a unique digital portfolio ID assigned at the start of their first capstone project, which persists across all eight projects in the sequence and beyond graduation. This seemingly simple design choice has significant consequences: rather than producing a collection of course outputs, students accumulate a longitudinal record of entrepreneurial development — a living document of how their thinking, their ventures, and their capabilities evolved over time.



The ecosystem operates through three interrelated components. The AI engine functions as a continuous monitoring layer, tracking student engagement, submission patterns, and activity levels, and using this data to flag at-risk students, predict completion timelines, and generate targeted alerts to supervisors before problems escalate. This transforms supervision from a reactive to a proactive function: rather than discovering that a student is struggling at the point of a missed deadline, supervisors receive early signals that allow for timely intervention. The student portfolio component tracks the full arc of each student's entrepreneurial journey — from initial idea generation and market analysis through operational planning, financial modeling, expansion strategy, and final implementation. An AI-generated skills heatmap visualizes competency development across domains including ideation, market research, marketing, operations, management, finance, and strategy, producing an evidence-based record that carries meaning beyond the academic context, serving as a credential for future employers and investors. The third component, the case study database, captures anonymized records of student projects — including the decisions made, the pivots undertaken, the failures encountered, and the lessons drawn — which become a searchable institutional resource available to future cohorts. As illustrated in Figure 1, student work flows continuously into this database, which in turn feeds forward into the learning environment of the next generation of students, creating a self-reinforcing knowledge loop in which each cohort contributes to the infrastructure inherited by the one that follows.

Several design features of the system deserve particular attention. The case study database does not record only successes. It deliberately documents failures, incorrect assumptions, and strategic pivots, and makes these visible to future students in anonymized form. This reflects a considered pedagogical choice: by normalizing iteration and reducing the stigma of failure, the system encourages students to take the kinds of risks that entrepreneurial learning requires. A “Case of the Trimester” feature recognizes the strongest project from each cohort as a teaching case in subsequent classes, creating peer recognition and an institutional culture of excellence without imposing a purely competitive dynamic. Students can also use the database proactively — searching for cases that match challenges they are currently facing, reading how previous students navigated similar situations, and applying those insights to their own work.

The ethical architecture of the system is equally deliberate. Student participation in the case study database is voluntary, governed by explicit consent, and subject to comprehensive anonymization protocols that remove all identifying information before any case enters the shared repository. Students retain the right to exclude sensitive sections — such as financial projections from a public view, and industry supervisors must provide written approval for any project involving company data. Participation carries no grade implications. These safeguards reflect an awareness that building institutional knowledge through student work creates genuine obligations of transparency and trust — obligations that become more rather than less important as AI systems become more deeply embedded in educational environments.

What this departmental ecosystem demonstrates, ultimately, is a working model of the principles this chapter has argued for at the institutional level. Learning is organized around real problems and cumulative development rather than discrete assessments. AI functions as infrastructure for support, monitoring, and knowledge generation rather than as a replacement for human judgment. Faculty and industry advisors contribute forms of guidance — ethical reasoning, critical evaluation, market validation — that remain distinctively human. And the institution itself becomes a learning system, not merely a delivery mechanism, accumulating and transmitting knowledge across cohorts in ways that compound in value over time.

Conclusion

The future of universities will not be determined solely by who adopts AI first, but by which institutions successfully redesign themselves around new realities of intelligence, work, and human development. The DIU case suggests that mid-tier universities can become significant sites of innovation, that constraint-driven institutions may evolve faster than expected, and that AI-native transformation is organizational, not merely technological.

Most importantly, universities must decide whether they will passively consume AI systems, or actively participate in shaping the educational, ethical, and societal futures those systems create. The next decade may determine not only the future of universities, but also the role universities play in preserving human-centered development within AI-mediated societies.

References

- Brynjolfsson, E., & McAfee, A. (2017). *Machine, platform, crowd: Harnessing our digital future*. W. W. Norton & Company.
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116. <https://hbr.org/2018/01/artificial-intelligence-for-the-real-world>
- Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial intelligence in education: Promises and implications for teaching and learning*. Center for Curriculum Redesign.
- McKinsey Global Institute. (2023). *The economic potential of generative AI: The next productivity frontier*. McKinsey & Company. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>
- Miao, F., & Holmes, W. (2023). *Guidance for generative AI in education and research*. UNESCO. <https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>
- OECD. (2023). *OECD digital education outlook 2023: Towards an effective digital education ecosystem*. OECD Publishing. <https://doi.org/10.1787/c74f03de-en>
- World Bank. (2020). *Realizing the future of learning: From learning poverty to learning for everyone, everywhere*. World Bank Group. <https://documents.worldbank.org/en/publication/documents-reports/documentdetail/250981606928190510>

Chapter 12. Continental impact and indigenous/African legal governance

Letlhokwa George Mpedi

We are undergoing seismic shifts globally. Africa stands at a crossroads. By 2050, one in every four people on the planet will be African (UN DESA, 2022), a demographic shift that presents both unparalleled opportunity and profound urgency. The continent's vast natural resources and youthful population sit alongside pressing challenges such as high unemployment, enduring inequality, the accelerating disruptions of climate change, and the risk of being left behind in the global race for artificial intelligence (AI) leadership. The Roman god Janus, who looks simultaneously backwards and forwards, offers a useful metaphor for Africa's current condition. In political theory, this has been described as a Janus-faced moment that simultaneously holds the promise of emancipation and the threat of deepening underdevelopment (Nairn, 1977). Against this backdrop, it is imperative to ask if African universities will be passive recipients of AI or active architects of it.

This chapter reflects on South Africa's University of Johannesburg's (UJ) evolving journey as what may be regarded as Africa's most deliberate, institutionally grounded experiment in building an AI hub within a higher education setting. This is not a triumphalist account of achievement (there is still too much to be done for that), but an honest reflection on what it means to pursue innovation with societal purpose at its center. Drawing on scholarly work at the intersection of AI and law, UJ's institutional initiatives, and the intellectual legacy of Professors Ihron Rensburg and Tshilidzi Marwala, this chapter examines what UJ has built and situates that work within the broader African context. The argument is that UJ's significance as an AI hub lies in the deliberate orientation of technological infrastructure towards measurable societal impact.

Institutional background

Any honest account of UJ's trajectory as a technology-forward institution must begin with the foundational work of Professor Ihron Rensburg, UJ's inaugural Vice-Chancellor and Principal. It was Rensburg who took on the audacious task of forging a coherent, ambitious university out of the 2005 merger of four historically distinct institutions (the Rand Afrikaans University, the Technikon Witwatersrand, and the Vista University campuses in Soweto and the East Rand) (UJ, 2005). That this merger produced, within two decades, an institution ranked among Africa's leading research universities is the consequence of deliberate, values-driven leadership rooted firmly in Pan-Africanism. Rensburg's tenure was defined by a commitment to transformation and access that gave UJ its distinctive character. The technological ambition that UJ subsequently pursued in the fourth industrial revolution (4IR) era was made possible by these institutional foundations that Rensburg painstakingly constructed.

During Marwala's tenure, UJ made several consequential institutional commitments that would define the university's identity as an AI-forward institution. Where Rensburg had built the organizational foundations, Marwala gave UJ an intellectual direction by positioning the university explicitly and publicly as Africa's leading institution in the 4IR (Marwala, 2020). A compulsory AI course was introduced for all students, regardless of their field of study. This was a curricular decision of considerable ambition that signaled, at an institutional level, that AI literacy was part of the common intellectual formation of every UJ graduate. The Institute for Intelligent Systems (IIS) was established as a dedicated research entity, providing an organizational home for the university's 4IR ambitions and a platform for sustained engagement with industry and government (IIS, 2016). UJ's national co-hosting of the Artificial Intelligence Institute of South Africa (AII-SA) alongside the Tshwane University of Technology (TUT) was a direct response to the recommendations of the Presidential Commission on the Fourth Industrial Revolution (PC4IR, 2019) and further cemented the university's standing as a node of national AI strategy. These structural choices established UJ as a recognizable actor in continental and global AI conversation.

The current institutional leadership has sought to consolidate what has been built and deepen its roots in societal purpose. The Strategic Plan 2035 that subsequently emerged reflects this orientation (UJ, 2024).

Strategic architecture: the Plan 2035 and technology for the future

UJ's Strategic Plan 2035, launched in October 2024, provides the institutional scaffolding through which the university's AI ambitions are operationalized (UJ, 2024). The plan is organized around three defining pillars: Societal impact and sustainability; Global footprint and partnerships; and Technology for the future.

The Technology for the future pillar is the most directly relevant to UJ's AI hub ambitions but treating it in isolation would be a conceptual error. Technology at UJ is conceived as an instrument of societal transformation, not an end in itself. This reflects a substantive philosophical commitment to harnessing AI responsibly and ensuring its benefits reach those who have historically been excluded from technological progress. This commitment finds growing support in the global higher education literature (Miao et al., 2021; Schroeder, 2021).

The operationalization of this vision has taken several institutional forms. The IIS continues to serve as the primary research engine, working across machine learning, AI, Internet of Things, advanced communications, 3D printing, and blockchain (IIS, 2016). In 2025, the launch of UJ Digital as a comprehensive online learning platform reflected an understanding that extending AI's reach requires first extending educational access. The platform currently offers more than 80 Short Learning Programs (SLPs) and 17 fully online qualifications (UJ, 2025a). This year, UJ became the first South African university to implement AI-driven facial recognition for student identity verification at scale, demonstrating that institutional operations can themselves become sites of AI application and refinement (UJ, 2026).

Perhaps most representative of the innovation-to-impact aspiration is the UniPod, launched in partnership with the United Nations Development Program (UNDP) as part of its *timbuktoo* pan-African innovation initiative (UNDP, 2025). Located on UJ's Auckland Park campus, the UniPod provides aspiring entrepreneurs with access to AI laboratories, prototyping equipment, incubation programs, and investor networks. It was conceived not as a showcase of technological sophistication but as infrastructure for young

Africans whose ideas have historically lacked the environment in which to take root. As articulated at its launch, the conviction underpinning the UniPod is that “innovation does not begin with hardware. It begins with hope, access and the right environment to dream boldly” (Mpedi, 2025a). The UniPod is a node in a pan-African network and an early expression of the collaborative model that the global literature increasingly identifies as essential for sustainable AI development in low- and middle-income contexts (Cloete et al., 2015; Tandon & Chakraborty, 2021).

