Katz

Katz School
of Science and Health

M.S. in Artificial Intelligence

Course Descriptions

AIM 5001 Data Acquisition and Management

Data Acquisition and Management focuses on the data structures, data design patterns, algorithms, methods, and best practices for the pre-modeling phases of data science workflows, including problem formulation, gather, analyze, explore, model, and communicate, analytics programming focuses on the gather, analyze, and explore workflow steps. This comprises the "data wrangling" work which is where most data scientists spend the majority of their time. Because data science is iterative, this preparatory work informs the modeling phase. Often, the creation and validation of new models requires going back for additional data, different data transformations, and exploration of data distributions. In short, every effective data scientist needs to master analytics programming. Course topics include reading from or writing to databases, text files, and the web; shaping data into "tidy" data frames, exploratory data analysis, data imputations, feature engineering, and feature scaling.

Industry Applications: Data acquisition skills enable reliable analytics and modeling pipelines. Companies such as Amazon and Uber rely on data engineers to transform raw transactional and behavioral data into structured datasets for machine learning. Healthcare organizations including UnitedHealth Group integrate clinical, claims, and demographic data to support analytics and compliance reporting. Consumer goods firms like Procter & Gamble apply feature engineering and exploratory analysis to demand forecasting and marketing optimization.

AIM 5002 Computational Statistics and Probability

Arguably, most of data science is statistical learning, which requires strong foundational knowledge in probability and statistics. And applying computational methods such as direct simulation, shuffling, bootstrapping, and cross-validation to statistical problems is often more intuitive, and intuitive and can provide solutions where analytical methods would prove computationally intractable. This course introduces students to the statistical analysis of data using modern computational methods and software. Probability, descriptive statistics, inferential statistics and computation methods such as simulations sample distributions, shuffling, bootstrapping, and cross-validation will be covered.

Industry Applications: Computational statistical techniques are fundamental to experimentation, forecasting and risk assessment. Technology firms such as Meta and Netflix use bootstrapping and cross-validation to evaluate product changes and personalization algorithms. Financial institutions including JPMorgan Chase and BlackRock apply simulation-based probability models to quantify market risk and portfolio uncertainty. Public health organizations such as Centers for Disease Control and Prevention rely on probabilistic modeling and simulation to evaluate disease spread and intervention strategies.

AIM 5003 Numerical Methods

Algorithms in machine learning and neural networks are built upon a strong foundation of linear algebra. For example, modern recommendation systems may have sparse matrices with millions of users and millions of items; matrix factorization methods make the underlying calculations tractable say this course builds a foundation of linear algebra concepts such as matrices, determinants, vectors and eigen values. Then it deepens it into data science applications around network analysis and logistic algorithms. In addition, some multi-variate calculus and graph theory topics are covered.

Industry Applications: Numerical methods make large-scale analytics computationally feasible. Financial data providers like Bloomberg rely on matrix operations and optimization for real-time analytics. Industrial firms such as Siemens use numerical optimization and network analysis for systems modeling. Recommendation systems across technology and retail sectors depend on matrix factorization techniques to handle massive user-item datasets.

AIM 5004 Predictive Models

Prerequisites: Data Acquisition and Management; Computational Statistics and Probability

Predictive modeling answers the question, “What will happen next?” Linear regression and logistic regression are foundational predictive modeling methods, used to predict continuous and categorical output respectively. The main topics covered in this course include simple and multiple linear regression, variable selection and shrinkage methods, binary logistic regression, count regression, weighted least squares, robust regression, generalized least squares, multinomial logistic regression, generalized linear models, panel regression, and nonparametric regression.

Industry Applications: Predictive modeling expertise is widely used in analytics, forecasting, and risk assessment roles. Banks such as Wells Fargo apply regression models to credit risk and regulatory analysis. Healthcare systems like Kaiser Permanente use predictive models to forecast patient outcomes and resource utilization. Marketing teams across industries rely on regression for pricing strategy and campaign effectiveness analysis.

AIM 5005 Machine Learning

Prerequisites: Data Acquisition and Management; Computational Statistics and Probability

In classical programming, answers are obtained from rules and data. In machine learning, rules are obtained from data and answers. The widespread availability and sharing of data, and improvements in computing capacity, processing methods, and algorithms have given machine learning the power to deliver game-changing systems and technologies to organizations that compete on predictive, prescriptive, and/or autonomous analytics. In this course, we’ll look at methods for using, tuning, and comparing machine learning algorithms, based on measures of performance, accuracy, and explainability. We'll also look at recent advances and trends in automated machine learning.

