

I. Vision Statement

UMBC will critically examine how Artificial Intelligence (AI) is influencing its mission as a higher education institution, in Research/Scholarship/Creative Achievement, Education, Pedagogy, and Service to the state. It will strategically pilot new approaches to improve university operations and train faculty and staff to use AI tools where appropriate. As in all our work, the foundation of our AI Strategy and Implementation will be based on ethical principles and institutional values.

Our AI strategy is being developed during a time of accelerating change and AI adoption in both higher education and general society. AI has been described as a force that is "coming for the campus" [Alexander], with multiple possible consequences depending on whether institutions respond with strategic clarity or institutional drift. This moment presents both an institutional responsibility and a distinctive chance for UMBC as we craft an AI strategy that builds on the foundational strengths of the university while embracing educational, research, and administrative opportunities.. This context matters for how we frame our strategy. Any strategy that treats AI primarily as a set of tools to be governed, rather than as a structural force with implications for the institution's long-term viability and social role, will underestimate what is at stake. Managed thoughtfully and transparently, AI becomes an integral part of our evolution. By developing genuine institutional capacity in AI, not just as users of vendor-provided tools, but as leaders providing critical AI scholarship, workforce preparation, and ethical leadership, we will offer unique capacity and capability to the state that the commercial AI industry cannot replicate.

Our vision will strive to provide strategic clarity grounded in who we are at our core. As a public research university, we have deep strengths in core areas that undergird AI research across our three colleges, including but not limited to computer science, computer engineering, information systems, mathematics, psychology, philosophy, and public policy. Units like these have always worked on AI as a part of their scholarship. With the advent of Generative AI, we also have a growing list of units across our campus exploring how AI tools influence their disciplinary scholarship and pedagogy. Well-attended workshops on these themes run by the Faculty Development Center attest to this interest. We should not approach AI as an outside force to be passively accommodated or resisted. UMBC has already built significant institutional capacity: the Center for Artificial Intelligence (<https://ai.umbc.edu/>) provides an interdisciplinary hub for research and education; the Division of Information Technology (DoIT) has made tools available for applications in administrative, academic, and research settings; and a substantial cohort of faculty is advancing the frontiers of AI scholarship across multiple departments in all colleges. The development of new academic programs, including undergraduate and graduate degrees in artificial intelligence, signals that we are not merely a consumer of AI innovation but a producer of the knowledge, talent, and critical frameworks this rapidly growing field requires.

At the same time, our identity as a public institution committed to inclusion, access, and social impact shapes how this vision must be pursued. The transformative potential of AI is real, but it is not uniformly distributed, and its risks are not hypothetical. Peer institutions across American higher education, from public R1 universities to Ivy League research centers, have learned

through their own efforts that the most consequential questions about AI are not primarily technical. They concern governance, pedagogy, labor and intellectual property, and compliance and data governance, to name a few. Maryland has now enacted statutory constraints on AI use that create direct obligations for public institutions, and the federal government is discussing similar constraints, as recent debates about Mythos [Regulation of AI] show.

We will respond to this complexity with neither uncritical enthusiasm nor reflexive resistance. The word “critically” in the vision statement with which this report leads is not rhetorical; it signals a commitment to evidence-based evaluation of AI applications before and during deployment, to genuine faculty and staff participation in governance, and to disciplinary differentiation. This means recognizing that the guidelines for the appropriate use of AI in a chemistry research lab, a writing-intensive humanities seminar, the student advising office, and a financial aid process are not the same and should not be governed by identical policy. The word “strategically” similarly reflects an understanding that the pace of widespread AI adoption should not outpace the institutional deliberation necessary to sustain it responsibly based on our workgroups’ review of peer approaches.

The ethical commitments that anchor our vision are not constraints on our ambition. They are, in fact, conditions for its sustainability. An AI strategy that fails to address, among other things, data privacy, equity of access, the integrity of academic work, the fair treatment of employees, and treats ethics as merely a compliance exercise will generate distrust and institutional friction that undermine innovation. Conversely, an AI strategy overly consumed by caution will leave us unable to develop the AI capabilities that our research mission and state mandates require. We will also be unable to support our students’ futures and desires, and deal with the reality that the generation about to enter college is already using AI tools. Our vision is to navigate that tension with integrity: to be an institution that contributes to the reinvention of higher education in the age of AI, grounded in the conviction that how we engage with AI matters. The proposed strategy supports the distinctive contribution UMBC is positioned to make: not only applying AI but helping to define, through research, pedagogy, and public engagement, what responsible AI looks like in higher education. We aim to be a university that others look to as an institution that has thought carefully about what it means to be a public university in the age of artificial intelligence, and that has built its AI practices on that foundation.

II. Guiding Principles

The principles that follow emerge directly from the deliberative process of our workgroups, each of which engaged broad cross-sections of faculty, staff, and shared governance representatives during the period of November 2025 through May 2026. Many distinct and occasionally competing perspectives were expressed in these deliberations. That they converged independently on a coherent ethical framework is itself significant. It reflects the seriousness with which UMBC's community has engaged these questions and the degree of institutional consensus underlying them. The workgroups focused on the following four topics:

- Academic and Instructional AI – John Schumacher & John Fritz, Co-Chairs

- Research & Creative Achievement AI – Karl V. Steiner, Chair
- Administrative and Operational AI – Jack Suess, Chair
- Workforce Development AI – Tim Finin, Chair

