Katz
Katz School
of Science and Health

M.S. in Data Analytics and Visualization

Course Descriptions

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 Application: The iterative, customer-driven product development approach in this course mirrors how leading technology companies actually build AI products, especially those powered by large language models and other generative systems. Students apply hypothesis-driven development, rapid experimentation, and customer discovery to design and validate features before full-scale deployment, similar to practices at Google, Microsoft, and venture-backed startups. Advanced applications include LLM tools and generative AI as core building blocks in chat-based assistants, domain-specific copilots, RAG-style knowledge bots, AI product copilots for customer support and sales, workflow and automation agents, and AI-powered analytics and decision-support tools.

DAV 5000 Business Modeling and Data Analysis

While data analysts need to be competent with a variety of tools, they will most often work with stakeholders who only use spreadsheets. Therefore, deep and broad skills working with spreadsheets and fluency in moving data between spreadsheets, business intelligence applications, and relational databases are critical for data analysts to be effective and credible. In this project-based course, students will build and apply key spreadsheet skills in the service of organizational data management, modeling, and analysis.

Industry Application: Despite the proliferation of specialized analytics tools, spreadsheet fluency remains the most universally required skill for data professionals because spreadsheets

are how business stakeholders actually work. Investment banking analysts at firms like Goldman Sachs and Morgan Stanley build complex financial models entirely in Excel—leveraging advanced functions, scenario analysis, and sensitivity tables that inform billion-dollar decisions. At consulting firms like McKinsey and Bain, the ability to rapidly prototype analyses in Excel and then communicate findings to executives who are themselves Excel users is what separates effective consultants from those who struggle to gain client trust. The critical skill isn't just spreadsheet manipulation—it's the ability to translate between spreadsheets, databases, and BI tools, ensuring data integrity throughout the workflow. This versatility is why analysts who can seamlessly move data between Excel, SQL databases, and visualization platforms are consistently more effective than those who specialize in only one environment.

DAV 5100 Structured Data Management

Organizations require reports and analyses that are both accurate and useful. This course emphasizes the skills that database developers rely on to 1) translate organizational requirements into information architectures, 2) create operational and reporting databases, 3) manage data sources, 4) perform data integration into reporting databases, and 5) create ad-hoc analytics reports using business intelligence reporting tools. Students will gain both essential theory and hands-on practice, enabling them to build the database systems, supporting data workflows, and reporting architectures required to produce accurate and useful information in support of organizational decisions.

Industry Application: The ability to translate business requirements into working database systems is where many analytics initiatives succeed or fail, making these skills exceptionally valuable. Organizations don't need reports—they need *accurate* reports built on properly designed data architecture. At retail companies like Walmart and Target, database developers build the information systems that reconcile point-of-sale data across thousands of locations, ensuring inventory and sales reports reflect reality rather than garbage data. Financial institutions face regulatory requirements where database accuracy isn't optional—incorrect regulatory reporting at firms like Citigroup or Charles Schwab can result in millions in fines. The critical skill this course develops is the end-to-end capability: understanding organizational needs, designing appropriate data models, implementing ETL workflows, and creating BI reporting layers. Healthcare organizations particularly struggle with data integration across clinical, billing, and operational systems—professionals who can actually build these integrated reporting databases rather than just query them are consistently in short supply.

DAV 5200 Visual Design and Storytelling

Analysts must present their data in effective and compelling visualizations. This course combines the best heuristics for data presentation with hands-on experience in creating spreadsheet charts and data visualizations from a variety of source data. Students will learn how to combine text and visualizations to craft stories that promote deeper engagement with data analyses and conclusions.

Industry Application: The difference between presenting data and telling a compelling data story often determines which projects get funded and which recommendations get implemented. Effective visualization isn't about making pretty charts—it's about combining visual and textual elements to drive specific decisions and actions. At media organizations like The New York

Times and Reuters, data visualization specialists make complex topics accessible to millions of readers, influencing public discourse on policy and social issues. In consulting, the ability to create executive-ready visualizations that clearly communicate insights is what justifies premium fees—McKinsey and BCG consultants are essentially selling clarity of insight, which requires mastery of visual communication. The storytelling component addresses a critical reality: data rarely speaks for itself. Marketing teams proving campaign ROI, analysts recommending budget allocations, and researchers seeking funding all face the same challenge—they must craft narratives that help stakeholders understand not just what the data shows, but why it matters and what to do about it.

DAV 5300 Computational Math and Statistics

Deeper math literacy and computational thinking are essential for deeper data literacy. Probability, statistics, and mathematics—especially fundamental linear algebra—are critical to the success of data analysts as they implement increasingly complex solutions. This course gives the non-mathematician practice using mathematical and statistical computational methods in the service of data analysis.