AI in service of society: from hub to heartbeat

The designation of UJ as an AI hub risks, if one is not careful, becoming a badge of technological prestige rather than a description of developmental purpose. The more demanding question is not whether AI research and innovation concentrate at UJ, but whether it radiates outward into communities and into policy. This is the distinction between an AI hub as an institutional asset and AI as a societal force.

Research on the role of universities in national innovation systems consistently finds that translating research outputs into societal benefits is neither automatic nor guaranteed; it requires deliberate institutional mediation (Etzkowitz & Leydesdorff, 2000; Maassen & Stensaker, 2011). UJ’s approach acknowledges this directly. The AI-driven facial recognition system introduced for student registration was pursued because identity verification delays were creating tangible access barriers, disproportionately affecting students from disadvantaged backgrounds (UJ, 2026). Moreover, UJ Digital was launched because physical campus capacity cannot accommodate the full demand for higher education in a country where youth unemployment exceeds 45% and where geography, cost, and life circumstances prevent millions from accessing contact learning (Stats SA, 2024). In each case, the AI application exists when read against the social problem it addresses.

UJ’s Master of AI program, offered through the IIS, further reflects this orientation. Designed to develop graduates who work at a high level of competence across engineering, commerce, the natural sciences, and the law, the program is explicitly framed around AI research relevant to real-world practice (UJ, 2025b). The intent is not simply to produce technically capable graduates, but graduates who are intellectually equipped to ask whose problems their AI systems serve, and by what values they are governed. This is increasingly recognized in the literature on transformative higher education as the essential disposition for graduates who will operate in societies marked by deep inequality (Walker & McLean, 2013).

This is a recognition that the most sophisticated research often emerges from engagement with complex real-world problems. The most transformative education occurs when students learn in and with communities, and universities are strongest when deeply rooted in the societies they serve. Community-engaged learning should thus become standard practice. Johnson (2020) argued that community engagement professionals often refer to this work as “the stepchild of higher education” because it is not valued or funded in the same way as teaching and research. This, however, is precisely what must change if we are serious about societal impact. Every program at UJ thus includes opportunities for students to apply their knowledge to real local challenges such as unemployment, service delivery, small business development or social cohesion. This serves a dual purpose. It develops graduate capabilities while contributing to immediate community needs. Our students learn by doing, and communities benefit from their emerging expertise.

No example illustrates this more concretely than UJ's long-running societal impact project in the remote villages of Gwakwani, Matatani, and Mbodi in the Vhembe district of Limpopo, situated between Musina and Thohoyandou. Initiated in 2013 through the Faculty of Engineering and the Built Environment (FEBE) in partnership with Schneider Electric South Africa, the project with an acknowledgement that without electricity, there could be no sustainable clean water supply, and without clean water and connectivity, education remained out of reach for children who were walking hours each day simply to reach a school (UJ, 2023). What followed was a decade-long program of integrated technological interventions in solar energy, water infrastructure, and internet connectivity, explicitly designed around community needs rather than institutional convenience. The measurable outcomes speak directly to the innovation-to-impact thesis of this chapter. Hanyani Secondary School, located outside Thohoyandou and serving learners from these villages, recorded a matric pass rate of 57% in 2023. Following UJ's sustained educational engagement as part of the project, that figure rose to 92.1% in 2024 (UJ, 2025d). More recently, the project has expanded into economic empowerment, with a sewing initiative introduced in Gwakwani to equip women in the community with practical skills and pathways to income generation (Mpedi, 2025b). The Gwakwani project illustrates what responsible AI-era university engagement with marginalised communities should look like.

The AI and law nexus: filling the governance gap

One of UJ's most distinctive contributions to continental AI conversation lies at the intersection of artificial intelligence and law. This is a domain that remains largely undertheorized in the African context. The collaborative volume *Artificial Intelligence and the Law* (Marwala & Mpedi, 2024), published by Springer Nature, addresses a striking gap in the existing literature. While a voluminous global discourse on AI governance exists, much of it is generated from the perspective of European or North American legal traditions. The book examines how AI intersects with a wider range of legal systems operative in South Africa and across Africa, including Roman-Dutch law, English common law, customary and indigenous law, and religious personal law. The intersection of AI with customary law is particularly underexplored.

The urgency of this work cannot be overstated. Scholars of technology governance have observed that regulatory frameworks for emerging technologies tend to be captured early by the jurisdictions with the greatest technical capacity and the loudest institutional voice, and that late-arriving jurisdictions often find themselves adopting frameworks designed for different social, economic, and legal contexts (Bradford, 2020; Pistor, 2019). Africa is at risk of precisely this dynamic with respect to AI. The continent's legal systems embody distinct conceptions of personhood, community, obligation, and justice that have no straightforward equivalent in the training data or the regulatory assumptions of the AI systems now being deployed at scale across African markets. Customary law, for instance, governs significant aspects of land tenure and dispute resolution for millions of Africans, yet it is almost entirely absent from the datasets on which legal AI tools are trained, and from the governance frameworks being developed elsewhere. If African universities do not generate the scholarship that makes this absence visible and the frameworks that begin to address it, it is unlikely that anyone else will. This is not a peripheral academic concern.

The establishment of UJ's AI and the Law Institute, launched in October 2025 within the Faculty of Law, gives this scholarly program an institutional home (UJ, 2025c). The

Institute is designed to generate African-grounded analysis of regulatory developments, including the EU AI Act (2024), thereby ensuring that Africa's governance frameworks are not simply adapted from European templates but are built on African conditions, legal traditions, and developmental priorities. A forthcoming volume on *Artificial Intelligence and Employment Law* (Marwala & Mpedi, forthcoming) extends this program into one of the most consequential social domains of the AI age. In South Africa, where youth unemployment exceeds 45% (Stats SA, 2024), the transformation of labor markets by AI is particularly urgent. UJ's contribution is to bring interdisciplinary rigor to these questions before the policy window closes.

Africa must balance AI's benefits and risks through appropriate governance and ethical regulation. This requires African scholars to engage seriously with the technical dimensions of AI, and African legal academics and policymakers to develop governance frameworks grounded in African constitutional values and customary legal traditions. The EU AI Act, the OECD AI Principles, and the UN Secretary-General's AI Advisory Body are among the sites where global norms are being forged (OECD, 2019; United Nations, 2024). However, Africa's voice in these processes remains too quiet in relation to the continent's size and its stake in the outcome. Part of UJ's contribution must be to produce the scholarship, the trained graduates, and the policy engagement that amplifies African voices in these global conversations.

Reflections on the innovation-to-impact journey

In less than two decades, UJ has moved from a contested merger of historically distinct institutions to a recognized actor in continental and global AI conversation. However, achievement in a continent facing the challenges Africa faces cannot be measured by internal metrics alone. Though rankings matter and are attended to, the real test of UJ's AI hub is whether AI research and education at UJ generate tangible improvements in the lives of South Africans and Africans more broadly. We must interrogate whether the innovation ecosystem generates decent work and whether legal and governance research produces frameworks that protect citizens.

Against these tests, the journey from innovation to impact remains incomplete. Computing capacity and reliable connectivity remain constraints at many African universities (Gillwald, 2020). The global competition for AI expertise means that African graduates with strong technical skills face considerable incentives to pursue careers outside the continent. This is a dimension of the brain drain problem that is well-documented in the higher education literature (Ndulu, 2004). AI research is expensive, and the gap between what African universities can resource internally and what their global competitors spend remains significant (Cloete et al., 2015). Finally, there is the challenge of ensuring that the AI systems built and researched at UJ address genuine African priorities rather than reproducing problems imported from other contexts.

These challenges are real and acknowledging them is a clarion call for deliberate institutional strategy, stronger Pan-African partnerships, more sustained government investment in university-based research, and a frank assessment of where AI can best serve African development.

Conclusion

African universities are tasked with generating research, training talent, building governance frameworks, and advocating for policies that ensure AI serves African development rather than undermining it. Janus, looking backwards, sees a UJ that has travelled an extraordinary distance in a short time. Looking ahead, Janus sees terrain that is less mapped but no less compelling. The road from innovation to impact is long, and UJ has been walking it with deliberate purpose. Yet there is much further to go, and Africa's Janus-faced moment demands precisely this.

References

- Bradford, A. (2020). *The Brussels effect: How the European Union rules the world*. Oxford University Press. *Change: Book title should be in sentence case in APA 7, not title case.*
- Cloete, N., Bailey, T., & Maassen, P. (2015). *Universities and economic development in Africa*. African Minds. *Correct as is.*
- Etzkowitz, H., & Leydesdorff, L. (2000). The dynamics of innovation: From national systems and "Mode 2" to a triple helix of university–industry–government relations. *Research Policy*, 29(2), 109–123. *Change: Title and journal name formatting corrected to sentence case per APA 7; en-dash used between page numbers.*
- European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union. *Correct as is.*
- Gillwald, A. (2020). *Beyond access: Addressing digital inequality in Africa*. Research ICT Africa. *Correct as is.*
- Institute for Intelligent Systems. (2016). *About the IIS*. University of Johannesburg. <https://iis.uj.ac.za/> *Change: Remove "(IIS)" from author name — parenthetical abbreviations are not used in APA 7 author fields.*
- Johnson, B. J. (2020). Community engagement: Barriers and drivers in South African higher education. *South African Journal of Higher Education*, 34(6), 87–105. *Correct as is.*
- Maassen, P., & Stensaker, B. (2011). The knowledge triangle, European higher education policy logics and policy implications. *Higher Education*, 61(6), 757–769. *Correct as is.*
- Marwala, T. (2020). *Closing the gap: The fourth industrial revolution in Africa*. Pan Macmillan South Africa. *Change: Title to sentence case; remove line break between "Pan" and "Macmillan."*
- Marwala, T., & Mpedi, L. G. (2024). *Artificial intelligence and the law*. Springer Nature Singapore. <https://doi.org/10.1007/978-981-97-2827-5> *Change: Title to sentence case.*
- Marwala, T., & Mpedi, L. G. (in press). *Artificial intelligence and employment law*. Springer Nature. *Change: APA 7 uses "in press" not "forthcoming" for works accepted but not yet published.*
- Miao, F., Holmes, W., Huang, R., & Zhang, H. (2021). *AI and education: Guidance for policy-makers*. UNESCO. *Correct as is.*
- Mpedi, L. G. (2025a). *Remarks at the launch of the UniPod*. University of Johannesburg. *Correct as is.*

