UNIVERSITY OF MARYLAND
OFFICE OF THE PRESIDENT

1101 Thomas V. Miller, Jr. Administration Building
College Park, Maryland 20742
301.405.5803 TEL
301.314.9560 FAX

April 25, 2025

Dr. Sanjay Rai
Secretary
Maryland Higher Education Commission
217 East Redwood Street, Suite 2100
Baltimore, MD 21202

Dear Secretary Rai:

I am writing to request approval for a new Master of Science program in Artificial Intelligence. The proposal for the new program is attached. I am also submitting this proposal to the University System of Maryland for approval.

The proposal was endorsed by the appropriate faculty and administrative committees. I also endorse this proposal and am pleased to submit it for your approval.

Sincerely,

Darryll J. Pines
President
Glenn L. Martin Professor of Aerospace Engineering

DJP/mdc

cc: Candace Caraco, Associate Vice Chancellor
Jennifer King Rice, Senior Vice President and Provost
Amitabh Varshney, Dean, College of Computer, Mathematical, and Natural Sciences

Office Use Only: PP#

Cover Sheet for In-State Institutions
New Program or Substantial Modification to Existing Program

Institution Submitting Proposal	University of Maryland, College Park

Each action below requires a separate proposal and cover sheet.

- (●) New Academic Program
- ( ) New Area of Concentration
- ( ) New Degree Level Approval
- ( ) New Stand-Alone Certificate
- ( ) Off Campus Program
- ( ) Substantial Change to a Degree Program
- ( ) Substantial Change to an Area of Concentration
- ( ) Substantial Change to a Certificate Program
- ( ) Cooperative Degree Program
- ( ) Offer Program at Regional Higher Education Center

Payment Submitted: (●) Yes ( ) No

Payment Type: ( ) R*STARS # JJ551994 ( ) Check # JJ551994

Payment Amount: 850

Date Submitted: 4/25/2025

Department Proposing Program	College of Comptuer, Mathematics, and Natural Sciences	
Degree Level and Degree Type	Master's; Master of Science	
Title of Proposed Program	Artificial Intelligence	
Total Number of Credits	30	
Suggested Codes	HEGIS: 079904	CIP: 11.0102
Program Modality	( ) On-campus ( ) Distance Education (fully online) (●) Both	
Program Resources	(●) Using Existing Resources ( ) Requiring New Resources	
Projected Implementation Date (must be 60 days from proposal submisison as per COMAR 13B.02.03.03)	(●) Fall ( ) Spring ( ) Summer	Year: 2025
Provide Link to Most Recent Academic Catalog	URL: https://academiccatalog.umd.edu/	
Preferred Contact for this Proposal	Name: Michael Colson	
	Title: Senior Coordinator for Academic Programs	
	Phone: (301) 405-5626	
	Email: mcolson@umd.edu	
President/Chief Executive	Type Name: Darryll J. Pines	
	Signature: Darryll J. Pines	Date: 04/25/2025
	Date of Approval/Endorsement by Governing Board:	

Revised 1/2021

A. Centrality to the University's Mission and Planning Priorities

Description. The University of Maryland, College Park (UMD) proposes a **Master of Science (M.S.) in Artificial Intelligence (AI)**. This program will combine technical courses in the fundamentals of AI and courses that address the interaction between AI and humans and society. During their coursework, students will build solid foundations in mathematics, statistics and computing and also obtain a broader view of human centered AI and its societal implications. Students will gain expertise in machine learning, deep learning, and AI-driven decision-making while exploring areas such as AI ethics, human-computer interaction, explainable AI, and policy considerations. The program prepares graduates to develop AI solutions that enhance human well-being, promote fairness, and integrate seamlessly into social and professional contexts. The program consists of 30-credit course work and is a non-thesis MS program. Students will be prepared for careers across disciplines and they will develop skills to be collaborative, adaptable problem solvers in a rapidly changing field. **Please note: The program will be offered both in-person and through a fully online modality.**

Relation to Strategic Goals. UMD is the state's flagship campus and a national leader in higher education. UMD is ranked in the top 20 public universities in the nation and in the top 50 among all national institutions by US News and World Report.[1] Through our strategic partnership with the University of Maryland, Baltimore, we rank among the top 30 institutions for research funding according to the National Center for Science and Engineering Statistics.[2] UMD ranks in the top 10 for research expenditures in computer and information sciences. As written in our mission statement, "UMD embraces its flagship status and land-grant mission to share its research, educational, cultural, and technological strengths to bolster economic development, sustainability, and quality of life in Maryland and beyond." The proposed program will equip students with advanced technical skills in machine learning, data science, and AI systems, while fostering ethical, human-centered, and socially responsible innovation. Graduates will be prepared for leadership roles in AI across industry, government, and research, with the ability to develop impactful and trustworthy AI solutions.

Funding. This MS program will be self-supporting with tuition revenue. The College of Computer, Mathematical, and Natural Sciences, through its Science Academy, already offers self-support master's programs in Applied Machine Learning, Data Science, Quantum Computing, and Bioinformatics and Computational Biology. The Science Academy already has the administrative and physical infrastructure to offer the program.

Institutional Commitment. UMD is committed to being a higher education leader in AI and will leverage its instructional, research, and administrative capabilities toward this end. On April 9, 2024, the University of Maryland launched the Artificial Intelligence Interdisciplinary Institute at Maryland (AIM) (see https://aim.umd.edu/about), a collaborative hub advancing responsible

[1] U.S. News & World Report: https://www.usnews.com/best-colleges/university-of-maryland-2103.
[2] National Center for Science and Engineering Statistics: https://ncses.nsf.gov/surveys/higher-education-research-development/2023.

and ethical AI education, research, and innovation across all disciplines. Building on UMD's existing AI expertise and over 100 faculty engaged in AI scholarship, AIM supports faculty research, experiential learning, and workforce development while coordinating new academic programs, high-performance computing resources, and partnerships with government and industry. Aligned with national and state executive orders on trustworthy AI, AIM prepares students across all majors to apply AI in their fields and drive public good in a technology-rich world.

