

MASTER OF SCIENCE PROGRAM IN APPLIED DATA SCIENCE

The Master of Science in Applied Data Science gives students a thorough knowledge of techniques in the field of analytics and data science, and the ability to apply them to real-world business scenarios. Building from a core in applied statistics, math and programming, students are provided with advanced analytical training to develop their ability to draw insights from big data. This includes: machine learning and predictive analytics, deep learning, reinforcement learning, data engineering platforms, time series analysis, linear and non-linear models, statistical methods, and other sophisticated techniques for analyzing complex data.

The program is highly applied in nature, integrating business strategy, project-based learning, simulations, case studies, and specific electives addressing the analytical needs of various industry sectors. Through partnerships with key employers, the program also provides students with a client based, 2 quarter Capstone experience (<https://datascience.uchicago.edu/education/masters-programs/ms-in-applied-data-science/capstone-projects/>) (1 quarter for MBA/MS in Applied Data Science joint degree students) as well as access to career networks and employment pathways upon graduation.

- Program type: Masters of Science degree program
- Program structure, courses, requirements, and application (<https://datascience.uchicago.edu/education/masters-programs/ms-in-applied-data-science/>)
- Location: Cityfront Campus (Downtown) - NBC Tower and the Gleacher Center
- Time to completion: 1-4 years
- Only courses with a grade of B- or better will count toward degree requirements
- Minimum cumulative GPA for satisfactory academic progress: 3.0

PROGRAM MODALITIES

MS in Applied Data Science - In-Person

The In-Person program is ideal for those who learn best in an in-person classroom environment, with access to in-person student services, and prefer to be in Chicago for the duration of the degree program. To view a complete list of course offerings, please visit the in-person program (<https://datascience.uchicago.edu/education/masters-programs/in-person-program/>) website.

MS in Applied Data Science - Online

Leverage the latest tools in data analytics with our online applied data science Master's degree. The online program gives you maximum flexibility, while broadening your network of peers. To view a complete list of course offerings, please visit the online program (<https://datascience.uchicago.edu/education/masters-programs/online-program/>) website.

Joint MBA/MS in Applied Data Science

The joint degree with UChicago's Booth School of Business is ideal for ambitious students looking to supplement their MBA studies with a cutting-edge education in data science. To view a complete list of course offerings, please visit the joint MBA/MS program (<https://www.chicagobooth.edu/mba/joint-degree/mba-ms-applied-data-science/>) website.

PROGRAM REQUIREMENTS FOR MS IN APPLIED DATA SCIENCE

- For students admitted to MS in Applied Data Science In-Person or Online prior to Autumn 2024 (1200 units required for completion):
 - Foundational courses (4 courses – Zero Unit)
 - Core Courses (7 courses – 100 Units each)
 - Electives (3 courses – 100 Units each)
 - Capstone Project (2 courses – 100 Units each)
- For students admitted to MS in Applied Data Science In-Person or Online in Autumn 2024 and beyond (1200 units required for completion):
 - Foundational courses (4 courses) – Optional
 - Core Courses (6 courses – 100 Units each)
 - Electives (4 courses – 100 Units each)
 - Capstone Project (2 courses – 100 Units each)
 - Career Seminar Requirement (Zero Units)

For information on the program requirements for the joint MBA/MS in Applied Data Science programs, please visit the joint MBA/MS program (<https://www.chicagobooth.edu/mba/joint-degree/mba-ms-applied-data-science/>) website.

Career Seminar Requirement: Our career seminar requirement focus on how to identify and develop industry skills, explore the broad array of jobs that data scientists obtain, and determine what types of companies may be the best fit for the next stages of our students' careers. Our seminar series helps prepare students for the data science marketplace, to present themselves confidently, and assist students in receiving an internship and/or full-time job offer upon graduation. Students with significant full-time eligible work experience may be eligible to waive this course. For more information on the waiver submission process and the career course seminar sequence, please visit the Curriculum (<https://applieddatascience.psd.uchicago.edu/curriculum/>) section of the current student website.