- Mpedi, L. G. (2025b). *Vice-Chancellor message, 25 July 2025*. University of Johannesburg. <https://news.uj.ac.za/> *Correct as is.*
- Nairn, T. (1977). *The break-up of Britain: Crisis and neo-nationalism*. New Left Books. *Change: Title to sentence case.*
- Ndulu, B. J. (2004). Human capital flight: Stratification, globalisation, and the challenges to tertiary education in Africa. *Journal of Higher Education in Africa*, 2(1), 57–91. *Correct as is.*
- OECD. (2019). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). OECD Publishing. *Change: The document identifier should appear in parentheses after the title, and publisher should be “OECD Publishing” rather than omitted.*
- Pistor, K. (2019). *The code of capital: How the law creates wealth and inequality*. Princeton University Press. *Change: Title to sentence case.*
- Presidential Commission on the Fourth Industrial Revolution. (2019). *Report of the Presidential Commission on the Fourth Industrial Revolution*. The Presidency of South Africa. *Change: Remove “(PC4IR)” abbreviation from author field.*
- Schroeder, R. (2021). *Social theory after the internet: Media, technology and globalization*. UCL Press. *Change: Title to sentence case.*
- Statistics South Africa. (2024). *Quarterly Labour Force Survey Q3 2024*. Statistics South Africa. *Change: Remove “(Stats SA)” abbreviation from author field.*
- Tandon, R., & Chakraborty, D. (2021). Community engagement and higher education: Exploring the nexus. *Journal of Higher Education Outreach and Engagement*, 25(1), 1–18. *Correct as is.*
- United Nations Department of Economic and Social Affairs. (2022). *World population prospects 2022: Summary of results*. United Nations. <https://population.un.org/wpp/> *Change: Remove “(UN DESA)” abbreviation from author field; title to sentence case.*
- United Nations Development Programme. (2025). *UNDP and UJ unveil South Africa’s first UniPod, accelerating innovation and entrepreneurship*. <https://www.undp.org/south-africa/press-releases/> *Change: Remove “(UNDP)” abbreviation from author field.*
- United Nations. (2024). *Governing AI for humanity: Interim report*. UN Secretary-General’s Advisory Body on Artificial Intelligence. *Change: Title to sentence case.*
- University of Johannesburg. (2005). *UJ founding and merger history*. <https://www.uj.ac.za/about/> *Correct as is.*
- University of Johannesburg. (2023). *UJ’s multinational outreach project connects engineering, technology and community*. UJ News. <https://news.uj.ac.za/> *Correct as is.*
- University of Johannesburg. (2024). *Strategic Plan 2035*. <https://www.uj.ac.za/about/strategic-plan/> *Correct as is.*
- University of Johannesburg. (2025a). *UJ Digital: Online learning platform*. <https://www.uj.ac.za/> *Correct as is.*
- University of Johannesburg. (2025b). *Master of Artificial Intelligence (Coursework)*. <https://www.uj.ac.za/university-courses/master-of-artificial-intelligence/> *Correct as is.*
- University of Johannesburg. (2025c). *Launch of the AI and the Law Institute*. UJ News. <https://news.uj.ac.za/> *Correct as is.*

University of Johannesburg. (2025d). *UJ's societal impact project helps Limpopo school achieve record matric results*. UJ News. <https://news.uj.ac.za/> Correct as is.

University of Johannesburg. (2026). *UJ drives innovation to meet record demand for study space through digital registration and online learning*. UJ News. <https://news.uj.ac.za/> Correct as is.

Walker, M., & McLean, M. (2013). *Professional education, capabilities and the public good: The role of universities in promoting human development*. Routledge. Change: Title to sentence case.

Chapter 13. The university as regenerative ecosystem: a planetary mission experiment in Mexico

Victoria Haro

We are observing how AI is calling into question the knowledge-transmission function of universities and intensifying a longstanding question: what universities are for beyond information delivery. This is not the first time their social purpose has been called into question. In recent decades the growing awareness of global socioenvironmental crises highlighted the need for new ways of learning, relating, and acting on the planet.

Some universities responded more substantially, rethinking their missions and curricula, while others did so more superficially by adding some sustainability-related courses or programs. A few were already fully oriented around a socioenvironmental mission, and new institutions also emerged with this mission at their core. Universidad del Medio Ambiente in Mexico (UMA) is one such example.

A university with a socioenvironmental mission is oriented toward cultivating the capacities and conditions that foster synergistic human and planetary wellbeing, what Luksha and Taddei (2025) describe as a *planetary mission* for the university. This purpose cannot be achieved through new knowledge content alone. It also requires transformations in pedagogy, practices, processes, and institutional design. In this context, formation, meaning-making, ecological relationship, and community become the primary organizing principles of the institution, rather than peripheral concerns.

The current AI disruption is calling into question the role of content delivery and weakening the value proposition of credentials, bringing universities back—once again—to the question of purpose. Dominant responses—skills training, AI integration, credential reform—largely remain within the frame that the university exists to serve the knowledge economy.

Universities with a planetary mission, however, point to another path. These institutions are organized around purpose-driven, relational, and place-based commitments—not as a reaction to AI, but because they were addressing a different question from the outset. As a result, they are less dependent on the information-transfer model that AI is now challenging, and therefore less compelled to “pivot” in response.

These universities may offer insight into a different model of higher education—one that engages in the challenges posed by AI and takes up the socioenvironmental. UMA is one such institution. Our experience offers a situated example of an exploration around formation rather than information transfer, while continuing to evolve in practice.

Why a university as a regenerative ecosystem?

Universidad del Medio Ambiente (University of the Environment, UMA), was founded in Mexico in 2009, with the mission of *fostering a regenerative, sustainable, and ethical future by forming collective-change agents capable of promoting initiatives that transform socioenvironmental systems*.

In practice, this means that UMA is continually exploring how to create the conditions for people to cultivate regenerative ecosystems—social and ecological—to emerge and thrive over time. This involves nurturing relationships, practices, and forms of organization that restore and sustain the health of interconnected human and natural systems.

UMA is the first formal university in Latin America focused on socioenvironmental studies — by *formal universities* we refer to institutions that grant accredited degrees and are organized around a physical campus. It is one of around ten in the world. This kind of university—entirely focused on sustainability and regeneration across disciplines—is still an emerging idea. UMA is a small university set in a forest village two hours from Mexico City. Each year, we host on average 150 graduate students, along with 150 participants in shorter, non-degree courses.

We present UMA as a question, not as an answer, recognizing that humanity is still in the early stages of understanding what a regenerative and sustainable civilization may look like, let alone how to achieve it. This is uncertain ground, but also an exciting one to inhabit.

From this perspective, it seems that universities with a planetary mission cannot be anything other than spaces of collective inquiry, rather than spaces of learning alone. For this reason, UMA has explored the idea of a university as a *community of practice*: a community that collectively inquires and learns by putting ideas, methods, and projects into practice; learning from achievements, surprises and excellent errors.

Inspired by Gandhi's (1909) powerful ethos of reflecting the ends in the means, if our community of practice aims to foster regenerative ecosystems in the world, then we are called to experiment with how the university itself may be a regenerative ecosystem—one that enables renewal and allows new potential to flourish for people, relationships, community, place, knowledge, and professional practice.

Gandhi's idea that “the means are the ends” is often understood as a moral or ethical call for congruence — which it is — but it also carries a deeply practical insight. On the one hand, the means are the seed, or DNA, of the ends: we cannot harvest what has not been sown. On the other hand, and perhaps more interestingly, Gandhi teaches that to understand what an end is, we must attempt to express it in the means. It is a research proposal: use the means to discover the ends. Be the change you want to see in the world. UMA can therefore be seen as a kind of lab where we are experimenting with what it means to be a regenerative ecosystem that fosters regenerative ecosystems.

UMA as a learning experiment

We founded UMA out of deep concern for the socioenvironmental future, and with a recognition of the emerging potential of sustainability and regenerative movements. We saw higher education as a powerful way to contribute to this transformation and were inspired to explore new processes for human creativity and collective endeavor.

We sensed that a different kind of university was needed, though we did not yet know what it should look like. To explore this, we visited institutions around the world working at the intersection of sustainability and regeneration. We then convened a group of 40 people in Mexico, with experience in education and socioenvironmental systems change across diverse disciplines, to codesign UMA.

Together, we co-created the university's mission and imagined a space to engage with questions of sustainability and systems change across disciplines and worldviews—

while developing real-world projects, learning how to learn, exploring diverse thinking frameworks, and critically inquiring into professional and personal development as interconnected dimensions of a single process. We envisioned not only a place for learning, but also one for exploring purpose, meaning, nature, and human connection.

We began with three master's degrees not offered in Latin America at the time and which we saw as key nodes of societal transformation: Environmental Law and Public Policy, Socioenvironmental Business, and Sustainable Architecture. UMA was launched as a graduate school with hybrid programs combining short, intensive in-person workshops with virtual learning, enabling participation from experienced professionals engaged in real-world projects.

Over time, we expanded our offerings to include Education for Sustainability, Agroecology and Regenerative Food Systems, Sustainable Tourism, and Socioenvironmental Projects. We also developed non-degree programs and, in 2021—after 12 years focused on graduate education—launched our first bachelor's degree in Entrepreneurship and Socioenvironmental Projects: a fully residential program designed to prepare young collective-change agents.

To advance our mission and give form to this vision, we embedded a set of core strategies into the design of our programs and learning processes. These have been developed and refined in practice over the past 17 years, operating across multiple levels, from pedagogical design to relational conditions and learning environments. Below, we describe each strategy and share the insights and challenges that have emerged through their implementation.

