



AI Adoption in Research Administration at Emerging Research Institutions

Dylan Ruediger
Ruby MacDougall
Stefanie Brachfield
Doug Dechow
Jonathan Parker
Jana Remy

March 30, 2026



ITHAKA S+R

Through research, collaboration, and innovation, Ithaka S+R informs policy, guides strategies, and builds evidence-based tools to broaden access to postsecondary education, improve student and workforce outcomes, and advance scholarship.

Ithaka S+R is part of ITHAKA, a nonprofit with a mission to improve access to knowledge and education for people around the world. We believe education is key to the wellbeing of individuals and society, and we work to make it more effective and affordable.

Copyright 2026 ITHAKA. This work is licensed under a Creative Commons Attribution Non-commercial 4.0 International License. To view a copy of the license, please see <https://creativecommons.org/licenses/by-nc/4.0/>.



This material is based upon work supported by the National Science Foundation under Grant. No. 2437518. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.



Introduction

Research administration, an essential component of a university's research enterprise, is growing more complex, costly, and cumbersome each year.¹ According to the Council on Government Relations (COGR), the federal government has issued over 200 new or revised policies related to the administration of research funded by federal agencies in the past 10 years.² Funding from private philanthropy brings with it additional compliance requirements that may differ substantively from those of the federal government. Though one effect of these regulations is a web of red tape, at their core are principles that help ensure research is ethical, legal, trustworthy, and makes good use of funders' resources.

Research universities of all sizes depend on well-trained staff and well-organized procedures and workflows to manage the pre-award tasks required to effectively compete for research dollars and the post-award tasks associated with expending them. Maintaining the staff and infrastructure to support a robust research enterprise is challenging for all universities, but especially for Emerging Research Institutions (ERIs), where limited staffing resources may restrict opportunities to grow their research portfolios—and along with it the capacity to expand their staffing. As generative AI transitions into an everyday technology, university research offices are exploring its potential to reduce administrative burden and increase operational efficiency. Vendors who provide enterprise level research information management systems are asking similar questions

¹ National Academies of Sciences, Engineering, and Medicine, *Simplifying Research Regulations and Policies: Optimizing American Science* (Washington DC: National Academies Press, 2025), <https://doi.org/10.17226/29231>; for estimated costs of specific types of compliance activities see: "Research Security and the Cost of Compliance: Phase I Report – Results from COGR's Phase I Survey on the Costs of Complying with Research Security Disclosure Requirements for the Fiscal Year 2022-2023," Council on Governmental Relations, November 2022, <https://www.cogr.edu/sites/default/files/Version%20Dec%205%202022%20research%20security%20costs%20survey%20FINAL.pdf>; "Data Management and Sharing (DMS) and the Cost of Compliance: Results from the COGR Survey on the Cost of Complying with the New NIH DMS Policy," Council on Governmental Relations, May 11, 2023, https://www.cogr.edu/sites/default/files/DMS_Cost_of_Compl_May11_2023_FINAL%20%281%29.pdf.

² "Changes in Federal Research Requirements Since 1991," *Council on Governmental Relations Blog*, January 29, 2026, <https://www.cogr.edu/blog/changes-federal-research-requirements-1991>.

and building AI capabilities into their platforms. While AI seems to present real opportunities to automate aspects of research administrators' workflows, to date the actual return on investment on AI tools is unclear, and widespread adoption presents change and risk management problems.

A substantial scholarly literature about generative AI use by instructors, researchers, and students provides insights and analysis to institutions making strategic decisions about AI implementation within the academic enterprise.³ In contrast, the literature on AI implementation in administrative contexts, and specifically in research administration, is thin.⁴ Universities have been slow to develop policy guidelines for administrative staff.⁵ Cross-institutional information sharing is occurring at

³ For a few entry points, see Shan Wang et al, "Artificial Intelligence in Education: A Systematic Literature Review," *Expert Systems with Applications* 252, Part A (October 15, 2024): Article 124167, <https://doi.org/10.1016/j.eswa.2024.124167>; Dylan Ruediger, Melissa Blankstein, and Sage Jasper Love, "Generative AI and Postsecondary Instructional Practices: Findings from a National Survey of Instructors," *Ithaka S+R*, June 20, 2024, accessed March 13, 2026, <https://doi.org/10.18665/sr.320892>; Claire Baytas and Dylan Ruediger, "Making AI Generative for Higher Education: Adoption and Challenges Among Instructors and Researchers," *Ithaka S+R*, May 1, 2025, accessed March 13, 2026 <https://doi.org/10.18665/sr.322677>; Vicent Maberizi et al, "A Systematic Review of the Impact of Generative AI on Postsecondary Research: Opportunities, Challenges, and Ethical Implications," *Discover Artificial Intelligence* 5 (2025): article no. 238, accessed March 13, 2026, <https://doi.org/10.1007/s44163-025-00495-3>.

