

# Data Science Institute

Brown University's Data Science Institute (DSI) serves as a campus hub for research and education in data science. Through our research and academic programs, we strive to ensure that those most in need are not the last to benefit from fundamental research in data science or data-driven applied research. Brown's DSI is unique because we:

- Equally value both domain-driven and fundamental methodological research in data science;
- Increase data fluency and educate the next generation of data scientists, through our master's program and providing outreach to students and researchers at a variety of career stages;
- Explore the impact of the data revolution on culture, education, health and genomics, society, and social justice by engaging with research partners within the University and beyond.

The DSI offers multiple academic programs:

- Residential master's degree (<https://dsi.brown.edu/academics/masters-programs/masters-data-science/>) in Data Science, including a fifth-year master's (<https://dsi.brown.edu/academics/masters-programs/masters-data-science/fifth-year-masters-degree/>) option for Brown undergraduates
- Online master's (<https://dsi.brown.edu/academics/masters-programs/online-masters-data-science/>) in Data Science: Policy, Governance, and Society
- Doctoral certificate (<https://dsi.brown.edu/academics/doctorsal-certificate/>) for Brown graduate students
- Undergraduate certificate (<https://dsi.brown.edu/academics/certificate-data-fluency/>) in Data Fluency

The DSI also supports two research centers, the Center for Computational Molecular Biology (<https://ccmb.brown.edu/>) (CCMB) and the Center for Technological Responsibility, Reimagination, and Redesign (<https://cntr.brown.edu/>) (CNTR).

To support data science research and education across Brown's campus, the DSI hosts seminars and public lectures, offers workshops in data science skills, and offers small grants to Brown researchers in all disciplines.

For additional information, please visit the institute's website: <http://dsi.brown.edu/>.

## Data Fluency Certificate Requirements

The Certificate in Data Fluency provides a formal pathway for undergraduates in concentrations other than applied mathematics, computational biology, computer science, math, and statistics (see details below) who wish to gain fluency and facility with the tools of data science. The driving intellectual question motivating certificate students is how we can infer meaning from data whilst avoiding false predictions. The required experiential learning component provides you with the opportunity to apply your data-science skills in applied settings, engage in research that uses data science, teach data science as an undergraduate teaching assistant, or undertake an internship that has a substantive data-science component.

As with all undergraduate certificates (<https://www.brown.edu/academics/college/degree/undergraduatecertificates/>), the certificate has the following requirements:

- Students may not earn more than **one certificate** and may only have **one declared concentration**.
- Students must be enrolled in or have completed at least **two courses** toward the certificate at the time they declare in ASK.
- No more than **one course** may count toward your concentration and the certificate.
- Students may declare in ASK **no earlier** than the beginning of the fifth semester and must declare **no later** than the last day of classes of the antepenultimate (typically the sixth) semester, in order to facilitate planning for the capstone or other experiential learning opportunity.
- Students must submit a proposal for their experiential learning opportunity by the end of the **sixth semester**.

Excluded Concentrations: Applied Mathematics, Computational Biology, Computer Science, Mathematics, and Statistics. This includes joint concentrations in these areas; for example, Applied Mathematics-Economics is also excluded. According to the certificate guidelines, a student's concentration and certificate cannot have substantial overlap.

For more information on the Certificate in Data Fluency, please visit the Data Science Institute (<https://dsi.brown.edu/academics/certificate-data-fluency/>) website.

## Certificate Requirements

### Certificate Requirements

Core Courses:

DATA 0080	Data, Ethics and Society	1
CSCI 0111	Computing Foundations: Data	1
or CSCI 0150	Introduction to Object-Oriented Programming and Computer Science	
or CSCI 0170	Computer Science: An Integrated Introduction	
or CSCI 0190	Accelerated Introduction to Computer Science	
or CPSY 0950	Introduction to programming	
DATA 0200	Data Science Fluency	1
Elective Course: Select one follow-up Applied Math, Biostatistics, Computer Science or domain-specific course with a significant data component from the following list (or another course with approval from the certificate advisor):		1
ANTH 1201	Introduction to Geographic Information Systems and Spatial Analysis	
APMA 1650	Introduction to Probability and Statistics with Calculus	
BIOL 0495	Statistical Analysis of Biological Data	
BIOL 1555	Methods in Informatics and Data Science for Health	
BIOL 1565	Survey of Health Informatics	
CPSY 0900	Statistical Methods	
CPSY 1291	Computational Methods for Mind, Brain and Behavior	
CPSY 1580C	Visualizing Information	
CSCI 1411	Foundations of AI	
CSCI 1420	Machine Learning	
CSCI 1450	Advanced Introduction to Probability for Computing and Data Science	
CSCI 1470	Deep Learning	
CSCI 1951A	Data Science	
DATA 1030	Hands-on Data Science	
DATA 1150	Data Science Fellows <sup>1</sup>	
DATA 1500	Data Visualization & Narrative	
ECON 1620	Introduction to Econometrics	
ECON 1660	Big Data	
EDUC 1230	Applied Statistics for Ed Research and Policy Analysis	
ENVS 1105	Introduction to Environmental GIS	
EEPS 1320	Introduction to Geographic Information Systems for Environmental Applications	
EEPS 1330	Global Environmental Remote Sensing	
EEPS 1340	Machine Learning for the Earth and Environment	
HIST 1825J	History of Artificial Intelligence	
MATH 1210	Probability	
MUSC 1210	Seminar in Electronic Music: Real-Time Systems	
PHP 1501	Essentials of Data Analysis	
PHP 1510	Principles of Biostatistics and Data Analysis	

PHP 1560	Using R for Data Analysis	
SOC 1020	Methods of Social Research	
SOC 1100	Introductory Statistics for Social Research	
SOC 1340	Principles and Methods of Geographic Information Systems	
Capstone:		0-1
The required experiential learning component provides students with the opportunity to apply their data-science skills in their concentration, engage in research that uses data science, teach data science as UTAs, or undertake an internship that has a data-science component. The capstone may be completed for credit via an independent study course or not for credit. <sup>2</sup>		
Options for fulfilling this requirement include:		
1. Participate in a Brown University credit experience (i.e. independent study).		
2. Participate in a non-credit experience: summer Internship; TA for data-related course; work with a local organization on a data-related project. A 10-12 page reflective paper is required for this option.		
3. Be a Data Science Fellow. <sup>1</sup>		
<b>Total Credits</b>		<b>4-5</b>

<sup>1</sup> Students may complete DATA 1150 and the concurrent Data Science Fellows project to fulfill both the elective and experiential components of the certificate.

<sup>2</sup> Students must submit a proposal for their experiential component by the end of the sixth semester.

## Data Science Graduate Program

### Master of Science in Data Science

The Data Science Institute at Brown offers a master's program (ScM) that prepares students from a wide range of disciplinary backgrounds for distinctive careers in Data Science. With connections to departments across campus, in particular Brown's Division of Applied Mathematics (<https://appliedmath.brown.edu/>) and Department of Computer Science (<https://cs.brown.edu/>), the master's program offers a unique and rigorous education for people building careers in data science. The program is designed to provide a fundamental understanding of the methods and algorithms of data science, to be achieved through a study of relevant topics in mathematics, statistics, and computer science, including database engineering, visualization, machine learning, and deep learning. The program also provides experience in important, frontline data-science problems in a variety of fields, and introduces students to ethical and societal considerations surrounding data science and its applications.

The program's course structure, including the capstone experience, ensures that students meet the goals of acquiring and integrating foundational knowledge for data science, applying this understanding in relation to specific problems, and appreciating the broader ramifications of data-driven approaches to human activity.

The program can be completed in 12 months (September to August). All students begin the program in September; **there is no option for starting in the spring semester**. Students may elect to complete the program over 16, 21, or 24 months, and most do so. In some cases, exceptionally well-prepared students complete their work in nine months.

The curriculum for the Data Science Master's Program consists of nine credits: eight required courses, one of which is the experiential project course, and one elective. The nine credit-units divide as follows:

- 3 credits in mathematical and statistical foundations
- 3 credits in data and computational science
- 1 credit in societal implications and opportunities
- 1 elective credit to be drawn from a wide range of focused applications or deeper theoretical exploration
- 1 credit capstone experience.