Industry Applications: Machine learning (ML) techniques drive predictive and prescriptive analytics across industries by leveraging statistical patterns in data to automate complex decisions. Financial institutions use ML for credit scoring, underwriting, and fraud detection; retailers and logistics firms apply forecasting and optimization for demand and inventory management; and healthcare organizations employ clustering and survival analysis for patient risk and resource planning. Emerging applications include generative and agentic AI systems such as Claude Code, ChatGPT, and Gemini, embodied AI in robots, drones, humanoids, and autonomous vehicles, and vision-language models and coding agents that integrate perception, reasoning, and adaptive decision-making.

AIM 5006 Artificial Intelligence

Prerequisites: Data Acquisition and Management; Computational Statistics and Probability

Artificial Intelligence (AI) is an interdisciplinary field, integrating knowledge and methods from computer science, mathematics, philosophy, psychology, economics, neuroscience, linguistics, and biology. Intelligent agents mimic cognitive functions to implement intelligent behaviors such as perception, reasoning, communication, and acting in symbolic and computational models. AI is used in a wide range of narrow applications, from medical diagnosis to speech recognition to bot control.

The autonomous single, multiple, and adversarial agents that students build in this course will support fully observable and partially observable decisions in both deterministic and stochastic environments. Topics covered include search, constraint satisfaction, Markov decision processes, planning, knowledge representation, reasoning under uncertainty, graphical models, and reinforcement learning. The techniques and technologies mastered here will provide the foundational knowledge for the ongoing study and application of AI in other applications across practice areas.

Industry Applications: Professional practice in artificial intelligence focuses on the design and deployment of AI systems. Applications span large-scale optimization at companies like Google and Microsoft, generative and reasoning models such as ChatGPT, Claude Code, and Gemini, and embodied AI in robots, drones, humanoids, and autonomous vehicles. In aerospace, defense, and healthcare, organizations like Boeing, Lockheed Martin, and Mayo Clinic apply planning, constraint satisfaction, and probabilistic reasoning for mission planning and clinical decision support. Professionals collaborate closely with domain experts, systems engineers, and risk or compliance teams to ensure reliable and ethical AI deployment. Advanced applications include text and image generators, and Agentic AI systems that are goal-oriented and proactive—capable of planning and executing multi-step tasks autonomously in domains like software development agents, research copilots, autonomous vehicles, robotic systems, and enterprise workflow automation.

AIM 5007 Neural Networks and Deep Learning

Prerequisites: Machine Learning

Data scientists have been able to leverage better algorithms on faster hardware optimized with graphical processing units to deliver improved performance and accuracy in whole classes of applications that had been previously commercially unviable. The biggest beneficiaries are applications that require unstructured data, such as audio and/or video processing. Deep neural networks have also provided gains for other complex applications, from recommendation systems to natural language processing. This course builds on the concepts in machine learning to train multi-layered neural networks.

Industry Applications: Deep learning powers high-dimensional pattern recognition and autonomous perception across sectors by leveraging GPU-accelerated architectures to derive insights from complex data. In healthcare, convolutional neural networks automate anomaly detection in MRI, CT, and X-ray imaging. Transformers and large language models such as ChatGPT, Claude Code, and Gemini enable advanced NLP, code generation, multimodal reasoning, and conversational systems. Emerging embodied and agentic AI applications include vision-language models, coding agents, and autonomous robots, drones, humanoids, and vehicles. Deep learning also drives retrieval-augmented generation (RAG) with fine-tuned code assistants, next-generation search, marketing, and creative content production using generative and diffusion models.

AIM 5008 Capstone for Artificial Intelligence and Machine Learning

The Capstone integrates prior coursework, research, colloquia, and professional experience, and provides the opportunity to synthesize theory with practice in an applied project, thesis, approved internship, or equivalent activity. Examples include developing an AI application or methodology, publishing a research paper at a peer-reviewed conference, or creating a startup company through YU's Innovation lab-though students may propose other related projects based on their interests. The Capstone will include four components: a brief proposal and project schedule; the main deliverable (e.g. thesis, conference paper, working system with analysis/code/data); and a final presentation to the student and faculty body. Faculty will provide students with mentorship and feedback at each stage of the work.

