
REPORT OF THE UNIVERSITY OF MARYLAND

Report of the AI Commission



AUGUST 2024

CONTENTS

Executive Summary.....	1
I. Charge	2
II. AI Background	4
III. AI Commission Recommendations	5
Rules and Guiding Principles	5
Equitable Access to AI	8
AI Accelerator for UMD Initiatives	9
IV. Additional Pilot Project Recommendations	11
V. Appendix	19
Oversight Committee	19
Teaching and Learning overview	21

EXECUTIVE SUMMARY

Artificial intelligence (AI) has the potential to transform the way we work, teach, learn, and conduct research. This report outlines the strategic implementation of AI to ethically and fairly enhance the operational efficiency, teaching and learning objectives, and research mission of the University.

Key Recommendations

1. An AI oversight committee to ensure ethical, transparent, privacy-respecting use of AI
2. An AI Accelerator to support exploration, community connection, and organizational use of AI
3. A repository of AI tools for campus use

Implementation Strategy

Phase 1: Assessment and Planning

- Develop a detailed AI implementation roadmap specifying which tools to adopt and license to the university.
- Create plans to establish new committees and organizations on campus to support AI objectives.
- Work with units across campus to assess their needs and test whether existing AI technology can be beneficial.
- Create training materials for faculty, staff, and students about the use of AI in teaching, learning, research, and administration.

Phase 2: Pilot Projects

- Launch pilot projects in select units to test AI applications and gather performance data.
- Adjust strategies based on pilot outcomes and feedback.

Phase 3: Full-Scale Deployment

- Scale successful AI initiatives campuswide.

Risk Management

We recommend the creation of an AI Oversight Committee that will create guidelines and review AI tools to ensure they are fair, equitable, and privacy-respecting. This committee will create rules around AI adoption and create policies to ensure the rules are followed across campus. We should consider how this group may relate to campus efforts to work with the new state AI law.

I. CHARGE

The Presidential Commission on AI for Maryland was established to address the transformative potential and ethical considerations surrounding artificial intelligence. This multidisciplinary commission is tasked with exploring, evaluating, and providing recommendations for the responsible development, deployment, and ethical use of AI technologies within the university community. The commission aims to ensure that AI is leveraged to enhance education, research, service, and administrative operations while addressing key issues such as bias, transparency, accountability, privacy, and security.

President Pines asked the AI Commission to address the following areas of interest:

1. **Assessment of Current Landscape:** Review AI applications across academic, research, and administrative functions; identify AI tools for enhancing university operations.
2. **Ethical Considerations:** Investigate issues such as bias, transparency, accountability, privacy, and security.
3. **Curriculum Integration:** Ensure students have access to relevant AI education and training.
4. **Research and Innovation:** Identify opportunities for interdisciplinary AI research and foster innovation.
5. **Community Engagement:** Engage faculty, students, staff, and external stakeholders for diverse perspectives on AI.
6. **Policy and Governance:** Recommend policies and governance structures for ethical AI use.
7. **Collaboration with Industry and Government:** Explore partnerships to advance AI research, application, and policy.
8. **Accessibility and Inclusivity:** Ensure AI advancements do not exacerbate disparities.
9. **Continuous Monitoring and Reporting:** Establish mechanisms for ongoing monitoring and regular reporting.

Given the range of questions asked, the AI Commission decided to organize its activities under five subgroups:

- **MARKETING**: led by Amir Mansour and included Chris Carroll, Craig Kier and Ruibo Han
- **RESEARCH**: led by Katie Shilton and included Autumn Perkey, Naeemul Hassan, Quynh Nguyen and Stephen Dowdy
- **RULES**: led by Courtney Brown and included Aiden Wechsler, Frauke Kreuter, Leslie Krafft and Sheena Erete
- **STUDENT AFFAIRS**: led by Kevin Pitt and included Alexandra DeBus and Gus Sam
- **TEACHING**: led by Jim Zahniser and included Marcio Oliveira, Min Wu, Chester Harvey and Rance Cleaveland (deceased)

These groups held their own meetings and reported back on a regular basis to the working group.

AI COMMISSION MEMBERS

AXEL PERSAUD

Co-Chair

JENNIFER GOLBECK

Co-Chair

COURTNEY BROWN

CHRIS CARROLL

RANCE CLEAVELAND

ALEXANDRA DEBUS

STEPHEN DOWDY

MICHELE EASTMAN

SHEENA ERETE

RUIBO HAN

CHESTER HARVEY

NAEEMUL HASSAN

CRAIG KIER

LESLIE KRAFFT

FRAUKE KREUTER

ALAMIR (AMIR) MANSOUR

QUYNH NGUYEN

MARCIO OLIVEIRA

BALAJI PADMANABHAN

AUTUMN PERKEY

KEVIN PITT

AUGUSTUS (GUS) SAM

KATIE SHILTON

AIDEN WECHSLER

MIN WU

JAMES (JIM) ZAHNISER

II. AI BACKGROUND

The rapid development of open-source and publicly available artificial intelligence frameworks has significantly democratized access to AI technology. This democratization allows for widespread adoption and integration of AI in various sectors, including higher education. The potential of AI in higher education may help address some of the challenges educational institutions face. For example, ChatGPT access currently costs \$20 monthly, so students would face \$240 per year in fees to access this tool on their own.

While the hope is that AI makes institutions more engaging, efficient, effective, and equitable, it is crucial to approach this promise cautiously, ensuring that the implementation of AI technologies is grounded in the realities of educational needs and outcomes. For instance, it is important to note that two years after the public release of an open AI tool (ChatGPT), generative AI is still not trustworthy or mainstream. Generative AI can save time and boost productivity in some cases but it can also create more work and doesn't always deliver the expected results and return on investment. There are tasks it is helpful with, but also substantial hype that is not yet substantiated.