These guiding principles are also well-aligned with external frameworks — and, critically, with our legal responsibilities as a public higher education institution. For example, Maryland's Artificial Intelligence Governance Act of 2024 (SB 818) establishes binding requirements for public higher education institutions regarding the governance of AI and risk-based stratification. The University System of Maryland (USM) is currently formulating rules that will govern how UMBC and its sister USM institutions comply. EDUCAUSE's 2024 Action Plan for AI Policies and Guidelines¹, developed from a survey conducted across more than 900 higher education institutions, also identifies governance, operations, and pedagogy as the three domains any comprehensive AI policy must address, and grounds those domains in the same values of shared governance, faculty autonomy, equitable access, and transparency that shaped our strategy. The practitioner literature on public-sector AI governance, synthesized in a 2026 UC Berkeley study [Chirikov], reaches many of the same conclusions, namely that risk differentiation rather than uniform constraint is the right governance architecture, that principles should be stable while operational guidance remains adaptive, and that governance should enable responsible innovation rather than function as a barrier. The principles below are grounded in who we are as UMBC. The fact that they align with what the state requires, what the broader scholarly community recommends, and what peer institutions have independently learned should give the campus community confidence that these principles provide a durable foundation.

II.a Human Agency First

We affirm that AI systems serve a consultative and supportive role, augmenting human judgment rather than supplanting it across all institutional functions: research, teaching, and administration. Faculty retain full authority over their courses and research practices; staff own and are accountable for AI-assisted work product, in accordance with applicable university policies; students are in control of their own learning. Human review is required before any AI-assisted output carries institutional consequences. This principle is fundamental to our commitment to developing institutional capacity and defining responsible AI.

II.b A Governance Structure to Guide AI Adoption

We will establish a permanent, adaptive governance framework for AI. We recommend a standing cross-functional body to support the Chief AI Officer that has a clear nexus with existing governance structures. This body will have the agility and ability to respond to the pace of technological change. This structure will support disciplinary differentiation, recognizing that appropriate AI use varies significantly across fields and purposes, from

¹<https://library.educause.edu/resources/2024/5/2024-educause-action-plan-ai-policies-and-guidelines>

a research lab to a writing-intensive seminar to an alumni outreach event. It will also operate within a nested regulatory environment (including applicable federal requirements, Maryland statute, and USM system policy) and will be transparent about where institutional discretion ends and external obligation begins.

II.c Inclusive Excellence

Our AI strategy will be designed with inclusive excellence as a foundational commitment, not an afterthought. We commit to mitigating the risks of differential access to AI tools and training, ensuring that the transformative potential of AI is distributed across all student, staff, and faculty populations, and that they have the literacy needed to use AI critically. We will evaluate and monitor AI systems to ensure alignment with the University's values. If we determine that any AI systems undermine institutional values (for example, by introducing bias in decision-making), we will, if possible, attempt to address concerns with those systems and will discontinue their use, as appropriate. We will pursue AI applications that strengthen access, enhance student success, and reflect the diversity of our community. These obligations extend beyond our campus. As a community-engaged Research 1 institution in Maryland, we will play a leading role in creating AI Tech Extensions statewide envisioned under SB 597, which creates an AI Partnership between the State and USM. We have a responsibility to bring AI capacity to communities across the state that have been historically underserved by both higher education and the technology industry.

II.d Data Privacy and Regulatory Compliance

We commit to comprehensive data privacy, security, and regulatory compliance. We treat these not as administrative requirements but as ethical obligations foundational to institutional trust. Certain uses of AI are categorically excluded without clearly identified safeguards: AI will not be used in contexts governed by FERPA or HIPAA in ways that compromise protected records. It will also not be used in behavioral intervention or counseling contexts without meaningful human review and recourse. We will evaluate and monitor AI systems to ensure alignment with the University's values. In the event we determine that any AI systems undermine institutional values (for example, by introducing bias in decision making), we will, if possible, attempt to address concerns with those systems and will discontinue their use, as appropriate. Our compliance obligations reflect applicable federal requirements, Maryland statutes, and USM system policy.

II.e Recognition of Faculty and Staff Labor

We acknowledge the significant intellectual, pedagogical, and operational contributions that responsible AI integration requires of faculty and staff. Critically examining AI tools, redesigning courses and workflows, developing new literacies, and maintaining accountability for AI-assisted work all demand real investment of time and expertise. We commit to providing the resources necessary to make that work possible and

sustainable. We emphasize that faculty retain full authority over their pedagogical choices, including how, when, and whether AI tools are incorporated into their courses. Staff will not be penalized for discussing with their managers how, when, and whether to use AI tools, where doing so reflects a reasoned professional judgment related to their job duties.

III. Strategic Pillars/Goals

III.a Academic and Instructional AI

Our approach to AI in teaching and learning is based on a foundational conviction: that education is a fundamentally human endeavor. Learning requires effortful engagement, intellectual struggle, and the kind of relational trust between teacher and learner that no automated system can replicate or replace. This conviction does not lead to a default resistance to AI tools; it defines the terms on which those tools earn a place in our classrooms and beyond.