Project lab and communities of practice

Learning is project-based across all programs. Students develop real-world projects and are organized into small communities of practice, where they share progress and reflect through an action research approach: strengthening collective learning, community, and collaborative capacity.

Working with real projects raises the stakes of learning. Curiosity and demand for applicable knowledge become highly visible. This creates a high-energy environment, but also stress and frustration as projects meet real-world complexity.

A recurring learning process is the gradual recognition of what constitutes a “realistic” project. While challenging—particularly in early stages—this becomes one of the most formative aspects of the UMA experience, shifting learning from abstraction toward situated judgment and responsibility.

We have also observed that connecting collective-change agents within communities of practice significantly expands students' willingness to experiment and reimagine what is possible. These communities often function as relational safety nets, enabling forms of creative and intellectual risk-taking that would be unlikely in isolation.

System change

Projects are oriented toward socioenvironmental systems change and developed through codesign processes with those involved, fostering situated and multi-perspective understandings. Students work at different scales—from family to neighborhood to industry—choosing what is most feasible, relevant, and meaningful. Moving from linear change perspectives to systemic ones takes time, but it does develop. We work at two

complementary levels: learning to weave perspectives through the codesign process, and learning to make sense of the collective understanding of the system.

System mapping is one of the key tools we use. We observe students improve their mapping capacity over time, but in final (graduation) projects, approximately 30% of maps still contain inaccuracies. Rather than interpreting this solely as a limitation of learning, we see it as evidence of the inherent difficulty of working with complex systems representation. For this reason, we continue to use system mapping, while recognizing that its value lies as much in the learning process it activates.

What matters is the capacity to understand the system sufficiently to design meaningful pathways for change. To support this, the Project Lab takes students through a developmental progression: beginning with their personal systems, moving to organizational systems, and eventually engaging with broader socioenvironmental systems. This staged approach allows students to build systemic sensitivity through progressively more complex contexts.

Almost all students are able to foster valuable systemic changes in their personal and organization systems; but socioenvironmental change requires multiple cycles of experimentation, reflection, and adaptation that extend beyond the timeframe of formal education. While our intention is that students graduate with a well-founded theory of change supported by collective intelligence, we also recognize that not all projects reach full implementation during or after the program. In many cases, graduation marks the transition into a longer trajectory of learning and action.

For this reason, we increasingly understand systems change at UMA as the cultivation of the capacity to remain engaged in iterative processes of understanding and transforming complex systems over time.

Thinking frameworks

Projects function as a space for experimenting with what is learned in the program. At the same time, they are sustained by it. The foundation is provided by a program-wide process of “thinking classes”, focused on learning and practicing different ways of thinking.

Systems thinking provides interconnectedness; critical thinking supports evidence evaluation in socioenvironmental contexts; evolutionary thinking frames nature as ongoing change; and ethical thinking invites reflection on what “doing good” means in context, avoiding fixed external prescriptions.

Engaging with these frameworks is both intellectually demanding and, at times, uncomfortable—but, from our perspective, central for the development of change agents. We have observed initial resistance in learning to assess scientific evidence, shaped in part by fragmented disciplinary culture (“if we are not scientists, we cannot understand”). However, once the range of controversies and evidentiary structures becomes more visible, change agents recognize how easily effort may be directed toward problems that are not well founded in evidence or are incorrectly framed. We have also come to see the value of methodological ethical inquiry for change agents, as it enables them to make explicit and revise assumptions about what constitutes “the right thing” in complex socioenvironmental contexts—especially questioning the assumption that ethical intentions imply ethical results.

Disciplinary strength

Each program offers advanced knowledge in its field, contextualized within a socio-environmental perspective and taught by experienced practitioners actively engaged in sustainability and regeneration within their disciplines.

The current socioenvironmental situation requires a fundamental rethinking of how we produce, consume, feed ourselves, educate ourselves, build, and more broadly organize collective life. This implies not only applying existing knowledge differently, but revisiting the purpose and practices of each discipline.

In our experience, one of the key challenges is holding together two movements that are often in tension: developing sufficient depth within a discipline, while also questioning and transforming its foundations. Without depth, proposals risk becoming superficial; without critical revision, they risk reproducing the same patterns that contribute to current challenges.

We have found that working with practitioners actively engaged in their fields is essential in navigating this tension. Rather than offering finalized knowledge, they bring processes that are still unfolding—marked by constraints, trade-offs, and experimentation. This exposes students to the reality that transforming a discipline is not a linear process, but one that involves navigating uncertainty and contradiction in practice.

At the same time, students are often preparing to act in professional contexts that are themselves in transition. In some cases, the fields they are entering are only beginning to integrate socioenvironmental perspectives; in others, the forms of practice they seek do not yet fully exist. This can be both disorienting and generative, as it requires students not only to learn a discipline, but to participate in its ongoing redefinition. For this reason, disciplinary strength at UMA is not understood as mastery of a fixed body of knowledge, but as the capacity to engage a field critically and creatively—grounded in practice, informed by multiple perspectives, and oriented toward transformation.

Sustainability and regeneration

This curricular axis addresses current socioenvironmental problems, as well as the history, movements, theoretical frameworks, and diverse perspectives within sustainability and regeneration. Its aim is to situate disciplinary knowledge and practice within a broader understanding that integrates ecological, economic, historical, and cultural dimensions. Through this process, one of our central learnings has been the need to hold together two interrelated perspectives: the (uneven) expansion of key dimensions of human wellbeing over the past two centuries, and the simultaneous deterioration of planetary health.

While wellbeing is subjective, certain conditions, such as health, food, water, housing, and literacy are broadly recognized as foundational. In our experience, many students are surprised to discover that these have improved significantly on average, albeit with profound inequalities in their distribution.

Working with these two trajectories together can be disorienting. It challenges simplified narratives of progress or collapse, and calls for a more demanding form of attention: to recognize that we are inheritors of both care and damage. Many of the conditions that sustain our lives today are the result of accumulated efforts, knowledge, and struggles across generations. This recognition can open a sense of gratitude toward our ancestors—those who came before us—not as idealization, but as acknowledgment of what has been made possible.

At the same time, it becomes evident that the dominant ways in which these gains have been produced are generating ecological degradation and social inequities that threaten their continuity. In practice, this tension is not resolved conceptually, but lived as a condition of the present: we are able to receive, but not assured of passing on.

From here, the question shifts. It is no longer only how to preserve what has been achieved, but how to transform the systems through which it is generated—so that well-being can endure, deepen, and be more justly shared. This perspective also expands the sense of who “we” are. It situates human action within a longer continuity of interdependence—an extended human tribe across time. We are not only individuals or cohorts, but part of an ongoing process of receiving and giving. In practice, this framing has proven to be a source of orientation and energy for action. It invites a shift from reacting to crisis toward engaging with the possibility of acting as regenerators—of participating, however modestly, in the renewal of the systems that sustain life.

Agency Lab

A program-long personal development process supports the integration of professional work with personal growth and wellbeing, grounded in action research and in communities of practice, as in the Project Lab.

Our learning process seeks to dissolve the separation between professional and personal development. This is valuable for integral formation in general, but particularly important for change agents, since many of the qualities required for effective action for change—such as patience, resilience, resourcefulness, and empathy—belong to the personal domain.

In our experience, there is an additional reason why this integration matters. Change agents often engage with situations they find deeply problematic or unacceptable. Working toward transformation requires sustained attention to these conditions, which can involve frustration, fatigue, or loss of motivation over time. For this reason, developing the capacity to remain connected to sources of meaning, beauty, and personal enthusiasm becomes essential. This involves learning to move energy between what one seeks to change and what sustains one—holding together a sense of nonconformity with a capacity for appreciation and renewal.

The Agency Lab invites students to engage in action research on their own sources of enthusiasm and how to activate them, drawing insight from others within their community of practice. They also explore key capacities such as self-observation, presence, and self-directed learning, strengthening their ability to sustain engagement in processes of change over time. Although this approach is coherent within our model, many students initially experience discomfort in engaging in this type of work within a university setting, particularly in the shared and sometimes intimate space of their communities of practice. However, bringing the personal into the learning process is a core dimension of the work. It reflects a broader form of interconnectedness: between inner and outer transformation, between individual and collective processes, and between the different scales at which change unfolds.

Relationality

The curricula extend beyond classes to include facilitated dialogues, shared experiences such as documentaries and bonfires, extended breaks, and informal gatherings.

These create conditions for relationships to emerge and for ideas to be exchanged, challenged, and transformed. In our experience, these spaces are not peripheral to learning, but constitutive of it. Understanding does not develop only through engagement with content, but through sustained interaction with others who are also making sense of their work, testing ideas, and encountering uncertainty. In this sense, relationships become a medium through which thinking evolves. However, relational learning is neither automatic nor inherently harmonious. The same diversity that enriches perspectives also generates friction, misunderstanding, and, at times, conflict. We have learned that creating space for exchange is not sufficient; it requires ongoing attention to how conversations are held, how differences are navigated, and how trust is built over time.

Another insight has been the importance of informality. Some of the most meaningful shifts in understanding do not occur within structured sessions, but in the interstitial spaces around them—during breaks, shared meals, or unplanned conversations. These moments allow for a different quality of attention, where ideas can be explored with less pressure and relationships can deepen more organically. At the same time, not all students engage in these spaces in the same way. For some, entering into open dialogue and relational exposure can be uncomfortable. This reveals a limitation: relationality cannot be fully designed or required; it can only be invited and cultivated through conditions that make participation possible without forcing it. For this reason, we increasingly understand relationality not simply as a set of activities, but as an ongoing practice of creating and sustaining the conditions in which collective inquiry can take place. It is through these relationships that ideas are not only shared, but transformed— and where learning becomes, in a deeper sense, a shared process.

Sustainable campus in nature

The campus itself is part of the learning environment. Intentionally simple, beautiful and designed for positive impact, it offers a setting that fosters both focus and inspiration, as well as learning through and with nature. The campus is conceived as a pedagogical medium. Built with earth walls, locally certified wood, and green roofs, it mirrors local architecture and integrates rainwater harvesting, water treatment, and solar energy generation. The land was originally barren; today, it includes a growing edible forest cultivated by the UMA community, which produces food, regenerates soil, and has brought back birds and ducks.