We must strategically commit to the responsible adoption of AI technologies on our campus. We recommend that we leverage these new technologies to accelerate the goals already established under the commitment of the strategic plan.

We can ensure that AI integration is effective and contextually relevant by empowering faculty, staff, and students with the training and resources necessary to actively involve them in decision-making processes. A collaborative approach amongst the university community could enhance the quality of education and foster a culture of innovation and continuous improvement. Faculty expertise will guide AI's ethical and responsible use, ensuring that these technologies augment rather than undermine the critical human elements of teaching and learning, ultimately leading to more meaningful, impactful, human-centered educational experiences for our students.

Integrating AI technologies at UMD must align with our strategic commitment to reimagining learning, taking on humanity's grand challenges, and achieving our institutional goals. By embracing responsible and informed use of AI, we can develop innovative and inclusive teaching approaches, expand experiential learning opportunities, improve administrative efficiency, and foster multidisciplinary collaboration. Through these efforts, we aim to create an educational environment that is more engaging, effective, equitable, and supportive of all learners.

III. AI COMMISSION RECOMMENDATIONS

AI has the potential to have impacts across campus, in teaching, research, administration, and student life. To realize these benefits, we must ensure that the adoption and use of AI on campus is equitable and preserves privacy. Within teaching and research applications, we must recognize the diversity of needs and support when making decisions about what AI tools to adopt and support. To foster innovation, we should support programs to bring the campus and community together to collaborate and explore uses of AI.

RULES AND GUIDING PRINCIPLES

Equity is a fundamental principle in developing and integrating AI technologies. AI has the potential to magnify disparities and exploit personal information if implemented without conscious efforts to prevent such outcomes. To avoid this requires analysis and auditing of AI tools to ensure they are not re-entrenching social biases. Protecting the personal information of our community and ensuring data security is also essential to maintaining trust and safeguarding against misuse. AI systems in education must be designed with robust privacy protections, including data encryption, secure storage, and strict access controls. It is vital to implement transparent data governance policies that inform students, parents, and educators about what data is being collected, how it is being used, and who has access to it. Additionally, users should control their data in ways consistent with university policies and Maryland Higher Education Data Privacy Law, allowing them to opt in or out of data collection and usage.

Oversight Committee

We recommend the formation of a dedicated oversight body responsible for monitoring AI usage across the campus. This body will comprise multidisciplinary experts in AI, ethics, law, and relevant industry domains. We recommend a committee of 5-7 to include members from:

- DIT
- The AI Institute
- TLTC
- VPR
- Student Affairs

The oversight body will be responsible for reviewing and approving the use of AI tools as well as addressing major issues and concerns that arise from the use of the technology. The committee should have the ability to provide feedback and comments about technologies before development begins, during deployment, and periodically post-deployment.

The oversight body must provide periodic reports on the status of AI systems, including audit findings, incident reports, and compliance status. Furthermore, the committee must provide clear accountability mechanisms outlining the responsibilities of individuals and departments involved in AI projects, ensuring that any breaches of rules or guidelines are addressed promptly and effectively.

The oversight body should foster a culture of continuous learning and adaptation, encouraging the

incorporation of new insights, technologies, and best practices. The oversight body also organizes periodic meetings, surveys, and public consultations to ensure stakeholder perspectives are considered.

In connecting with the research and educational mission of the university, the Oversight Committee should:

- Encourage researchers to create and follow field-specific frameworks for evaluating potential AI tools for their research needs. This could include:
 - Providing a model template based on existing guidelines.
 - Working with associate deans for research (ADRs) in each unit to create field-specific frameworks and share them with research faculty and staff.
- Clarify the chain of responsibility around using new AI tools for research and education. This should include:
 - Clarifying who, if anyone, researchers need to ask before using new tools.
 - Consulting with IRB about research ethics responsibilities.
 - Identifying AI expertise on campus.
 - Clarifying how both students and researchers can find out about AI resources on campus. This might take the form of a recommendation that DIT create a central site (e.g. within TerpWare) to house information about accessing campus AI resources.

Transparency and Documentation

There must be a document that defines and describes how UMD handles, maintains, uses, and develops AI systems that is available to the entire UMD community. This document must be updated annually and should contain information about decision-making criteria, data sources, and algorithms used. The document should include information about the detailed function of the oversight committee and how it makes decisions about AI technologies. The oversight committee and other relevant stakeholders should be able to access the documentation to ensure transparency and accountability. This responsibility should be handled by DIT.

Developing guidelines for AI use that are easy to understand and practical to implement is critical. Like our recording guidelines for lectures and meetings, we should document best practices and establish norms for our campus community. Faculty should be transparent about how, when, and why these tools may be used in instructional settings. Faculty should make expectations for student use of AI clear for each assignment. The Teaching and Learning Transformation Center (TLTC) provides sample language that can be included in a course syllabus. Faculty should be strongly encouraged to include this language. This responsibility should be handled by TLTC.

Given the quickly evolving nature of AI technologies, our practices and guidelines must be revisited frequently.

AI Literacy

It is important that faculty, staff, and students develop literacy around AI tools, awareness of their ethical and responsible use, and an understanding of their capabilities and limitations. We recommend that the university take the following steps:

- Incorporate fundamental information about AI in the TerrapinSTRONG campus onboarding activity by leveraging the TLTC's and University Libraries' AI and Information Literacy Module.
- Widely promote existing TLTC and library resources (especially syllabus language and AI training for instructors and students).
- AI ethics training should be incorporated into universitywide and field-specific responsible-conduct-of-research training.
 - Work with UMD libraries and TLTC, or existing contractors like CITI, to develop or find appropriate training.

