# Course Change Request

## New Course Proposal

Date Submitted: 03/13/26 12:14 pm

## Viewing: COS 503 : Introduction to AI Applications in Science

Last edit: 04/10/26 1:38 pm

Changes proposed by: jbazaz

In Workflow

1. **SC Curriculum Committee**
2. SC Assistant Dean
3. Assoc Provost-Graduate
4. Registrar-Courses
5. Banner

**Are you completing this form on someone else's behalf?**

Yes

**Requestor:**

| Name | Extension | Email |
|---|---|---|
| Arie Croitoru | 5302 | acroitor@gmu.edu |

**Effective Term:** Summer 2026

**Subject Code:** COS - College of Science

**Course Number:** 503

**Bundled Courses:**

**Is this course replacing another course?** No

**Equivalent Courses:**

**Catalog Title:** Introduction to AI Applications in Science

**Banner Title:** Intro to AI Application in Sci

**Will section titles vary by semester?** No

**Credits:** 3

**Schedule Type:** Lecture

**Hours of Lecture or Seminar per week:** 3

**Repeatable:** May only be taken once for credit (NR)
*GRADUATE ONLY*

**Default Grade Mode:** Graduate Regular

**Recommended Prerequisite(s):**

**Recommended Corequisite(s):**

**Required Prerequisite(s) / Corequisite(s) (Updates only):**

**Registrar's Office Use Only - Required Prerequisite(s)/Corequisite(s):**

| And/Or | ( | Course/Test Code | Min Grade/Score | Academic Level | ) | Concurrency? |
|---|---|---|---|---|---|---|
| | | | | | | |

**Registration Restrictions (Updates only):**

**Registrar's Office Use Only - Registration Restrictions:**

**Field(s) of Study:**

**Class(es):**

**Level(s):**

**Degree(s):**

**School(s):**

**Catalog Description:**

This course explores how artificial intelligence (AI) is transforming scientific research across the natural, mathematical, computational, and data sciences through a critical examination of the research literature. Students will review and analyze peer-reviewed studies demonstrating AI applications in these and other scientific fields while developing expertise in evaluating and synthesizing AI-driven research. Each week focuses on a scientific domain, with students learning to identify, synthesize, and critically evaluate the AI methods employed by researchers in that field. The course emphasizes conceptual understanding of AI approaches, ethical considerations in AI-driven research, and the ability to conduct rigorous literature reviews that contextualize AI methods within scientific problems. No programming is required.

**Justification:**

What: Creating a new course.

Why: To offer a graduate course focusing on scientific AI applications.

**Does this course cover material which crosses into another department?**

No

**Learning Outcomes:**

**Will this course be scheduled as a cross-level cross listed section?**

**Attach Syllabus**

COS503_Apr2026.pdf

**Additional Attachments**

**Staffing:**

Dr. Billy Lamberti with guest lectures from COS departmental faculty.

**Relationship to Existing Programs:**

Currently none.

**Relationship to Existing Courses:**

N/A

**Have you reached out to the Libraries to determine whether there are adequate resources to support your course? If not, please email Meg Meiman, Associate University Librarian for Learning, Research, and Engagement at mmeiman2@gmu.edu.**

No

**Additional Comments:**

**Reviewer Comments**

Key: 19283

# Course Proposal: Advanced AI Applications in Science

**Author:** Dr. William Franz Lamberti (wlamber2@gmu.edu)

**Date Created:** 3/4/2026

**Date Last Updated:** 4/10/2026

**Reviewer:** Arie Croitoru

**Date Last Reviewed: 3/9/2026**

**Proposed Course Number:** COS 503

**Course Level:** Graduate

---

## Course Description

This course explores how artificial intelligence and machine learning are transforming scientific research across disciplines. Students will examine real-world AI applications in biology, chemistry, physics, environmental science, neuroscience, forensic science, and other fields while developing practical skills in R programming. Each week focuses on a scientific domain, with students learning relevant AI methods in the context of solving authentic problems from that field. The course emphasizes understanding how and why AI tools are applied in different scientific contexts, ethical considerations in AI-driven research, and hands-on implementation of machine learning workflows using scientific datasets.