The most significant long-term opportunity in academic AI is ensuring that all UMBC graduates possess the critical AI literacy that their futures require, regardless of their major or discipline. The evidence for this position has accumulated rapidly and now spans federal policy, employer surveys, and higher education research. The Department of Labor's AI Literacy Framework², released in February 2026, identifies AI literacy as a foundational priority to expand across the workforce and education system as a whole. This federal framework reflects what employers are already communicating directly: while 86% of organizational leaders continue to deem data literacy essential for daily workflows, 69% now identify AI literacy specifically as essential. Among students themselves, the largest AI survey ever undertaken in higher education, drawing nearly 100,000 responses from across the California State University system³ found that 69% of students (and 78% of *faculty* and 82% of *staff*) believe AI will be an essential part of most professions, yet only 37% said their curriculum offers adequate AI exposure. EDUCAUSE frames this as an institutional obligation rather than a curricular option, explicitly calling on universities to develop students' AI competencies as a core element of workforce preparation rather than a supplement to it. The conclusion these sources converge on is one that UMBC has reached independently: AI competency is no longer a specialized technical credential but a foundational element of professional readiness that the university has an obligation to provide across its full range of disciplines and degree programs. Our ongoing General Education Program (GEP) reform provides the structural vehicle for making this universal. Embedding AI literacy into the GEP, covering the principles underlying AI, practical application, critical evaluation of outputs, and

² <https://www.dol.gov/agencies/eta/advisories/ten-07-25>

³ <https://www.calstate.edu/csu-system/news/Pages/CSU-Releases-Largest-and-Most-Comprehensive-Survey-on-Artificial-Intelligence-in-Higher-Education.aspx>

responsible and ethical use, ensures that these competencies are foundational to every undergraduate's education rather than siloed in technically oriented programs.

Faculty authority over course-level decisions about AI is not merely a matter of academic freedom. It reflects a recognition that appropriate AI use varies so significantly across disciplines, course types, and student developmental stages that no institution-wide mandate, whether permissive or restrictive, can substitute for informed faculty judgment. For instance, a chemistry course built around hands-on laboratory reasoning, a writing-intensive seminar in the humanities, and an advanced data science elective present categorically different questions about where AI assistance serves learning and where it undermines it. Our AI strategy places those questions where they belong: with the faculty who design and deliver those courses, supported by institutional resources and a clear framework of expectations rather than top-down prescription.

This faculty authority carries a corresponding obligation of transparency to students. Our students are entitled to explicit, advance notice of the rules related to AI governing their academic work (ideally in the syllabus) and, where appropriate and possible, the pedagogical reasoning behind them. Whether permitting AI use for certain assignments, requiring its use, or excluding it entirely, ambiguity in AI expectations disadvantages students who have the least access to institutional guidance.

Academic integrity remains the most frequently cited barrier to faculty adoption of AI-assisted pedagogy, and it deserves a structural response rather than only a policy response. When faculty cannot trust the integrity of high-stakes assessments, they are rationally reluctant to permit AI use in lower-stakes, formative work where it offers the most genuine pedagogical benefit. Providing a secure, institutionally controlled testing environment as a pedagogical resource addresses this barrier directly and concretely, allowing the high-stakes and low-stakes functions of assessment to be clearly distinguished.

Faculty attention to AI is also increasingly a matter of assessment design. As AI tools become capable of producing competent work products across many disciplines, the question shifts from whether students are using AI in violation of stated rules to whether assessments are designed in ways that meaningfully evaluate the competencies they are meant to evaluate. Our AI strategy will support faculty in this work through the Faculty Development Center's ongoing programming on AI-aware assessment design and through sharing best-practice examples across disciplines. The aim is not to AI-proof assessment but to ensure that what we certify our graduates can do reflects what they have actually learned.

Academic integrity policies, currently being reviewed by the appropriate deans and shared governance, must consider the role of AI in Academic Integrity violations. At the same time, policies must be able to stand the test of time, even as AI technology and tools evolve. Finally, UMBC's approach to academic AI requires ongoing institutional

stewardship. The technological landscape will continue to evolve faster than any static policy document can anticipate. The governance and community-building infrastructure described in Section IV is the mechanism by which the academic AI strategy remains a living framework rather than a filed document.

Actionable Steps

- Establish a standing AI and Teaching Community of Excellence (AIT-CoE), facilitated by the Faculty Development Center and the Office of the Vice Provost and Chief AI Officer, to serve as a central hub for peer learning, quality assurance, and continuous monitoring of the evolving landscape.
- Current institutional guidelines around AI use in class are voluntary. Establish that faculty must use those guidelines to provide students with explicit, advance notice of AI expectations in their courses, including what is permitted, what is excluded, and where appropriate, the pedagogical rationales. Faculty should also consider transparently disclosing to their students how they are using AI in their pedagogy.
- Leverage the GEP reform process to embed critical AI literacy as a formal, universal undergraduate requirement, progressing from baseline AI awareness integrated into the UNIV 101 seminar sequence through discipline-spanning AI micro-credentials to a formal GEP designation.
- Begin planning for a centralized UMBC Testing Center that provides a secure, institutionally controlled environment for both face-to-face and online summative assessment, with the explicit aim of expanding faculty confidence in permitting AI use in formative and preparatory contexts.
- Task the Chief AI Officer, the Faculty Development Center, the academic Colleges, and DoIT jointly with curating and disseminating evolving subject matter expertise on AI in instruction, including developing an AI resource tool enabling community members to submit and access relevant references and examples as an ongoing community of practice.
- Conduct continuous benchmarking of UMBC's instructional AI approaches against peer institutions, with annual reporting through the CoE to the governance structure described in Section IV.

III.b AI Research

Our research and creative achievement (RCA) mission has long been central to UMBC's identity, and AI is increasingly central to how that mission is conducted across disciplines. Our RCA strategy in AI rests on four observations from the research subcommittee's working groups:

- AI tools and infrastructure available to UMBC researchers have grown substantially but require sustained curation and investment.
- AI research and AI-enabled research are increasingly interdisciplinary.

- The external funding landscape for AI is large but evolves quickly, requiring institutional capacity to identify and pursue opportunities at scale.
- Ethical use of AI in research carries distinctive obligations beyond those in administrative or instructional contexts.

