RESEARCH BRIEF

Artificial Intelligence in Graduate Education:
Implications and a Decision-making Framework for Action

Debra W. Stewart | NORC
Stewart-Debra@norc.org

Karen Grigorian | NORC
Grigorian-Karen@norc.org

Executive Summary

Artificial intelligence (AI) is rapidly impacting higher education. Across U.S. institutions, adoption is already widespread but uneven, governance structures are still emerging, and graduate education leaders are making consequential decisions without clear precedent.

This Research Brief synthesizes findings from a national survey and convening of graduate education leaders to provide a concise, decision-oriented interpretation of the current landscape. It focuses on what matters most for graduate deans: understanding key dynamics, identifying central challenges, and acting effectively under conditions of uncertainty.

Five core insights define the current moment:

1. AI adoption is broad but inconsistent. Practices vary significantly across institutions and within them.
2. Practice is outpacing policy. AI is already embedded in operations even where formal guidance is incomplete or absent.
3. Efficiency is the primary driver for deans…initially. Adoption often begins with administrative gains but is expanding into strategic domains.
4. The central challenges are value-based, not technical. Leaders are navigating tensions between academic freedom, learning, efficiency, and fairness.
5. Institutions are operating in an experimental mode. Most are testing tools and learning iteratively rather than implementing fixed strategies.

These dynamics require action. Decisions cannot wait for certainty; flexible frameworks are more effective than rigid policies; and variability must be managed without stifling innovation.

This Research Brief offers a practical framework to guide decision-making for senior officers for graduate education, emphasizing clarity of use cases, evaluation of trade-offs, human-centered implementation, and continuous adaptation.

Introduction: A Leadership Moment for Graduate Education

Artificial intelligence (AI) is no longer an emerging issue in graduate education; it is an active and accelerating force shaping daily practice in most institutions. Many graduate schools across the United States are already using AI in administrative operations, student communication, recruitment, and admissions. What remains unsettled is not whether AI will be integrated, but how.

Graduate deans are facing a moment defined by uneven adoption, incomplete governance, and competing institutional values. Unlike earlier technological shifts, AI is advancing rapidly and simultaneously affecting multiple domains: administrative work, academic practice, and the preparation of students for a changing workforce. Institutions are therefore being asked to make consequential decisions without stable norms or clear precedent.

This Research Brief brings together findings from a national survey and a convening of graduate education leaders and reframes them for practical use.[1] NORC at the University of Chicago (NORC) conducted a comprehensive national survey of graduate deans at leading STEM PhD-granting institutions across the United States in the Spring/Summer of 2025. The survey participants were senior officers for graduate education on their campuses, a role that encompasses a variety of titles; the most common of these is graduate dean, the term used throughout this Research Brief to refer to this position. The survey explored the current status and implementation of AI-related policies and processes across all domains of graduate education, including administrative operations, curriculum development, research practices, and student conduct standards.

Subsequently, NORC hosted a convening in November 2025 that brought together over 100 graduate education leaders from across the nation to explore how AI is being integrated into graduate STEM education and to examine its implications for innovation and access in the field. The convening featured both in-person and virtual participation, creating a diverse forum for dialogue among institutional leaders grappling with similar challenges and opportunities. A comprehensive White Paper reporting the results of this research is available on the NORC website.[2] Building off the full report which provides details on all observed practices, this Research Brief offers a concise interpretation of what matters most for decision-making and proposes a framework to guide action by Senior Officers for Graduate Education on US campuses.

[1] This research and the resulting report "AI in Graduate Education" issued July 2026 were funded by the Alfred P. Sloan Foundation.
[2] See https://www.norc.org/research/projects/artificial-intelligence-in-graduate-stem-education.html.

The Current Landscape: Broad Use, Uneven Development

AI is already embedded in graduate education, regardless of whether formal policies exist. The 2025 survey showed institutions fall into three broad categories: those with established policies, those actively developing them, and those that have yet to begin as shown in figure below. Yet these categories do not reliably predict practice.