Required Resources:

- Faculty/staff support, such as through TLTC, in developing resources, such as syllabus language and instructional case studies of how and why AI is used and the best practices and pitfalls to avoid.
- Resources to develop and publicize webinars and other training material for various groups of the campus communities and make periodic updates as technologies evolve. (The DIT security training “Defend Your Shell” may be one of the examples to model when developing the training material for the responsible use of AI.)
- Resources to hold an annual or bi-annual workshop/retreat (by colleges and major divisions and at the campus level) to review the use of AI and reflect and update guidelines, best practices, and/or resources.
- TLTC and Academic Technology & Innovation
 - AI Specialist | 1 FTE \$120,000 | Enhanced LMS with AI capabilities
 - AI Data Analyst | 1 FTE \$120,000 | Data & Prompt management/analysis
 - UDL specialist | 1 FTE | \$100,000 | Digital Accessibility, testing and training
 - Inst. Designer | 1 FTE | \$100,000 | Course Content Design and Development
 - Faculty Coach | 1 FTE | \$100,000 | Faculty Training, Curriculum Development
 - Professional development for existing staff | \$150,000
 - Design and deployment of training and workshop sessions | \$50,000/year

EQUITABLE ACCESS TO AI

AI tools have quickly become common on campus, leaving members of the UMD community who may not have access to them for financial or technical reasons at a potential disadvantage. UMD should work to ensure that all campus community members have access to a baseline set of intuitive AI tools that are on par with the capabilities of prevailing technologies and offer privacy protections consistent with other institutions. This will allow instructional faculty to easily incorporate AI tools into curricula and enable all students to learn about the opportunities and limitations of AI technologies. Universal access to mainstream tools will ensure that students, faculty, and staff can effectively collaborate across campus and with broader communities of academics and practitioners, inspiring innovation and affording widespread impact.

Among the AI tools to prioritize for baseline access are Universal Access AI technologies. These can democratize education by making learning resources accessible to a broader audience. This includes creating AI-driven platforms that provide educational materials in various formats (text, audio, video, interactive simulations) and languages, accommodating users with different needs and preferences. Moreover, AI can help remove barriers to education for students with disabilities by offering assistive technologies such as speech-to-text, text-to-speech, and real-time captioning.

In addition to baseline access, the campus must recognize and support the needs of certain community members for specialized AI tools. This may include tools supported centrally but useful for only a subset of community members. It may also include extremely specialized tools—used only in a specific course, research, or administrative setting—and can be procured and supported in limited and particular ways. The wide breadth of potential AI tooling, similar to the breadth of software used by different disciplines or even uniquely by small groups or individuals, suggests a need for flexible procurement and support of AI systems to reach different portions of the campus community equitably.

There may also be AI tools that UMD discourages or even precludes the use of due to known faults or prohibits because they pose substantial risks. The campus should identify problematic tools, educate community members about their flaws, and, in extreme cases, implement policies or systems to restrict their use on campus.

The rise of AI may increase inequities across multiple fronts, given that groups with resources will have access to AI first and most, which can increase economic, health, and social inequities. Without subsidies or inclusion plans, certain groups may be left behind. In the university context, special attention should be paid to student researcher access, as well as accessibility to disciplines with historically less access to computing resources.

This responsibility should be handled within DIT.

Short Term:

- Enable universal campus access to a generative AI tool with an interface and capabilities comparable to current versions of the prevailing generative AI technology (e.g. ChatGPT).
- Enable universal campus access for AI tools built into existing licensing agreements (e.g., Microsoft Office, Google Workspace).

Required Resources:

- Cloud computing resources to support campuswide access to a generative AI tool (Cost Estimate: Infrastructure \$600,000/year)
- Staff and consulting resources to support and consistently enhance a campuswide generative AI tool (Cost Estimate: 2 FTE - \$300,000/year, \$100,000/year consulting)
- Staff resources to curate and maintain lists of AI tools that are supported, specialized, discouraged, and forbidden for use on campus (Cost Estimate: .5 FTE - \$75,000/year)
- Staff resources to provide robust technical and application assistance for universal and supported tools (Cost Estimate: 2 FTE - \$300,000)

AI ACCELERATOR FOR UMD INITIATIVES

The AI Accelerator is envisioned as a central, agile AI team focused on leveraging AI to transform key operations at the University of Maryland in partnership with various UMD divisions. It can be seen as a parallel entity to the Artificial Intelligence Interdisciplinary Institute at Maryland (AIM), which is focused on research and learning. The AI Accelerator would complement AIM's work with a focus on enhancing operational efficiency, improving organizational outcomes, and bringing opportunities for AI exploration to campus and community members.

An AI Accelerator (e.g., <https://ai.asu.edu/openAI>) could serve as an incubator for ideas, and a sandbox for AI experimentation. The barriers to entry should be low, and the resources should be moderate. The goal is to encourage AI innovation with guardrails to protect data and privacy. Campus community members can develop an idea, gain access to higher-level tools and expertise to pilot that idea into a solution, and then share the results. Great results could scale up for more impact. We can learn from less-than-great results as well. Developing these mini-case studies of what people have tried will foster a culture of collaboration.

The AI Accelerator would serve as a focal point for a community of practitioners to build up knowledge about what is working on this campus with the constantly evolving tools. This would provide ongoing reinforcement of the culture of collaboration and experimentation. Ideas from the AI Accelerator could be scaled up with central professional resources to help bring the best ideas to larger audiences for a more significant impact.