Industry Application: Mathematical literacy is what enables analysts to move from executing predefined analyses to designing novel solutions and understanding why methods work or fail. At quantitative finance firms like Two Sigma and Citadel, every trading strategy rests on probability theory and linear algebra—analysts who lack this foundation simply cannot understand or improve the models they work with. The recent explosion of machine learning has created a problematic divide: many analysts can run algorithms but cannot diagnose when they're inappropriate, overfitting, or producing nonsensical results. Organizations increasingly recognize this gap—tech companies like Google and Netflix specifically seek candidates who understand the mathematical foundations of their recommendation and optimization systems rather than treating them as black boxes. This computational approach to mathematics is particularly valuable because it builds intuition through practice: students learn probability by running simulations, understand linear algebra by manipulating matrices, and grasp statistical concepts by implementing them computationally rather than through pure theory.

DAV 5400 Analytics Programming

Code-based solutions can be richer, more accurate, and more flexible than those that rely on off-the- shelf software and analytic packages. This course teaches the programming skills that data analysts need to prepare structured and unstructured data for downstream analysis. Students will learn to use high-level programming languages to create rich data analysis workflows.

Industry Application: Programming skills fundamentally change what is possible in data analysis—enabling automation, flexibility, and sophistication that point-and-click tools cannot achieve. Code-based workflows at companies like Amazon and Google are version-controlled, peer-reviewed, and deployed across thousands of datasets, whereas manual spreadsheet analyses remain error-prone and hard to audit. Financial institutions such as JPMorgan and Goldman Sachs require programming proficiency to ensure that risk calculations are reproducible and compliant, while e-commerce companies like Shopify and Etsy rely on

programmatic data preparation to handle millions of daily transactions. As generative models and agentic AI systems become embedded in analytics workflows, programming also becomes the way analysts design, control, and audit AI-driven pipelines—wrapping LLM calls in reproducible code, building agents that orchestrate multi-step data tasks, and enforcing guardrails around data access, privacy, and model behavior.

DAV 6500 Capstone

In this course, students will integrate the skills developed in the previous classes into a comprehensive body of knowledge and will provide tangible evidence of analytic and visualization competencies. The capstone will include four components: 1) a brief proposal and project schedule; 2) the main project deliverable; 3) a final presentation; and 4) a reflection on the student's data analytics and visualization skills and competencies.

Industry Application: The capstone creates tangible proof of end-to-end analytical and AI product capability—the evidence hiring managers want to see beyond credentials and coursework descriptions. Capstone projects that incorporate LLM-powered analysis, RAG-based knowledge retrieval, code assistants (e.g., Cursor, Claude Code), or autonomous agents for data workflows stand out as cutting-edge demonstrations of production-ready skills, showing not only that students can build models but that they can design, implement, and govern AI-augmented analytics in realistic settings. When technology companies like Amazon, Google, Spotify, and Airbnb evaluate candidates, they assess whether someone can scope ambiguous problems, integrate LLM and agentic components into data and software pipelines, execute sophisticated analyses, and communicate findings to non-technical stakeholders—precisely the integrated skills demonstrated through capstone work. The portfolio component aligns with hiring realities in tech firms and AI-focused consultancies, where managers explicitly request work samples that show practical capability with modern AI tools rather than only academic performance. The presentation requirement further develops the ability to defend analytical and design choices, explain the behavior and limitations of LLM- and agent-driven systems, and respond to stakeholder questions—skills that determine whether AI-driven recommendations and products are trusted, adopted, and successfully deployed.

DAV 6000 Talent Analytics

To survive and prosper, organizations must make good use of data and analytics to improve their workforce-related processes. This is particularly critical at times of low employee engagement and high turnover. In this course, students will learn the key processes, measures, and tools that enable data- driven workforce analysis to deliver competitive organizational advantage.

Industry Application: Workforce analytics addresses organizations' largest controllable expense and most valuable asset—people—making these skills strategically important despite being relatively specialized. The impact potential is substantial: reducing turnover by even a few percentage points at large organizations like Walmart or Amazon saves tens of millions annually, while improving hiring quality compounds over years as better employees make better decisions. Technology companies including Google and Microsoft maintain dedicated people analytics teams precisely because data-driven workforce decisions provide competitive advantage—better hiring, compensation, and retention strategies

directly impact innovation capacity. What makes talent analytics particularly valuable is the combination of statistical rigor with sensitivity to organizational politics and employee privacy—analysts must balance data-driven insights with ethical considerations and change management. HR technology companies like Workday and ADP increasingly need professionals who can both build workforce intelligence platforms and help clients interpret sensitive employee data responsibly.