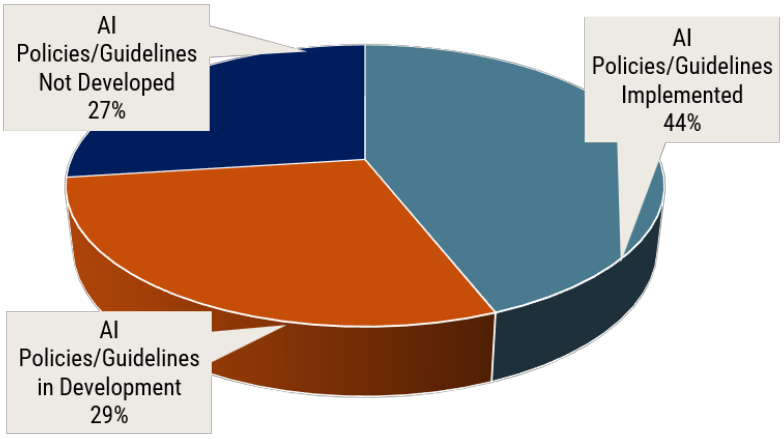

Across all three groups, AI is being used in core administrative and operational functions, especially:

- Office operations and routine administrative tasks
- Communication with graduate students and applicants
- Recruitment and admissions processes

This creates a landscape in which practice is advancing ahead of policy. Graduate schools are not waiting for fully developed governance frameworks; they are acting in real time.

At the same time, there is substantial variation within institutions. Departments, programs, and individual faculty members often determine how AI is used. This decentralized approach enables innovation but produces inconsistencies in expectations, experiences, and outcomes.

For graduate leaders, the landscape is therefore not simply evolving, it is uneven, fragmented, and already operational.

Adoption without Alignment

A defining feature of the current environment is the gap between adoption and alignment. AI is being used broadly, but not always coherently.

This disconnect appears in several ways:

- Institutional policies often lag behind practice
- Faculty approaches vary widely, even within the same department
- Students encounter inconsistent expectations across courses and programs

Decentralization enables experimentation while complicating coordination. Graduate deans are not simply facilitating adoption, they are managing divergence.

Without intentional alignment, local decisions accumulate into institutional consequences, introducing risks of confusion, inequity, and inefficiency.

What is Driving AI Adoption?

For the graduate schools and their leadership, the most immediate driver of AI adoption is efficiency. Institutions are using AI to:

- Reduce administrative workload
- Improve responsiveness in communication
- Enhance data analysis and reporting

These uses respond directly to operational pressures, including limited staffing and growing complexity.

At the same time, broader motivations are emerging:

- Improved strategic decision-making
- Enhanced student engagement and services
- Alignment with workforce expectations

AI is therefore evolving from an operational tool to a strategic one. Institutions are not only increasing efficiency but reconsidering how work is structured and outcomes are defined.

The Real Challenges: Values, Not Technology

The most significant challenges presented by AI are not technical but value-based. Graduate leaders must navigate tensions among:

- Academic freedom and institutional consistency
- Strategic agility and deliberate, values-grounded institutional infrastructure
- Efficiency and values-based human oversight

AI amplifies these longstanding tensions. For example, decentralized academic structures can produce dramatically different expectations for AI use across courses or programs.

Similarly, efficiency gains must be weighed against concerns about skill development and overreliance on automation.

The value of equity presents an additional challenge and opportunity. While AI can expand access, uneven implementation may reinforce disparities.

These are not problems that can be solved through uniform policy. They require sustained judgment and leadership.

The Role of Human Judgement

Across institutions, a consistent principle is clear: AI should support human decision-making, not replace it.

For example in areas such as admissions or evaluation, institutions emphasize maintaining human oversight. AI may assist with analysis or pattern recognition, but it cannot assume responsibility for consequential decisions.

This reflects both practical and ethical considerations. AI outputs may be unreliable or biased, while graduate school practices on rely fundamentally on professional judgment and contextual expertise.

Maintaining human oversight is therefore a defining feature of responsible AI use.

An Experimental Phase, Not a Settled One

Graduate schools are not implementing fully developed AI strategies, they are experimenting and innovating.