AI Tools and Research Infrastructure

We have built a strong foundation of AI tools available to researchers through a combination of high-performance computing infrastructure and access to leading commercial and open-source generative AI systems. The High Performance Computing Facility (HPCF) maintains GPU-equipped machines that run open-source large language models — including Llama, Qwen, gpt-oss, and Sarvam — at no additional cost to researchers, supporting research that requires the kind of control over data, models, and provenance that commercial AI services cannot provide. Through DoIT's enterprise contracts and gateway infrastructure, UMBC researchers also have access to ChatGPT, Claude, Google Gemini, Microsoft Copilot, and other commercial systems under data-protection terms that prohibit using UMBC data for model training. UMBC's LiteLLM AI Gateway and OpenWebUI provide unified, secure access points to multiple commercial and open-source models, with the gateway architecture allowing researchers to move between models without changing their underlying code while centralizing billing, usage analytics, and budget controls.

This existing infrastructure is necessary but not sufficient. Faculty surveys conducted by the RCA subcommittee identified several persistent gaps. The most consistent finding was the need for additional specialized expertise on campus, particularly AI/ML research professionals and software development engineers who can support faculty across disciplines. This is especially important for faculty who are considering the role of AI in non-AI research. They need support in understanding the capabilities of existing AI tools available on campus, including how to use them and the ethical and data privacy concerns presented by each. Researchers identified the need for support in navigating data privacy, security, research compliance, and licensing agreements with AI vendors, including assurances against intellectual property loss and clarity around how research data is handled when commercial models are used. Across the survey responses, faculty consistently indicated that the availability of such support directly affects their willingness to pursue grant opportunities and new lines of research. The implication is that UMBC's AI research capability is currently constrained more by the staffing and administrative infrastructure surrounding the tools than by access to the tools themselves. This support could come from a variety of sources like DoIT, FDC, the Library, the CAIO, and the Colleges.

Interdisciplinary AI Research

AI is, by its nature, a deeply interdisciplinary field. The ethical, social, cultural, and domain-specific dimensions of AI research require collaboration across computer and information science, engineering, the social sciences, the arts and humanities, public policy, the health sciences, and beyond. UMBC has established institutional assets that

support this kind of work — including the Center for AI (CAI), the Imaging Research Center (IRC) , the UMBC Cybersecurity Institute (UCI), the Center for Innovation, Research, and Creativity in the Arts (CIRCA), the Dresher Center for the Humanities, and various cross-college collaborations. However, much of UMBC's interdisciplinary AI work currently depends on informal networks and individual faculty relationships, which are not institutionally sustainable as faculty leave, retire, or rotate through administrative roles.

The broader question of how UMBC supports interdisciplinary research, including matters of faculty appointment, promotion, tenure recognition for cross-college work, and graduate program structure, extends well beyond AI and is not for this strategy to resolve. However, what this AI strategy can commit to is making interdisciplinary AI collaboration easier to start and sustain within existing institutional structures. This means systematic visibility into who is doing what AI-relevant work across the campus, mechanisms to connect researchers across disciplinary boundaries around shared problems, and consulting infrastructure (statistical, computational, and AI-specific) that supports disciplines without deep computing expertise. Where UMBC's broader institutional structures present friction for interdisciplinary AI work, the AI Governance body described in Section II.b is well-positioned to surface those issues for resolution by the appropriate university decision-making bodies.

The External Funding Landscape

The funding environment for AI research has expanded substantially across federal agencies, state programs, philanthropic foundations, and corporate partners. Federal priorities shifted markedly in 2025 toward use-inspired and agentic AI — exemplified by the DOE Genesis Mission, NSF's agentic ecosystems program, and NIH's strategic plan for data science — while the NIST AI Risk Management Framework is moving from voluntary guidance toward effectively mandatory operational practice. At the state level, Maryland's AI Enablement Strategy, the AI IQ initiative, SB 597, and HB 1057 collectively position the state as actively seeking university partnerships in AI workforce development and K–12 collaboration. Foundation support and corporate funding have grown substantially as well. For instance, Schmidt Sciences offers grants up to \$5M+ for trustworthy AI research, the Humanity AI coalition has committed \$500M to people-centered AI, and major cloud providers (AWS, NVIDIA, Azure, IBM) offer significant cloud-credit awards.

This landscape is too large and too fast-moving for individual faculty members or departments to navigate comprehensively. The RCA subcommittee piloted a process using large language models to identify, score, and disseminate AI funding opportunities matched to UMBC faculty expertise. The pilot demonstrated feasibility and identified a path toward a sustained, semi-automated system that would scan federal, state, foundation, and corporate sources, match opportunities to faculty interests using internal data from Kuali, Watermark, Academic Analytics, and faculty CVs, and route targeted opportunities to faculty and research administration. Such a system would meaningfully

expand UMBC's capacity to compete for AI funding without requiring proportional growth in research administration staffing. However, UMBC should consider additional support for developing grants to non-federal funders, which would also support access to the diverse funding opportunities in AI.

Ethics of AI in Research

The RCA subcommittee developed an ethical framework for AI in research grounded in transparency, accountability, and integrity. Its core principle is that AI use in research is ethical when it augments human judgment, creativity, and efficiency, and becomes problematic when it substitutes for the core intellectual contributions on which authorship, reasoning, and accountability rest. The framework introduces several analytical concepts that have value across our research and education activities: cognitive debt (the accumulated burden of deferred sense-making), cognitive surrender (the abdication of judgment to AI), and reward hacking (the substitution of metrics for genuine judgment). It identifies specific ethical risks distinctive to research contexts — including data provenance, hallucination, bias amplification, and the emerging risk of adversarial or misleading training practices — and emphasizes that ethical evaluation of AI in research, scholarship, and creative efforts should work through existing structures (i.e. the Institutional Review Board, the Office of Research Protections and Compliance, the academic misconduct reporting process, and disciplinary scholarly communities) rather than create parallel oversight mechanisms specific for AI.