The local community forest is our neighbor, and students systematically visit it as part of the curricula. The forest is healthy and extensive. There is a specific field learning process situated there, through which students learn about, from and with nature, as well as explore new, deeper ways to relate to it and to self-discover as part of it. The facilitated, experiential learning process brings them from perceiving it differently, to taking care of it, to understanding the forest ecosystemic dynamics. Although other ecosystems are also part of field learning, the neighboring forest is the most visited. This is an intentional design so that students may experience knowing and relating to a specific natural place, as if it was a friend one knows well. At the same time, we have observed that the capacity to relate to nature is unevenly developed. While some students quickly form a deep and meaningful connection, others experience the forest more distantly, or even with discomfort. This points to the need to actively cultivate this relationship as part of the learning process.

Another important challenge lies in transferability. While the campus offers a privileged environment for this kind of learning, most students will not carry out their pro-

fessional work in such contexts. This raises an ongoing question: how can the depth of relationship developed with a specific place translate into ways of perceiving, relating, and acting in very different environments?

Scholarship program

A high proportion of scholarships, including specific support for students from Indigenous communities, contributes to the cultivation of a diverse cultural and socioeconomic community, enabling a richer and more complex social fabric. Students come mainly from Mexico, but also from Latin America. We foster geographical, economic, and cultural diversity as an explicit design to facilitate relationships across customary borders. A planetary mission requires that we learn not only to encounter diversity, but to weave it in ways that are respectful, generative, and capable of holding difference without erasing it. In practice, diversity does not automatically translate into integration or interaction. While the presence of different backgrounds creates the potential for richer learning, it does not by itself generate exchange or understanding. Students often arrive with distinct frames of reference, expectations, and ways of relating to knowledge, which may initially coexist without deeply interacting.

We have also observed that broader social inequalities do not disappear within the university space. Differences in educational background, economic conditions, and cultural capital can shape participation, confidence, and the ability to engage in dialogue. This requires ongoing attention to the conditions that enable more equitable forms of participation and mutual recognition. At the same time, weaving diversity involves both enrichment and friction. Intercultural encounters, particularly when engaging with Indigenous perspectives, can challenge dominant assumptions and open new ways of understanding, while also generating moments of tension or misunderstanding. In this sense, weaving is not a smooth process, but a practice that requires time, attention, and the capacity to remain in relationship across difference. In our experience, this process is not only valuable in itself but also develops key capacities for project work. Learning to listen across difference, to integrate multiple perspectives, and to navigate tension constructively are essential for codesign processes in complex socioenvironmental contexts. In this way, the diversity of the community becomes a formative space for the development of weaving capacities that extend beyond the university.

Evaluation: findings and insights

Last year, we conducted the first systematic evaluation of our mission through a survey of alumni from our master's programs over the past 15 years. The aim was to understand the extent to which the capacities cultivated at UMA are being applied in practice, and what kinds of projects and trajectories are emerging from them.

The design of the survey itself proved valuable. By prioritizing a small set of core capacities and defining clear levels of application, it required a process of synthesis that sharpened our understanding of what we are actually trying to cultivate. It became, in a sense, a moment of collective clarification—of looking more closely at what we have been building. While the sample was representative, it may still reflect a bias toward more engaged alumni. Even so, the results offer a meaningful window into what is emerging from this learning experiment in practice.

1) Engagement and orientation: Nearly all graduates (93%) have a central socioenvironmental professional activity, though fewer than half (45%) derive most of their income from it.

2) Projects and pathways: The majority are building on their UMA experience — 60% are working on projects related to those developed there, and 63% have taken on founder or director roles.

3) Collaboration and community: Community ties remain strong: 75% collaborate professionally with a fellow UMA graduate, and 90% are co-designing projects — though only 43% are doing so in direct partnership with beneficiaries.

4) Impact and measurement: This is an area with room to grow. Just 30% have defined impact indicators, and perceived impact averages 3.5 out of 5.

5) Application of capacities: Adoption is broad but depth varies — 90% report applying the core capacities developed at UMA, yet only 35% do so at the highest of four defined levels. Notably, 95% are engaged in explicit personal development processes.

These findings point to clear strengths: alumni are actively engaged in the socioenvironmental sector; there is continuity between learning and practice; strong entrepreneurial agency; dense collaboration networks; and a high degree of application of core capacities.

At the same time, they reveal important tensions inherent to this kind of educational model. First, there is a tension between purpose-driven work and financial sustainability: while most alumni are engaged in socioenvironmental work, fewer are able to fully sustain themselves through it. This often means holding a second job, which limits the time and focus they can dedicate to their socioenvironmental activities.

One possible interpretation is that UMA is preparing people to operate in fields that are still emerging and structurally incomplete, where roles, funding models, and professional pathways are not yet fully established. In this context, it is perhaps not coincidental that a significant proportion of alumni are founders or take on leadership roles, actively contributing to shaping the very systems in which they work.

Second, while a strong culture of collaboration and codesign is evident, deeper engagement from the outset with those most directly affected remains a practical and context-dependent challenge.

Third, although many projects are underway, the development of indicators and evidence of systemic impact remains limited.

Responses also reveal a range of interpretations of sustainability and regeneration. While this diversity reflects openness, it also points to a lack of shared criteria for understanding and evaluating what “regenerative” practice entails. Finding the right balance between openness and clarity remains an ongoing challenge for the community. Taken together, these tensions suggest that cultivating purpose-driven, project-based learning is not sufficient on its own. It must be accompanied by deliberate attention to economic viability, the practical conditions required for sustained codesign with directly affected actors, and the development of shared criteria and evidence of impact. These are not peripheral concerns, but core dimensions of the work.

In this context, UMA appears to be supporting the development of socioenvironmental projects with a systems change perspective. At the same time, we are still in the process of understanding in what ways and to what extent their effects contribute to sustainability or regeneration. In the next phase of evaluation, we will analyze projects with defined indicators.

UMA as a regenerative ecosystem

We refer to regeneration as both a system's capacity to sustain itself over time (Lyle, 1994) and to renew itself by evolving into new forms through the activation of inherent potentials (Reed, 2007). In this sense, the capacity to transform depends on the capacity to sustain.

The notion of ecosystem is used here metaphorically to describe communities characterized by diversity and interdependence, which behave as dynamic “living” systems that evolve through collaboration, exchange, and continuous adaptation (Kinchin, 2022). A regenerative ecosystem, therefore, may be understood as a diverse and interdependent community that not only sustains its basic conditions of existence, but also continuously expands its capacities over time. It does so across multiple levels—individual, relational, epistemic, and systemic—enabling ongoing renewal of people, relationships, knowledge, and practices in service of broader forms of collective wellbeing. To reflect on how UMA expresses this form of regenerative ecosystem, we turn to the five regenerative dimensions that we developed in the UMA Project Lab to invite students to explore through their own project processes:

Personal regeneration: *Renewing oneself through the development of new capacities and sources of enthusiasm.* At UMA, this unfolds through self-directed learning, action research on sources of enthusiasm, and participation in communities of practice that provide continuity, support, and shared reflection. These processes generate cycles in which reflection, challenge, and renewed motivation are continuously reorganized into expanded capacities and forms of appreciation. In this sense, personal regeneration is not a stable condition, but an ongoing practice of re-engagement. At the same time, the intensity and immersion of the learning model can lead to overload and emotional fatigue, making individual and collective strategies for balance, care, and pacing a necessary condition for sustaining this form of renewal.

Relational regeneration: *Renewing relationships through collaborative work and learning.* This is expressed through codesign processes, communities of practice, and structured and informal spaces for dialogue where ideas are shared, tested, and transformed collectively. Codesign implies developing projects *with* others rather than *for* others, while collective learning involves constructing understanding together rather than individually. Both require weaving diverse perspectives and allowing relationships to evolve through shared work. Through these interactions, relationships evolve from coordination toward deeper forms of trust, mutual recognition, and shared inquiry. Diversity plays a central role by expanding perspectives and enabling learning across difference. However, diversity is also shaped by tension. Differences in worldview, communication styles, and positionality can generate friction, misunderstanding, or fragmentation—even within a community oriented toward care and collaboration. Developing a regenerative approach to conflict—one that can hold and transform disagreement when it arises—requires sustained attention.

Socioenvironmental regeneration: *Renewing ecosystems and their embedded communities, through how we work in and with place.* At UMA, this takes shape across three nested scales: the campus ecosystem, the surrounding territory, and student-led projects in diverse contexts. At the campus level, regeneration is expressed in the renewal of soil and biodiversity through the edible forest, which is cared for collectively by the UMA community and provides both ecological restoration and daily sustenance. At the territorial level, it is present in the relationship with the neighboring community forest, where students and educators engage in sustained observation, care, and learning with nature. At a broader scale, students develop socioenvironmental initiatives in diverse

contexts beyond UMA. While UMA participates in and influences local socioenvironmental efforts, its direct territorial impact remains limited, as most student projects are implemented elsewhere. This is gradually shifting as more alumni and students remain in the region to develop local initiatives, and through the bachelor's program, which includes extended engagement with local organizations.

Knowledge regeneration: *Renewing how systems are understood and transformed.*

At UMA, this is most evident in the action research processes embedded in the Project and Agency Labs, where students experiment with change across personal, professional, and socioenvironmental systems. Following Kurt Lewin's (1947) insight that the best way to understand a system is to try to change it, knowledge is regenerated through engagement with transformation itself. These processes often generate shifts in paradigms and ways of thinking. However, while projects produce valuable insights, these are not yet consistently documented, synthesized, or disseminated, which limits their broader contribution beyond group learning contexts.