⁴ There are several excellent case studies: Amber Hedquist et al., "Reflections on AI Implementation in Research Administration: Emergent Approaches and Recommendations for Strategic and Sustainable Impact," *SRA International blog*, November 25, 2025, accessed March 13, 2026, <https://www.srainternational.org/blogs/srai-ira2/2025/11/25/reflections-on-ai-implementation-in-research-admin>; Lisa A. Wilson, Benn Konsynski, and Tubal Yisrael, "Advancing Research Administration with AI: A Case Study from Emory University," *SRA International blog*, May 22, 2025, accessed March 13, 2026, <https://www.srainternational.org/blogs/srai-ira2/2025/05/22/advancing-research-administration-with-ai-emory>; and useful best practice documents: "A Practical Guide to Using AI in Research Administration: Guide Only, Artificial Intelligence for Research Administration (AI4RA)," University of Idaho, July 2025, accessed March 13, 2026, <https://ai4ra.uidaho.edu/wp-content/uploads/2025/07/A-Practical-Guide-to-Using-AI-in-Research-Administration-GUIDE-ONLY.pdf>;

⁵ Yunjo An, Ji Hyun Yu, and Shadarra James, "Investigating the Higher Education Institutions' Guidelines and Policies Regarding the Use of Generative AI in Teaching, Learning, Research, and Administration," *International Journal of Educational Technology in Higher Education* 22 (2025): Article 10, accessed March 13, 2026, <https://link.springer.com/article/10.1186/s41239-025-00507-3>; Jenay Robert, "The

professional meetings and professional development resources created by Society of Research Administrators International (SRAI), National Council of University Research Administrators (NCURA), National Organization of Research Development Professionals (NORDP) (among others), but many research offices are making critical decisions more or less on their own.

With funding from the National Science Foundation's GRANTED program (grant #2437518), Ithaka S+R, Chapman University, and Montclair State University organized two workshops to help research administrators consider how to leverage AI to build research capacity at ERIs. Our first workshop, held at Montclair State in September 2025, brought together 31 participants from 13 academic and medical institutions in the New York/New Jersey/Pennsylvania region. Our second workshop, hosted by Chapman University on December 5, 2025, included 32 participants from 13 colleges and universities in Southern California. The approximately 2,600 ERIs in the United States receive a disproportionately small amount of federal research awards, but ERIs have been identified by Congress and leading scientific organizations as strategically important for the growth of the research enterprise and the STEM workforce.⁶

AI adoption presents distinctive challenges at ERIs, who already face systematic barriers to improving and supporting their research capacities.⁷ Integration of generative AI into their research infrastructure

Impact of AI on Work in Higher Education," *Educause*, January 12, 2026, accessed March 13, 2026, <https://www.educause.edu/research/2026/the-impact-of-ai-on-work-in-higher-education>.

⁶ H.R. 4346, Legislative Branch Appropriations Act, 2022, 117th Cong. (2022), enacted as Pub. L. No. 117-167 (Aug. 9, 2022), accessed March 13, 2026, <https://www.congress.gov/bill/117th-congress/house-bill/4346>; National Academy of Engineering and National Research Council, *Partnerships for Emerging Research Institutions: Report of a Workshop* (Washington, DC: The National Academies Press, 2009), accessed March 13, 2026, <https://www.nationalacademies.org/read/12577/chapter/1#viii>; Anna M. Quider and Gerald C. Blazey, "How to Keep Emerging Research Institutions From Slipping Through the Cracks," *Issues in Science and Technology* 39, no. 3 (Spring 2023): 50–53, accessed March 13, 2026, <https://issues.org/emerging-research-institutions-quider-blazey/>; Sara Partridge, Victor Santos, and David K. Sheppard, "Bolstering the Role of HBCUs in Federal Research and Development," *Center for American Progress*, September 24, 2025, accessed March 13, 2026, <https://www.americanprogress.org/article/bolstering-the-role-of-hbcus-in-federal-research-and-development/>.