We would like to see the AI Accelerator host an AI @ Maryland day to showcase the ongoing efforts around AI, new activities, and available resources. This event will educate and connect people with similar interests/needs. It will encourage a community of practice around AI on Campus. We may need multiple years of support for this event to establish the community.

The AI Accelerator should exist within DIT.

Short Term:

- Create a mechanism for people to gain access to AI resources and tools. As part of this expanded access, there should be a commitment to share lessons learned with the community. This would be a community of practice with resources to support it.
- Create an AI @ Maryland event to jump-start the Campus AI community.

Long Term:

- Curate a list of AI tools to help community members identify AI tools for which they have varying degrees of access and support. These may include the following categories:
 - Universal (available at no additional cost to the entire UMD community)
 - Supported (available on-demand and with potential additional cost)
 - Specialized (may be deployed for specialized applications requiring substantial expertise and/or cost)
 - Discouraged (use discouraged due to technical or ethical risk)
 - Forbidden (use prohibited on UMD systems due to known security risk)
- Develop an accounting system (e.g., software allowance) to provide equitable access to supported tools customizable to individual needs.

Required Resources:

- Staff support for AI Accelerator to include organizational and technical resources to assist the community. (Cost Estimate: \$300,000 for 2 FTE)
- Create a pool of higher-level resources for members of the AI Accelerator to use. This pool would need to include funding for emerging and existing tools. (Cost Estimate: \$120,000 tool pool for ~500 accounts)
- Event support for AI @ Maryland Day (Cost Estimate: \$40,000 for 300-person event)

IV. ADDITIONAL PILOT PROJECT RECOMMENDATIONS

The commission identified the following projects as being of potential interest for pilot testing. These are areas where there is speculation that AI may bring a benefit but where it is, as yet, unproven. Pilot studies would use existing tools and university data and analyze the impact AI has on efficiency and quality of results.

Development Officer AI Copilot

STRATEGIC VALUE: The AI copilot may be able to analyze and summarize donor interactions, philanthropic history, and interests to help development officers create personalized cultivation and solicitation strategies with their assigned prospects. By leveraging philanthropic databases like DonorScape and social platforms like LinkedIn, the AI copilot may be able to learn donor behavior and interests. If successful, this would provide predictive insights, identify high-potential donors, and suggest the best approaches for engagement, leading to more successful proposals and interactions.

BENEFITS: The key benefit of this AI initiative would be improved fundraising performance, with higher rates of proposal success, higher yield, and overall dollars raised. Additionally, it would enhance productivity by allowing development officers to focus on high-impact activities and improve portfolio coverage by optimizing donor relations and understanding donor preferences and behaviors.

IMPLEMENTATION: Currently, UR's Business Application team is building a proof of concept using API integration with OpenAI to provide prospect briefings with actionable summaries for development officers. The proof of concept launched in October 2024. The AI copilot could further benefit from a customGPT architecture, especially for enhanced internet capabilities

Constituent Life-Cycle Management

STRATEGIC VALUE: Implementing AI-driven constituent life-cycle management may be able to deepen engagement and personalization. By analyzing data from multiple touchpoints and performing sentiment analysis on UMD social accounts, the AI system could create tailored strategies to enhance affinity and increase constituent lifetime value. This would help understand each constituent's journey and identify key moments for intervention to maintain and strengthen relationships.

BENEFITS: The key benefit would be promoting a smooth transition towards higher engagement and affinity with UMD. Personalized strategies foster a deeper connection, driving sustained support and involvement. Proactive engagement addresses constituent needs before they arise, resulting in meaningful interactions, higher satisfaction, and a stronger university reputation.

IMPLEMENTATION: Implementing this AI system would require the digestion, analysis, and updating of large datasets. The CustomGPT architecture would analyze data from various sources to ensure better interoperability and comprehensive understanding of constituent behavior. Additional resources are needed, especially with anticipated challenges with data integration and management, which would need to be addressed to ensure the system's success.

Content Optimization for Email Communication

STRATEGIC VALUE: AI tools could enhance email campaigns by addressing low email conversion rates. By analyzing audience preferences and behaviors, AI could create more engaging and impactful email content.

BENEFITS: The primary benefit would be improved email performance, resulting in higher open rates, click rates, and overall conversion rates. AI-driven content optimization ensures emails are relevant and compelling, fostering a stronger connection with the audience. Continuous performance monitoring allows for real-time adjustments, maximizing the effectiveness of email marketing efforts.

IMPLEMENTATION: Using A/B testing, AI may be able to automate the testing process by sending different variations of an email to segmented audiences and analyzing the results. The AI system could then recommend the most effective elements to include in future emails, ensuring continuous improvement in campaign performance. The implementation is involved and lengthy in comparison to other opportunities, however.

Research COI and Protocol Review

STRATEGIC VALUE: AI tools could help researchers and the university identify and manage potential conflicts of interest, such as financial interests or relationships that may influence research objectivity. It could also analyze research protocols, consent forms, and other study documents to identify potential issues related to ethical considerations, participant safety, data privacy, and regulatory compliance.

BENEFITS: This tool could help flag areas of concern, such as unclear language, missing information, or potential risks, and provide recommendations for improvement. It could also help with regulations analysis to determine what requirements exist that pertain to specific items. These capabilities would help ensure UMD is in compliance with the law and best practices.

IMPLEMENTATION: Tools like Convercent, SMART COI, and Cayuse COI for COI identification, and ProtocolBuddy, Advarra INSIGHT, and Forte Research Systems for protocol review currently exist to help with some of this management. This pilot would investigate their existing AI capabilities as well as how AI might be integrated with these tools for improved compliance.