Industry Applications: Capstone projects replicate professional practice in consulting, research, and product development, producing deliverables comparable to those in corporate R&D labs, healthcare analytics teams, financial institutions, and startups. Advanced applications include large language and video models, Claude Code, ChatGPT, Gemini, embodied and agentic AI systems, vision-language models, coding agents, RAG architectures, and next-generation autonomous and generative technologies.

AIM 5013 Advanced Data Engineering

Prerequisites: Data Acquisition and Management

As both the volume and the velocity of data increase exponentially, problems in both commerce and research become increasingly reliant on environments with distributed data storage and data processing capabilities. This requires rethinking how our entire approach to distributed environments. This course provides students with the concepts, data structures, and algorithms needed to implement data science applications in distributed computing environments. In this course, we will implement and apply distributed algorithms, data frames, and streaming. You will also learn how to choose appropriate distributed algorithms based on the characteristics of the problem and the system.

Industry Applications: Distributed data engineering supports real-time and large-scale analytics. Organizations such as Databricks and Walmart rely on distributed processing for supply chain analytics and personalization. Telecommunications firms like AT&T use streaming systems for network monitoring and customer analytics.

AIM 5009 Bayesian Methods

Prerequisites: Data Acquisition and Management; Computational Statistics and Probability

Bayesian inference provides powerful tools to model random variables. While Bayesian methods often yield the most accurate theoretical results, historically analytical complexity made it challenging to apply Bayesian methods against less trivial problems. Now, the confluence of more powerful computing resources and improved computational algorithms make Bayesian methods the best choice for tackling some of the most complex data science problems. Bayesian analysis is increasingly important in academic research, and research and is the preferred standard statistical analysis tool in data science practice. In this course, we'll build from Bayes' probability foundations to first applying Bayesian methods to infer binomial probabilities, then hierarchical models, and finally generalized linear models. We'll provide comparisons between frequentist approaches and Bayesian approaches. We'll build basic algorithms from scratch, as well as using high-performance Markov Chain Monte Carlo (MCMC) methods.

Industry Applications: Bayesian approaches are commonly applied within research, experimentation, and risk-focused roles. Companies such as Stripe use Bayesian inference for experimentation and forecasting, while pharmaceutical organizations like Pfizer apply hierarchical Bayesian models in clinical research. Government and policy institutions also employ Bayesian methods for uncertainty-aware decision analysis. These roles are often situated within analytics research or decision science teams.

Complex Systems: Financial Time Series Analysis

Prerequisites: Data Acquisition and Management; Computational Statistics and Probability

Provides a rigorous introduction to modeling and prediction of financial time series. The goals are to learn basic characteristics of financial data, understand the application of financial econometric models, and gain experience in analyzing financial time series. We begin with the basic concepts of linear time series analysis such as stationarity and autocorrelation function, introduce regression models with time series errors, seasonality, unit-root non-stationarity, and long-memory processes. We provide methods of analysis in the presence of conditional heteroscedasticity and serial correlations of asset returns. The course introduces heavy-tailed distributions, and their application to financial risk management. In particular, we discuss modern valuations of credit risk. We introduce multivariate time series analysis and apply the concept of co-integration to investigate arbitrage opportunity in pairs trading. The course places great emphasis on empirical data analysis. We use real examples and exercises in R and Python are included. The course aims to broaden the horizons of students in applied mathematics and to provide conceptual background to students who are interested in a career in financial industry.

Industry Applications: Financial time series expertise supports quantitative roles in trading, asset management, and risk control. Financial time series analysis is foundational in quantitative finance. Asset managers such as BlackRock and investment banks like Goldman Sachs use these models to manage portfolio risk and identify trading opportunities. Regulatory institutions apply time series methods to monitor systemic financial stability. Work emphasizes empirical modeling, validation, and real-time data analysis.

AIM 5012 Data Visualization

Data scientists depend on data visualizations for their own exploratory analysis to support their modeling decisions--the mind can process visual information must faster than numbers. Data visualization is also important to inform—and often to persuade—other people about what can be inferred from the data. These explanatory visualizations often require higher production values, interactivity, and guiding text. In this course, students apply the concepts, methods, and best practices of data visualization to create reproducible, code-based exploratory and explanatory data visualizations.