---

## Prerequisites

- STAT 250 or STAT 344 , or equivalent (introductory statistics)
- Graduate standing or permission of instructor
- No prior programming experience required (R fundamentals covered in Week 3)

---

## Course Objectives

Upon successful completion of this course, students will be able to:

- Recognize diverse applications of AI across scientific disciplines
- Identify how different scientific fields leverage AI to address domain-specific challenges
- Implement machine learning algorithms in R for scientific data analysis
- Evaluate the appropriateness of AI tools for various scientific problems
- Identify ethical considerations and societal implications of AI use in scientific research
- Preprocess and prepare scientific datasets for machine learning analysis
- Interpret and communicate results from AI models in scientific contexts
- Apply supervised and unsupervised learning methods to real scientific problems
- Critically evaluate AI-driven research in peer-reviewed literature

---

## Required Textbook

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). *An Introduction to Statistical Learning with Applications in R* (2nd ed.). Springer.

**Available free as PDF from:** https://www.statlearning.com/

Additional domain-specific readings will be provided throughout the semester, including peer-reviewed articles demonstrating AI applications in each scientific field.

---

## Software

**R** (freely available) will be used extensively throughout the course.

Students are expected to install R on their personal computers. Installation instructions will be provided during Week 3.

## Common Policies Addendum

Students must abide by the **Common Policies Addendum**.

## Academic Standards: Use of AI Tools

All work submitted in this course must be your own original work. The use of generative AI writing tools, including but not limited to ChatGPT, Claude, Gemini, Copilot, and similar technologies, is prohibited for completing assignments, exams, reflection posts, and presentations unless explicitly authorized by the instructor for a specific activity.

This policy exists because the learning objectives for COS 503 require you to develop foundational understanding of AI concepts, applications, and ethical considerations through your own critical thinking and analysis. Work produced by generative AI does not constitute original work and will not support your learning in this course. Mastering these foundational concepts independently is essential for evaluating AI tools, understanding their limitations, and engaging meaningfully with the scientific applications explored throughout the semester.

Use of AI tools in violation of this policy constitutes a violation of academic standards, and all violations will be reported using the Academic Standards Referral Form. Student work may be analyzed using originality detection tools, including those designed to identify content produced by generative AI.

Note that this policy applies to using AI to complete coursework; the course itself will involve guided exploration and discussion of AI tools as part of the curriculum, which is distinct from using such tools to produce your submitted work. You may be asked to use a generative AI tool in an assignment. Only when told in instructions for a particular question/assignment, may you use generative AI tools. If you are uncertain whether a specific tool or practice is permitted for a particular assignment, consult the instructor before proceeding.

## Course Content and Schedule

### Week 1: Introduction to AI in Scientific Research

- What is artificial intelligence? Definitions and historical context
- Overview of AI subfields: machine learning, deep learning, natural language processing, computer vision
- How AI is transforming scientific discovery across disciplines
- The machine learning workflow in scientific contexts
- Survey of AI applications: from drug discovery to climate modeling

- Course structure, expectations, and assessment overview

## Week 2: Ethics of AI in Science

- Foundational principles of research ethics
- Bias in AI systems: sources, consequences, and mitigation strategies
- Data privacy and consent in AI-driven research
- Reproducibility and transparency in AI-assisted science
- Case studies: ethical failures and successes in scientific AI applications
- The role of institutional oversight (IRB considerations for AI research)
- Responsible AI development and deployment in scientific contexts

## Week 3: Foundations of Scientific Computing in R

**Reading:** ISLR Chapter 2; R Lab sections

- R installation and interface orientation
- Basic R syntax, data types, and data structures
- Exploratory data analysis workflow
- Preparing scientific datasets for analysis

**Lab:** Setting up R environment; exploratory data analysis of a scientific dataset

## Week 4: AI in Environmental and Earth Sciences

**Focus:** Climate and air quality prediction

- How environmental scientists use AI for monitoring and prediction
- Applications: climate modeling, pollution forecasting, ecosystem monitoring
- Remote sensing data analysis with machine learning

- Challenges unique to environmental data (spatial/temporal dependencies, missing data)
- Method Applied: Linear regression