B. Critical and Compelling Regional or Statewide Need as Identified in the State Plan

Need. AI technologies are rapidly evolving and being more integrated into various aspects of society and industry. AI's evolution will rival the internet itself in terms of its potential to be both extraordinarily beneficial and profoundly detrimental to society. In January 2024, Governor Moore issued an executive order recognizing that "Artificial intelligence (AI) is transforming society and work in myriad ways, and the pace of that change will continue to accelerate-unlocking new opportunities and risks for Maryland's residents, workers, and economy."[3] AI is rapidly becoming a cornerstone of innovation across both public and private sectors, transforming industries, reshaping workforce demands, and redefining decision-making in everything from national security to consumer services. As AI capabilities expand, Maryland's public institutions, federal agencies, and private enterprises will increasingly depend on a workforce equipped not only with technical AI skills but also with the ethical and contextual understanding required to apply AI responsibly. The proposed program responds directly to this need. The program provides rigorous technical training in machine learning, data science, and computing systems, equipping students with the skills needed to build and deploy advanced AI solutions. At the same time, the program emphasizes the ethical, societal, and policy dimensions of AI through specialized coursework in responsible and human-centered design. Drawing on UMD's world-class AI research enterprise—including over 100 faculty across disciplines and the Artificial Intelligence Interdisciplinary Institute at Maryland—students gain both cutting-edge knowledge and the broader perspective needed to lead AI innovation responsibly.

State Plan. The proposed program aligns with the 2022 *Maryland State Plan for Postsecondary Education,* specifically Priority 5, "Maintain the commitment to high-quality postsecondary education in Maryland," and its Action Item to "Identify innovative fields of study." Artificial intelligence represents one of the most transformative and fast-evolving domains of the 21st-century economy, with sweeping implications for Maryland's public and private sectors. By offering interdisciplinary, ethically grounded, and technically rigorous AI education, the MSAI program prepares students to meet the workforce demands of this innovative field while upholding the state's commitment to academic excellence, equity, and responsiveness to

[3] State of Maryland Executive Order Catalyzing the Responsible and Productive Use of Artificial Intelligence in Maryland State Government: https://governor.maryland.gov/Lists/ExecutiveOrders/Attachments/31/EO%2001.01.2024.02%20Catalyzing%20the%20Responsible%20and%20Productive%20Use%20of%20Artificial%20Intelligence%20in%20Maryland%20State%20Government_Accessible.pdf

emerging global challenges. Through this program, Maryland is poised to become a national leader in responsible AI talent development.

C. Quantifiable and Reliable Evidence and Documentation of Market Supply and Demand in the Region and State

There is strong and growing demand for professionals with advanced training in artificial intelligence, both nationally and in the state of Maryland. According to the U.S. Bureau of Labor Statistics, employment of computer and information research scientists—a role closely associated with AI professionals—is projected to grow by 26% nationally between 2021 and 2031, significantly faster than the average for all occupations.[4] The Maryland Department of Labor similarly projects a 21% increase in these occupations between 2022 and 2032.[5]

A labor market analysis conducted for the University of Maryland confirms this trend, showing that from January 2020 to January 2025 there were over 30,000 unique job postings nationally listing artificial intelligence as a required skill, with Washington, D.C. among the top ten U.S. cities for such postings. The analysis further reveals that 41% of these job postings required or preferred a master's degree, reinforcing the need for graduate-level training in AI.

AI will be particularly critical for Maryland and the surrounding region due to the high concentration of federal agencies, national laboratories, and private-sector firms engaged in AI research and application. The current downsizing of the federal workforce makes this need even more urgent, as agencies increasingly turn to AI to fulfill their missions and require personnel with advanced technical and analytical skills. The proposed program addresses this gap by producing graduates trained in both the technical foundations of AI and in its human-centered, ethical application. In doing so, it supports workforce development across key sectors—including healthcare, defense, finance, cybersecurity, and public policy—where AI is rapidly becoming integral to operations and strategic decision-making.

D. Reasonableness of Program Duplication

Demand for AI expertise is rapidly outpacing the supply of graduate-level programs—especially those that combine rigorous technical training with critical attention to ethical, social, and policy considerations. As AI becomes increasingly foundational to sectors such as defense, healthcare, finance, and education, Maryland must expand its capacity to educate professionals who are prepared to lead in this space.

National labor market data underscore this mismatch. From 2020 to 2025, over 30,000 unique job postings listed artificial intelligence as a required skill, yet in 2022, only 554 master's

[4] USBLS Occupational Outlook Handbook: Computer and Information Research Scientists: https://www.bls.gov/ooh/computer-and-information-technology/computer-and-information-research-scientists.htm
[5] Maryland Department of Labor: Maryland Occupational Projections 2022-2032: https://labor.maryland.gov/lmi/iandoproj/maryland.shtml

degrees were awarded in AI nationwide—yielding approximately 54 job postings per graduate. This highlights a clear shortage in the graduate talent pipeline, particularly for programs that equip students with both technical and ethical competencies. Rather than representing duplication, the proposed program addresses a quantifiable and urgent workforce need.

Only two master's programs in Artificial Intelligence are currently approved in Maryland: those at Capitol Technology University and Johns Hopkins University. Given the scale and urgency of AI's transformation across the public and private sectors, the demand for high-quality graduate programs is not only unmet—it is growing. Maryland should encourage multiple, complementary offerings to meet distinct student and workforce needs. UMD's program will help ensure the state remains at the forefront of responsible AI development while expanding access to education in a field central to its economic and civic future.

UMD's program is uniquely positioned due to its integration with the newly launched Artificial Intelligence Interdisciplinary Institute at Maryland (AIM), which unites over 100 faculty across disciplines to support innovative AI research, education, and policy development. The program also benefits from UMD's strategic location near the federal government and major employers driving AI adoption, making it an ideal hub for training the next generation of AI leaders.

E. Relevance to High-demand Programs at Historically Black Institutions (HBIs)

No HBI's currently offer a master's program in Artificial Intelligence.