UMD Capital Campaign AI Copilot

STRATEGIC VALUE: With the launch of a new capital campaign, the AI copilot could analyze performance, prospect potential, and donor analytics to set and track fiscal year goals for various units and individual development officers. It could also evaluate development activities and recommend scoring formulas, incentives, and policy changes, ensuring that strategies are data-driven and effective. Additionally, the AI copilot could identify where new resources should be placed based on need and potential, recommending areas that require additional staffing to maximize fundraising outcomes.

BENEFITS: The key benefit of this AI initiative would be improving goal prediction and attainment, ensuring accurate setting and achievement of fundraising targets. Other benefits include higher return on investment through increased donation amounts, enhanced agility to quickly adapt to changing

donor behaviors, and more informed decision-making, enabling leadership to make data-driven decisions for effective fundraising.

IMPLEMENTATION: The implementation plan leverages existing data, currently used for ad-hoc analysis, ensuring a starting point for tested insights and recommendations. The AI copilot could be built using OpenAI CustomGPT, providing a robust and flexible framework for AI-driven analysis. However, the architecture is not yet in place to securely connect CustomGPTs to the analytics database. Additional resources may be needed for a timely full implementation and ongoing support for the AI system.

AI Products for Researchers and Research Administrators

STRATEGIC VALUE: This project would investigate existing tools that could support research and research administrators. This pilot task would convene a group that reviews existing AI tools through this lens and offers back recommendations on which should be considered for adoption.

BENEFITS: This will provide the university with a comprehensive look at available tools and their benefits to our research mission. The final report will be of interest to AIM and the AI Accelerator, as well as to DIT, VPR, and research support units across campus.

IMPLEMENTATION: We recommend review of the following tools to identify their capabilities, costs, and the interest faculty and staff have in them.

- Literature review tools
 - **Elicit** (developed by Semantic Scholar) - This AI-based search engine uses natural language processing and machine learning to help researchers find relevant academic papers. It can identify key concepts, citations, and connections between papers, and also provides features like topic summaries and literature maps.
 - **ConnectedPapers** - This tool uses machine learning algorithms to analyze the citation network of academic papers and identify relevant literature based on the connections between papers. It can generate visual maps of research topics and highlight influential papers in a field.
- Data analysis tools
 - **Automated Machine Learning (AutoML) Tools:**
 - Tools like Google Cloud AutoML, Amazon SageMaker Autopilot, and DataRobot offer automated machine learning capabilities, allowing users to build and deploy models with minimal coding or machine learning expertise.
 - These tools automate tasks like data preprocessing, feature engineering, model selection, and hyperparameter tuning.
 - **Conversational Analytics and Chatbots**
 - Tools like ChatGPT, IBM Watson Assistant, Google Dialogflow, and Amazon Lex enable organizations to build and deploy conversational interfaces and chatbots for customer service, data querying, and knowledge management.
 - **Natural Language Processing (NLP) Tools:**
 - Tools like Google Cloud Natural Language, Amazon Comprehend, and SpaCy can analyze and extract insights from text data, such as sentiment analysis, named entity recognition, text classification, and topic modeling.

- **Computer Vision and Image Analysis Tools:**

- Tools like Google Cloud Vision, Amazon Rekognition, and OpenCV enable organizations to analyze images and videos for tasks like object detection, facial recognition, image classification, and video analytics.

The output of this pilot would be a comprehensive report on available tools, their attributes, and a list of concrete ways they could be applied toward research and research administration tasks.

AI-Enhanced Career and Academic Planning:

STRATEGIC VALUE: AI tools could be designed to personalize and refine career and academic advising processes. By analyzing a student's academic records, extracurricular involvement, and career preferences, AI may be able to provide customized recommendations for majors, courses, and career paths.

BENEFITS: This tailored guidance could be useful for aligning educational pursuits with career goals, thereby enhancing individual academic experiences and outcomes. This AI could also assist the university's Career Center to streamline the process of identifying suitable internships and job opportunities, reducing the time students spend searching for relevant positions.

IMPLEMENTATION: The Career Center and DIT would work together to consider how to develop AI advising tools. With tools properly identified and trained on existing data, a pilot program with a select group of students across various majors would allow the university to analyze the results.

To effectively support the deployment of AI-enhanced career and academic planning tools, we recommend evaluating the following software solutions, known for their capability to handle complex data analytics and personalized recommendation algorithms:

- **Salesforce Einstein AI:**
 - Ideal for creating personalized student engagement and academic advising experiences by leveraging machine learning to predict and recommend academic pathways.
 - *Estimated Cost:* Typically ranges from \$25,000 to \$50,000 annually, depending on the scale of deployment and specific feature utilization.
- **Google Cloud AI:**
 - Provides a suite of machine learning tools that can be customized for academic advising, including the integration of data across various student activities and academic records.
 - *Estimated Cost:* Costs can vary based on usage but generally start at around \$1,500 per month for access to all machine learning tools and necessary cloud storage.
- **Oracle AI Platform:**
 - Offers powerful data processing capabilities and machine learning models suitable for developing sophisticated advising tools that can adapt to changing educational needs.
 - *Estimated Cost:* Starting at \$10,000 per year for basic packages, with more advanced features and support increasing the cost.

These tools could allow DIT and the Career Center, along with academic departments, to develop and implement effective AI advising solutions. It's important to consider additional expenses related

to system integration, ongoing maintenance, training, and possible customizations to meet specific institutional needs.