**Homework 1 assigned** (due end of Week 6)

## Week 5: AI in Biology

**Focus:** Species classification and ecological modeling

- How biologists use AI for biodiversity research
- Applications: species identification, habitat modeling, population dynamics
- Image-based species recognition systems

- Conservation applications of machine learning
- Method applied: Logistic regression and classification metrics

## Week 6: AI in Chemistry and Biochemistry

**Focus:** Molecular property prediction and compound classification

- How chemists use AI for molecular discovery
- Applications: drug design, materials science, reaction prediction
- Representing molecules for machine learning
- AI in laboratory automation
- Method applied: Linear Discriminant Analysis and K-Nearest Neighbors

**Homework 1 due**

## Week 7: AI in Physics and Astronomy

**Focus:** Particle detection and experimental data analysis

- How physicists use AI for discovery

- Applications: particle physics, gravitational wave detection, cosmological simulations
- Handling massive experimental datasets

- Signal detection in noisy data
- Method applied: Cross-validation and model evaluation techniques

**Homework 2 assigned** (due end of Week 9)

## Week 8: Midterm Project Presentations

- Student presentations on AI applications in a scientific field of their choice
- Each student selects a domain, identifies a dataset, and applies methods from Weeks 4-7
- Peer feedback and discussion
- Synthesis of applications and methods covered in first half of course

**Midterm Project Presentations** (in-class)

**Homework 2 due**

## Week 9: AI in Neuroscience

**Focus:** Brain imaging analysis and cognitive prediction

- How neuroscientists use AI to understand the brain
- Applications: fMRI analysis, EEG signal processing, brain-computer interfaces
- Challenges of high-dimensional neuroimaging data
- AI in mental health diagnostics
- Method applied: Ridge regression and Lasso for high-dimensional data

## Week 10: AI in Forensic Science

**Focus:** Evidence analysis and pattern recognition

- How forensic scientists use AI for criminal investigation
- Applications: DNA analysis, fingerprint matching, digital forensics
- Legal and ethical considerations unique to forensic AI
- Reliability and admissibility of AI-generated evidence
- Method applied: Decision trees and random forests

**Homework 3 assigned** (due end of Week 12)

## Week 11: AI in Biochemistry and Structural Biology

**Focus:** Protein structure and function prediction

- How structural biologists use AI for molecular understanding
- Applications: protein folding (AlphaFold), drug-target interactions, enzyme function
- The revolution in protein structure prediction
- Sequence-to-function prediction
- Method applied: Support Vector Machines

## Week 12: AI in Geography and Geoinformation Science

**Focus:** Spatial pattern analysis and land use classification

- How geographers use AI for spatial understanding
- Applications: land use mapping, urban planning, natural disaster prediction
- Remote sensing and satellite image analysis
- Integrating spatial data with machine learning
- Method applied: K-means and hierarchical clustering

**Homework 3 due**

## Week 13: AI in Mathematical Sciences

**Focus:** Pattern discovery in high-dimensional data

- How mathematicians and statisticians develop and apply AI methods
- Applications: automated theorem proving, mathematical discovery, statistical inference
- Understanding the structure of complex datasets
- Visualization of high-dimensional scientific data
- Method applied: Principal Components Analysis and dimensionality reduction

**Homework 4 assigned** (due end of Week 14)

## Week 14: Integrative AI Applications and Final Project Workshop

- Combining methods across scientific domains
- Building complete machine learning pipelines for scientific research
- Best practices for reproducible AI-driven science
- Scientific communication of AI results
- Final project workshop: peer feedback on project progress

**Homework 4 due**

## Week 15: Final Project Presentations

- Final project presentations
- Peer evaluation and feedback
- Course synthesis: AI across scientific disciplines
- Emerging trends and future directions
- Resources for continued learning

**Final Project Presentations** (in-class)

---

## Teaching Methods

- Lectures and discussions on AI applications in scientific domains
- Live coding demonstrations in R
- In-class problem-solving activities
- Analysis of peer-reviewed literature demonstrating AI in science
- Individual projects applying AI methods to scientific problems