⁷ "Building America's STEM Workforce: Eliminating Barriers and Unlocking Advantages," American Physical Society Office of Government Affairs, January 22, 2021, accessed March 13, 2026, <https://www.aps.org/publications/reports/building-americas-stem->

offers opportunities to improve research competitiveness for federal and private dollars increasingly allocated to AI-enabled and AI-focused work, but ERIs will need to meet the challenges of governing and implementing these technologies with limited financial, staffing, and technical resources. Absent careful planning and coordinated action, the rapid evolution of generative AI risks exacerbating, rather than mitigating, existing inequities between ERIs and established research institutions. Yet, AI also presents new opportunities to strengthen their capacity to develop competitive project proposals and meet the milestones and requirements of funded projects.

Working in small groups, participants at both workshops assessed their existing capacities and needs, shared information about experiments with AI, and ideated implementation strategies and collaborative possibilities. Throughout the workshops, we also surfaced a number of important open questions about the future of an AI-enabled research enterprise.

Challenges

At both workshops, participants raised a series of interconnected challenges related to integrating AI into research administration, reflecting technical, organizational, and cultural barriers:

- Data security concerns are making some research administrators wary of using third-party AI tools.
- Many administrators and faculty are skeptical that AI is sufficiently reliable to be trusted to meet rigorous legal, ethical, and fiscal standards required for compliance.
- AI can add stress to already strained relations between staff and senior leaders, especially as these leaders are sometimes perceived as pushing impractical or counterproductive AI practices.
- Participants were cautiously optimistic that AI would eventually lessen their workload, but they were cognizant of the significant up-

[workforce](#); Rocio C. Chavela Guerra and Carolyn Wilson, “From Lack of Time to Stigma: Barriers Facing Faculty at Minority Serving Institutions Pursuing Federally Funded Research,” ASEE Annual Conference Proceedings, July 2021, accessed March 13, 2026, <https://par.nsf.gov/servlets/purl/10304434>; Anna M. Quider and Gerald C. Blazey, “Toward More Equitable Academic Research,” *Physics* 16, no. 30 (March 6, 2023), accessed March 13, 2026, <https://physics.aps.org/articles/v16/30><https://physics.aps.org/articles/v16/30>.

front costs of realizing those benefits and concerned about AI's broader impact on the labor market and talent pipeline.

- Uneven rates of AI literacy and the lack of clear institutional strategies are producing fragmented adoption.

Promising practices

Participants also shared a number of practices and areas of experimentation that point toward productive approaches to integrating AI into research administration workflows:

- AI shows promise for assisting with risk-based proposal and compliance reviews—such as transactions that may not align with funder requirements, research methods that require heightened IRB attention, problematic contract language, or early identification of proposals that will require specialized facilities.
- The individuals who are best positioned to find AI-enabled efficiencies are research administrators. Several institutions described holding “process-palooza” events or “use-case showcases” to facilitate the lateral and upwards flow of ideas and practices.
- Several enterprise-level, community-built AI solutions are in development or available. Two of these received significant attention and interest at our workshop: TritonGPT, developed by UCSD, which is now available as SaS, and The Vandalizer (funded by the NSF GRANTED program) and developed at the University of Idaho.

Description of activities

The workshop at Montclair State University and the workshop at Chapman University followed similar structures. We opened each convening with a plenary discussion, followed by small breakout groups focused first on identifying shared challenges. Subsequent sessions were dedicated to discussing examples of strategies and experiments that have shown promise and those that did not. Afternoon sessions included more expansive, “pie in the sky” thinking about potential implementation projects, and what institutions would need to get started. (See Appendix for workshop agendas).