F. Relevance to the identity of Historically Black Institutions (HBIs)

We do not anticipate any negative impacts on the unique identities of Maryland's HBIs. UMD already offers a master's program in Computer Science with an Artificial Intelligence concentration, as well as graduate programs in Data Science and Applied Machine Learning. The proposed MS in Artificial Intelligence builds on this foundation to meet rapidly growing demand in a field that is critical to the state. As global investment in artificial intelligence continues to increase each year, we believe the state should actively encourage the development of more highly specialized, technical graduate programs—particularly in diverse geographic areas—to promote inclusive economic growth and workforce development in this vital sector.

G. Adequacy of Curriculum Design, Program Modality, and Related Learning Outcomes

Curricular Development. When developing the proposed curriculum, several factors were considered to ensure the program's success. An advisory group was formed to identify necessary skills, capacities, and capabilities in job postings in the market. These in demand skills led the curriculum development to ensure that students will be qualified to satisfy the workforce gap at the time of graduation.

Faculty Oversight. The College of Computer, Mathematical, and Natural Sciences' Science Academy will work with the Mathematics and Computer Science department chairs, as well as the Artificial Intelligence Interdisciplinary Institute at Maryland (AIM) for oversight. A Computer Science faculty member will serve as the faculty director. This faculty director will provide the curriculum oversight, course evaluation, and advise students. In addition, the academic faculty director, in collaboration with the Assistant Dean for Professional Graduate Education, is responsible for all instructor selections and appointments. Appendix A is a list of faculty who will be teaching in the program.

Educational Objectives and Learning Outcomes. The main objective of the program is to blend highly technical, applied, experiential learning in AI with courses on AI's societal impact, human-AI collaboration, and responsible AI development. Graduates will be prepared to enter into industry at the completion of the program. Courses will combine a theoretical foundation from tenured professors and researchers on campus with applied learning from lecturers and industry practitioners. Students will build a solid foundation in mathematics, statistics, and computing, and pair this foundation with hands-on application to real world problems. They will learn different methods and approaches to solve complex problems using artificial intelligence, data analysis techniques, modern tools, and state-of-the-art technologies while also recognizing potential ethical implications of AI. The learning outcomes for the program are as follows:

1. Understand the theoretical foundations of artificial intelligence, including mathematics and statistics, machine learning, deep learning, and optimization techniques.
2. Develop advanced problem-solving and analytical skills in AI by gaining proficiency in scripting and programming, leveraging machine learning frameworks, and employing high-performance computing platforms to optimize AI solutions.
3. Apply AI technologies to address critical societal challenges and to solve complex problems in disciplines such as healthcare, finance, climate change, and public policy.
4. Design AI systems to work effectively alongside people in healthcare, education, business, and creative industries.
5. Analyze the impact of AI-driven automation in real-world settings by evaluating its effectiveness and efficiency.
6. Describe the societal implications of AI, including issues of safe and trustworthy AI.
7. Evaluate the ethical and legal implications of AI by analyzing issues such as bias, privacy, accountability, and transparency, and developing strategies for responsible AI design, governance, and policy compliance.
8. Explain practical strategies for implementing and scaling AI solutions within enterprise and societal contexts.
9. Communicate AI technologies and applications in a variety of fields to technical and non-technical audiences.

Institutional assessment and documentation of learning outcomes. Assessments will be projects based using publicly available data when possible, to create innovative solutions to societal challenges. Please see Appendix B for information about assessing the program's learning outcomes.

Course requirements. This is a 30-credit non-thesis master's degree, with no thesis option. Students complete 21 credits of required core coursework and select 9 credits of electives from a list of specialized AI courses.

Curriculum		
Course Number	**Course Title**	**Credits**
Core		
MSAI601	Probability and Statistics	3
MSAI602	Principles of Data Science	3
MSAI603	Principles of Machine Learning	3
MSAI605	Computing Systems for AI	3
MSAI606	Human-Centered and Participatory Approaches to AI	3
MSAI630	Safe and Trustworthy AI	3
MSAI631	AI and Society	3
Elective Courses – Select 3 Courses (9 Credits Total)		
MSAI604	Introduction to Optimization for AI	3
MSAI612	Deep Learning for AI	3
MSAI632	Generative AI	3
MSAI633	AI Policy	3
MSAI634	AI in Engineering	3
MSAI635	Reinforcement Learning	3
MSAI636	Explainable and Interpretable AI	3
MSAI640	Computer Vision for AI	3
MSAI641	Natural Language Processing for AI	3
MSAI642	Robotics for AI	3
MSAI650	Cloud Computing for AI	3
MSAI651	Big Data Analytics for AI	3
MSAI660	Probabilistic Graphical Models and Bayesian Learning	3
MSAI661	Causal Inference and AI Decision Making	3
MSAI662	Adversarial Machine Learning and Robustness	3
MSAI663	Graph Neural Networks and Structured Data Learning	3
MSAI664	Meta-Learning and Few-Shot Learning	3
MSAI665	AI for Healthcare and Biomedical Applications	3
MSAI666	AI for Cybersecurity and Threat Detection	3
MSAI667	AI for Finance and Algorithmic Trading	3
MSAI670	Applied Ethics of AI	3

A list of courses, including those that count for the elective category, is included in Appendix C.

General Education. Not applicable for our graduate programs.

Accreditation or Certification Requirements. No accreditation or licensure is required for this program.

Other Institutions or Organizations. The offering unit is not planning to contract with another institution or non-collegiate organization for this program.

Student Support. The Science Academy in the College of Computer, Mathematics and Natural Science will provide administrative coordination for the program, in collaboration with the Office of Extended Studies. Students will be supported through the Science Academy for academic guidance and advising. They will also have access to the Graduate School Counseling and the Counseling Center resources. The Science Academy Program Manager will be the first point of contact for students, while the Office of Extended Studies, which provides administrative services for a host of professional programs, provides student and program services, such as admission support, scheduling, registration, billing and payment, graduation, and appeals. Students will see admission criteria, financial aid resources, costs, and complaint procedures on both the Science Academy website and the Extended Studies program page. For technical aspects of both the in-person and online versions of the program, specific technological competence and equipment will be included in the admission criteria. Learning management information will also be included in these materials.

Marketing and Admissions Information. Students will see admission criteria, financial aid resources, and costs on both the Science Academy website and the Extended Studies program page.