Process regeneration: *Renewing processes through positive feedback loops.*

At UMA, this operates through two interrelated feedback loops: the learning system itself and the project ecosystem it generates. The learning system is continuously renewed through institutional action research, where feedback, reflection, and experience inform ongoing redesign of curricula and processes. At the same time, a project-based feedback loop emerges: the value generated by student projects increases visibility and attracts new participants, who in turn generate new projects and further learning. While this dynamic produces regenerative momentum, it is not yet sufficient to fully sustain the institution. External resources remain necessary due to the extent of the scholarship program, the quality and intensity of the learning model, and the emerging nature of the fields in which UMA operates. This reveals both a regenerative dynamic and a structural limitation: while learning processes are highly iterative, institutional sustainability is still not fully self-reliant.

Taken together, these five dimensions offer a way of understanding UMA as a regenerative ecosystem in practice. They show how regeneration is not located in a single domain but unfolds across nested levels: internal experience (personal and relational), engagement with context (socioenvironmental), knowledge production (epistemic), and systemic learning dynamics (process). Rather than operating as separate domains, these dimensions interact as a system through which transformation continuously emerges.

Artificial Intelligence and regenerative learning ecosystems

At the beginning of this article, we stated that universities with planetary missions, like UMA, are less dependent on the information-transfer model that AI is now challenging. So what does UMA's model illuminate about what AI cannot replace?

The rapid development of artificial intelligence is transforming how knowledge is produced, accessed, and applied. In educational contexts, large language models can already perform many functions traditionally associated with teaching, such as explaining concepts, structuring information, and supporting individualized learning. In this sense, AI expands access to high-quality instructional interaction and increasingly functions as a continuously available learning support system. However, this transformation also clarifies a key distinction: while AI can support information processing and accelerate understanding, it does not generate vocational purpose, situated meaning, or shared responsibility. It does not ask why knowledge matters in a specific context, nor does it participate in the lived consequences of learning over time.

UMA's model makes this distinction visible in practice. It is not organized around information delivery, but around collective inquiry situated in real projects, relationships, and socioenvironmental contexts. Knowledge is not only accessed or applied but continuously constructed through dialogue and shared reflection. In this setting, learning becomes a process of co-creating meaning and co-responsibility. Students are not only engaging with concepts such as sustainability or regeneration; they are continuously negotiating what these terms signify in different contexts. This negotiation is inseparable from attention, trust, affect, and the capacity to remain engaged in uncertain and evolving processes. Within this experience, several dimensions of learning emerge that cannot be fully replicated through individual interaction with an AI system.

First, there is the dimension of **collective relational learning**. Meaning emerges through sustained engagement with others who are also forming questions, making decisions, and navigating uncertainty. Learning is not only cognitive, but intersubjective and shared.

Second, there is the dimension of **situated and embodied engagement**. UMA's projects are embedded in real socioenvironmental contexts, where knowledge is continuously tested through constraints, feedback, and consequences that unfold over time.

Third, there is the dimension of **purpose formation**. In a context where AI can already provide answers, a central educational shift lies in the capacity to formulate meaningful questions: why something matters, for whom, and with what implications. This process cannot be externalized, as it emerges through lived experience and responsibility.

Fourth, there is the dimension of **ecological and place-based learning**. Engagement with the campus and surrounding forest introduces forms of attention and understanding shaped by direct interaction with living systems, extending learning beyond human-only interaction.

Taken together, these dimensions help explain why the relevance of UMA—and of other universities with a planetary mission—may increase rather than diminish in the context of AI. As informational access becomes increasingly automated, the distinctive value of education shifts toward what cannot be substituted.

This does not position AI as external to UMA's work. On the contrary, it opens the possibility of more intentional forms of collaboration between human, natural, and artificial intelligences. Such collaboration, however, requires clarity about what each form of intelligence contributes, and what remains irreducible in collective learning processes. UMA may have some experience in collaboration between human and natural intelligences through the exploration of nature-based practices, but has little to say in terms of how IA will or should enter that intelligence network. It is in any case untrue to say that AI is not challenging UMA and other planetary-mission universities in any way. The challenge ahead is not only to integrate artificial intelligence into these processes, but to ensure that it expands rather than displaces the capacities that make regenerative learning possible in the first place.

In general, universities are far from understanding how to guide students and teachers to engage artificial intelligence in ways that enhance human creativity, judgment, and collective inquiry, rather than reducing learning to increasingly efficient forms of information processing and analysis. This is not only a matter of how AI is used within educational processes, but of how it subtly reshapes what is considered valuable in learning itself.

From this perspective, AI introduces a triple shift for education. First, it reduces the centrality of information-based learning. Second, it highlights the importance of those dimensions that cannot be substituted. Third, it calls for inquiry into how learning sys-

tems can expand through collaboration with artificial intelligence while remaining grounded in lived contexts, relational processes, and real-world transformation.

An experiment in collaborative inquiry with Artificial Intelligence

As a small prototype to explore pathways for meaningful human-AI interaction in processes of learning and inquiry, this article was developed through an experiment in collaboration with artificial intelligence. Rather than treating AI primarily as a tool for efficiency, it was engaged as a dialogical interlocutor in the formation, refinement, and testing of arguments, embedded within a broader reflection on knowledge, learning, and intelligence. The process unfolded through iterative cycles of writing, reflection, and revision. What gradually became clear is that there are at least three modes of interaction. These modes were not fully anticipated at the beginning; they emerged with how different types of prompts shaped the relationship between thinking and writing.

A first **instrumental mode** represented the most typical interaction. I noticed that when my prompts were oriented toward **efficiency**—such as asking AI to summarize, translate, restructure, or clarify existing text—the interaction remained essentially external to the thinking process. Typical prompts in this mode were of the form: *“Can you summarize this section into its key arguments without changing the meaning?”* *“Can you improve the clarity and structure of this paragraph while preserving its original logic?”* What I discovered is that, while useful, this mode does not really transform the underlying structure of thought; it mainly accelerates it.

A second **dialogical mode** emerged when I began to treat AI less as a tool and more as a **conversational partner** in shaping ideas. Here, my prompts shifted from output-oriented requests to conceptual interrogation. Instead of asking for better versions of what I already knew, I started asking questions such as: *“What assumptions are embedded in this framing?”*, *“How would this argument change if I changed its underlying concept of system?”*, or *“Where is this idea weak or overextended?”* In writing terms, this became a process of iterative co-writing. In research terms, I later realized, it resembled using AI to test conceptual coherence and explore alternative theoretical framings. What I discovered in this mode is that AI begins to participate in shaping the boundaries of thought, even if it still does not define its direction.

A third **inquiry mode** only became visible after working through the previous two. Here, the interaction stopped being primarily about improving text or refining ideas, and started to function as a space for **exploring how meaning itself is produced**. I noticed that when I formulated prompts that deliberately altered the conceptual architecture of an argument, something different happened. For example: *“What changes if I replace ‘education system’ with ‘regenerative ecosystem’ across this entire section?”*, or *“What becomes visible and what disappears if I shift from thinking in terms of institutions to thinking in terms of relational fields?”* or *“What new questions emerge if I assume that knowledge is something that is continuously co-created rather than transferred or applied?”*

In writing terms, these prompts did not simply improve the text—they destabilized it productively, revealing tensions in the underlying structure of thought. In research terms, this became even more evident: I was no longer asking AI to analyze an idea, but to explore how the framing of a problem generates different forms of knowledge. What I discovered in this mode is that writing and research begin to overlap, and that AI can function as a medium through which conceptual worlds are actively reconfigured, not just expressed.

This third mode is also where the most important shift became visible. The question was no longer how to use AI to think better or write faster, but what kinds of thinking become possible when meaning is continuously co-constructed through iterative interaction between human and artificial intelligence. This does not replace human inquiry; it amplifies inquiry in which assumptions, framings, and conceptual boundaries are actively tested within the process itself.

Across these three modes, I came to see a progression: from using AI to process thought, to using it to refine thought, to using it to explore how thought is formed. This progression is not strictly linear, and the modes coexist, but they are qualitatively distinct. The inquiry mode stood out in particular, as it was experienced as genuine co-creation.

What this experiment revealed is not a fixed method, but a shift in orientation: that the most generative potential of AI in learning and research may not lie in producing better answers, but in enabling better questions—and, more importantly, in making visible how questions themselves shape the worlds we are able to think and build.

Regenerative dimensions of collaboration with AI

This experiment in collaborative inquiry also resonates with the five regenerative dimensions that structure UMA's learning model, offering a concrete example of how they may be extended into emerging forms of intelligence and inquiry.

At the level of **personal regeneration**, the shift between the three modes of interaction revealed how deeply inquiry depends on attention and intentionality. Instrumental use of AI tends to reinforce efficiency-oriented habits of thinking, while dialogical and exploratory modes require slowing down, questioning assumptions, and tolerating uncertainty. In practice, I noticed that moving toward deeper forms of engagement was less a technical shift than a personal one: it required cultivating curiosity, restraint from premature closure, and a willingness to let ideas remain open longer than usual.

At the **relational level**, the experiment made visible that even when working individually with AI, inquiry remains fundamentally relational. The quality of the interaction changed significantly depending on how I positioned myself in relation to the system: as a user requesting outputs, as a conversational partner testing ideas, or as a co-participant in a shared space of exploration. What became clear is that relationality is not limited to human-to-human interaction, but also includes how we enter into dialogue with non-human forms of intelligence—and how that dialogue shapes our thinking in return.

At the level of **knowledge regeneration**, the most significant insight was that AI does not only process existing knowledge, but can actively participate in revealing how knowledge is structured. In the third mode of interaction in particular, shifting conceptual framings (rather than refining content) exposed how different epistemic worlds emerge depending on how a problem is defined. This aligns with UMA's emphasis on knowledge as something that is produced through engagement with transformation, rather than simply applied to it.

At the level of **process regeneration**, the experiment itself functioned as a micro-cycle of iterative learning: articulation, feedback, reframing, and re-articulation. What became visible is how AI can accelerate the timing of iteration, but also how important it is to maintain reflective pauses that allow meaning to consolidate rather than constantly optimize. In this sense, the regenerative potential of the process does not lie in speed, but in the quality of the feedback loops that are created.

At the level of **socioenvironmental regeneration**, the connection is more indirect but still relevant. The key shift lies in how conceptual frames influence action. The third mode of interaction showed that changing the underlying framing of a system (for example, from institutional to relational, or from static to ecological) fundamentally alters what kinds of interventions become imaginable. In this sense, collaborative intelligence is not only about thinking differently, but about expanding the range of possible relationships between thought and action in real-world contexts.