The opening plenary discussions at both workshops brought together participants from different positions, units, and institutional types. The Montclair State University panel was moderated by Stefanie Brachfeld, vice provost for research of the graduate school at Montclair State University, and included:

- Libby Barak, Assistant Professor, Linguistics and School of Computing, Montclair State University
- Kevin Cooke, Director of Research Policy, Association of Public and Land-grant Universities
- Eric Hetherington, Associate Vice Provost, Sponsored Research Admin, Office of Research & Development, New Jersey Institute of Technology
- Todd Slawsky, Assistant Vice President of Finance and Administration, Rutgers University

The Chapman University panel was moderated by Essraa Nawar, assistant dean for library engagement and program initiatives and chair of the libraries' AI committee, Chapman University. This panel included:

- Brett Pollak, Executive Director of IT Services, University of California San Diego
- Ed Clark, CIO, California State University System
- Sylvia Bradshaw, Executive Director of Sponsored Programs, Southern Utah University

The diversity of vantage points and experiences represented on each panel underscored the fact that AI integration, like every major technological introduction before it, is not a one-size-fits-all undertaking.⁸ Many of the issues raised during the plenaries became jumping-off points for the later breakout sessions, and as conversations moved deeper into institutional contexts, campus dynamics, and professional responsibilities, it became increasingly clear that institutions will need to take a holistic approach to integrating AI into research administration workflows. Indeed a strong throughline across both workshops was that institutional context impacts governance structures, and in an era dominated by “scalable” thinking, meaningful integration of AI into institutional research

⁸ Sheila Jasanoff, *The Ethics of Invention: Technology and the Human Future* (New York: W. W. Norton, 2016), <https://www.norton.com/books/The-Ethics-of-Invention/>.

administration systems may, in fact, prove antithetical to facile notions of scale.

Montclair's plenary panelists repeatedly emphasized that instead of relying on vague rhetoric and sweeping promises, institutions should evaluate specific tools for specific use cases (document comparison, timeline conflict detection, funding opportunity matching, portfolio analysis, contract review) while carefully considering end users, implementation pathways (commercial versus in-house), cost, staff time, benefits, risks, and data security and governance. As the discussion unfolded, the different perspectives of each panelist brought deeper structural challenges into focus. Heterogeneous data sets across colleges, escalating power demands, rapidly aging technologies, and the need for computational expertise within administrative units complicate implementation, particularly for institutions already stretched thin. Data security and stewardship remain paramount, especially when contracts contain proprietary information. Panelists described needing to build rigorous processes to ensure that data are not stored or cached by vendors, emphasizing that the least expensive component of implementation is often the tool itself, not the expertise required to use it responsibly.

With regard to staffing and employee job security, panelists acknowledged that AI may eventually reshape roles, but that the current roll-outs they are involved in are intended to support, not replace, staff by freeing time for higher-value negotiation, relationship-building, and strategic work. The most constructive path forward, they suggested, will be sustained and disciplined self-assessment: identifying institutional pain points, strengthening data foundations, understanding disciplinary needs, learning the strengths and limits of different models, and perhaps most importantly, cultivating the human capacity to engage AI intentionally rather than reactively.

The Chapman University opening plenary discussion touched on many of these themes and even embodied, by way of who was on the panel, how strongly institutional type, scale, and resource availability can shape the structure of AI initiatives on campuses. Panelists described markedly different entry points for AI integration that were tied to specific conditions of their institutional context. At Southern Utah University (SUU), a smaller teaching focused institution, AI is used by individuals in targeted and operationally focused ways. For example, the campus developed a budget

availability report that automates the flagging of high-risk expenditures for further review. By contrast, University of California San Diego (UCSD), a large, research-intensive university, described a much more expansive ecosystem approach, including the development of a locally-branded GPT environment running on an alumni-supported platform and combining commercial and bespoke tools. Within its contracts and grants office, UCSD is developing risk-based proposal review processes to identify grants involving facilities needs or human subjects/IRB research and is measuring return on investment carefully. One example of this is automating NDA redlining in ways that reduce drafting time by roughly 70 percent. Zooming further out to the system level, Ed Clark of the California State University (CSU) system discussed their 94,000-response AI sentiment survey, the largest of its kind, as a way to establish baseline metrics and define institutional success.

Contextual differences also reflected the current governance structures Chapman University's panelists described: SUU relies on specific individuals to move initiatives forward, and AI implementation is largely driven by individual initiative rather than formal committee oversight. This has enabled rapid experimentation but has also concentrated responsibility within an already small group. At UCSD, academic governance and administrative governance have been separated, with the administrative body focused explicitly on impact and ROI before handing initiatives to IT for management. At the CSU system level, shared governance was framed as essential infrastructure, necessary to build consensus, coordinate priorities, and reduce siloing across campuses. While this is what we heard in the workshop, panelists also emphasized that their governance structures, or lack thereof, are dynamic and in flux, necessarily responsive to changes in the AI ecosystem. They also agreed that, even if it is not currently the norm across all units, siloing ownership of AI initiatives poses a significant risk to meaningful integration, regardless of institutional structure.