H. Adequacy of Articulation

Not applicable for this graduate program.

I. Adequacy of Faculty Resources

Program faculty. Appendix A contains a list of faculty members who will teach in the program. Instructional resources for the program will comprise current tenure track faculty, professional track faculty, and adjunct instructors. These instructional personnel will come from the Computer Science Department and the Mathematics Department, UMD's Artificial Intelligence Interdisciplinary Institute, and outside the university (e.g., ARLIS, NASA, federal agencies, and industry). Instructors may come from adjacent federal agencies, which will increase the exposure of students to real-world problems as part of the program curriculum.

Faculty training. Faculty teaching in the program will use the university's learning management system along with its extensive electronic resources. They will have access to instructional development opportunities available across the College Park campus, including those offered as part of the Teaching and Learning Transformation Center, many of which are delivered in a virtual environment. Instructors will work with the learning design specialists on campus to incorporate best practices when teaching in the online environment.

J. Adequacy of Library Resources

The University of Maryland Libraries assessment concluded that the Libraries are able to meet, with current resources, the curricular and research needs of the program.

K. Adequacy of Physical Facilities, Infrastructure, and Instructional Resources

No additional physical facilities, infrastructure and instructional equipment is required for this program. Existing facilities (e.g., general purpose classrooms) and resources (e.g., instructional equipment) will be used, and these are demonstrably adequate for the proposed program. For the online components of the coursework, UMD maintains an Enterprise Learning Management System (ELMS). ELMS is a Web-based platform for sharing course content, tracking assignments and grades, and enabling virtual collaboration and interaction. All students and faculty have access to UMD's electronic mailing system.

L. Adequacy of Financial Resources

Tables 1 and 2 contain the details of resources and expenditures.

Table 1 Resources:
The program will be self-supported through tuition revenue.

1. Line 1 shows no reallocated funds since the program is supported by tuition from existing students.
2. Graduate students pay tuition by the credit.
3. Students will complete 24 credits in the first year, and are shown in this chart as full-time students. Part-time students reflect those finishing the program in the second year.
4. No external sources of funding are assumed.
5. No other sources of funding are assumed.

Table 2 Expenditures:

1. Faculty salaries are based on cost per course. We assume an annual increase of 3% in salaries with a corresponding 35.6% benefits rate.
2. Administrative staff represents the program director salary and benefits (.2 FTE).
3. Support staff represents program manager salary and benefits (.33 FTE).

4. Other expenditures include campus administrative fees, travel and recruitment, marketing, hourly grader wages, and director stipend.

M. Adequacy of Program Evaluation

Formal program review is carried out according to the University of Maryland's policy for Periodic Review of Academic Units, which includes a review of the academic programs offered by, and the research and administration of, the academic unit (http://www.president.umd.edu/policies/2014-i-600a.html). Program Review is also monitored following the guidelines of the campus-wide cycle of Learning Outcomes Assessment (https://irpa.umd.edu/Assessment/loa_overview.html). Faculty within the department are reviewed according to the University's Policy on Periodic Evaluation of Faculty Performance (http://www.president.umd.edu/policies/2014-ii-120a.html). Since 2005, the University has used an online course feedback survey instrument for students that standardizes course feedback across campus. The course survey has standard, university-wide questions and allows for supplemental, specialized questions from the academic unit offering the course.

N. Consistency with Minority Student Achievement goals

Recruitment for the Master of Science in Artificial Intelligence will be led by the college's Science Academy, which employs a targeted, inclusive digital strategy focused on UMD alumni, graduating seniors, and working professionals in the Washington, D.C. metropolitan area. The admissions process evaluates not only academic readiness but also diversity in experience, background, and professional goals to ensure a well-rounded and inclusive student body.

To attract a diverse applicant pool, the program will be represented at educational fairs and conferences such as the National Society of Black Engineers Leadership Conference and GEM Grad Labs. Outreach efforts will include advertising through organizations like NSBE, SWE, AWM, and AWC; targeted email campaigns to partner institutions; engagement with UMD student organizations and military veterans; and robust digital marketing including virtual open houses and career panels.

Once enrolled, students benefit from a supportive, inclusive environment fostered by Science Academy staff and faculty. Students are encouraged to participate in diversity and inclusion programs such as TerrapinSTRONG, Cultivating Community Conversations, and the Graduate School's Spring Speaker Series. Faculty bring a variety of academic backgrounds and career experiences, offering students multiple mentorship and career development pathways. The Academy provides academic advising, access to counseling and funding resources, and ongoing student support to promote persistence and timely degree completion.

Retention strategies include hosting seminars like "Women in Engineering, Computing, and STEM," requiring regular academic advising sessions, and implementing an early warning system to identify and assist students facing academic challenges. These initiatives are designed

to ensure that all students—particularly those from underrepresented groups—are supported, empowered, and well-prepared to succeed in the program and beyond.

O. Relationship to Low Productivity Programs Identified by the Commission

N/A

P. Adequacy of Distance Education Programs

The distance-education version of the program will be entirely online. This will allow the program to reach a wider audience, including those in the Washington, DC area whose professional commitments may not allow for regular travel to College Park. The online curriculum will be the same as the in-person curriculum. Learning outcomes, academic rigor and program curricula will be exactly the same for the online program as it is for the on-campus program. The program will go through periodic evaluations, at least every three years, by the Science Academy leadership and academic department chairs. Students will have access to the same services that online students and will be advised by both the Science Academy and the Office of Extended Studies.