Providing Wellness and Mental Health Front-End Administrative Support

STRATEGIC VALUE: An AI-driven chatbot designed to provide front-end administrative support for students seeking mental health resources could help reduce barriers to finding and accessing in-person on-campus resources.

BENEFITS: This chatbot could serve as an initial point of administrative contact, offering basic administrative guidance on connecting students to human staff and resources at the Counseling Center.

IMPLEMENTATION: This project would begin with collaboration with the university Counseling Center to assess student needs and challenges accessing counseling and mental health resources. With these tasks in mind, DIT would work with the Counseling Center on training an AI chatbot to answer queries about scheduling, available resources, and other issues identified as information challenges for students seeking services. A pilot test could then be run to assess how helpful the AI is for prospective student connections.

To support the development of an AI-driven chatbot for front-end administrative support in mental health services, we recommend the following software solutions:

- **IBM Watson Assistant:**
 - This AI platform specializes in creating conversational interfaces for customer service, making it suitable for a chatbot that provides mental health administrative support. It can be integrated into the university's systems to offer seamless interaction with students.
 - *Estimated Cost:* Pricing starts at approximately \$120 per month for the standard edition, with additional costs based on the number of queries processed and the level of customization required.
- **Microsoft Azure Bot Service:**
 - Utilizes Microsoft's cognitive services to build, connect, deploy, and manage intelligent bots that interact naturally wherever your users are talking, from your app to Microsoft Teams and other channels.
 - *Estimated Cost:* The cost can range from \$0.50 to \$1 per 1,000 messages, depending on the services used and the computational requirements for handling queries.
- **Dialogflow by Google Cloud:**
 - This platform enables the development of chatbots and other conversational AI applications with natural language understanding. It is ideal for creating a chatbot that can understand and respond to student inquiries regarding mental health scheduling and support.
 - *Estimated Cost:* The standard plan is free, but enterprise-grade solutions that offer more advanced analytics and support start at approximately \$20,000 annually.

These tools are designed to provide the required capabilities for developing a responsive and sensitive chatbot that prioritizes confidentiality and ethical practices. Additional considerations for the budget

should include the integration with existing student affairs portals, ongoing maintenance, and possible scalability to handle more complex queries or a larger user base.

Dynamic Campus Engagement Platform

STRATEGIC VALUE: An AI-powered platform that provides personalized recommendations for events, clubs, and activities could help enhance campus life. This system would analyze interests that students choose to share to suggest the most relevant and engaging opportunities, thereby fostering a vibrant campus community and enhancing the overall university experience.

BENEFITS: Personalized recommendations could help students engage more deeply with campus life, enriching their college experience by connecting them with events and activities that align with their interests. The platform may make it easier for students to find and participate in events they care about, fostering a stronger, more inclusive community.

IMPLEMENTATION: This project would begin with a collaboration with student affairs, academic IT services and student life coordinators to develop the AI platform. Students can express interest in events and the AI would use data that other students share with the platform as well as data about events and organizers to help recommend similar events. This could be pilot tested with small user groups to assess any benefits it offers before deciding on roll out to the full campus

For the development of an AI-powered platform to enhance campus engagement, the following software solutions are recommended:

- **Salesforce Community Cloud:**
 - This platform can be tailored to create a dynamic engagement environment that integrates seamlessly with other university systems, providing personalized event and activity recommendations.
 - *Estimated Cost:* The cost typically starts around \$10,000 annually for basic setups but can increase based on custom features and the number of user licenses.
- **Adobe Experience Platform:**
 - Known for its robust data management capabilities, this tool can be used to analyze student interaction data effectively and deliver personalized content.
 - *Estimated Cost:* Pricing starts at approximately \$50,000 per year, depending on the scale of deployment and specific configuration requirements.
- **Microsoft Azure AI:**
 - Offers a comprehensive suite of AI tools that can be integrated into the development of a campus engagement platform, capable of processing complex datasets to provide personalized recommendations.
 - *Estimated Cost:* Monthly costs can start at around \$1,000, varying with the extent of usage and the complexity of data processing needed.

These platforms will support the integration of diverse data sources, from student schedules to past event participations, enabling the AI to make accurate and relevant recommendations. Additional

budget considerations should include costs for initial setup, ongoing maintenance, custom development, and potentially, costs related to scaling up the platform as student participation increases.

Academic Integrity and Support for the Office of Student Conduct:

STRATEGIC VALUE: The rising use of artificial intelligence (AI) tools in academics at the University of Maryland presents new challenges for maintaining academic integrity. The Office of Student Conduct has noted a significant uptick in AI-related integrity referrals, highlighting an urgent need for enhanced staffing and resources to tackle these challenges effectively. This pilot would introduce AI tools to detect the use of AI in academically dishonest ways and create training for staff and faculty around managing AI and academic integrity.

BENEFITS: With improved resources and continuous education efforts, the university will better equip both faculty and students to understand and navigate the ethical challenges posed by AI, ensuring the integrity of the educational process.

IMPLEMENTATION: This project will develop and implement targeted training programs for faculty and staff on the ethical use of AI in education. This initiative will include creating materials that outline best practices for integrating AI tools into teaching in a manner that upholds academic integrity. It will also roll out workshops for students that focus on the ethical utilization of AI in their educational processes. These workshops will emphasize AI as a supportive tool that enhances learning and will guide students in understanding the limitations and responsible use of AI technology. Finally, it will identify and deploy AI-driven tools designed to detect academic integrity violations effectively. These tools will enhance the ability of the Office of Student Conduct to monitor and respond to incidents swiftly.