Following the plenary discussions at both Montclair State University and Chapman University, participants spent the remainder of the convening working primarily in small-group breakout sessions. Institutional contexts and overall levels of acceptance, comfort, and fluency with AI varied to some degree between the two workshops. However, many concerns, opportunities, ideas, and experiences were shared across campus types, roles, and resource environments. The findings below synthesize insights from these breakout sessions into a composite view drawn from research

administrators, librarians, IT leaders, and research development professionals grappling in real time with how to adopt AI tools responsibly, strategically, and sustainably within the research enterprise. Although presented as distinct findings, many of these contain themes that overlap and layer upon one another. Questions of trust, for example, run through nearly all of them; workforce implications are embedded across multiple areas of concern; and confusion over AI literacy and responsibility are relevant across findings.

Key findings

Data security, quality, and governance are enduring concerns

Participants consistently described data security and quality as key concerns in whether and how AI can be integrated into research administration. Participants noted that the quality of AI outputs depends on the quality of their inputs, and several voiced concerns that their institutional data was too fragmented, spotty, and diverse to yield accurate results. Others noted that their institutional policies and standard operating procedures were sometimes contradictory or that it would be difficult to ensure that AI would have access to the most recent versions of them as they evolved over time.

We heard that “data, data, data” is both the starting point and the pain point for many efforts, and that while campuses are exploring AI-enabled services, differences in how institutions define, structure, and govern their data limit interoperability and reduce the feasibility of cross-institutional solutions. As a result, participants pointed to the evergreen need to focus on shared data standards and aligned governance frameworks. Without greater standardization at the data layer, successfully syncing AI capacity across institutions felt like a distant reality for quite a number of participants.

Participants see data security as a serious barrier to AI adoption. A number of them described, in particular, the risk of exposing institutional data to external vendors, especially as AI functionality is embedded, often

unexpectedly, into commercial platforms accompanied by shifting terms and conditions. For some, this concern prompted moving sensitive workflows out of the cloud and back onto their premises so that they could deploy local models that provide greater control over protected data (including HIPAA-covered information and materials related to patent filings). Questions around intellectual property added another layer of complexity: if research materials, draft proposals, or patent language are run through AI systems, who owns the resulting work product, and where does that IP reside? These unresolved questions about data governance and ownership, surfaced through multiple discussions, demonstrate that AI integration is not about an isolated technological adoption; it is inseparable from broader institutional data strategy. Workshop participants stressed that responsibility for data security does not lie solely with universities; companies developing foundational models and those developing AI tools also bear accountability for data stewardship and transparency. Over the course of this project, we engaged with representatives from two widely used research information and compliance platforms. They agreed with the concerns articulated by participants and are taking their responsibility for creating secure AI environments seriously.

Trust, transparency, and reliability are non-negotiable

Participants all agreed that without trust, grounded in clarity about how tools work, what they can and cannot do, how data are used, and how decisions are made with regard to AI use across campus, AI adoption will stall or provoke resistance rather than strengthen the research enterprise. They voiced a wide-ranging spectrum of concerns covering everything from the accuracy of AI outputs, to faculty fears that their research data will be handed over to vendors for LLM training, to worries about their roles being eliminated due to AI-driven automation.

Research administration requires meticulous attention to legal, ethical, and fiscal standards. Mistakes can have significant consequences for research subjects, researchers, and institutions. Many participants were concerned that currently available AI tools are not accurate or reliable enough to be trusted with such high-stakes workflows.

Research administrators often work closely with faculty, so faculty

skepticism about AI directly affects them. Participants described building relationships with faculty as a core part of their job and noted that they had worked hard to build researchers' trust. AI tested that trust. How AI tools and models are trained, how research administrators deploy them, and how companies developing and hosting these systems respect data privacy are questions faculty are deeply worried about, and yet, as the workshop participants acknowledged, remain unresolved. Participants who shared those concerns felt they could not assuage them without being disingenuous.