Table 1: Resource Table

Resources Categories	Year 1	Year 2	Year 3	Year 4	Year 5
1.Reallocated Funds					
2. Tuition/Fee Revenue (c+g below)	298512	418362	439800	489240	550470
a. #FT Students	9	10	10	11	12
b. Annual Tuition/Fee Rate (based on 24 credits)	33168	34152	35184	36240	37320
c. Annual FT Revenue (a x b)	298512	341520	351840	398640	447840
d. # PT Students	0	9	10	10	11
e. Credit Hour Rate	1382	1423	1466	1510	1555
f. Annual Credit Hours	6	6	6	6	6
g. Total Part Time Revenue (d x e x f)	0	76842	87960	90600	102630
3. Grants, Contracts, & Other External Sources	0	0	0	0	0
4. Other Sources	0	0	0	0	0
TOTAL (Add 1 - 4)	298512	418362	439800	489240	550470

Table 2: Expenditure Table					
Expenditure Categories	Year 1	Year 2	Year 3	Year 4	Year 5
1. Faculty (b+c below)	155685	200696	206716	212918	219305
a. #FTE	2	2	2	2	2
b. Total Salary	119850	154500	159135	163909	168826
c. Total Benefits	35835	46196	47581	49009	50479
2. Admin. Staff (b+c below)	35047	36098	37182	38298	39446
a. #FTE	0.2	0.2	0.2	0.2	0.2
b. Total Salary	25846	26621	27420	28243	29090
c. Total Benefits	9201	9477	9762	10055	10356
3. Total Support Staff (b+c below)	22374	23045	23737	24449	25182
a. #FTE	0.33	0.33	0.33	0.33	0.33
b. Total Salary	16500	16995	17505	18030	18571
c. Total Benefits	5874	6050	6232	6419	6611
4. Graduate Assistants (b+c)	0	0	0	0	0
a. #FTE	0	0	0	0	0
b. Stipend	0	0	0	0	0
c. Tuition Remission	0	0	0	0	0
d. Benefits	0	0	0	0	0
5. Equipment	1500	1545	1591	1639	1688
6. Library	0	0	0	0	0
7. New or Renovated Space	0	0	0	0	0
8. Other Expenses: Operational Expenses	55307	113068	116694	123164	130862
TOTAL (Add 1 - 8)	269913	374452	385920	400468	416483

Appendix A: Faculty Information

The following faculty members are projected to teach in the program. All faculty are full-time unless otherwise indicated.

Name	Highest Degree Earned, Program, and Institution	UMD Title (indicate if part-time)	Courses
Bahar Asgari	PhD, Electrical and Computer Engineering, Georgia Tech	Assistant Professor, Computer Science Affiliate Professor, UMIACS	MSAI 605: Computing Systems for AI
Neda Atanasoski	PhD, Literature and Cultural Studies, University of California San Diego	Professor and Chair, Harriet Tubman Department of Women, Gender and Sexuality Studies Associate Director of Education, AIM	MSAI631: AI and Society
Behtash Babadi	PhD, Engineering Sciences, Harvard	Associate Professor and Associate Chair for Grad Studies, Electrical and Computer Engineering	MSAI 630: Safe and Trustworthy AI
Abhinav Bhatele	PhD, Computer Science, University of Illinois Urbana-Champaign	Associate Professor, Computer Science and UMAICS Affiliate Professor, AIM and AMSC Director, PSSG	MSAI 605: Computing Systems for AI
Margrét Bjarnadóttir	PhD, Operations Research, MIT	Associate Professor of Management Science and Statistics, DO&IT group, Smith School	MSAI631: AI and Society
Jordan Boyd-Graber	PhD, Computer Science, Princeton University	Associate Profession, Computer Science, UMIACS, and iSchool	MSAI 632: Generative AI MSAI 641: Natural Language Processing for AI
Holly Brewer	PhD, American History, UCLA	Burke Chair of American Cultural and Intellectual History, History Director of Undergraduate Studies, History Associate Professor, History	MSAI631: AI and Society
Maria Cameron	Ph.D., Mathematics, UC Berkeley	Prof & Associate Chair, Mathematics. Affiliate Professor with Computer Science.	MSAI 612: Deep Learning for AI
Sandra Cerrai	Ph.D., Mathematics, Scuola Normale Superiore of Pisa	Prof & Assoc Chair, Mathematics	MSAI 601: Probability and Statistics
Yizheng Chen	Ph.D., Computer Science, Georgia Institute of Technology	Assistant Professor, Computer Science	MSAI 630: Safe and Trustworthy AI

Sanghamitra Dutta	PhD, Electrical and Computer Engineering, Carnegie Mellon University	Assistant Professor, Electrical and Computer Engineering	MSAI 630: Safe and Trustworthy AI
Sue Dwyer	Ph.D., Philosophy, MIT	Associate Professor, Philosophy	MSAI 670: Applied Ethics of AI
Sheena Erete	PhD, Technology and Social Behavior, Northwestern (joint degree in computer science and communication)	Associate Professor, College of Information Founder and Director, Community Research and Design Collective	MSAI 606: Human-centered and Participatory Approaches to AI
Soheil Feizi	PhD, EECS, MIT	Associate Professor, Computer Science Director, Reliable AI Lab	MSAI 612: Deep Learning for AI MSAI 660: Probabilistic Graphical Models and Bayesian Learning MSAI 663: Graph Neural Networks and Structured Data Learning
Naomi Feldman	Ph.D., Cognitive Science, Brown University	Professor, Linguistics and UMIACS	MSAI 632: Generative AI
Jonathan Fernandes	Ph.D., Mathematics, University of Maryland	Senior Lecturer, Mathematics	MSAI 601: Probability and Statistics
Nancy Gallagher	Ph.D., International Relations and National Security Studies, University of Illinois Urbana-Champaign	Research Professor, Public Policy Director, CISSM	MSAI 633: AI Policy
Tom Goldstein	PhD, Applied Mathematics, UCLA	Associate Professor, Computer Science Director, Maryland Center for Machine Learning	MSAI 662: Adversarial Machine Learning and Robustness
Elias Gonzales	M.Ed, Curriculum and Instruction, UMD	Lecturer and Curriculum Innovation Lead, Computer Science	MSAI631: AI and Society
Charles Harry	Ph.D., Public Policy, UMD	Director, GoTech Associate Research Professor, Public Policy Operations Director, MaGIC Senior Research Associate, CISSM	MSAI 633: AI Policy; MSAI 633: AI Policy
John Horty	Ph.D., Philosophy, University of Pittsburgh	Distinguished University Professor, Philosophy Affiliate Professor, UMIACS	MSAI 670: Applied Ethics of AI
Furong Huang	PhD, Electrical and Computer Engineering, University of California Irvine	Associate Professor, Computer Science	MSAI 630: Safe and Trustworthy AI MSAI 635: Reinforcement Learning MSAI 664: Meta-Learning and Few-Shot Learning