We also offer an option as a 5-th Year Master's Program if you are an undergraduate at Brown. This allows you to substitute maximally 2 credits with courses you have already taken. **5th-Year students must complete the degree in one year.**

DATA 1030	Hands-on Data Science	1
DATA 1050	Data Engineering	1
APMA 1690	Computational Probability and Statistics	1
DATA 2020	Statistical Learning	1
CSCI 1951Z	Fairness in Automated Decision Making	1
CSCI 2470	Deep Learning	1
DATA 2060	Machine Learning: from Theory to Algorithms	1
DATA 2050	Data Science Practicum	1

The practicum experience is a hands-on thesis project that entails an in-depth study of a current problem in data science. Students will synthesize their knowledge of probability and statistics, machine learning, and data and computational science. Students will work in teams on projects with Brown faculty members or with external companies. The project will be completed as part of a course that includes additional career-oriented skills development.

<b>One elective:</b>		<b>1</b>
Domain knowledge relevant to individual interest, 1 credit, must be a graduate level course with 4-digit course number starting with a non-0 digit. Most graduate level CSCI and APMA courses qualify. Please contact the DGS if you plan to take a course from a different department.		

<b>Total Credits</b>		<b>9</b>
----------------------	--	----------

## Data Science Policy, Governance & Society Graduate Program

### Master of Science in Data Science: Policy, Governance & Society

The acceleration of artificial intelligence, machine learning, and advanced analytics is rapidly reshaping industries and societies across the globe. With this immense potential for transformation comes the need to develop robust governance frameworks that ensure data and AI systems are deployed responsibly, ethically, and legally. Brown's online Master's in Data Science: Policy, Governance & Society program prepares students to be leaders and stewards of their organization's data assets and analytics capabilities.

This fully online program brings a world-class education to a globally diverse student body while maintaining the same academic rigor and excellence of an Ivy League University. A flexible format delivered asynchronously gives working professionals global access to graduate education from one of the country's leading universities.

The curriculum for the online ScM in Data Science consists of eight required courses.

For more information about the online Master's Program curriculum and when courses are offered, please visit the department website (<https://dsi.brown.edu/academics/masters-programs/online-masters-data-science/>) or Courses@Brown (<https://cab.brown.edu/>).

DSIO 2000	Technical Foundation for Data Science Success	1
DSIO 2100	Basic AI & Policy Ethics	1
DSIO 2110	Evidence-Driven Policy Making	1
DSIO 2010	Data Engineering in Disguise	1
DSIO 2020	Machine Learning/DL/LLM	1
DSIO 2120	Fairness & Bias	1
DSIO 2130	Advanced Topics of AI Governance	1
DSIO 2030	Applied Learning Experience	1

## Courses

### Data Science

#### DATA 0080. Data, Ethics and Society.

A course on the social, political, and philosophical issues raised by the theory and practice of data science. Explores how data science is transforming not only our sense of science and scientific knowledge, but our sense of ourselves and our communities and our commitments concerning human affairs and institutions generally. Students will examine the field of data science in light of perspectives provided by the philosophy of science and technology, the sociology of knowledge, and science studies, and explore the consequences of data science for life in the first half of the 21st century. Fulfills requirement for Certificate in Data Fluency

Fall	DATA0080	S01	18081	MWF	9:00-9:50(09)	(D. Turner)
------	----------	-----	-------	-----	---------------	-------------

#### DATA 0200. Data Science Fluency.

As data science becomes more visible, are you curious about its unique amalgamation of computer programming, statistics, and visualizing or storytelling? Are you wondering how these areas fit together and what a data scientist does? This course offers all students regardless of background the opportunity for hands-on data science experience, following a data science process from an initial research question, through data analysis, to the storytelling of the data. Along the way, you will learn about the ethical considerations of working with data, and become more aware of societal impacts of data science. Course does not count toward CS concentration requirements.