Importantly, the specifics of disclosure — when AI assistance must be acknowledged, in what form, and with what level of detail — are appropriately governed by funding agencies, journals, and the disciplinary norms of professional societies. The IEEE, ACM, APA, AHA, AMA, and ABA, among other professional bodies, have developed and continue to refine guidance on AI disclosure in scholarly work, and NIH, NSF, DOE, and major journals have issued their own policies on AI use in submitted work. Our role is to ensure that researchers know these obligations exist, can find and comply with them, and have institutional support for doing so. It is not to substitute our own disclosure rules in place of the bodies that govern scholarly publication and grantmaking. Where the RCA subcommittee or individual UMBC units have developed particularly thoughtful examples of disclosure practices or AI-resilient assessment design, we can highlight these as best-practice references for the broader campus community.

Actionable Steps

- Continue and expand our investment in HPCF GPU infrastructure to support research that requires institutional control over models, data, and provenance — including open-source LLM work that cannot rely on commercial AI services.
- Maintain and publicly document our curated portfolio of AI tools through DoIT's Software Catalog, the LiteLLM Gateway, OpenWebUI, and HPCF resources, updating it on a recurring basis as the tool landscape evolves and ensuring researchers have clear guidance on which specific tools are appropriate for which kinds of work.

- Invest in research support infrastructure specifically scaled to AI research needs: dedicated AI/ML research software engineers, statistical and AI consulting services available to faculty across disciplines, and streamlined procurement and legal review processes for AI vendor agreements that protect institutional and faculty intellectual property.
- Develop institutional capacity to make interdisciplinary AI collaboration easier within existing structures, including systematic visibility into AI-relevant work across the campus, mechanisms for connecting researchers across disciplinary boundaries, and channels for the AI Governance body described earlier to surface structural friction issues to the university bodies responsible for resolving them.
- Develop and deploy the semi-automated funding-opportunity identification and matching system piloted by the research subcommittee. Begin with a focused twelve-month pilot, scale toward a sustained service integrated with Kualī, Watermark, and faculty research profiles such as those in Academic Analytics, and extend the model to research areas beyond AI as it matures.
- Accelerate and strengthen existing research partnerships with the University of Maryland Baltimore, the University of Maryland Medical Center, and the Maryland Psychiatric Research Center for AI-driven health research; with regional and international universities for cross-institutional AI work; and across the broader Baltimore–Washington research corridor.
- Adopt the RCA subcommittee's ethical framework as institutional guidance for AI in research, integrating it with existing IRB, ORPC, and academic integrity processes rather than creating new oversight structures. Ensure that researchers can readily access current disclosure requirements from funding agencies, journals, and relevant professional societies, and that UMBC research administration provides support for complying with them.
- Position UMBC RCA strategy to align with the major federal funding shifts identified by the RCA subcommittee, such as the NIST AI Risk Management Framework as operational practice, the DOE Genesis Mission's twenty-six technology challenges, and the NSF agentic ecosystems and NIH data science strategic priorities.

III.c Administrative and Operational AI

UMBC's administrative and operational AI strategy is grounded in a clear institutional position: AI serves as a utility, not a decision-maker. In every operational context - from student services to financial workflows to human resources — the individual deploying an AI tool or agentic workflow retains full responsibility for reviewing, validating, and owning the output. This principle is not a constraint on operational ambition; it is the condition of responsible deployment. AI tools that inform administrative judgment can genuinely improve the quality and efficiency of UMBC's operations. AI tools that substitute for that judgment introduce accountability gaps that we as an institution are not prepared to accept.

Our strategy is organized around four interconnected areas: workforce preparation and values alignment, operational applications, data governance, and access to AI tools and infrastructure. Operational AI deployments will be evaluated for their effect on the staff who currently do that work, so that it might supplement rather than supplant them.

Workforce Preparation and the Training Framework

No operational AI strategy succeeds without a workforce that understands not just how to use AI tools but when and whether to use them, and what remains their own responsibility regardless. The Administrative and Operational AI Subcommittee, through a dedicated Training Workgroup, developed a values-based framework that addresses this need directly. The framework rests on five institutional principles:

- The human always remains accountable for AI-assisted work.
- AI informs decisions but does not make them.
- What matters is how AI is used, not whether it is used.
- Responsibility is shared across the individual, the supervisor, and the institution.
- Some work should not involve AI, particularly work in which the human relationship or judgment is the service itself.

These principles are operationalized through a three-tiered training architecture. The first tier establishes a universal, mandatory institutional foundation — comparable in standing to Title IX training — covering the five principles, clear ethical boundaries, and a practical self-assessment framework. Critically, staff members can complete this training without requiring any staff member to use AI tools, providing space for those who choose not to. The second tier provides resources, guidance, and accountability structures, including a decision framework to help staff reason through novel situations, and formalizes the institution's commitments regarding sanctioned tools, paid training time, and protections from retaliation. The third tier is unit-specific and voluntary, enabling departments to define AI use cases, role-based expectations, and any additional work that is locally off limits, provided all decisions honor the institutional floor established in tier one.