Heng Huang	PhD, Computer Science, Dartmouth College	Brendan Iribe Endowed Professor, Computer Science, UMIACS, ECE, and CBCB	MSAI 665: AI for Healthcare and Biomedical Applications
Jia-Bin Huang	Ph.D., Electrical and Computer Engineering, University of Illinois Urbana Champaign	Capital One Endowed Associate Professor, Computer Science	MSAI 632: Generative AI
Hal Daume III	Ph.D, Computer Science, University of Southern California	Professor, Computer Science, UMIACS; Director, AIM	MSAI 600: Human-centered and Participatory Approaches to AI MSAI 630: Safe and Trustworthy AI;
Mohit Iyyer	PhD, Computer Science, UMD	Associate Professor, Computer Science	MSAI 641: Natural Language Processing for AI
David Jacobs	PhD, Computer Science, MIT	Professor, Computer Science and UMIACS	MSAI 640: Computer Vision for AI MSAI 632: Generative AI
Leonid Koralov	Ph.D., Mathematics, SUNY at Stony Brook	Prof & Assoc Chair, Mathematics	MSAI 601: Probability and Statistics
Frauke Kreuter	Ph.D., Social Science Research Methods; Survey Methodology, University of Konstanz	Co-Director of the Social Data Science Center Professor, Joint Program in Survey Methodology	MSAI631: AI and Society
Vince Lyzinski	Ph.D., Applied Mathematics & Statistics, Johns Hopkins	Associate Professor	MSAI 651: Big Data Analytics for AI
Kevin McGarry	MA, Political Science and Government, University of California Berkeley	Clinical Professor, Smith School	MSAI 633: AI Policy
Abdirisak Abdullahi Mohamed	PhD, Mathematics, University of Karlsruhe (KIT), Germany	Adjunct Faculty, College of Information Development Expert and AI Ambassador, SAP	MSAI 661: Causal Inference and AI Decision Making
Louiqa Raschid	PhD, Electrical Engineering, University of Florida Gainesville	Deanâ€™s Professor of Information Systems, Smith School Professor, UMIACS and Computer Science	MSAI 667: AI for Finance and Algorithmic Trading
Philip Resnik	PhD, Computer and Information Science, University of Pennsylvania	Professor, Linguistics and UMIACS Affiliate Professor, Computer Science	MSAI 641: Natural Language Processing for AI
Paul Rodrigues	Ph.D., Linguistics, Indiana University Bloomington	Chief Artificial Intelligence Officer, Microsoft: National Security Group	MSAI 651: Big Data Analytics

Rachel Rudinger	PhD, Computer Science, Johns Hopkins University	Assistant Professor, Computer Science, UMIACS, and Linguistics	MSAI 641: Natural Language Processing for AI
Zoltan Safar	Ph.D., ECE, University of Maryland	Director, Telecommunications	DATA/MSML 650: Cloud Computing
Craig Schlenoff	Ph.D. Computer Science, UniversitÃ© de Burgundy	Deputy Associate Director of Laboratory Programs (acting) at NIST Lecturer, MATH (Part-time)	MSAI 631: AI and Society
Katie Shilton	PhD, Information Studies, UCLA	Professor, College of Information	MSAI 606: Human-centered and Participatory Approaches to AI
Abhinav Shrivastava	PhD, Artificial Intelligence, Carnegie Mellon University	Associate Professor, Computer Science and UMIACS	MSAI 640: Computer Vision for AI
Ido Sivan-Sevilla	PhD, Public Policy and Governance, The Hebrew University of Jerusalem	Assistant Professor, College of Information Affiliate Professor, Public Policy Founder, UMD Tech Policy Hub	MSAI 633: AI Policy
Shabnam Tafreshi	Ph.D., Computer Science, George Washington University	Machine Learning Senior Advisor - NLP Researcher at EviCore by Evernorth	MSAI 641- Natural Language Processing
Mohammad Teli	PhD, Computer Science, Colorado State University	Senior Lecturer, Computer Science	MSAI 605: Computing Systems for Machine Learning
Pratap Tokekar	PhD, Computer Science, University of Minnesota	Assistant Professor, Computer Science and UMIACS	MSAI642: Robotics for AI
Mumu Xu	PhD, Mechanical Engineering, California Institute of Technology	Associate Professor, Aerospace Engineering	MSAI642: Robotics for AI
Yun Yang	Ph.D., Statistics, Duke University	Associate Professor, Mathematics	MSAI 603: Principles of Machine Learning for AI
Haizhao Yang	PhD, Mathematics, Stanford University	Associate Professor, Mathematics, Affiliated Associate Professor (UMIACS & CS)	MSAI 603: Principles of Machine Learning for AI
Tianyi Zhou	PhD, Computer Science, University of Washington	Assistant Professor, Computer Science, UMIACS, and AIM	MSAI 632: Generative AI MSAI 612: Deep Learning for AI

Appendix B: Plan for Assessing Learning Outcomes

To maintain the credibility of the MS in Artificial Intelligence, the student learning outcomes will be assessed using a combination of formative and summative assessments throughout the semester and at the completion of each course. These assessments will focus on the direct application of AI technologies to ensure that students can enter the workforce with the skills necessary for success in their future career. For example, many of the elective courses will include final projects, presentations, and assignments where students have to work with real data sets. Students will be expected to process the data, perform tasks and analysis, and make recommendations as if they are entry-level AI professionals. The projects may include hands-on coding assignments, performance evaluations on AI solutions, and other assignments that utilize various machine learning frameworks. These projects may also allow students to explore different applications or areas of AI and can serve as a portfolio for future job searches. Through these assessments, students will demonstrate their problem-solving capabilities, their proficiency in scripting and programming, and their ability to leverage machine learning and computational frameworks and high-performance computing platforms by solving real-world problems in artificial intelligence.