Spr	DATA0200	S01	26334	TTh	1:00-2:20(08)	'To Be Arranged'
-----	----------	-----	-------	-----	---------------	------------------

#### DATA 0250. Applied Statistics in Python.

As more students engage in data science there is a need to provide guidance on conducting basic statistical analysis in Python. This course will provide a non-specialist approach to applied statistics, specifically linear models Python. Students will learn how to conduct linear modules using the Statsmodels package in Python. Students should have good working knowledge of descriptive statistics (equivalent to a high school AP level). Python coding experience is helpful but not required. Student learning would be assessed through hands-on Python coding activities and written interpretation of statistical reports. Students from the humanities and social sciences are particularly encouraged to enroll in this course. Pre-req- Basic High School A.P. knowledge of statistics.

#### DATA 1010. Probability, Statistics, and Machine Learning.

An introduction to the mathematical methods of data science through a combination of computational exploration, visualization, and theory. Students will learn scientific computing basics, topics in numerical linear algebra, mathematical probability (probability spaces, expectation, conditioning, common distributions, law of large numbers and the central limit theorem), statistics (point estimation, confidence intervals, hypothesis testing, maximum likelihood estimation, density estimation, bootstrapping, and cross-validation), and machine learning (regression, classification, and dimensionality reduction, including neural networks, principal component analysis, and unsupervised learning).

#### DATA 1030. Hands-on Data Science.

Develops all aspects of the machine learning pipeline: data acquisition and cleaning, handling missing data, exploratory data analysis, visualization, feature engineering, modeling, interpretation, presentation in the context of real-world datasets. Fundamental considerations for data analysis are emphasized (the bias-variance tradeoff, training, validation, testing). Classical models and techniques for classification and regression are included (linear and logistic regression with regularization, support vector machines, decision trees, random forests, XGBoost). Uses the Python data science ecosystem (e.g., sklearn, pandas, matplotlib). Prerequisites: Python coding experience suggested. Familiarity with linear algebra and calculus. A course equivalent to CSCI 0050, CSCI 0150 or CSCI 0170 is strongly recommended.

Fall	DATA1030	S01	18611	MW	3:00-4:20(10)	(A. Zsom)
------	----------	-----	-------	----	---------------	-----------

#### DATA 1050. Data Engineering.

The course will cover the storage, retrieval, and management of various types of data and the computing infrastructure (such as various types of databases and data structures) and algorithmic techniques (such as searching and sorting algorithms) and query languages (such as SQL) for interacting with data, both in the context of transaction processing (OLTP) and analytical processing (OLAP). Students will be introduced to measures for evaluating the efficacy of different techniques for interacting with data (such as 'Big-Oh' measure of complexity and the number of I/O operations) and various types of indexes for the efficient retrieval of data. The course will also cover several components of the Hadoop ecosystem for the processing of 'big data.' Additional topics include cloud computing and NoSQL databases. Introduction to concepts and techniques of computer science essential for data science will also be covered.

Fall	DATA1050	S01	18397	TTh	9:00-10:20(05)	(A. Ashraf)
------	----------	-----	-------	-----	----------------	-------------

#### DATA 1080. Large Language Models and Generative AI.

The primary emphasis of this course will be on the technical foundations of large language models (LLMs) and generative AI (Gen AI). The course will cover the use of agents and tools to combine the power of large language models and Gen AI with a wide range of traditional computer algorithms. Additionally, the course will cover the social impacts of this technology and the safety concerns it raises. Students will be required to complete a project to build a multi-agent system in which each agent uses both the power of generative AI and traditional algorithms and in which the collaboration of the agents produces a system which can perform a complex task such as planning a trip or building a predictive model or creating a database or a creating a website. Programming assignments will be done in Python or PyTorch.

#### DATA 1150. Data Science Fellows.

DATA 1150 for juniors and seniors possessing data science skills, seeking to apply these skills and collaborate with faculty to integrate data science content into Brown courses. The course teaches communication, teaching and learning strategies, and determining project requirements. Qualified students have a combination of programming experience (intermediate level or above in R or Python), statistical knowledge (intermediate level or above) and knowledge of how data and computing can be used in applied fields. Students in the data fluency certificate must have DATA 0200 prior to DATA 1150. Students need to complete the application (url below) no later than August 1st for consideration. Qualified students must participate in an interview with the instructor and override requests will be granted only to students by instructor approval. <https://forms.gle/Je3Przs3NDEo4eG9>

Fall	DATA1150	S01	18408	TTh	1:00-2:20(06)	(L. Clark)
------	----------	-----	-------	-----	---------------	------------

#### DATA 1200. Reality Remix - Experimental VR.

This course pursues collaborative experimentation with virtual and augmented reality (AR and VR). The class will work as a team to pursue research (survey of VR/AR experiences, scientific and critical literature review), reconnaissance (identifying VR/AR resources on campus, in Providence and the region), design (VR/AR prototyping). Research findings are documented in a class wiki. The course makes use of Brown Arts Initiative facilities in the Granoff Center where an existing VR laboratory will be expanded through the course of the semester based on student needs. Class culminates in the release the class wiki as a resource for the Brown community.