Please see our prospective website for admissions requirements and deadlines (<https://datascience.uchicago.edu/education/masters-programs/ms-in-applied-data-science/how-to-apply/>) and our current student website for information on policies, resources, and program activities (<https://www.applieddatascience.psd.uchicago.edu/>).

APPLIED DATA SCIENCE COURSES

ADSP 31000. Introduction to Statistical Concepts. 000 Units.

This course provides general exposure to basic statistical concepts that are necessary for students to understand the content presented in more advanced courses in the program. The course covers theoretical distributions and the way these distributions are used to assign probabilities to events in some depth. The course also introduces students to descriptive statistical methods to explore and summarize data, methodologies for sampling units for measurement or analysis, drawing inferences on the basis of knowledge gained from samples to populations, assessing relationships between variables, and making predictions based upon relationships between variables.

Terms Offered: Autumn Spring

ADSP 31006. Time Series Analysis and Forecasting. 100 Units.

Time Series Analysis is a science as well as the art of making rational predictions based on previous records. It is widely used in various fields in today's business settings. For example, airline companies employ time series to predict traffic volume and schedule flights; financial agencies measure market risk via stock price series; marketing analysts study the impact of a newly proposed advertisement by the sales series. A comprehensive knowledge of time series analysis is essential to the modern data scientist/analyst. This course covers important issues in applied time series analysis: a solid knowledge of time series models and their theoretical properties; how to analyze time series data by using mainstream statistical software; practical experience in real data analysis and presentation of their findings in a logical and clear way to various audiences.

Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): MSCA 31007: Statistical Analysis

ADSP 31008. Data Mining Principles. 100 Units.

Drawing on statistics, artificial intelligence and machine learning, the data mining process aims at discovering novel, interesting and actionable patterns in large datasets. This class will introduce the student to the fundamentals of data mining: association and sequence rules discovery, memory-based reasoning, classification and regression decision trees, comparison of data mining models, logistic models, scorecard models, and neural network models. The class follows a learn-by-doing approach in which the student will complete bi-weekly assignments using real world datasets. The student will also propose and complete a data mining research project of their own design.

Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): MSCA 31007: Statistical Analysis

ADSP 31009. Machine Learning & Predictive Analytics. 100 Units.

This course in advanced data mining will provide a practical, hands-on set of lectures surrounding modern predictive analytics and machine learning algorithms and techniques. It will emphasize practice over mathematical theory, and students will spend a considerable amount of class time gaining experience with each algorithm using existing packages in R, Python, and Linux libraries. The course will cover the following topics: regression and logistic regression, regularized regression including the lasso and elastic net techniques, support vector machines, neural networks, decision trees, boosted decision trees and random forests, online learning, k-means and special clustering, and survival analysis.

Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): MSCA 31008: Data Mining Principles, MSCA 31010: Linear and Non-Linear Models; MSCA 37014: Python for Analytics

ADSP 31010. Linear and Non-Linear Models. 100 Units.

This course concentrates on the following topics: Review of statistical inference based on linear model, extension to the linear model by removing the assumption of Gaussian distribution for the output (Generalized Linear Model), extension to the linear model by allowing a correlation structure for the model residuals (mixed effect models), and extension of the linear model by relaxing the requirement that inputs are combined linearly (nonparametric regression, regime switches). Course emphasizes applications of these models to various fields and covers main steps of building analytics from visualizing data and building intuition about their

structure and patterns to selecting appropriate statistical method to interpretation of the results and building analytical model. Topics are illustrated by data analysis projects using R. Familiarity with R at some basic level is not a requirement but recommendation. Students can pick up the programming language by following the descriptions of the examples.

Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): MSCA 31007: Statistical Analysis

ADSP 31012. Data Engineering Platforms for Analytics. 100 Units.

Effective data engineering is an essential first step in building an analytics-driven competitive advantage in the market. Modern data engineering platforms reduce manual data preparation by automating processes, which in turn, enable companies to focus on deriving efficiencies in data processing to develop impactful business insights. This course provides students with a thorough understanding of the fundamentals of data engineering platforms, for both operational and analytical use cases, while gaining hands-on expertise in building these platforms in a way to develop analytical solutions effectively. Students will have the opportunity to construct both relational and analytical databases on the cloud or on premise from real-life datasets while using programmatic or configuration driven data pipelines. By the end of the course, students will be able to design and implement an end-to-end data engineering platform capable of supporting sustainable analytics solutions.