Operational Applications

In several areas that demonstrate the practical value of operational AI, we have already moved beyond strategy into deployment. The myUMBC Answers system has processed over 58,000 inquiries from more than 11,500 unique users during the 2025-26 academic year, providing a concrete baseline for measuring continued improvement in service responsiveness. The next phase of development will move this system toward a more agentic architecture, tailoring responses to the user's specific role and question type. A graphical workflow integration platform is in active development, enabling AI-assisted validation and process improvement across administrative workflows.

Looking ahead, AI-assisted redesign of the student residency petition process will enable students to receive real-time feedback on whether they meet residency

requirements before submitting a formal petition, substantially reducing administrative burden and improving the student experience. A software and cloud contract review tool, planned for the next academic year, will apply AI to the review of legal contracts against applicable Maryland, USM, and UMBC regulatory requirements, with planned expansion to include cybersecurity, privacy, and accessibility compliance assessment.

Data Governance

The quality of AI-assisted administrative work is entirely dependent on the quality, structure, and security of the underlying data. UMBC's data governance strategy leverages AWS DataZone capabilities via the Internet2 NET+ contract, providing enterprise-grade security for all institutional data and automated metadata extraction supporting a comprehensive institutional data catalog, supplemented by a secondary catalog through HelioCampus. This infrastructure is being developed as the major focus of the Data Management Committee through FY27. Where cost or contractual constraints require it, UMBC will deploy and leverage on-site, open-source AI models rather than commercial large language models.

Access to AI Tools and Services

UMBC provides access to a structured, tiered portfolio of AI tools through a dedicated Generative AI category in our Software Catalog⁴. Basic Gemini access is available at no cost to all Google users. Premium licenses for Microsoft Copilot and Google Gemini Pro are available through DoIT at standard institutional rates. Access to OpenAI and Claude is provided through a Libre chat interface routed through UMBC's AI gateways with Microsoft and AWS, respectively, with institutional contracts ensuring data protection and initial free usage allocations. The LiteLLM AI Gateway provides a single unified API interface enabling users to move between AI models without altering their underlying code, while centralizing billing, usage analytics, and budget controls. An internal OLLAMA service provides additional model access for faculty at no charge.

Actionable Steps

- Complete and publish the institutional Tier 1 AI training for all employees within the first year of our strategy's implementation. Discuss with shared governance whether it should be treated as mandatory. The training will be slightly different for faculty and staff. Contractors, especially in auxiliary support areas such as food or janitorial service, are not assumed to be covered.
- Similar to the faculty-focused AIT-COE, establish a staff-focused CoE supported by DoIT and the CAIO, charged with monitoring adoption patterns, identifying friction points, formulating policy recommendations, and maintaining the training framework as a living document with version control and annual review.
- Convene a structured conversation with appropriate employee groups and governance bodies about how AI affects the conditions of staff work, and how the

⁴ <https://softwarecatalog.umbc.edu>

principles in this strategy apply in operational settings. Such a conversation could explore questions including: how to distinguish good-faith reliance on sanctioned AI tools from negligent over-reliance; how to ensure AI is used to supplement rather than supplant staff; and define what “meaningful review” of AI outputs means in different operational contexts. The aim is shared understanding of these issues, not a predetermined outcome.

- Complete the AWS DataZone data architecture buildout, including the HelioCampus catalog integration, as the foundational data governance infrastructure prerequisite to deploying AI in consequential administrative processes.
- Launch the residency petition AI redesign and the software and cloud contract review tool in the 2026-27 academic year, with measurement frameworks designed from the outset to capture before-and-after staff time, error rates, and user experience.
- Expand marketing and awareness of the myUMBC Answers system and the LiteLLM AI gateway to ensure all eligible users understand available tools and access pathways, and initiate the rearchitecture of myUMBC Answers toward a more agentic, role-aware interface.

III.d Workforce Development

Preparing students and the broader workforce for careers in the context of an AI-pervasive economy requires more than awareness that AI exists. It requires critical AI literacy including the judgment of when and if to use AI tools, the ability to apply AI tools effectively, when indicated, to evaluate their outputs critically, to direct them purposefully, and to use them responsibly within the ethical and regulatory frameworks that govern professional practice across workplaces and industries. Our workforce development strategy must emerge from our broader UMBC curriculum, building specialized pathways for students seeking deeper AI expertise, and extending that capacity beyond campus to serve Maryland's workforce, communities, and public institutions.

AI Across the Curriculum

The most consequential workforce development investment we can make is ensuring that AI competency is not confined to a handful of technical programs. Employers across virtually every sector — healthcare, public administration, law, education, social services, the arts — increasingly require graduates who can engage with AI tools critically, ethically, and responsibly in their specific professional contexts. For example, a sociology graduate who understands how algorithmic systems shape social discourse, a public health professional who can evaluate AI-assisted diagnostic tools, a policy analyst who can reason about the governance of automated decision-making, a computer scientist who improves the power efficiency of AI algorithms — these are not niche

specializations. They are the baseline expectations of an economy where AI use is routine.

To that end, we are engaged in a campus-wide conversation about what AI-relevant courses can and should look like across disciplines, driven by faculty expertise and disciplinary perspective rather than a centrally prescribed template. Whether, how, and when to cover AI will be decided by academic units. This effort recognizes two distinct kinds of courses, both valuable, which different disciplines will approach differently. The first kind uses AI tools to support and enrich disciplinary inquiry — such as courses that help students in chemistry, political science, public health, or the visual arts develop practical fluency with AI methods as they apply to the problems their disciplines care about. The second kind interrogates AI from the perspective of the discipline — such as courses in philosophy that examine the relationship between AI and theories of mind and cognition; in psychology, courses that explore how AI interfaces with models of human behavior, learning, and emotion; in sociology or gerontology, courses that examine how AI is reshaping social relationships, amplifying or mitigating isolation, or restructuring labor markets and care systems. Both kinds of courses develop workforce-relevant AI literacy. Both efforts reinforce that literacy is in disciplinary depth rather than treating AI as a generic technical skill to be bolted onto any degree program.