Taken together, this experience suggests that artificial intelligence can be meaningfully integrated into a regenerative learning ecosystem, not as a substitute for human or natural intelligence, but as an additional layer that intensifies certain aspects of inquiry. Its most relevant contribution is not the automation of knowledge work, but the amplification of reflective and conceptual shifts that already lie at the core of regenerative learning processes.

Conclusion: Artificial Intelligence and Planetary Mission

UMA's broader experiment is about testing what becomes possible when a university is organized as a relational system of inquiry for planetary wellbeing rather than a structure for information transmission.

This matters now because the conditions that sustained the 20th-century university model—scarcity of information, centralized expertise, and linear knowledge transfer—are being fundamentally altered by artificial intelligence. As these conditions dissolve, the question is no longer only how universities should update their tools or curricula, but what kind of institutions they are called to become, and in service of what forms of life.

In UMA's case, that question has always been inseparable from a socioenvironmental and planetary orientation. From the outset, the university has been guided by the challenge of how to contribute to more regenerative relationships between people and planet, across present and future generations. This planetary dimension is not an add-on to its educational model, but the very field in which its learning processes are situated and tested.

The collaboration experiment with artificial intelligence developed in this article can be understood as a situated instance of this broader orientation. Rather than positioning intelligence as located in a single agent, it explores how meaning can emerge through distributed, dialogical processes in which human intention, contextual knowledge, and computational language generation interact. In this sense, the writing process itself becomes a small-scale enactment of the broader transformation this article describes: a shift from knowledge as transmission toward knowledge as co-produced, situated inquiry across multiple forms of intelligence.

UMA's experience suggests that one direction for higher education lies in strengthening the relationship between knowledge transmission and conditions for collective inquiry oriented toward planetary wellbeing: where internal experience, relational life, engagement with place, and epistemic practice co-evolve with the socioenvironmental systems they seek to transform.

Looking ahead, an emerging possibility lies in the intentional use of artificial intelligence within communities of regenerative practice—not as individual assistants, but as shared inquiry companions embedded in collective processes of design, reflection, and experimentation. In such settings, groups of practitioners could co-construct knowledge with AI systems across cycles of action and reflection: testing assumptions, reframing

problems, exploring alternative theories of change, and making visible how different framings shape possible interventions.

In this context, AI becomes less a tool for producing answers, and more a relational layer within a broader ecology of inquiry—one that can amplify iteration, surface hidden assumptions, and support the exploration of complexity. Its value lies not in replacing human or natural intelligence, but in expanding the space within which regenerative practices are designed, tested, and evolved over time.

In this sense, the notion of a planetary mission names not only how a university is organized, but an evolving response to a deeper question: how can regenerative learning ecosystems collaborate with different intelligences to contribute meaningfully to the flourishing of life within planetary boundaries, while remaining open to continuous transformation themselves?

References

- Gandhi, M. K. (1997). *Hind Swaraj and other writings* (A. J. Parel, Ed.). Cambridge University Press. (Original work published 1909)
- Lewin, K. (1947). Frontiers in group dynamics: Concept, method and reality in social science; Social equilibria and social change. *Human Relations*, 1(1), 5–41. <https://doi.org/10.1177/001872674700100103>
- Luksha, P., & Taddei, F. (2025). *Manifesto for the planetary mission of the university*. Learning Planet Institute.
- Lyle, J. T. (1994). *Regenerative design for sustainable development*. Wiley.
- Reed, B. (2007). Shifting from “sustainability” to regeneration. *Building Research & Information*, 35(6), 674–680. <https://doi.org/10.1080/09613210701475753>

Editorial conclusion: Toward an agentic future

Sholpan Tazabek and Pavel Luksha

*Admit that the waters around you have grown,
and accept it that soon you'll be drenched to the bone.
If your time to you is worth savin', then you better
start swimmin', or you'll sink like a stone.*
— **Bob Dylan**

Artificial intelligence is no longer a horizon event. It is here, embedded in research workflows, student writing processes, administrative systems, and the labor markets into which universities send their graduates. The question posed to universities is therefore not whether to engage with AI or not, since that question has already been answered by the world beyond their walls. The question is whether universities will act as authors of the transition or merely its subjects. Integrating tools is only one part of the task. The larger challenge is to shape the social, intellectual, and human purposes toward which these technologies are directed.

The scale of the shift demands clarity about what is actually changing. In little more than a decade, AI has moved from a peripheral research topic to a ubiquitous layer of the global knowledge economy. Language models capable of producing fluent text, working code, competent legal analysis, and functional research summaries are now accessible to any student with a smartphone. By 2025, the World Economic Forum estimated that more than 40 percent of workplace tasks across major economies were susceptible to AI augmentation — a transformation unfolding substantially faster than most institutional planning cycles anticipated (World Economic Forum, 2025). At the same time, global investment in AI infrastructure, models, and applications reached approximately \$375 billion in a single year, approaching the scale of the entire research enterprise of the world's universities (International Data Corporation, 2024). Some leading AI companies now command market valuations measured in trillions of dollars, exceeding the annual GDP of most countries. Employer surveys across North America, Europe, and East Asia suggest that degree requirements are quietly disappearing from many job postings as organizations discover that task capability can often be assessed more directly, and that fluency in working with AI increasingly matters as much as, or more than, disciplinary credentials.

These are not signals at the edge. They describe a structural shift in the cognitive economy that universities were built to serve. Yet much of the institutional response remains focused on adaptation at the margins: pilot projects, faculty workshops, AI guidelines, and revised assessment policies. Necessary as these measures may be, they do not match the scale of the transformation now unfolding.

The chapters assembled in this volume — from institutions operating in contexts as diverse as Central Asia, the Caucasus, South Asia, Africa, Latin America, China, North America, and Western Europe — converge on a harder claim: that genuine institutional transformation requires not just the adoption of AI tools but the construction of an ecosystem capable of sustaining full-spectrum change. Curriculum, pedagogy, governance, assessment, faculty development, institutional identity, and relationships with regional communities must move together. Notably, many of the most ambitious attempts documented here emerge not from the world's most resource-rich universities, but from institutions operating under conditions of constraint, where the pressure to redefine

educational purpose is often felt most acutely. Partial reforms that leave the industrial university's operating logic intact are likely to produce surface modernization while the deeper mismatch between institutional form and civilizational need continues to grow.

And yet, in the race to integrate AI, something far older deserves to be named with equal urgency. History suggests a pattern. Each major technological rupture, from the printing press and steam engine to electrification and the digital revolution, has eventually produced not only economic reorganization but a return to fundamental questions about what it means to be a human. These are not retreats from progress; they are progress recognizing its own incompleteness. The printing press made the Reformation possible, but it also catalyzed a Renaissance of Humanism — a renewed insistence on the full formation of the person as the proper aim of education. Industrial modernity produced not only scientific positivism and mass production, but also Dewey, Freire, and a century of pedagogical traditions insisting that democracy depends on educated persons, not merely trained workers.

We are in such a moment again. And this time the technological rupture is more intimate than previously, because it reaches into the cognitive interior, i.e. the very processes of reasoning, writing, analysis, and creativity that modern education has treated as the defining marks of human formation. When these processes can be replicated at scale by machine systems, the question that returns is: what are we actually cultivating?

The contributors to this volume have named what they believe endures. Sidorkin names *relational labor*: the mentoring, judgment, and sustained human contact that massification made economically impossible and that AI now makes possible to recover. Chen et al. name *AI collaboration fluency* — the capacity to work productively with machine intelligence while retaining human judgment, purpose, and metacognitive agency. Taddei and Costambeys-Kempczynski name *planetary agency* — the formation of people capable of acting for the world, not merely succeeding within it. Haro names *ecological embeddedness* — learning that is situated within living systems rather than abstracted from them. Kozhakhmetov et al. name *Tolyq Adam* — a holistic person formed through the integrated cultivation of mind, heart, and will — a vision grounded not in imported templates but in the philosophical inheritance of Central Asian civilization.

These contributions represent different vocabularies, different intellectual lineages and traditions, and different institutional starting points. Nevertheless, they all point to same underlying movement: the transition from education as transmission toward education as formation. As the procedural, analytical, and communicative dimensions of cognitive labor are increasingly delegated to machines, the challenge of the AI era is not simply preserving the legacy functions of the industrial university but determining what should be cultivated in their place.

Emphasizing these human dimensions including but not limited to moral orientation, relational depth, and ecological awareness, is not a retreat from intellectual seriousness, but is its deepest expression. If AI can produce competent text, the question is no longer whether students can write. The question that remains provocative and open is whether they have something worth saying, and whether they possess the ethical and relational capacities to say it well, straightforwardly, and in the service of something larger than themselves. If AI can synthesize existing knowledge, the question is no longer whether students can recall information, but whether they can navigate genuine uncertainty, form authentic relationships, take responsibility for their judgments, and act as citizens of a planet under extraordinary stress.

The university that takes this seriously is not competing with AI, but doing something AI cannot do: *growing human beings*.

Unfinished questions: inequality and alternative priorities

This volume has examined the university's encounter with AI across multiple dimensions, but the editors acknowledge, with some humility, that certain questions remain inadequately addressed. These questions remain open not because they are less important, but because they represent the frontier of what the field has yet to fully articulate.

Chief among these is the question of inequality and inclusion. The AI transition is not a neutral event. It arrives into a world already marked by deep asymmetries of access, capability, and infrastructure. AI systems trained predominantly on English-language data perform significantly worse in the languages spoken by the majority of the world's students. Institutions in high-income countries are building AI governance frameworks, technical capacity, and curriculum infrastructure that are not being developed at the same pace in the Global South. The labor market consequences of AI-driven automation are likely to fall most heavily on workers in middle-skill, routine cognitive occupations — precisely the professional pathways that higher education in emerging economies has been building toward.

In this sense, AI does not create new inequalities so much as it acts as a litmus test that makes visible the existing fractures: between rich and poor, between the formally educated and those excluded from education systems, between the employed and those structurally displaced, and between the Global North and the Global South. It reveals with unsparing clarity which institutions, which communities, and which populations have enough resources to adapt, and which do not.