#### DATA 1340. Machine Learning for the Earth and Environment (EEPS 1340).

Interested students must register for EEPS 1340.

#### DATA 1450. Text Analytics.

This course will first cover techniques for compiling textual corpora from web pages, pdfs, scanned pdfs, images, audio clips, etc. Secondly, it will look at processes for extracting some common types of information from these corpora. In particular, we will cover extracting named entities (persons, locations, organizations, etc.), relations between entities, events, transactions, topics, document summaries, abstracts, legal clauses, etc. This course is different from standard courses in Natural Language Processing and Computational Linguistics in that we will spend significant amount of course time on compiling textual corpora from documents in a variety of formats and our emphasis will be on extracting information that can be fed to analytics pipelines.



**DATA 1491. Fairness in Automated Decision Making (CSCI 1491).**  
Interested students must register for CSCI 1491.

**DATA 1500. Data Visualization & Narrative.**

Data visualization is an essential tool in both discovering and communicating key analytical findings. However, data practitioners and developers can sometimes undervalue the visual polish that goes into creating the most effective graphics. This course will act as a technical primer for building data visualization using code, but a core focus will be the graphic design decisions – color, hierarchy, font selection, labeling – that elevate visualizations. Additional topics will include cartography, web design, and interaction.

Spr DATA1500 S01 26336 Th 4:00-6:30(17) (R. Fischer-Baum)

**DATA 1720. Tackling Climate Change with Machine Learning (EEPS 1720).**

Interested students must register for EEPS 1720.

**DATA 2020. Statistical Learning.**

A modern introduction to inferential methods for regression analysis and statistical learning, with an emphasis on application in practical settings in the context of learning relationships from observed data. Topics will include basics of linear regression, variable selection and dimension reduction, and approaches to nonlinear regression. Extensions to other data structures such as longitudinal data and the fundamentals of causal inference will also be introduced.

**DATA 2040. Deep Learning and Special Topics in Data Science.**

A hands-on introduction to neural networks, reinforcement learning, and related topics. Students will learn the theory of neural networks, including common optimization methods, activation and loss functions, regularization methods, and architectures. Topics include model interpretability, connections to other machine learning models, and computational considerations. Students will analyze a variety of real-world problems and data types, including image and natural language data.

**DATA 2050. Data Science Practicum.**

This course is a requirement for master's students in Data Science and is only open to them. The course includes a semester-long capstone project as well as instruction on topics that prepare students for working as data scientists, such as requirements gathering, version control, bug tracking, software deployment, and other professional development. Capstone projects will be sourced from entities in Brown (departments, labs, researchers, etc.) as well as external entities (companies, nonprofits, etc.) and will engage most of the core components of data science: Data sculpting (data cleaning, formatting, feature selection, etc.), exploratory data analysis, data modeling, and data visualization. Students will also address any social and ethical issues raised by their project. Students will usually work in teams of 3 - 4.

Fall DATA2050 S01 18418 Arranged (A. Ashraf)

**DATA 2060. Machine Learning: from Theory to Algorithms.**

Data science techniques and tools are all around us. Machine learning is a term used across many different disciplines, and often people use machine learning tools without a thorough understanding of how and why the tools work. This course will provide a foundation of machine learning grounded in the mathematical models behind the techniques. We will cover the theory, computational methods, and visualization inherent in the application of machine learning models. Students will learn the statistical learning framework, common assumptions in the data generation process, the mathematics behind machine learning models, including supervised and unsupervised techniques, as well as how to implement machine learning models in Python from scratch. Strong python skills (preferably object-oriented programming) are required. All DSI and CS ScM students will receive overrides as this is a required course for them.