This curricular development work will be coordinated through the governance and community of practice structures described in Section III.a and Section IV, ensuring that new courses reflect shared institutional expectations regarding ethics, critical evaluation, and responsible use, while preserving the disciplinary autonomy that makes them meaningful.

Specialized Pathways and Credentials

For students seeking deeper AI expertise, we offer a growing set of dedicated pathways. The BS and MS programs in Artificial Intelligence are designed to train students to construct AI systems — grounded in mathematical foundations, core AI theory, domain applications, and mandatory ethics requirements — rather than merely use them. The online MS in Information Systems with an AI track, designed for working professionals, extends this capacity to mid-career learners without requiring them to leave their positions while continuing their education.

Beyond degrees, AI micro-credentials developed through the UMBC Institute for eXtended Learning (IXL) — aligned with the U.S. Department of Labor's foundational AI literacy framework — provide verifiable, employer-recognized evidence of specific competencies for students who are not pursuing AI degrees. These credentials allow a political science graduate, a public health student, or an arts professional to document meaningful AI literacy for prospective employers in terms that those employers can evaluate. They also serve as stackable pathways: students can build toward deeper credentials over time without committing to a full degree program.

Statewide AI Literacy and Workforce Capacity

Our workforce development obligations extend beyond our enrolled students. As a Carnegie classified community-engaged Research 1 institution within Maryland, and in our role leading the AI Technology Extension Hubs envisioned under SB 597, we have both a mandate and a distinctive capacity to bring AI literacy to Marylanders who will not come to a university classroom. This includes state employees navigating the adoption of AI tools in public agencies, workers in industries undergoing AI-driven transformation, and residents in communities — particularly rural and historically underserved ones — who risk being left behind as AI reshapes the economy around them.

We are exploring with state partners, including the Maryland Department of Labor and the Maryland Higher Education Commission, how to develop AI literacy programs that can reach these populations at scale. A potential vehicle for this work is a cloud-hosted, browser-based AI Range, built on our existing Cyber Range infrastructure, which can deliver hands-on AI content to any learner in the state without requiring specialized hardware, physical computer labs, or high-speed broadband. Device-agnostic and designed for low-bandwidth environments, the AI Range is the technical infrastructure that makes genuine statewide reach possible — allowing a rural high school teacher, a county government employee, or a small business owner to access the same quality of AI learning environment as an on-campus UMBC student. Completion of AI Range modules will be recognized through the same IXL micro-credential infrastructure that serves our enrolled students, providing portable, employer-legible credentials regardless of where a learner first encounters UMBC's programs.

Actionable Steps

- Work with Academic Deans to identify faculty, facilitated through the AIT-CoE described in Section III.a, to develop AI course offerings across colleges, including both AI-as-tool and AI-as-subject courses, with the goal of ensuring that every one of our Colleges and Schools has relevant AI curriculum options available to its students.
- Work with the GEP re-envisioning group to discuss the integration of AI literacy requirements into the General Education Program reform, with the aim of ensuring that foundational AI competency, including critical evaluation of AI outputs, responsible use, and awareness of AI's societal implications, is a universal hallmark of a UMBC undergraduate education.
- Develop and launch a suite of AI micro-credentials through IXL, aligned with the U.S. Department of Labor's AI Literacy Framework, available to both enrolled students and external learners through non-credit pathways, and designed to stack toward degree programs where appropriate.
- Build and deploy the cloud-hosted AI Range, extending our existing Cyber Range infrastructure to support AI literacy delivery across the state, with explicit attention to accessibility for low-bandwidth and device-limited environments.

- Establish formal partnerships with Maryland's Department of Labor and Department of Commerce, with MHEC, community colleges, and local workforce development boards to coordinate AI literacy programming for state employees, community learners, and workers in industries undergoing AI-driven change.
- Ensure that AI literacy programming delivered through the AI Range and IXL credentialing infrastructure is recognized in statewide workforce development frameworks, including those governing the SB 597 Technology Extension Hubs, so that UMBC-developed training contributes directly to Maryland's measurable AI workforce readiness goals.

IV. Implementation Roadmap

IV.a Governance Structure

The President and the Provost shall appoint a small group called the AI Advisory Board (AIAB), comprised of approximately 6-8 individuals drawn from faculty, staff, and administration, to work with the CAIO. We recommend that the CIO, the VP RCA, and the General Counsel become ex-officio members of the AIAB. The other members of this group should serve renewable two-year terms.

The AIAB will carry out the governance functions described in various sections of this document. These include providing a review and assessment of the risk level of various AI uses and working with the risk management framework defined in IV.c, formulating policy recommendations, and bringing faculty and staff concerns regarding AI usage to the notice of the leadership. Once a semester, this group will present a brief update to the University Steering Committee.