DAV 6050 Data-Driven Organizations

The best data analysis projects are implemented in the context of an organization's business model, culture, key strategic initiatives, and processes. Data analysts who understand these contexts are more likely to see their efforts lead to improved organizational processes and decision-making. This course examines three important organizational-level analytical frameworks and emphasizes using data, analysis, and experimentation within each. Students will also study centralized data warehouses.

Industry Application: Understanding organizational context separates analysts who deliver insights that get ignored from those who drive actual change and decision-making. The harsh reality is that most analytics projects fail not because of poor analysis but because analysts don't understand the business model, strategic priorities, or organizational culture. At consulting firms like McKinsey and Bain, the ability to quickly diagnose a client's strategic context and design analytics initiatives that align with existing processes is what makes consultants valuable—technical skills alone are insufficient. Technology companies like Salesforce face a similar challenge when helping enterprise customers adopt analytics: the software works, but customers struggle with organizational change management and data governance. This course's focus on analytical frameworks and centralized data warehouses addresses the infrastructure reality—without proper data architecture and governance, even sophisticated analyses lack credibility. Private equity firms increasingly value professionals who can transform portfolio companies into data-driven organizations, recognizing that operational improvement through analytics directly impacts investment returns.

DAV 6100 Information Architectures

Organizations combine data from many different sources, including spreadsheets, databases, and data warehouses. As the volume, variety, and velocity of data increases, more enterprise data is stored in cloud-based distributed data stores. In this course, students will learn to design, populate, and report on these enterprise data architectures.

Industry Application: The shift from traditional databases to cloud-based distributed data systems represents one of the most significant infrastructure transitions in decades, creating urgent demand for professionals who understand both worlds. Organizations aren't starting fresh—they're migrating legacy systems while maintaining business operations, requiring architects who can design hybrid solutions that bridge traditional and modern approaches. Cloud platforms like Amazon Web Services and Google Cloud actively seek solutions architects who can help enterprise clients navigate this complexity. What makes these skills particularly valuable is the combination of theoretical knowledge with practical implementation: companies like Uber and Netflix have built cloud-native data architectures processing billions of events, but most organizations are still figuring out how to migrate decades of data from on-

premise systems. Financial institutions face especially challenging migrations—firms like Goldman Sachs and Capital One must modernize infrastructure while maintaining regulatory compliance and data security. The course emphasis on understanding both traditional and distributed systems prepares graduates for this transition phase rather than assuming a pure cloud-native world.

DAV 6150 Data Science

Frequently, analysts use data to describe the current state of an organization. Data science extends the analyst's reach into the future. Data science has been almost exclusively the domain of people who have STEM degrees and especially those with a quantitative background. Recent fast-paced tool development and abstraction now allow motivated data analysts to perform useful and rigorous predictive analyses using high-level languages and their rich scientific ecosystems. This course will cover classification, regression, and clustering methods, and students will apply these methods in designing, modeling, and building model applications that use natural language processing and recommender systems.

Industry Application: The democratization of data science—making predictive analytics accessible to analysts without traditional STEM backgrounds—represents a significant career opportunity for motivated professionals. Recent advances in high-level languages and abstraction layers mean that rigorous predictive modeling no longer requires PhD-level mathematics, expanding who can perform sophisticated analyses. E-commerce companies like Amazon and Netflix demonstrate this shift: their recommendation engines generate billions in revenue using classification and clustering methods that are now accessible through modern tools. The practical applications are immediate—fraud detection at companies like PayPal and Stripe, churn prediction at subscription businesses, and demand forecasting across retail organizations. What distinguishes successful data science practitioners isn't just technical proficiency but the ability to frame business problems as prediction tasks, validate models appropriately, and communicate uncertainty to stakeholders. Healthcare organizations particularly need analysts who can apply predictive methods to clinical and claims data for care management, where the impact isn't just revenue but patient outcomes.

MAN 5580 Project Management

Most data analysis and visualization work is project-based, and successful data analysts are effective at managing projects and collaborating as members of project teams. This course teaches project management using several tools from the leading methodologies for managing software projects. The most effective project managers will combine methods to create a "right-sized" methodology appropriate to the organizational culture and project team members' background and experience.

Industry Application: Project management capabilities are often the determining factor in whether analysts advance into leadership roles or remain individual contributors. Most analytics work happens in project contexts—data warehouse implementations, reporting system upgrades, dashboard rollouts—and these initiatives fail not due to technical problems but due to poor coordination, unclear requirements, and missed deadlines. Organizations like Google and Microsoft explicitly promote data analysts to project management roles when they demonstrate the ability to coordinate cross-functional teams and deliver results on

schedule. In consulting at firms like Accenture and Deloitte, project management skills directly impact billable utilization and client satisfaction, making them essential for career progression. The challenge this course addresses is "right-sizing" methodology—using heavyweight processes for small projects wastes time, while insufficient structure on complex initiatives leads to chaos. Learning to adapt methodologies to team experience and organizational culture is what makes project managers effective rather than bureaucratic.