Fall DATA2060 S01 18612 TTh 2:30-3:50(12) (A. Zsom)

**DATA 2080. Data and Society.**

A course on the social, political, and philosophical issues raised by the theory and practice of data science. Explores how data science is transforming not only our sense of science and scientific knowledge, but our sense of ourselves and our communities and our commitments concerning human affairs and institutions generally. Students will examine the field of data science in light of perspectives provided by the philosophy of science and technology, the sociology of knowledge, and science studies, and explore the consequences of data science for life in the first half of the 21st century.

**DATA 2110. Topics in Econometrics.**

This course will begin with a survey of the literature on identification using instrumental variables, including identification bounds, conditional moment restrictions, and control function approaches. The next part of class will cover some of the theoretical foundations of machine learning, including regularization and data-driven choice of tuning parameters. We will discuss in some detail the canonical normal means model, Gaussian process priors, (empirical) Bayes estimation, and reproducing kernel Hilbert space norms. We will finally cover some selected additional topics in machine learning, including (deep) neural nets, text as data (topics models), multi-armed bandits, and data visualization.

**DATA 2980. Research in Data Science.**

Section numbers vary by instructor. Please check Banner for the correct section number and CRN to use when registering for this course.

## Data Science (Online)

**DSIO 2000. Technical Foundation for Data Science Success.**

Focused on practical applications, this course covers the foundational concepts from programming, linear algebra, calculus, probability, and statistics that are most relevant to technical data science skills in the curriculum. At the end of the course, all students will achieve a baseline level of technical knowledge to prepare them for the technical components of the program. Learners will have the opportunity to explore more in-depth topics to discover how foundational topics like derivatives, matrices, and distributions are applied in machine learning and data modeling. Examples are drawn from advanced courses to highlight how these tools can be applied to theories and practice examined in this course.

Fall DSIO2000	S01	18406	Arranged	(L. Clark)
Spr DSIO2000	S01	26338	Arranged	'To Be Arranged'

**DSIO 2010. Data Engineering in Disguise.**

This course introduces students to the core principles of data engineering, emphasizing the often-hidden ethical choices that shape how massive datasets are managed. Students will learn about the fundamentals of data architecture, storage, and processing, while exploring critical values issues such as privacy, biases, data provenance, ownership, and copyright. This course interweaves the ethical considerations with the technical mechanics of data engineering, which exemplifies the real-world choices data engineers make as well as their broader societal implications. By the end of the course, students will understand not just the technical foundations of data engineering, but also the value-laden decisions involved in handling large-scale data.

Spr DSIO2010	S01	26339	Arranged	'To Be Arranged'
--------------	-----	-------	----------	------------------

**DSIO 2100. Basic AI & Policy Ethics.**

This course provides an introduction to the ethical and policy considerations surrounding artificial intelligence (AI) in today's society. Students will explore key ethical concerns, such as data privacy, bias, and accountability, as well as the societal and historical contexts that have shaped current AI governance. In addition, the course will offer a high-level overview of how AI systems are developed, including the basics of data collection, usage, and the training process for machine learning models. By examining these topics, students will gain a foundational understanding of the complex interactions between AI technologies and societal impacts, preparing them for deeper discussions in future courses on AI governance and responsible innovation.

Fall DSIO2100	S01	18424	Arranged	'To Be Arranged'
Spr DSIO2100	S01	26340	Arranged	'To Be Arranged'

**DSIO 2110. Evidence-Driven Policy Making.**

This course explores the role of artificial intelligence (AI) and machine learning (ML) in shaping evidence-driven policy decisions across various sectors. Using case studies from various AI/ML sectors, students will critically examine how AI/ML tools influence policy outcomes. Rather than delving into the technical intricacies of ML, this course emphasizes a "black box" approach, where data inputs lead to predictions. Students will learn to distinguish between prediction and intervention, recognize the limitations of AI/ML, and develop a transparent, precise language for discussing these technologies. By fostering healthy skepticism, this course equips students to make informed decisions about AI's role in evidence-based policy making.

Spr	DSIO2110	S01	26341	Arranged	'To Be Arranged'
-----	----------	-----	-------	----------	------------------