Additionally, each course in the program will have homework assignments, quizzes, and/or other assessments that will be graded with constructive feedback to help assess the student's learning. These smaller, more formative assessments will strengthen students' understanding of the theoretical foundations of AI within each specific course. The formative assessments will be evaluated for their accuracy of the foundational knowledge and reasoning students need for more advanced applications of the materials. Summative and cumulative assessments, such as midterms and final exams or projects, will be used to determine if and to what level the student mastered the core AI concepts and specific learning outcomes for each course.

Many of the program courses will also address concerns of emerging AI, including the ethical, legal, and societal implications of AI. Assessments in these courses may include projects and presentations where students will be expected to effectively communicate AI technologies and applications to both technical and non-technical audiences. Students will also be expected to review case studies and articles in the field and to synthesize the information therein.

Lastly, students will be challenged to complete reflective assessments to apply knowledge and skills in their future professional work. This work will provide students with the skills that will assist them in the job search process and enable them to identify, apply to, and earn positions in artificial intelligence. The assessments will all follow best practices for adult and professional students. By the end of the program, students will have gained the knowledge and experience illustrated through the program-level outcomes, and their performance on the assessments will reflect how well they have achieved these goals.

Appendix C: Course Descriptions

Please note that the MSAI is not yet an existing course prefix and so there are no MSAI courses listed in our Graduate Catalog: https://academiccatalog.umd.edu/.

Core Courses

MSAI601 Probability and Statistics for AI (3 Credits)
The course provides a foundational understanding of concepts in probability theory and statistics tailored for artificial intelligence. The course covers the basic probabilistic concepts such as probability spaces, random variables and vectors, expectation, covariance, correlation, probability distribution functions, and hypothesis testing. etc. Conditional probabilities, the Bayes formula, limit theorems, and properties of jointly distributed random variables are also covered. Students will explore practical applications of probabilistic and statistical methods within the field of artificial intelligence through hands-on exercises and real-world problems.

MSAI602 Principles of Data Science for AI (3 Credits)
This course provides an introduction to the data science pipeline, including the processes of data collection, cleaning unstructured and messy data, data visualization, and statistical analysis. Students will also explore ethical considerations such as fairness, transparency, and bias mitigation. The course will offer students a broad overview of data science and the common tools and systems used in data science problems. Through case studies, students will consider different AI systems through the lens of data science.

MSAI603 Principles of Machine Learning for AI (3 Credits)
This course offers an introduction to the core concepts of machine learning. Students will learn fundamental ML techniques, including supervised and unsupervised learning, neural networks, decision trees, clustering, and PCA. The course will also discuss recent applications of machine learning in AI solutions, such as computer vision, data mining, autonomous navigation, and speech recognition. Students will also gain a basic understanding of ethical AI development and AI for social good.

MSAI605 Computing Systems for AI (3 Credits)
This course will focus on the programming, software and hardware design, and implementation issues of computing systems for machine learning and artificial intelligence applications. Students will explore a variety of topics, including basic Python program structure, functions and modules, basic I/O, object-oriented programming, database access, computer architecture, CPUs and GPUs, memory and I/O systems, virtual memory, and different processing architectures. The course will also cover AI model deployment, edge computing, and scalability challenges in large-scale AI systems.

MSAI606 Human-centered and Participatory Approaches to AI (3 Credits)

This course will cover a broad range of issues in developing human-centered AI with a focus on participatory approaches. We will look at approaches to building AI systems that expand human capabilities, and the interplay between human and AI skills. We will explore how to make use of expertise in those communities impacted by AI systems to design them better. Topics include the fundamentals of HCI and AI, interpretability and explainability in machine learning, human-centered design for AI, adaptive user interfaces, and conversational agents. The course will teach students to design machine learning systems that are well integrated with human capabilities and concerns.

MSAI630 Safe and Trustworthy AI (3 Credits)
Recent advances in AI have created powerful new models, but these models are not easily understood, and it is difficult to guarantee that they will behave in safe and predictable ways. In this course we will examine several key aspects of these models ranging from data privacy, secure code generation, bias and fairness, memorization and copyright infringement, poisoning and adversarial attacks on machine learning systems, reliability, robustness and safety.

MSAI631 AI and Society (3 Credits)
This course is an interdisciplinary exploration of the social impacts and ethical implications of AI. It examines the histories, social values and power dynamics shaping AI technologies, as well as how AI is reshaping culture, politics, and society. Students will develop a sociotechnical understanding of AI related to policy, education, labor, economic systems, and culture. Using approaches from the humanities and social sciences, students will develop frameworks to address ongoing challenges including digital inequality, bias, and surveillance. Students will also learn how AI has and can be used to foster positive social change.

Electives

MSAI604 Introduction to Optimization for AI (3 Credits)
This course introduces fundamental optimization techniques essential for artificial intelligence and machine learning. Students will start with an overview of linear algebra techniques, including vector spaces, linear ransformations, and eigen-decomposition, before moving to techniques in unconstrained and constrained optimization. The course will also explore global search methods, such as simulated annealing, with a focus on AI applications. Students will develop the skills to formulate and solve optimization problems, improving the efficiency and performance of AI models.

MSAI612 Deep Learning for AI (3 Credits)
This course provides a comprehensive introduction to deep learning, a key driver of modern artificial intelligence, with a focus on the main features in deep neural nets and their applications in AI. Students will explore a variety of topics, including backpropagation and its importance, coding tools and their use of parallelization, autoencoders, convolutional neural networks, recurrent and recursive neural networks, and attention-based models. Students will also apply deep learning techniques to real-world problems in computer vision, natural

language processing, and classification/clustering questions, gaining practical experience in building AI models.

MSAI632 Generative AI (3 Credits)
The course will explain the fundamental principles and important techniques in building large language models (LLMs), multi-modal LLMs, and image and video generation models. The class will study Transformer architectures and their use in pretraining, and discuss methods of fine-tuning models including the use of reinforcement learning. The class will study methods of data cleaning, including efficient methods of duplicate detection. And the class will examine computing methods for large scale models that are efficient and that can run in parallel. We will also discuss image and video generation methods, such as the use of stable diffusion.