Institutional initiatives and operational reality are often misaligned

Integrating AI into staff workflows and processes, as many participants described, is a stress test for relationships between employers and employees that, in many cases, are already strained. Some expressed concern that senior leaders do not fully understand the scope, complexity, and compliance-sensitive nature of their work, even as they advance sweeping AI initiatives. Leadership interest in AI was at times perceived as driven by the fear of missing out, resulting in AI implementation efforts that lack clear roadmaps, playbooks, or strategic plans. Participants, both those working in research administration and those working in libraries, expressed feeling like they are the only ones who fully understand their operational processes, and they felt that leadership decisions left them responsible for translating broad and vague ambitions into workable practice.

Opinions about the workforce implications of AI are simultaneously fraught with worry and hope

Participants and panelists acknowledged that AI will replace some roles in the broader workforce and emphasized the need for intentional workforce development over passive adaptation. Administrative burden and insufficient staff capacity were described as universal pressures, and these have led to interest in tools that can translate dense federal guidance into usable formats or accelerate contract review during periods of staffing shortages and budget contraction.

In logistical terms, AI is currently most effective at tasks that mirror entry-level research administration work. Participants expressed concern over how institutions, and the field of research administration, will cultivate expertise if foundational learning experiences are automated away. If early-career professionals no longer develop judgment through document review, compliance checks, and routine portfolio analysis, institutions may inadvertently weaken the human pipeline that sustains the profession. If AI reduces the administrative burden of entry level tasks, can institutions redesign training, career pathways, and role expectations in ways that preserve professional development and sustain the talent pipeline for the field?

If early-career professionals no longer develop judgment through document review, compliance checks, and routine portfolio analysis, institutions may inadvertently weaken the human pipeline that sustains the profession.

Uneven AI literacy and lack of institutional strategy are producing fragmented adoption

AI literacy and comfort levels vary widely across and within institutions, roles, and disciplines. Some participants described themselves as novice users with limited time to invest in learning how to use AI tools or in understanding the ethical, technical, and legal issues surrounding their deployment. Others reported greater comfort experimenting with AI, yet remained uncertain about how to integrate these tools into coordinated workflows across units. This mismatch in understanding and ability creates an overarching dynamic of confusion, worry, and frustration, as expectations for adoption outpace shared knowledge and institutional guidance.

In most cases, staff use of AI is left to individual discretion and occurs without a broader strategic framework or institutional scaffolding. As a result, experimentation is ad hoc, unevenly distributed, and difficult to scale or transfer across units. Many participants indicated that their

institutions have not yet articulated a systematic approach to AI adoption within research administration, leaving individuals to navigate implementation—and its potential consequences—largely on their own.

Promising practices

As described above, the integration of AI into research administration is a complex technical and social challenge. Some institutions represented in the workshop were just beginning to consider how they might leverage it for administrative tasks; others were working to scale new workflows and practices. Most were somewhere in-between, with a few individuals actively using AI tools on their own initiative, while formal implementation and/or strategic planning about AI remain in the early stages of development. Even so, we were able to identify several types of practices that institutions might consider adopting.

Most commonly, participants were using AI to assist with identifying high-risk items for human review. These types of use cases can pop up at all stages of the research lifecycle. For example, participants described using AI to review grant proposals to identify which would need to pass an ethics review or would require investments in labs or other facilities. One discussed using AI to conduct an initial review of contracts for problematic language and in some cases to automate redlining. A few compliance officers have used it to flag expenses that might violate a funder's regulations. At their core, these activities share an emphasis on using AI as an extra layer of review that directs human attention to potential problems. They may or may not make work more efficient—in the cases described by participants, staff still conducted line by line review—but they can contribute to making oversight more effective.

Participants also described additional types of use cases for different aspects of the research lifecycle. Research offices often work with proposals that come to them from faculty, but they are also involved in identifying potential faculty collaborators who could serve as co-principal investigators on multi-disciplinary projects. Several participants indicated that AI was helping them to locate individuals who could be recruited to compete for these types of opportunities. Participants expressed interest in AI's potential to help with other aspects of proposal development such

as organizing timelines and the distribution of labor and responsibility across complex projects. On the post-award side, one representative from a large research university noted that they had some success using AI to scan grants that had received funding to identify projects with the potential for commercialization. Others were intrigued by tools from major vendors that are designed to help institutions communicate compelling narratives about the impact of the research conducted on their campus.