Terms Offered: Autumn Spring Summer Winter

ADSP 31013. Big Data and Cloud Computing. 100 Units.

This course teaches students how to approach Big Data and large-scale machine learning applications. While there is no single definition of Big Data and multiple emerging software packages exist to work with Big Data, we will cover the most popular approaches. Students will learn the Big Data infrastructure, including Linux, Massive parallelization and Distributed Computing, and how to apply both Hadoop and Spark map-reduce concepts for clustering, similarity search, web analytics and classification. During the course, we will cover the applications of NoSQL systems, such as JSON stores, object storage and Elasticsearch. The cloud computing section of the course will focus on virtualization and container orchestration, including virtual machines, dockers and Kubernetes. During the course students will gain hands-on expertise leveraging Hive, Pig, Python and PySpark for Big Data applications in client-server environment. For MScA students in the 12-course curriculum who wish to take Big Data Platforms as a core course, instead of MSCA 31012 Data Engineering Platforms: Certain technical skills and knowledge are required to be successful in this course. Required skills may be found for MSCA 31013 Big Data Platforms at <https://professional.uchicago.edu/find-your-fit/masters/master-science-analytics/curriculum>. Contact the program administration for further details.

Terms Offered: Autumn Spring Summer Winter

ADSP 31014. Statistical Models for Data Science. 100 Units.

In a traditional linear model, the observed response follows a normal distribution, and the expected response value is a linear combination of the predictors. New methods based on probability distributions other than Gaussian appeared only in the second half of the twentieth century. These methods allowed working with variables that span a broader variety of domains and probability distributions. Besides, methods for the analysis of general associations were developed that are different from the Pearson correlation. This course begins in linear normal models. We will visit the foundations of generalized linear models (GLM) and take a detour to see the survival models. This journey ends at some nonlinear model lookout posts at the instructor's discretion. This course will prepare students to be ready for and capable of the statistical analysis process. Students will first discover the insights, formulate the propositions, validate the evidence, and finally build the solutions for solving business problems. Following the process properly raises credibility and increases the impact of the results. Besides developing Python codes for carrying out the process, students will learn to tune the software tools for the most efficient implementation and optimal performance. At the end of this course, students will have built their inventory of data analysis codes and their confidence in advocating their propositions to the business stakeholders.

Terms Offered: Autumn Spring Summer Winter

ADSP 31016. Leadership & Consulting in Data Science. 100 Units.

Professional organizations see value in data science when it helps them to achieve their strategic goals, and the current job market likewise rewards data scientists who can use data to advance organizational interests, either as an external consultant or within internal operations teams. Data scientists can become successful (and highly marketable) leaders in today's professional world, but they require an uncommon skill set: the strategic awareness to align data requirements with business requirements, the technical proficiency to choose a methodology appropriate to each new problem, and the communication skills to both execute the plan as part of a broader team and persuade others of their findings. The Leadership and Consulting in Data Scientist course is focused on: •

Terms Offered: Autumn

Prerequisite(s): ANLT-MS, ANLP-MS Students Only

ADSP 31017. Machine Learning I. 100 Units.

This course is aimed at providing students an introduction to machine learning with data mining techniques and algorithms. It gives a rigorous methodological foundation in analytical and software tools to successfully undertake projects in Data Science. Students are exposed to concepts of exploratory analyses for uncovering and

detecting patterns in multivariate data, hypothesizing and detecting relationships among variables, conducting confirmatory analyses, and building models for predictive and descriptive purposes. It will present predictive modeling in the context of balancing predictive and descriptive accuracies.

Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): ADSP 31014 Statistical Models for Data Science

ADSP 31018. Machine Learning II. 100 Units.