MSAI633 AI Policy (3 Credits)
How can regulatory strategies promote innovation while safeguarding public interest? This course provides an examination of national and international regulatory and legal frameworks governing artificial intelligence. Students will learn about topics in policy considerations, including copyright, data privacy, bias and discrimination, and the explainability and accountability of AI systems in sectors finance, healthcare, and national security. Students will also learn about contemporary developments in AI governance, including through international AI regulations, national policies, and the advocacy of standards organizations.

MSAI634 AI in Engineering (3 Credits)
This course explores the role of artificial intelligence in engineering disciplines. Students will examine how AI technologies, including machine learning and neural networks, can help solve complex engineering problems and optimize processes. The course will cover a variety of topics and applications such as predictive modeling, automation, intelligent systems design, computer-aided design optimization, and environmental engineering. Students will learn how AI tools can be integrated into engineering workflows across various fields.

MSAI635 Reinforcement Learning (3 Credits)
This course covers both model-free and model-based reinforcement learning (RL), and it explores Markov decision processes, dynamic programming, Q-learning, policy gradient methods, and deep RL. Applications include robotics, game AI, and real-world decision-making systems.

MSAI636 Explainable and Interpretable AI (3 Credits)
As AI becomes more ubiquitous, interpretability is critical. This course explores techniques for understanding deep learning models, such as feature attribution, model distillation, LIME, SHAP, and counterfactual explanations, with an emphasis on ethical and regulatory considerations.

MSAI640 Computer Vision for AI (3 Credits)

This course provides an in-depth introduction to computer vision, a key field in artificial intelligence that enables machines to interpret and analyze visual data. Students will explore fundamental concepts such as image filtering, correlation, object detection, image segmentation, and scene reconstruction. This course will also include discussion on facial recognition, motion tracking, and ethical considerations in vision-based AI. Students will apply computer vision techniques to real-world AI problems.

MSAI641 Natural Language Processing for AI (3 Credits)
This course provides students with the fundamental concepts related to computers generating and processing natural language, including morphological analysis, phrase structure, word sense disambiguation, word embedding models, and advanced deep learning architectures used in NLP. With a focus on the applications of NLP, students will explore topics related to question answering, sentiment analysis, machine translation, text summarization, and chatbot creation.

MSAI642 Robotics for AI (3 Credits)
This course introduces the design and programming of robotic systems with a focus on AI-driven applications. Students will explore core concepts such as kinematics, differential motion, velocity, dynamics, and forces, along with the integration of sensors, actuators, and drive systems. The course covers trajectory planning, motion control, and the implementation of open-loop and closed-loop controllers. Key AI techniques, including state estimation and Kalman filters, will be examined in the context of robotics. Additionally, students will study recent advancements in machine learning for motion planning, grasping, manipulation, and other AI-powered robotic applications.

MSAI650 Cloud Computing for AI (3 Credits)
This course provides an in-depth exploration of state-of-the-art cloud computing technologies and their applications in artificial intelligence. Studenst will explore topics, including telecommunication needs, architectural models, cloud computing platforms and services, and network and storage virtualization technologies. The course will also include a discussion of key concerns in cloud computing such as security, privacy, and trust management. Students will gain practical experience in utilizing cloud-based tools and services to enhance AI workflows, ensuring robust and scalable AI applications.

MSAI651 Big Data Analytics for AI (3 Credits)
This course explores the challenges, tools, and techniques for designing and implementing machine learning algorithms at scale, with a focus on AI applications. Students will learn how to configure and operate distributed computing platforms to efficiently process massive datasets. Key topics include scalable learning techniques, data streaming, data flow analytics, and machine learning on large graphs. The course covers massively parallel computing models such as MapReduce, along with methods to optimize memory, storage, and communication in parallel machine learning algorithms. Additionally, students will gain hands-on experience with

SQL and NoSQL databases, distributed file systems, key-value stores, document databases, graph databases, and large-scale data visualization.

MSAI660 Probabilistic Graphical Models and Bayesian Learning (3 Credits)
This course focuses on the representation and inference of uncertainty in AI using probabilistic graphical models, such as Bayesian networks and Markov random fields. It also covers variational inference, sampling methods, and applications in decision-making and reasoning.

MSAI661 Causal Inference and AI Decision Making (3 Credits)
Unlike traditional correlation-based learning, causal AI seeks to understand cause-and-effect relationships. This course explores Pearl's causal inference framework, causal discovery, and interventions for AI decision systems.

MSAI662 Adversarial Machine Learning and Robustness (3 Credits)
This course explores vulnerabilities in AI models, covering adversarial attacks, defenses, and the study of robustness in deep learning models against perturbations.

MSAI663 Graph Neural Networks and Structured Data Learning (3 Credits)
Graph neural networks (GNNs) enable AI models to work with non-Euclidean structured data. This course covers graph representation learning, message passing, and applications in social networks, bioinformatics, and knowledge graphs.

MSAI664 Meta-Learning and Few-Shot Learning (3 Credits)
This course explores learning-to-learn approaches, including model-agnostic meta-learning (MAML), few-shot classification, and applications in fast model adaptation.

MSAI665 AI for Healthcare and Biomedical Applications (3 Credits)
Applications of AI in medicine, including medical imaging, genomics, drug discovery, and personalized healthcare solutions.

MSAI666 AI for Cybersecurity and Threat Detection (3 Credits)
Examines AI-driven cybersecurity measures, including anomaly detection, malware analysis, adversarial robustness, and security threats in machine learning systems.

MSAI667 AI for Finance and Algorithmic Trading (3 Credits)
Covers AI applications in finance, including reinforcement learning for trading strategies, risk modeling, fraud detection, and AI-driven market forecasting.

MSAI670 Applied Ethics of AI (3 Credits)
As artificial intelligence tools are increasingly used in high stakes scenarios throughout our lives, it is increasingly important to understand the ethical considerations behind their use. This course will introduce students to applied ethics, a major subfield of contemporary Philosophy, as a way of making sense of how AI tools can be built and used ethically. The course will

consider a broad range of topics, ranging from human-robot interaction to algorithmic bias, from autonomous weapon systems to algorithmic accountability and opacity. The tools students learn in this course will be broadly applicable to all forms of artificial intelligence, both present and future.