Common approaches include:

- Piloting tools
- Testing use cases
- Learning through iteration

This reflects the pace of technological change. However, without coordination, experimentation can remain fragmented and fail to produce institutional learning.

The challenge is to enable experimentation while ensuring it contributes to coherent progress and innovation which is mission-aligned.

Implications for Graduate Deans

Several implications follow:

- Decisions must be made under uncertainty
- Frameworks are more effective than fixed policies
- Variability must be actively managed
- External pressures, including the skills and competencies employers expect from future graduates, must be considered
- Leadership is fundamentally a coordinating function

Graduate deans must align stakeholders and provide direction while preserving flexibility.

A Framework for Institutional Decision-making and Action

Institutions can adopt a simple framework to guide decisions:

1. Define the use case clearly
2. Clarify the Purpose
3. Evaluate the Tradeoffs
4. Clarify the role of AI
5. Select an appropriate governance approach
6. Implement iteratively and adapt over time

This framework supports structured decision-making while allowing implementation flexibility and is illustrated on page 7.

Conclusion: Leading Through Complexity

AI is transforming graduate education through interconnected changes affecting operations, learning, and decision-making.

Senior officers for graduate education must lead under conditions of uncertainty while balancing competing values. The goal is not uniformity, but thoughtful and adaptive integration grounded in institutional mission.

Assessing AI Use in Graduate Education: A Six-Step Decision Framework

STEP 1
Define the Use Case
Is it administrative, academic, support services, or other?

→

STEP 2
Clarify the Purpose
Is it for efficiency, quality, access, or something else?

→

STEP 3
Evaluate Trade-Offs
What do you gain or lose with AI: efficiency, learning, fairness, integrity, etc.?

→

STEP 4
Clarify the Role of AI
What will AI do: provide information, analysis, decision support, or other?

→

STEP 5
Determine Appropriate Governance
Where does authority reside: local, institutional, hybrid?

→

STEP 6
Implement Iteratively & Adapt
Action: pilot, evaluate, share, adjust

ABOUT STEP 3. EVALUATE TRADE-OFFS

Dimension	Key Question
Efficiency	Does this improve processes?
Learning	Does this affect skill development?
Fairness	Who benefits or is disadvantaged?
Integrity	Are ethical risks introduced?
Other dimensions	What is gained or lost?

ABOUT STEP 4. CLARIFY THE ROLE OF AI – GUIDING PRINCIPLE

AI should not determine high-stakes outcomes independently.

ABOUT STEP 5. DETERMINE GOVERNANCE APPROACH

- **Local flexibility** — decisions made at the program or unit level
- **Institutional guidance** — centralized policy applied campus-wide
- **Hybrid model** — shared institutional principles with local discretion in application

ABOUT STEP 6. IMPLEMENT & ADAPT

- **Pilot** — test the use case in a limited setting
- **Evaluate** — assess outcomes against the purpose and trade-offs identified
- **Share lessons** — disseminate findings across programs and units
- **Adjust practices** — monitor outcomes and revise and refine the use case, AI's purpose, or governance approach as needed

Disclaimer

Artificial intelligence tools were used in a limited, supportive role to assist with synthesizing and summarizing primary findings from the *AI in Graduate STEM Education* project for this Research Brief. All substantive analysis, interpretation, conclusions, and data verification were conducted by the authors.

About NORC at the University of Chicago

NORC at the University of Chicago is an objective, nonpartisan research organization that delivers insights and analysis decision-makers trust. NORC has been a pioneer in research and data science for more than 80 years. Our rigorously derived insights have informed the program, policy, and personal decisions of government agencies, businesses, organizations, and members of the public. And our methodological innovations have helped our clients ask new questions, include new populations, and further their missions. To learn more about NORC visit www.NORC.org.

Acknowledgements

The authors are grateful to *AI in Graduate STEM Education* project team for their excellent work, which was foundational to this Research Brief; the higher education institutional leaders who participated in the research project; the project's expert advisory committee members who inspired us to create this Research Brief; and the Alfred P. Sloan Foundation who generously supported the foundational research project.