The objective of this course is three-folds - first, to extend student understanding of predictive modeling with machine learning concepts and methodologies from Machine Learning I into the realm of Deep Learning and Generative AI. Second, to develop the ability to apply those concepts and methodologies to diverse practical applications, evaluate the results and recommend the next best action. Third, to discuss and understand state-of-the machine learning and deep learning research and development and their applications. This course clarifies concepts such as Artificial Intelligence, Machine Learning and Deep Learning and distinguishes between Expert and Learning systems. It introduces different types of learning systems such as error-based, information-based, representation, active, and generative modeling. The course expands on recommender systems and regularized regression from Machine Learning 1 using different types of data encoding for learning systems. Additional topics include model selection, fairness, and ethics. Students learn state-of-the-art machine and deep learning applications in the industry along with their pros and cons via case studies, assignments, and a class project. The course is taught in Python. The class covers popular machine learning library APIs and implementations using examples from public github repositories.

Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): ADSP 31017 Machine Learning I

ADSP 32001. Financial Analytics. 100 Units.

This course concentrates on the following topics: review of financial markets and assets traded on them; main characteristics of financial analytics: returns, yields, volatility; review of stochastic models of market price and their statistical representations; concept of arbitrage, elements of arbitrage pricing approach; principles of volatility analyses, implied vs. realized volatility; correlation, cointegration and other relationships between various financial assets; market risk analytics and management of portfolios of financial assets. The course puts special emphasis on covering main steps of building analytics from visualizing data and building intuition about their structure and patterns to selecting appropriate statistical method to interpretation of the results and building analytical models. Topics are illustrated by data analysis projects using R. Basic familiarity with R is a requirement.

Terms Offered: Summer Winter

Prerequisite(s): MSCA 31007: Statistical Analysis

ADSP 32007. Data Visualization Techniques. 100 Units.

In today's data driven enterprise, data storytelling using effective visualization strategies is an essential skill for analytics practitioners in almost every field to explore and present data. This course focuses on modern data visualization technologies, tools, and techniques to convert raw data into actionable information. Modern data visualization tools are at the forefront of the "self-service analytics" architectures which are decentralizing analytics and breaking down IT bottlenecks for business experts. Moreover, with its foundations rooted in statistics, psychology, and computer science, data visualization shows you how to better understand the data, present clear evidence of your findings to your intended audience and tell engaging data stories through charts and graphics. This course is designed to introduce data visualization as a medium of effective communication using strategic storytelling, and the basis for interactive information dashboards.

Terms Offered: Winter

ADSP 32009. Data Science in Healthcare. 100 Units.

Given the breadth of the field of health analytics, this course will provide an overview of the development and rapid expansion of analytics in healthcare, major and emerging topical areas, and current issues related to research methods to improve human health. We will cover such topics as security concerns unique to the field, research design strategies, and the integration of epidemiologic and quality improvement methodologies to operationalize data for continuous improvement. Students will be introduced to the application of predictive analytics to healthcare. Students will understand factors impacting the delivery of quality and safe patient care and the application of data-driven methods to improve care at the healthcare system level, design approaches to answering a research question at the population level, become familiar with the application of data analytics to impacting care at the provider level through Clinical Decision Systems, and understand the process of a Clinical Trail.

Terms Offered: Winter

ADSP 32013. Optimization and Simulation Methods for Analytics. 100 Units.

This course introduces students to how optimization and simulation techniques can be used to solve many real-life problems. It will cover two classes of optimization methods. First class has been developed to optimize real, non-simulated systems or to find the optimal solution of a mathematical model. The methods that belong to this class include liner programming, quadratic programming and mixed-integer programming. Second class of methods has been developed to optimize a simulation model. The difference with the classical mathematical programming methods is that the objective function (which is the function to be minimized or maximized) is not known explicitly and is defined by the simulation model (computer code). The course will demonstrate multiple

approaches to build simulation models, such as discrete event simulations and agent-based simulations. Then, it will show how stochastic optimization and heuristic approaches can be used to analyze the simulated system and design a sequence of computational experiments that allow to develop a basic understanding of a particular simulation model or system through exploration of the parameter space, to find robust plausible behaviors and conditions and robust near-optimal solutions that are not prone to being unstable under small perturbations.

Terms Offered: Summer Winter

Prerequisite(s): MSCA 31007: Statistical Analysis

ADSP 32014. Bayesian Machine Learning with Generative AI Applications. 100 Units.