IV.b Resource Allocation

REDACTED FOR EXTERNAL CIRCULATION

IV.c Risk Management Framework

The risks that AI poses to individuals and to institutional operations are real, and they vary considerably by use. Maryland's Artificial Intelligence Governance Act of 2024 (SB 818) recognizes this, and directs the University System of Maryland to adopt AI policies that are functionally compatible with those the Maryland Department of Information Technology has issued for state agencies, while preserving the latitude that academic and research work requires. That policy remains in draft and under active discussion between USM and the state. UMBC was instrumental in working with UMCP, UMB, and others in its development. In its current form it would require UMBC to:

- Adopt an institutional policy governing the procurement, development, deployment, ongoing assessment, and eventual retirement of AI systems across administrative, academic, and research operations.
- Designate an AI Risk Officer and a standing AI governance group.
- Inventory our high-risk and prohibited-use AI systems, review that inventory annually, and publish an aggregated summary, with exclusions where disclosure would compromise institutional security or where systems handle confidential, health, financial, clinical, or sponsored research data.
- Complete and document a risk assessment, together with mitigating controls, before any high-risk system is placed in use, and to revisit each assessment at least every three years.
- Maintain a process for receiving and acting on reports that an AI system has adversely affected an individual, including notification and a means of opting out.
- Operate a waiver process for academic and research uses of AI that would otherwise be prohibited, reviewed annually and reported to the USM CIO.

We expect specifics to change as USM completes its deliberations, and we will follow the policy as adopted. The functions it describes correspond closely to the governance structure proposed in Section IV.a. We recommend that an AI Risk Officer position be created and that the standing body described above serve as the institutional AI governance group, so that the campus is positioned to comply when the USM policy is finalized rather than beginning the work then.

IV.d Assessment and Evaluation Metrics

The measures below are the ones we can collect with some effort. The connection between the resources requested in Section IV.b and the numbers these measures produce is indirect, however. A faculty member who submits an AI-related proposal two years from now may have been assisted by the consulting staff, by access to compute resources added by this investment, by a seminar or a brown bag conversation, by a departmental colleague, or by none of these. Course development, publication, and external funding operate on timelines that do not align with the fiscal year in which resources are committed, and each responds to forces outside of our control, including disciplinary trends, federal funding priorities, and the market for our graduates. These measures should therefore be read as indicators of direction and of relative change, not as a return calculation on the investment. Taken together and over several years, they answer the question of whether AI capability at UMBC is broadening beyond the faculty, units, and disciplines already engaged with it. Several of these measures require a baseline captured in the first year of implementation, without which subsequent readings cannot be interpreted. We recommend that the governance body described in Section IV.a assemble these measures annually and report them as part of its update to the University Steering Committee.

Academic and Instructional

- Courses that use AI as a tool within a discipline and courses that take AI as their subject, tracked separately, with the number of departments and colleges represented and total student seats.
- Share of undergraduates completing AI literacy requirements if and when they are introduced through GEP reform, and completions of AI micro-credentials by non-AI majors.
- Share of syllabi carrying an explicit statement of AI expectations, as course-level notice moves from voluntary to expected.
- Participation in the AI and Teaching Community of Excellence and in Faculty Development Center AI programming, by college.

Research

- Student and faculty research output in AI and AI-enabled work (papers, presentations, dissertations, patents), with the understanding that publication lags might make this uninformative for at least two to three years.
- Proposals submitted in AI-related areas, and awards and award dollars, tracked separately, with particular attention to submissions from faculty who are not AI researchers.
- Proposals and publications with collaborators from more than one college, as a measure of whether interdisciplinary AI work is becoming easier to start.
- Opportunities disseminated through the funding-matching service and proposals that can be traced to it.
- Use of the new compute resources, measured by distinct users and their home departments rather than by node-hours alone.
- Volume, disciplinary spread, and wait times for the AI helpdesk and consulting services.

Administrative and Operational

- Completion of Tier 1 AI training across employee groups, and unit-level uptake of Tier 3 training.
- Usage of AI tools, disaggregated where possible by employee type and unit, reported as the share of eligible employees who are active users rather than by volume alone, since a small number of heavy users will otherwise dominate the totals.
- Measures for tools UMBC builds. For instance, before-and-after staff time, error rates, and user experience for the residency petition redesign and the contract review tool.

Workforce Development

- Enrollment and completions in degrees relevant to AI such as the BS and MS in Artificial Intelligence and the online MS in Information Systems AI track.
- Micro-credentials issued through IXL, separated by enrolled students and external learners.
- AI Range learners and completions, with geographic distribution across the state, including counties served through the SB 597 extension activity.
- Formal agreements in place with state agencies, community colleges, and workforce development boards.

V. Stakeholder Engagement Plan

This strategy will only succeed as a shared undertaking, and the engagement it requires runs in two directions. Upward, this report goes to the Provost and the President, and the governance body described in Section IV.a will maintain that connection after adoption through its semiannual updates to the University Steering Committee. The constituent senates were represented in the workgroup deliberations that produced this report, and that involvement should not end with its submission. The senates are the appropriate venues for the questions the strategy leaves deliberately open, such as whether Tier 1 training should be mandatory and how AI figures in academic integrity policy

Outward, the strategy must reach the academic and administrative units where its commitments will actually be carried out. The Chief AI Officer, working with the governance body, will present the adopted strategy to colleges, department chairs, and administrative division leaders, with two requests. The first is that units help collect and report the measures in Section IV.d that originate with them, including course offerings, syllabus practices, proposal activity, and training uptake, since no central office can observe these reliably. The second is that units tell us what is working and what is not. The measures in Section IV.d will show direction, but the reasons behind the numbers will surface in departments, and the Communities of Excellence described in Section III provide a standing channel for that feedback. The Chief AI Officer and the governance body will treat this as a continuous improvement cycle: an annual review of the measures and the feedback together, with adjustments to implementation reported through the same governance channels.