To effectively enhance the Office of Student Conduct's ability to manage academic integrity in the age of AI, we recommend the following software solutions:

- **Turnitin:**
 - Well-known for its plagiarism detection capabilities, this tool can be adapted to include AI-driven features that identify not only plagiarism but also other forms of academic dishonesty related to AI tool misuse.
 - *Estimated Cost:* Depending on the number of users and extent of features required, annual licensing can range from \$3,000 to \$15,000.
- **CopyLeaks:**
 - This AI-powered plagiarism detection software can scan text documents for academic integrity violations across a broad spectrum of data sources, including online content and databases.
 - *Estimated Cost:* Prices vary based on the volume of the submissions and the level of data access needed, typically starting at \$10 per month per user for educational institutions.
- **Canvas:**
 - Enhance the existing learning management system with integrated AI monitoring capabilities to observe student submissions and detect patterns indicative of unethical AI use.
 - *Estimated Cost:* Integrating advanced monitoring tools can increase the annual cost by \$5,000 to \$20,000, depending on the complexity of the integration and the number of users.

These tools will provide the necessary technological support to enhance the monitoring capabilities of the Office of Student Conduct, ensuring effective oversight of academic integrity. Additional budget considerations include the cost for training faculty and staff on these new tools, ongoing maintenance, and potential upgrades as AI technology evolves. This investment will ensure that the university remains at the forefront of ethical education practices, capable of addressing the new challenges presented by AI in academia.

APPENDIX

OVERSIGHT COMMITTEE

This section lays out detailed responsibilities of the proposed Oversight Committee.

Bias Auditing

The oversight committee should develop policies for bias auditing of AI in use on campus and serve as the body that holds, analyzes, and informs about rules and policies that relate to algorithmic bias. Audits should include the following:

- A complete AI impact assessment made up of:
 - The intended purpose for the AI and its expected benefit.
 - The potential risks of using AI, as well as mitigation measures the department is taking. This should include an analysis of stakeholders most likely to be impacted by the AI's decisions.
- The quality of the relevant data:
 - The data used in the AI's design, development, training, testing, and operation and its fitness to the AI's intended purpose.
- The quality and representativeness of the data for its intended purpose:
 - How the data is relevant to the task being automated.
 - Whether the data contains sufficient breadth to address the range of real-world inputs the AI might encounter.
 - how data gaps and shortcomings have been addressed.
- The AI's real-world performance:
 - Performance should consider both the specific technology used and feedback from human operators.
 - Whether the AI model could foreseeably use other attributes as proxies and its use would significantly influence model performance
 - Whether there are significant disparities in the model's performance across demographic groups
- Analyze disparate-impact risks of algorithmic eligibility determinations and algorithmic information availability determinations based on:
 - Actual or perceived race, color, religion, national origin, sex, gender identity or expression, sexual orientation, familial status, income, or disability
- Notify negatively affected individuals.
 - The department must notify individuals when use of the AI results in an adverse decision or action that specifically concerns them.

Transparency and Documentation

There must be a document that defines and describes how UMD handles, maintains, uses, and develops AI systems that is available to the entire UMD community. This document must be updated annually

and should contain information about decision-making criteria, data sources, and algorithms used. The document should include information about the detailed function of the oversight committee and how it makes decisions about AI technologies. The oversight committee and other relevant stakeholders should be able to access the documentation to ensure transparency and accountability.

Regular Review and Update of Guidelines

AI guidelines must be regularly reviewed and updated to reflect technological advancements and emerging ethical considerations. The oversight body is responsible for ensuring guidelines remain current and relevant.

Ethical and Legal Compliance

AI systems must adhere to ethical principles, such as fairness, transparency, and respect for user privacy. Compliance with local, national, and international laws and regulations is mandatory. The oversight body will be responsible for monitoring compliance and recommending corrective actions for any violations.

Incident Reporting and Management

A formal incident reporting process must be established for identifying, documenting, and addressing any issues or unintended consequences arising from the use of AI systems. All incidents must be reported to the oversight body within 72 hours of discovery. A thorough investigation must be conducted, and appropriate measures must be taken to mitigate risks and prevent future occurrences. The board documents will outline clear procedures for addressing and rectifying reported issues promptly and effectively.

Stakeholder Engagement

It is imperative that the full UMD campus is aware of what is happening. Regular engagement with stakeholders, including employees, users, and external partners, is essential to gather feedback and address concerns related to AI systems. The oversight body should organize periodic meetings, surveys, and public consultations to ensure stakeholder perspectives are considered.

Accountability and Whistleblower Protection

A secure and anonymous channel must be provided for employees and stakeholders to report unethical practices, regulatory violations, or any concerns regarding AI systems without fear of retaliation. The oversight body must ensure that whistleblower reports are taken seriously and investigated promptly.

Training and Awareness

Regular training programs on ethical AI practices, regulatory requirements, and accountability measures must be conducted for all employees involved in AI development and deployment. Resources should be provided to the campus community to understand and implement these guidelines.

Awareness campaigns should be launched to educate stakeholders about their rights, the potential impacts of AI systems, and the avenues available for raising concerns.

Sandbox Environment

A sandbox environment should be established where individuals can experiment and innovate within the enterprise system without affecting live operations. Approvals must be obtained from the oversight body

before moving any developments from the sandbox environment to actual production, ensuring all projects meet the required ethical and regulatory standards. This should be considered in the context of the AI Accelerator recommended in this report.

Reporting and Accountability Mechanisms

By adhering to these rules, the organization can ensure that its AI systems are developed and operated in a manner that is transparent, accountable, and ethically sound, thereby fostering trust and integrity in AI technologies.