This course provides a strong theoretical and practical skillset for probabilistic machine learning applications. Bayesian inference and modeling methods are important for several areas including prediction, decision making, and risk assessment where modeling the uncertainty is needed. The course begins with an introduction to Bayesian statistical analysis, covering the foundations of Bayesian inference and the application of Bayes' theorem for statistical inference. We then introduce Bayesian networks, which offer a powerful graphical tool for modeling complex systems and making probabilistic inferences. The course then advances to cover more sophisticated topics such as Markov Chain Monte Carlo (MCMC) methods for sampling from complex probability distributions, hierarchical models, and model selection techniques. The final three weeks are dedicated to cutting-edge methodologies like Generative Deep Learning, Variational Autoencoders, and Bayesian Neural Networks, all rooted in Bayesian Machine Learning. Upon completion, students will be equipped to apply Bayesian methods to a wide range of real-world problems in fields such as engineering, business, finance, and public policy, addressing challenges like missing data or training AI models that are able to say 'I don't know'.

Terms Offered: Autumn Spring Summer Winter

Note(s): In order to register for this course you must be a Data Analytic or Applied Data Science Student

ADSP 32015. Digital Marketing Analytics in Theory and Practice. 100 Units.

Successfully marketing brands today requires a well-balanced blend of art and science. This course introduces students to the science of web analytics while casting a keen eye toward the artful use of numbers found in the digital space. The goal is to provide marketers with the foundation needed to apply data analytics to real-world challenges they confront daily in their professional lives. Students will learn to identify the web analytic tool right for their specific needs; understand valid and reliable ways to collect, analyze, and visualize data from the web; and utilize data in decision making for their agencies, organizations or clients. By completing this course, students will gain an understanding of the motivations behind data collection and analysis methods used by marketing professionals; learn to evaluate and choose appropriate web analytics tools and techniques; understand frameworks and approaches to measuring consumers' digital actions; gain familiarity with the unique measurement opportunities and challenges presented by New Media; gain hands-on, working knowledge of a step-by-step approach to planning, collecting, analyzing, and reporting data; utilize tools to collect data using today's most important online techniques: performing bulk downloads, tapping APIs, and scraping webpages; and understand approaches to visualizing data effectively.

Terms Offered: Autumn

ADSP 32017. Advanced Machine Learning & Artificial Intelligence. 100 Units.

This course delves into advanced and evolving topics in machine learning, with a focus that may vary each semester. Students will engage in intensive programming projects, gaining hands-on experience with cutting-edge tools and techniques. For this semester, the course will emphasize two key areas: transformers and reinforcement learning. Students will explore the internals of transformer models, build a transformer from the ground up, and apply transformers to a range of tasks including NLP, time series analysis, and image processing. The course will involve extensive use of PyTorch, which we shall study, and TensorFlow, with a strong reliance on Hugging Face libraries, models, and datasets. Additionally, students will learn the fundamentals of reinforcement learning, working with environments such as OpenAI Gym and leveraging tf-agents to build and train RL models. This course is ideal for students seeking to deepen their understanding of advanced machine learning topics through practical, project-based learning.

Terms Offered: Summer Winter

Prerequisite(s): Required: MSCA 31009: Machine Learning & Predictive Analytics Recommended: MSCA 37011 Deep Learning & Image Recognition

ADSP 32018. Natural Language Processing and Cognitive Computing. 100 Units.

Extracting actionable insights from unstructured text and designing cognitive applications have become significant areas of application for analytics. Students in this course will learn foundations of natural language processing, including: concept extraction; text summarization and topic modeling; part of speech tagging; named entity recognition; semantic roles and sentiment analysis. For advanced NLP applications, we will focus on feature extraction from unstructured text, including word and paragraph embedding and representing words and paragraphs as vectors. For cognitive analytics section of the course, students will practice designing question answering systems with intent classification, semantic knowledge extraction and reasoning under uncertainty. Students will gain hands-on expertise applying Python for text analysis tasks, as well as practice with multiple IBM Watson services, including: Watson Discovery, Watson Conversation, Watson Natural Language Classification and Watson Natural Language Understanding.