Industry Applications: Data visualization skills are integrated into analytics, reporting, and communication roles across sectors. Organizations such as Tableau and McKinsey & Company rely on visualization to support executive decision-making and client communication. In journalism and public institutions, including The New York Times, visualization professionals translate complex data into accessible narratives. These roles are typically embedded within analytics or communications teams.

AIM 5011 Natural Language Processing

Prerequisites: Machine Learning; Neural Networks and Deep Learning

Natural Language Processing lives at the intersection of machine learning, artificial intelligence, and linguistics. It is the key to unlocking vast amounts of human-generated, unstructured data. The increased availability of corpuses of text data, the wide availability of cheap distributed systems, improvements in neural network algorithms, and increased access to graphical processing units (GPUs) have improved the performance and accuracy of entire families of once computationally intractable problems, making these commercially feasible. This course explores a series of text and voice processing use cases, including sentiment analysis and topic modeling. It is the key to unlocking vast amounts of human-generated, unstructured data. Along the way, students gain experience working with supervised and unsupervised methods using both machine learning algorithms and deep neural networks.

Industry Applications: Natural language processing capabilities are deployed in search, document analysis, conversational systems, and domain-specific applications such as sentiment analysis, semantic search, summarization, information extraction, and question answering. Technology companies like IBM and Google integrate NLP and modern large language models (LLMs) such as GPT-style transformers, Claude-like code assistants, ChatGPT-class models, and Gemini-style multimodal systems into enterprise and consumer products for translation, code generation, and content creation. In non-technology domains, including law firms, healthcare providers, and financial institutions, NLP supports document review, compliance monitoring, clinical text analysis, and knowledge access via retrieval-augmented generation (RAG) and next-generation search. Emerging embodied and agentic AI applications leverage vision-language models, coding agents, and autonomous decision systems to connect language with perception and action across robots, drones, humanoids, and autonomous vehicles.

AIM 5010 AI Product Studio

What is needed to convert a promising idea or research into a viable product or service? Bringing successful products to market is an experiential discipline that requires hands-on practice working through iterative workflow of a customer-driven product development lifecycle. In this course, students work with mentors to design products, develop customers, and create product development roadmaps. Students create and communicate hypotheses around customers, cost and revenue streams, activities, and value propositions. Agile project management, data-driven product design and customer feedback, and technical constraint identification are all covered.

Industry Applications: Product-focused work emphasizes the translation of technical capability into market-ready solutions. Consulting firms such as IDEO and corporate innovation labs structure teams around iterative discovery, customer validation, and roadmap development. In both startup and enterprise contexts, roles bridge technical, design, and business functions, reflecting industry-standard product development practices.

AIM 5012 Special Topics in Artificial Intelligence and Machine Learning

This course provides the opportunity to take on emerging theory, phenomena, and technologies in the field of artificial intelligence, machine learning, data science, and big data generally. This will be an advanced class, whether seminar style or project based.

Industry Applications: Advanced topics are explored in research labs, innovation groups, and advanced analytics units, where teams at organizations such as Meta AI and industrial research centers pursue exploratory projects that shape future products. Work emphasizes rapid experimentation, technical evaluation, and dissemination of findings. Applications span LLMs like ChatGPT, Gemini, and Claude; embodied and agentic AI systems; vision-language models; coding agents; and RAG-based domain-expert assistants that integrate curated, up-to-date knowledge for more accurate, transparent, and maintainable real-world solutions.

AIM 5999 Independent Study in Artificial Intelligence and Machine Learning

The course provides flexibility to learn more about a topic of interest outside of the formal course setting. The subject should be chosen in consultation with a faculty advisor who acts as the student's supervisor, and with the permission of the program director. The student is required to submit a course contract describing the course of study and its specific learning objectives.

Industry Applications: Independent study aligns with apprenticeship-style models found in research institutions, consulting environments, and specialized industry roles, with mentorship that emphasizes depth, autonomy, and alignment with organizational or domain-specific objectives. Advanced agentic AI systems use large language models and related architectures to plan, decide, and act toward specified goals across applications such as virtual assistants and chatbots, code assistants and autonomous coding agents, research and knowledge copilots with RAG, creative content generation, robotics and autonomous vehicles, drones and humanoid robots, predictive maintenance and operations optimization, clinical decision support, and multi-agent systems that coordinate complex workflows.

Updated: January 2026