TEACHING AND LEARNING OVERVIEW

One of our primary motivations should be to equip our faculty and students with all the necessary skills and digital fluency that are becoming increasingly important in the modern professional landscape. As AI tools become ubiquitous in various industries, our students must be prepared to utilize these technologies competently and confidently.

AI also presents a unique opportunity to make learning more accessible and expand inclusivity and agency in learning. We can create personalized and customized learning experiences catering to each student's needs and preferences. This could include, but may not be limited to, providing tailored feedback, identifying areas where students may need additional academic support, and offering resources that align with their learning preferences. It is possible that AI-driven personalization can significantly help us advance our academic mission by augmenting opportunities to promote a culture of effective, engaging, and inclusive approaches to teaching and learning. AI technologies have substantial potential to help students organize and access information and engage with learning materials in differentiated ways. AI is adept at distilling and transforming information into various modes—text, speech, visual, and more—addressing a perennial challenge in teaching and learning. AI can help students better understand and retain information by offering multiple ways to interact with and interpret content, ultimately enhancing their learning outcomes.

With the rapid development of AI technologies, we must approach its integration in our educational context with cautious optimism. The potential for AI to address pressing educational challenges is real, but we must be wary of promises that are overstated for financial gains. The critical evaluation and implementation of AI tools must be done thoughtfully, with the goal of improving learning outcomes at the forefront of each decision.

History has demonstrated that technology adopted in academic settings without a deliberate educational purpose could lead to unintended negative consequences. To prevent this, our approach to developing and deploying educational AI tools must prioritize promoting critical thinking, impactful research, and authentic intellectual engagement. This requires us to consider the unique needs of educational contexts and the diverse ways students learn. Our motivations for integrating AI into higher education are driven by the desire to enhance student learning, prepare them for their future careers, and address educational challenges through innovative solutions.

INSIGHTS FROM RECENT NATIONAL RESEARCH:

A recent poll conducted by [Impact Research for the Walton Family Foundation](#) highlights AI's growing prevalence and approval among students and educators. In summary, students are already experimenting with innovative prompts and discovering effective ways to utilize AI tools. An open exchange of ideas and experiences can enhance AI literacy, allowing students to learn from each other's successes and challenges. Encouraging such discussions will promote a deeper understanding of AI and support its more effective and responsible integration into educational practices.

The survey reveals a significant increase in AI usage among K–12 students and teachers over the past year. Less than 20% of students report never using generative AI, while almost half of U.S. teachers and K–12 students use ChatGPT weekly. Familiarity with ChatGPT has also surged, rising from 55% to 79% among teachers and 37% to 75% among K–12 students in just over a year.

Most students and educators view AI chatbots positively. Seventy percent of K–12 students have a favorable opinion of AI chatbots, which increases to 75% among undergraduates. Among parents, 68% expressed favorable views. However, only 25% of teachers report receiving training on AI chatbots, and 32% cite a need for more training and professional development as a primary barrier to their use. Educators are utilizing AI for various purposes, including idea generation, lesson planning, creating quizzes and tests, and developing student worksheets. Black and Hispanic K–12 students and undergraduates are more likely to use AI for school-related tasks. Among parents, 47% support the increased use of AI chatbots in schools, compared to 36% who oppose it.

Finally, despite these positive trends, nearly 20% of teachers believe ChatGPT has had a negative impact, up from 7% last year. A significant portion of K–12 students (56%) use AI chatbots to write essays and assignments, while 52% use them to study for tests and quizzes. Although 66% of [teachers permit using AI for schoolwork](#), only 32% report that their schools have appropriate use policies. Among undergraduates, 52% say their institutions have such policies.

INSIGHTS FROM UMD STUDENTS AND INSTRUCTORS SURVEY:

The 2023–2024 academic year marked the first full year during which generative AI (genAI) tools were broadly accessible to the entire campus community. To establish a baseline understanding of students' and instructors' familiarity with and use of genAI tools in academic and instructional contexts, we administered two parallel surveys to both groups in the Spring of 2024. This survey initiative was led by the Provost's office and designed and administered by the Academic Technology and Experience team, part of the Academic Technology and Innovation Division of IT, in collaboration with the Office of Institutional Research, Planning, and Assessment. Here is the report which includes an [executive summary](#), which expands on the following main overall findings:

- Most instructors and students do not categorize themselves as routine users of genAI tools. Only 14% of instructors and about a third of students said they routinely incorporate genAI in some or many aspects of their instructional or academic work. Moreover, a higher percentage of graduate students report using GenAI for academic work than undergraduate students.

- There is no dominant instructional or academic purpose for which genAI tools are used, as students and instructors report utilizing these tools for various similar purposes. Instead, instructors and students who reported using genAI tools tend to use them for various purposes. The top three instructional or academic uses of genAI for instructors and students include generating ideas, improving content, and summarizing concepts.
- There is an opportunity for additional guidance from instructors, departments, and the university around genAI use. About half of instructors (48%) who use genAI for instructional purposes said they do so without guidance from their departments or the university, and over a third of students (38%) who use these tools in academic contexts do so without direction from their instructors. Additionally, most instructors (69%) and students (73%) say they have not discussed how to cite the use of genAI tools in any of their classes, and almost half of instructors (48%) say they have not included guidance about these tools on any of their syllabuses.
- Instructors' perceptions of student use differ from students' actual reported use. Routine users of genAI tools within both groups identified using the technology to summarize concepts, improve content, and generate ideas. Yet, most instructors reported that students use genAI to create new content. Moreover, there needs to be more connection between students and instructors regarding perceptions of students' ability to use genAI responsibly. Sixty-four percent of students agree that they understand how to use genAI responsibly, but only 18 percent of instructors believe that students do so.