Terms Offered: Autumn Spring

Prerequisite(s): MSCA 31008: Data Mining Principles

ADSP 32019. Real-Time Intelligent Systems. 100 Units.

Developing end-to-end automation and intelligent systems is now the most advanced area of application for analytics. Building such systems requires proficiency in programming, understanding of computer systems, as well as knowledge of related analytical methodologies, which are the skills that this course aims to teach to students. The course focuses on python and is tailored for students with basic programming knowledge in Python. The course is partially project based. During the first three sessions, we will review basic python concepts and then learn more advanced python and the ways to use Python to handle large data flows. The later sessions are project based and will focus on developing end-to-end analytical solutions in the following areas: Finance and trading, blockchains and crypto-currencies, image recognition, and video surveillance systems.

Terms Offered: Autumn

Prerequisite(s): MSCA 31007: Statistical Analysis Recommended: MSCA 37014: Python for Analytics

ADSP 32020. Reinforcement Learning. 100 Units.

This course is an introduction to reinforcement learning, also known as neuro-dynamic programming. It discusses basic and advanced concepts in reinforcement learning and provides several practical applications. Reinforcement learning refers to a system or agent interacting with an environment and learning how to behave optimally in such environment. An environment typically includes time, actions, states, uncertainty and rewards. Reinforcement learning combines neuro networks and dynamic programming to find an optimal behavior or policy of the system or agent in complex environment setting. Neuro networks approximations are used to circumvent the well-known 'curse of dimensionality' which have been a barrier to solving many practical applications. Dynamic programming is the key learning mechanism that the system or the agent uses to interact with the environment and improve its performance. Students will master key learning techniques and will become proficient in applying these techniques to complex stochastic decision processes and intelligent control.

Terms Offered: Spring Summer Winter

Prerequisite(s): MSCA 31007 Statistical Analysis

ADSP 32021. Machine Learning Operations. 100 Units.

The objective of this course is two-folds - first, to understand what Machine Learning Operations (MLOps) is and why it is a key component in enterprise production deployment of machine learning projects. Second, to expose students to software engineering, model engineering and state-of-the-art deployment engineering with hands-on platform and tools experience. This course crosses the chasm that separates machine learning projects/experiments and enterprise production deployment. It covers 3 pillars in MLOps: software engineering such as software architecture, Continuous Integration/Continuous Delivery and data versioning; model engineering such as AutoML and A/B experimentation; and deployment engineering such as docker containers and model monitoring. The course focuses on best practices in the industry that are critical to enterprise production deployment of machine learning projects. Having completed this course, a student understands the machine learning lifecycle and what it takes to go from ideation to operationalization in an enterprise environment. Furthermore, students get exposure to state-of-the-art MLOps platforms such as allegro (<https://allegro.ai/>), xpresso (<https://abzooba.com/xpresso-ai>), Dataiku (<https://www.dataiku.com/>), LityxIQ (<https://lityx.com/>), DataRobot (<https://www.datarobot.com/>), AWS Sagemaker (<https://aws.amazon.com/sagemaker/>), and technologies such as gitHub, Jenkins, slack, docker, and kubernetes.

Instructor(s): Arnab Bose Terms Offered: Autumn

Prerequisite(s): MSCA 31009: Machine Learning & Predictive Analytics

ADSP 32023. Advanced Computer Vision with Deep Learning. 100 Units.

Computer vision is the field of computer science that focuses on creating digital systems that can process, analyze, and make sense of visual data in the same way that humans do. Deep learning is a subset of machine learning and a branch of Artificial Intelligence (AI). It involves the training, deployment, and application of large complex neural network architectures to solve cutting-edge problems. Deep Learning has become the primary approach for solving cognitive problems such as Computer Vision and Natural Language Processing (NLP) and has had a massive impact on various industries such as healthcare, retail, automotive, industrial automation, and agriculture. This course will enable students to build Deep Learning models and apply them to computer vision tasks such as object recognition, detection, and segmentation. Students will gain an in-depth understanding of the Deep Learning model development process, tools, and frameworks. Although the focus of the course will primarily be computer vision, students will work on both image and nonimage datasets during class exercises and assignments. Students will gain hands-on experience in popular libraries such as Tensorflow, Keras, and PyTorch. Students will also learn to apply state of the art models such as ResNet, EfficientNet, RCNNs, YOLO, Vision Transformers, etc. for computer vision and work on datasets such as CIFAR, ImageNet, MS COCO, and MPII Human Poses.