Yet recognizing AI as a diagnostic is not sufficient. What is needed is a deliberate reorientation of the discourse around AI in education. The dominant lexicon of AI development, i.e. innovation, disruption, productivity, competitive advantage, technological leadership, is oriented toward actors who already possess significant resources and whose primary challenge is maximizing returns on existing capabilities. It is not a language designed to illuminate the needs of communities where the challenges are access, equity, dignity, and basic capability formation.

What is needed is not only a new lexicon, but a new set of priorities. A discourse that places inclusion, care, local knowledge systems, and democratic participation at the center of conversations about AI in education is not merely a normative aspiration. It is a precondition for ensuring that the AI transition in higher education does not reproduce and deepen the inequalities that the industrial university model already carried.

The contribution of institutions in the Global South to this volume is therefore more than a matter of representation. It is a reminder that some of the most important innovations in human-centered education may emerge not from the most resource-rich environments, but from institutions forced to confront questions of purpose, inclusion, and social relevance under conditions of constraint. In such contexts, the stakes are often highest and the margin for failure smallest. What emerges there deserves to be understood not as adaptation at the periphery, but as a source of insight for the field as a whole.

'No template will travel'

One of the most important findings of this collection is the one that its editors did not fully anticipate at the outset. Despite the universality of AI as a technological phenomenon, no two institutional cases in this volume describe the same transformation. Caucasus University's response is shaped by brain drain and a thin labor market. The

International University of Central Asia's response is shaped by a pedagogy of resource constraint and meta-learning. Almaty Management University's response is shaped by a post-pandemic student wellbeing crisis and the philosophical inheritance of Abai. Daffodil International University's response is shaped by leapfrogging ambition in the context of rapid digital expansion. Universidad del Medio Ambiente's response is shaped by a foundational commitment to regenerative ecology that predates AI disruption by more than a decade.

These differences are not incidental — *they are the point*. The appropriate institutional response to AI is not a universal template to be downloaded and implemented, but an ecosystem to be constructed organically from each institution's particular conditions, resources, traditions, and a sense of purpose. Just as universities emerged historically in multiple forms serving different social, economic, and cultural purposes, the AI transition may accelerate the emergence of distinct institutional archetypes rather than a universal endpoint. The cases in this volume demonstrate that this construction is possible, and that it requires not just technical or curricular innovation but a willingness to ask what the institution is fundamentally for, and to answer that question honestly from within its own context.

This conclusion resonates with a body of scholarship that deserves to be explicitly named here. Iveta Silova and Gita Steiner-Khamsi's (2008) foundational work on traveling policies and postsocialist transformation package documented, in the context of neoliberal reform packages of the 1990s and 2000s, how policy models developed in one context were transferred wholesale to others, and how the results consistently failed to account for the ways that local institutional histories, cultural practices, and social structures shaped what any given reform could actually accomplish. Their analysis demonstrated that the apparent adoption of a traveling policy often concealed a process of creative recontextualization in which local actors used imported frameworks as resources for addressing challenges that the frameworks had not been designed to solve.

The AI transition presents a structurally similar risk. Models of AI integration developed at MIT, ETH Zurich, or the National University of Singapore are developed in response to particular resource environments, research infrastructures, and student populations. Their direct transposition to universities in Bishkek, Tbilisi, Dhaka, or Johannesburg is not merely unnecessary, but epistemically impoverished. It forecloses the possibility that institutions operating in different conditions might generate fundamentally different and more appropriate responses. As the cases in this volume demonstrate, the innovations that emerge from constraint and from the necessity of doing more with less, of building coherence without abundant resources, of maintaining institutional identity under pressure, may carry lessons that better-resourced institutions have not had the necessity to learn.

The imperative, then, is not to find the right global model and implement it. The imperative is for each university and each national system is to become an active agent in constructing its own path, drawing selectively on global experience while remaining anchored in local conditions, purposes, and communities. The diversity of pathways documented in this volume is not a sign of the field's immaturity. It may be its most valuable feature.

Limitations and what comes next

This volume was conceived as a contribution to the International Association of University Presidents' (IAUP) Central Asia Regional Conference in Almaty in June 2026. It

reflects the thinking of university presidents, faculty, and institutional leaders who were willing to engage seriously with questions that do not yet have settled answers. The editors are grateful for that willingness.

However, we owe the reader and ourselves an honest accounting of what the volume does not accomplish. Its most significant limitation is evidentiary. The cases presented here are institutional narratives: accounts of what universities are attempting, why they are attempting it, and what early results suggest, but these are not research studies. They do not provide systematic empirical evidence about student outcomes, the long-term effectiveness of the interventions described, or the causal mechanisms through which specific changes produce specific results. The volume documents what is being tried but does not yet demonstrate what works.

There is a second, related limitation. By design, this collection privileges institutional perspectives. The contributors are presidents, vice-chancellors, provosts, professors, researchers, and policy leaders. Their chapters illuminate how universities are responding to AI, but they do not fully answer a larger question: how these diverse responses relate to one another, what broader patterns they reveal, and what trajectories for higher education may be emerging beneath the surface of individual cases.

There is also a third, deeper absence that the editors feel with particular force: the absence of the voices of those who are most directly affected by higher education transformations this volume describes. The people who live within this transformation are students, junior faculty, administrative staff, and the communities universities serve, and they are largely absent from this conversation. Students are not merely the recipients of institutional decisions about AI. They are the people whose formation is at stake, whose relationship to knowledge and learning is being transformed, and whose voices are essential for understanding whether the changes institutions are making are actually serving the purposes they claim. We do not yet have systematic evidence about how students experience AI-integrated learning environments, what they find meaningful and what they find alienating, what they need from universities that AI cannot provide, and what they believe universities are failing to give them.

A second volume, currently in preparation, seeks to address these challenges. Rather than presenting additional institutional cases, it steps back from the immediate responses documented here to examine the larger transformation they collectively reveal. Drawing on historical analysis, contemporary evidence, foresight methods, and student perspectives, it explores the forces reshaping higher education and the different futures toward which universities may be moving. Its purpose is not to predict a single outcome, but to provide conceptual tools for navigating a period in which multiple institutional pathways remain open.

Acknowledgments

A volume of this scope does not emerge from editorial vision alone. It emerges from the willingness of people to think carefully, write honestly, and trust that the questions being asked are worth pursuing.

The editors wish to express their deep gratitude to all contributing authors: Fernando León García, Vadiv Novikov, Alexander Sidorkin, Natalia Cebotari, Charles Fadel, Timur Schukin, Raphael Costambeys-Kempczynski, François Taddei, Ralph Wolff, Kakha Shengelia, Ivan Ninenko, Assylbek Kozhakhmetov, Assel Aryn, Aigul Sarenova, Md. Sabur Khan, Shawn Chen and colleagues at SIAS University, Letlhokwa George Mpedi, and Victoria Haro. Each of these authors contributed not only a chapter but a genuine

intellectual encounter with one of the most consequential questions in contemporary higher education. Their willingness to think in public before the answers are settled is precisely the kind of institutional courage this moment requires.

The International Association of University Presidents (IAUP) has been an essential institutional home for this project. IAUP brings together university presidents across six continents and has, under its current leadership, made the integrated response to AI in higher education a central institutional priority. The association's commitment to ensuring that AI serves civility, inclusivity, equity, and humanity — and not merely productivity or institutional competitiveness — reflects exactly the orientation this volume has sought to develop. We are grateful for the space it provided for this inquiry, and for the community of university leaders it has convened.

Global Education Futures brought to this project a decade of inquiry into the future of education. Many of the questions explored in this volume — the changing role of universities, the future of knowledge and expertise, human development in an age of intelligent machines, and the emergence of new institutional forms — have been part of an ongoing conversation within the network for years. That work helped provide the conceptual foundations and long-term perspective that shaped the volume as a whole.

Almaty Management University (AlmaU), as a host of International Association of University Presidents' (IAUP) Central Asia Regional Conference in Almaty 2026 and IAUP's institutional partner, provided not only logistical support but something rarer: a willingness to subject its own practice to examination. AlmaU's development of the Tolyq Adam program, its experimentation with human-centered curriculum, and its engagement with the broader questions of what a Central Asian university can contribute to global educational discourse are not merely institutional achievements. They are demonstrations that the renewal of higher education need not wait for permission from prestige hierarchies.

Finally, the editors wish to acknowledge the many people across all the institutions represented here and many not represented, who are doing this work without recognition: faculty members redesigning courses under institutional uncertainty, students navigating a learning environment that the institutions themselves have not fully figured out, administrators managing the practical complexity of transformation while keeping institutions running. The future of higher education will be built, if it is built well, by people like these — people who do not see AI as a threat to resist or a technology to exploit, but as a condition within which the deeper question of what education is for has become impossible to avoid.

That question deserves most serious, most honest, and most pluralistic response we can give it. This volume is one contribution, but the work continues.

References

- Dewey, J. (1916). *Democracy and education: An introduction to the philosophy of education*. Macmillan.
- Freire, P. (1970). *Pedagogy of the oppressed* (M. B. Ramos, Trans.). Herder and Herder.
- International Data Corporation. (2024). *Worldwide artificial intelligence and generative AI spending guide*. IDC. https://www.idc.com/getdoc.jsp?containerId=IDC_P33198
- Silova, I., & Steiner-Khamisi, G. (Eds.). (2008). *How NGOs react: Globalization and education reform in the Caucasus, Central Asia and Mongolia*. Kumarian Press.
- World Economic Forum. (2025). *The future of jobs report 2025*. World Economic Forum. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>

**Universities After AI:
Voices from the Global Frontier**

Edited by Pavel Luksha, Sholpan Tazabek

Computer layout D. Tynbayev
Cover design M. Bakeev

Signed for seal 10.06.2026. Format 60×84 1/8. Offset paper.
Digital printing. Volume 15,5 standard sheets. 200 copies.

Printed in the printing house of the publishing house
«Smart University Press».
050012, Almaty, BC Business City,
st. Jambyl, 114/85, block Zh.