A few institutions are also creating enterprise-level AI solutions. Two of these received significant attention and interest at our workshop. TritonGPT, developed at the University of California, San Diego, is a mature AI platform that is trained on institutional resources and policies. TritonGPT, which has been available to the UCSD community since 2023, is now available as an SaaS for other institutions.⁹ The University of Idaho is leading development of an open-source initiative funded by the NSF's GRANTED program. On the hardware side, participants also expressed interest in building on existing models for cross-institutional data and computing resources, such as the Massachusetts Green High Performance Computing Center, to enable researchers at emerging research institutions to conduct the AI-enabled research that would keep them competitive to funders. Though it did not feature prominently in our discussions, New York's Empire AI, a massive project to create dedicated AI computing environments that serve research institutions in the state, is another example of consortial resource pooling. These kinds of consortial or community-led and developed resources are especially attractive to emerging research institutions, who typically do not have the resources to develop enterprise software or cutting-edge computational environments on their own.

We also surfaced several useful ways to encourage experimentation and to capture good ideas developed by staff. Many participants emphasized that the individuals who were best positioned to find AI-enabled efficiencies were those who most deeply understood administrative workflows. Several institutions described holding regular "process-palooza" or "use-case showcase" events to facilitate the lateral and upwards flow of ideas and practices. Those who held these events or

⁹ Claire Baytas, "How Can Universities Create AI Tools for their Communities? An Interview with the Creators of UC San Diego's TritonGPT," *Ithaka S+R*, December 7, 2023, <https://sr.ithaka.org/blog/how-can-universities-create-ai-tools-for-their-communities/>.

similar types of activities also reported that they helped assure staff that their expertise was valued and that leaders saw AI as a way to support them rather than replace them.

Our overall impression, though, is that administrators at many ERIs are still working to establish an actionable sense of how and when AI can be useful at scale. In this respect, they are in the same place as many other organizations, companies, and universities.

Conclusion

AI is disrupting higher education at a moment already characterized by similarly complex and deeply-rooted demographic, financial, and political challenges. The inseparability of these converging disruptions was evident at both workshops. Participants routinely described themselves and their colleagues as overwhelmed, exhausted, and hard-pressed to find the time to conscientiously integrate AI into their workflows, even if it could lessen their workload over time. The difficult financial headwinds that many universities face, and the resulting pressure to do more with less, have contributed to low morale and, often, created suspicions among staff about the intentions of senior leaders. These contexts are a large part of why interpersonal trust featured so prominently in conversations about barriers to AI adoption.

The generative AI tools that dominate current discussions about AI are error prone, though also capable of processing and interpreting material at dazzling speed. Humans are also fallible, but can bring actual expertise and—perhaps as importantly—accountability to the table. Compliance is a field with little room for error. One of the recurring themes at our workshops was a strong interest in identifying uses of AI that focused rather than replaced human judgment. Like much talk about “keeping humans in the loop,” it’s hard to foresee how this imperative will fare against the pressures of automation. However, the shared disposition towards responsible use, in a field with a culture that promotes attentive responsibility, is especially valuable at this moment, when use cases and norms are just beginning to come into focus.

Future research questions

- How are IRBs responding to the spread of AI as a research tool and to what extent are researchers disclosing use in accordance with IRB guidelines?
- How can the return on investment for AI tools be assessed, and what are the appropriate metrics?
- What implementations of AI at emerging research institutions are associated with growth of their research revenue, impact, and competitiveness for grants?
- Does the use of AI in research administration at ERIs impact their institutional risk profiles? Are there specific uses of AI in research administration that pose higher risks for data sovereignty, export controls, and national security vulnerabilities? What guardrails need to be in place to ensure risk to research security is minimized?

Appendix A: Montclair State University workshop

PDF located here: <https://sr.ithaka.org/wp-content/uploads/2026/03/Advancing-AI-Governance-Implementation-and-Equity-at-Emerging-Research-Institutions-Montclair-WorkshopAgenda.pdf>

Appendix B: Chapman University workshop

PDF located here: <https://sr.ithaka.org/wp-content/uploads/2026/03/PRINT-Advancing-AI-Governance-Implementation-and-Equity-at-Emerging-Research-Institutions-Chapman-WorkshopAgenda.pdf>