Instructor(s): Ashish Pujari Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): Know your computer (Setting environment variables, Using the Mac/PC terminal, traversing applications/folders, updating security preferences). Familiarity with data engineering and cloud computing. MSCA 37010 Programming for Analytics, MSCA 37014 Python for Analytics, MSCA 31009 Machine Learning and Predictive Analytics. ANLT-MS, ANLP-MS, and GSCP students only.

ADSP 32024. Data Science for Algorithmic Marketing. 100 Units.

This course focuses on marketing science methods and algorithms. The course will expose and immerse students in the marketing science algorithms. The course would cover algorithms for undertaking competitive analysis in the digital landscape, market segmentation, mining databases for effective digital marketing, design of new digital and traditional products, forecasting sales and product diffusion, real time product positioning, intra omni-channel optimization and inter omni-channel resource allocation, and pricing across both omni-channel marketing effectiveness and ROI. The course will use a combination of lecture, in-class discussions, group assignments, and a final group project. The course lays special emphasis on algorithms. Hence it draws heavily from the fields of optimization, machine-learning based recommendation systems, association rules, consumer choice models, Bayesian estimation, experimentation and analysis of covariance, advanced visualization techniques for mapping brand perceptions, and analysis of social media data using advanced NLP techniques.

Instructor(s): Anil Chaturvedi Terms Offered: Autumn Winter

Prerequisite(s): MSCA 31007 Statistical Analysis

Note(s): ANLP-MS, ANLT-MS and GSCP-MS only

ADSP 32025. Supply Chain Optimization. 100 Units.

"Big Data" continues to grow exponentially in our large-scale transactional world where 100,000s of SKUs and millions of customers are interacting with 1:1 offers that include differential pricing, shipping timing/costs and even made to order "custom" product configurations. These consumer behaviors are quickly advancing the availability of new data and techniques within the discipline of Data Science. This elective course will give students the opportunity to apply their skills in data visualization, data mining tools, predictive modeling, and advanced optimization techniques to address Supply Chain challenges. The course focuses on the use of Advanced Predictive Modeling, Machine Learning, AI and other Data Science insight and activation tools are to automate and optimize the performance of the Supply Chain. Students will also learn how to optimize the performance of the Supply Chain from the lens of multiple related disciplines including: Sales Forecasting, Warehousing/Inventory Management, Promotion, Pricing, Logistics Network Optimization, Freight Cost Management, Manufacturing, Retail POS Information, Ecommerce, Consumer Data, and Product Design/Packaging. After completing this course, you will be prepared to work in any of the numerous specialty areas possible in the world of Supply Chain Management.

Instructor(s): Gizem Aydin Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): MSCA 31008, MSCA 37014, and for ANLT-MS, ANLP-MS, GSCP students only

ADSP 32026. Python for ML Engineering. 100 Units.

This course prepares data science students to go beyond programming constructs and data science libraries. This course provides an understanding of web applications so students can get a deeper understanding of how their models are deployed. They are taught python environment management, required for production work. They are taught advanced data structures such as trees and graphs to allow them to work with more complex modes and more advanced control structures, such as generators a recursive functions to help them write more expressive code.

Instructor(s): Shahbaz Chaudhary Terms Offered: Autumn Spring Summer Winter

Note(s): It is recommended that students have taken the Python for Data Science foundational course prior to registering for this course.

ADSP 32027. Generative AI Principles. 100 Units.

This course dives into the realm of Generative AI, offering a comprehensive look into the world of Large Language Models (LLMs), image generation techniques, and the fusion of vision and text through multimodal models. Drawing from core concepts in neural networks, transformers, and advanced techniques such as prompt engineering, vision prompting, and multimodality representation, students will explore the capabilities, applications, and ethical considerations of generative models. This course culminates in hands-on projects, allowing participants to apply theory to practical scenarios.

Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): ADSP 31017 and ADSP 31018

ADSP 34002. Capstone I. 100 Units.

The overarching goal of this course is to take students two steps closer to being "Complete Data Scientists". The first step is by letting students manage and solve a real data science project with real clients and real problems. Students will complete the design of their Capstone Projects, and begin the implementation. The second step is by exposing them to data science methodologies in the absence of pre-existing data - by exposing them to quantitative methodologies in optimally designing data collection tasks. This course covers the Business analytic process from the translation of business problems and opportunities into questions that can be addressed by using data science, development of analytical plans including methodologies and data to address these issues, and initial implementation of these analytical plan.

Instructor(s): Anil Chaturvedi, Gregory Green Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): Prerequisites: 6 Core Courses Completed from MSCA Curriculum. Restricted to MSCA students completing the 12-course program curriculum.

ADSP 34003. Capstone II. 100 Units.

The Capstone II class is designed to: 1) Provide students maximum flexibility in the latter stages of their Capstone project to work heavily with their Capstone advisors in concluding the execution of the analytic

methodology and any client / sponsor deliverables for the project. 2) Provide maximum support to students in the curation and delivery of key project communications: a) Formal research paper. b) Formal business presentation of project details, value, findings and recommendations. c) Live presentation by the team in Capstone Showcase including question / answer session with a judging panel.

Instructor(s): Donald Patchell Terms Offered: Autumn Spring Summer Winter

Prerequisite(s): Required: MSCA 34002 Capstone 1. Restricted to MScA students completing the 12-course program curriculum.

ADSP 37025. MS in Applied Data Science Career Seminar I. 000 Units.

This course will help you navigate your career in data science and land a job that fits your needs and desires. It will lead you through a deeper discovery into who you are, clarifying what you want to do with your career, and navigating the market to find the right company and job match. This noncredit, 4 course sequence is required for all students. Students with extensive work experience, may be eligible to waive this course.

Terms Offered: Autumn

Note(s): In order to register for this course you must be a Data Analytic or Applied Data Science Student

ADSP 37026. MS in Applied Data Science Career Seminar II. 000 Units.

This course will help you navigate your career in data science and land a job that fits your needs and desires. It will lead you through a deeper discovery into who you are, clarifying what you want to do with your career, and navigating the market to find the right company and job match. This noncredit, 4 course sequence is required for all students. Students with extensive work experience, may be eligible to waive this course.

Terms Offered: Winter

Prerequisite(s): ADSP 37025 - MS in Applied Data Science Career Seminar I

Note(s): In order to register for this course you must be a Data Analytic or Applied Data Science Student

ADSP 37027. MS in Applied Data Science Career Seminar III. 000 Units.

This course will help you navigate your career in data science and land a job that fits your needs and desires. It will lead you through a deeper discovery into who you are, clarifying what you want to do with your career, and navigating the market to find the right company and job match. This noncredit, 4 course sequence is required for all students. Students with extensive work experience, may be eligible to waive this course.

Terms Offered: Spring

Prerequisite(s): ADSP 37025 and ADSP 37026

Note(s): In order to register for this course you must be a Data Analytic or Applied Data Science Student

ADSP 37028. MS in Applied Data Science Career Seminar IV. 000 Units.

This course will help you navigate your career in data science and land a job that fits your needs and desires. It will lead you through a deeper discovery into who you are, clarifying what you want to do with your career, and navigating the market to find the right company and job match. This noncredit, 4 course sequence is required for all students. Students with extensive work experience, may be eligible to waive this course.

Terms Offered: Autumn Summer

Prerequisite(s): ADSP 37025, ADSP 37026, and ADSP 37027

Note(s): In order to register for this course you must be a Data Analytic or Applied Data Science Student

ADSP 40100. Data Science Practicum. 000 Units.

Data Science Practicum is part of the co-operative educational agreement between MS in Applied Data science program and employers that provides off-campus work authorization for international students to pursue internships. The internships must meet the requirement that students archive at least five learning objectives of the course. The learning objectives are about students developing or sharpening their skills in applying analytical tools to solve real life problems.

Terms Offered: Autumn Spring Summer Winter