DAV 6200 Data Product Design

Successful entrepreneurs and consultants create value. Data analysts who can work alongside or act as value architects create more organizational value, more quickly. Today, this means using data, analysis, and experimentation to better understand customer goals and preferences. In this course, students learn analytical frameworks for using data in the service of customer insight, customer development, value proposition refinement, and product development.

Industry Application: The shift toward product-led growth strategies means organizations increasingly need analysts who think like entrepreneurs—focused on customer value creation rather than just data processing. This mindset distinguishes analysts who become strategic partners from those who remain order-takers executing predefined analyses. Product analytics teams at companies like Slack and Zoom use rigorous experimentation and customer data to optimize every aspect of user experience, from onboarding flows to feature adoption, directly impacting retention and revenue. The analytical frameworks for customer development and value proposition testing enable analysts to help organizations avoid costly mistakes—building features customers don't want or solving problems that don't actually matter. Startups backed by firms like Sequoia and Andreessen Horowitz particularly value this capability because early-stage companies live or die based on finding product-market fit through rapid iteration. Even within established companies, the ability to use data for customer insight rather than just reporting makes analysts indispensable to innovation initiatives and new product development.

DAV 6300 Special Topics

This course provides the opportunity to offer boutique short-term courses on emerging phenomena, policies, processes, technologies, and techniques in data analysis and visualization. The expectation is that this will be an advanced class that requires an appropriate student project and deliverable in line with the number of credits awarded for the course.

Industry Application: Emerging analytics techniques evolve faster than standard curricula can adapt, making specialized exposure to cutting-edge methods a meaningful competitive advantage. The gap between what's taught in established courses and what's being deployed in industry creates opportunities for motivated students who proactively develop emerging expertise. Areas like real-time analytics, edge computing, and AI-powered automation are transitioning from experimental to production deployment at technology companies, creating demand for professionals with hands-on experience. The project-based format is particularly valuable because it creates portfolio evidence of depth beyond standard coursework—demonstrating initiative and specialized knowledge that distinguishes candidates in competitive hiring. Companies explicitly value candidates who stay current with methodological advances rather than relying solely on established techniques. The flexibility to explore specialized topics

also enables students to align their skills with specific career trajectories, whether that's financial time series methods, geospatial analytics, or natural language processing applications.

DAV 6400 Internship

This course consists of an off-campus internship experience supervised by a staff person at the internship site and overseen by a faculty advisor. The program director must approve the site and the overall duration of student work must be no less than 150 hours (based on a 3-credit course). At the start of the internship, the student and faculty advisor will jointly develop specific learning objective. Over the course of the internship, students will submit weekly reflections, and at the end of the internship, students write a final paper that represents the culmination of the work performed.

Industry Application: Internship experience provides professional exposure and relationship-building that cannot be replicated in classroom settings, frequently serving as the primary pathway to full-time employment. Major employers use internship programs as extended interviews—the performance evaluation period happens during the internship rather than through traditional hiring processes. Technology companies like Google and Microsoft convert significant portions of their intern classes to full-time offers because internships allow mutual assessment of fit that résumés and interviews cannot provide. The structured 150-hour minimum with faculty oversight ensures substantive project work rather than administrative tasks, developing both technical capabilities and critical professional skills like stakeholder communication and deadline management. Financial institutions and consulting firms particularly value internship experience because it demonstrates that candidates can operate in professional environments and deliver under business constraints—capabilities that academic performance alone doesn't guarantee. The weekly reflections and final paper components develop the self-assessment and documentation habits that characterize high-performing professionals.

DAV 6450 Independent Study

This independent study course provides the student with the flexibility to learn more about a topic of interest outside of the formal course setting. Students will consult with a faculty advisor on the subject of the independent study 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. Course credit is determined in advance of the course, by the instructor with the approval of the program director.

Industry Application: Independent study demonstrates intellectual initiative and depth of investigation that distinguish exceptional candidates from competent ones in competitive hiring processes. The ability to self-direct learning and pursue specialized knowledge beyond required coursework signals the continuous learning mindset essential for long-term career success. Technology companies value this initiative because professional work requires exactly this capability—identifying knowledge gaps, finding resources, and developing expertise without formal instruction. The faculty mentorship relationship provides individualized career guidance and professional connections while developing the advisor-mentee dynamic common in professional research environments. For students targeting specialized roles in areas like sports analytics, political polling, or environmental data science, independent study enables deep

exploration of domain-specific applications that standard courses cannot cover. This experience also serves as excellent preparation for graduate programs requiring independent research capabilities and provides compelling evidence of genuine passion for specific analytical domains when competing for positions where cultural fit and intellectual curiosity matter as much as technical skills.

Updated: January 2026