---

## Assessment

**Core Course (220 points)**

- **Homework Assignments (88 points)**
    - 4 assignments, 22 points each
    - Each assignment focuses on applying methods to a scientific domain
    - Due dates: end of Weeks 6, 9, 12, and 14
- **Midterm Project (44 points)**
    - Individual project applying methods from Weeks 4-7 to a scientific dataset
    - Students select a scientific domain and identify a relevant research question
    - Includes written report (3-5 pages) and brief oral presentation (8-10 minutes)
    - Presented during Week 8
- **Final Project (88 points)**
    - Comprehensive individual project applying multiple methods to a scientific problem

- Students select a dataset from their field of study or from provided options
- Includes written report (8-12 pages) and oral presentation (15-20 minutes)
- Presented during Week 15

---

## Expectations for Graduate Students

Graduate students enrolled in this course are expected to demonstrate mastery at a level appropriate for advanced study. The following additional requirements apply:

- **LaTeX/Overleaf:** All homework assignments and project reports must be prepared using LaTeX (Overleaf recommended).
- **Quality of Work:** Projects should exemplify the quality expected of an academic paper or conference presentation, including:
  - Professional formatting and organization
  - Thorough literature context (minimum 5 peer-reviewed sources for midterm; 10 for final)
  - Rigorous methodology and analysis
  - Clear scientific interpretation of results
- **Presentation Skills:** The presenter is expected to handle complications that may occur during their presentation professionally. Examples of potential complications include, but are not limited to:
  - The most important person that needs to hear the presentation is late. The presenter only has half of their allotted time to present.
  - The electricity goes out on the presenter. The presenter must present without the use of electricity.

- A faculty member is asking challenging questions about the quality of the presenter's work. The presenter needs to professionally respond.

---

## Example Rubric for Midterm Project

| Component | Points | Criteria |
|---|---|---|
| Scientific Context | 8 | Clear problem statement; scientific relevance established; domain appropriately introduced |
| Data Description | 6 | Dataset described; variables explained; exploratory analysis included |
| Analysis | 12 | Appropriate methods applied; implementation correct; results clearly presented |
| Scientific Interpretation | 10 | Results interpreted in scientific context; limitations acknowledged |
| Presentation | 8 | Clear delivery; professional slides; appropriate timing; handles questions well |
| **Total** | **44** | |

---

## Example Rubric for Final Project

| Component | Points | Criteria |
|---|---|---|
| Introduction & Scientific Motivation | 12 | Clear problem statement; scientific relevance established; thorough literature context |
| Data Description | 10 | Dataset thoroughly described; variables explained; comprehensive exploratory analysis |
| Methodology & Analysis | 24 | Multiple appropriate methods applied and compared; implementation correct; validation sound |
| Results & Interpretation | 20 | Results clearly presented; appropriate visualizations; scientific interpretation of findings |
| Discussion & Conclusions | 12 | Limitations acknowledged; implications discussed; future directions identified |

| Component | Points | Criteria |
|---|---|---|
| Presentation | 10 | Clear delivery; professional slides; appropriate timing; handles questions professionally |
| **Total** | **88** | |

---

## Grading Policies

Grades are determined using points. The equation is:

**Grade = (Points Earned + Extra Credit) / (Points Available)**

The table converting points to letter grades is provided below:

| Points Range | Letter Grade |
|---|---|
| 211 - 220 | A+ |
| 198 - 210 | A |
| 193 - 197 | A- |
| 187 - 192 | B+ |
| 176 - 186 | B |
| 171 - 175 | B- |
| 127 - 170 | C |
| 0 - 126 | F |

---

## Relationship to Existing Courses in COS

This course complements existing CDS offerings by providing an applications-focused introduction to AI methods. It differs from methodology-focused courses by emphasizing scientific context and domain applications. Students completing this course will be prepared for more advanced offerings such as:

- **CDS 502 (Computational Science Tools):** Focuses on the technical infrastructure of computational science—programming environments, version control, high-performance computing, and software engineering practices for scientific computing—while COS 501 uses R as a tool to apply machine learning methods to scientific problems without emphasizing computational infrastructure or software development practices.
- **CDS 301/501 (Scientific Information and Data Visualization):** Emphasizes the principles and techniques of visualizing scientific data, including graphical design, interactive visualization, and communicating complex information visually, while COS 501 incorporates visualization only as part of exploratory data analysis and results presentation rather than as a primary learning objective.
- BINF 760 (Machine Learning for Bioinformatics): This course is focused only on bioinformatics and has bioinformatics-specific prerequisites (recommended).
- BINF 761 (Artificial Intelligence and Deep Learning in Bioinformatics): This course is focused only on bioinformatics.
- EVPP 692/991 (Master's Seminar in Environmental Science and Public Policy/Advanced Seminar in Environmental Science): This course is focused only on remote sensing, analysis and effectors both accelerate and enhance conservation research.
- GGS 788 (Deep Learning for Geoinformation): Focuses on machine learning theory and practice as applied to geoinformation only.
- MATH-662 (Mathematics of Machine Learning and Industrial Applications I): Covers mathematical optimization and probability theory in the context of classical learning and gradient based methods. Includes sample applications in satellite imagery, physics, biology and engineering
- MATH-663 (Mathematics of Machine Learning and Industrial Applications II): Covers mathematical and probabilistic models for artificial intelligence. Includes sample applications in computer vision, image processing, and physics
- NEUR 592/BINF 739/BIOL 691 (Special Topics in Neuroscience - Transforming Academia with Generative AI): Covers various applications of Generative AI that range from quantum computing and virtual reality to well-being, with a particular focus on examining the use of AI in academic settings and ethical use of AI.

The course serves students from diverse scientific backgrounds who wish to understand how AI can enhance their research.

---

## Relationship to External Courses at GMU

This course differs from other AI-related courses at GMU in its emphasis on scientific applications across multiple disciplines:

- **AII 600 (Foundations and Practice of Machine Learning for AI):** Emphasizes the theoretical foundations and algorithmic mechanics of machine learning for students building AI systems, while COS 501 introduces the same methods (regression, classification, clustering) in the context of solving authentic scientific problems across disciplines, prioritizing application over algorithmic depth.

- **AII 601 (Planning and Decision Making for Intelligent Agents):** Focuses on autonomous agent behavior, search algorithms, and decision-theoretic planning for robotics and AI systems, while COS 501 does not cover agent-based AI or planning—instead concentrating on data-driven machine learning for scientific analysis.

- **AII 602 (Foundations and Practice of Deep Learning for AI):** Provides in-depth coverage of neural network architectures, backpropagation, and deep learning frameworks, while COS 501 surveys classical ML methods (linear models, trees, SVMs, PCA), keeping the technical barrier accessible for non-programmers.

- **AII 603 (Engineering AI Systems and Pipelines):** Focuses on MLOps, deployment, scalability, and production-level AI infrastructure, while COS 501 emphasizes the research workflow—from scientific question to analysis to interpretation—rather than engineering systems for deployment.

- **ECE 576 (AI: Ethics, Policy, and Society):** Offers a dedicated, policy-oriented examination of AI ethics and governance as a standalone course, while COS 501 integrates ethical considerations throughout each scientific domain (i.e., forensic admissibility, IRB concerns, bias in medical AI) rather than treating ethics as a separate topic.

- **GBUS 662 (Management of IT and The Digital Enterprise):** Addresses business strategy, IT management, and organizational digital transformation, while COS 501 has no business or management component—focusing entirely on scientific research applications.

This proposed course uniquely surveys AI applications across the natural sciences, making it appropriate for graduate students in any College of Science program.

---

## Potential Enrollment

Anticipated enrollment includes 15-25 graduate students per semester. The course is designed to attract:

- Graduate students from any scientific discipline seeking AI literacy
- Students whose research may benefit from machine learning approaches
- Students preparing for careers in data-intensive scientific fields

---

## Justification

Artificial intelligence is rapidly transforming scientific research across all disciplines. Graduate students increasingly encounter AI methods in the literature and may benefit from incorporating these tools into their own research. However, many students lack exposure to how AI is applied in scientific contexts beyond their own field.

This course addresses this gap by surveying AI applications across the College of Science while providing hands-on experience implementing methods in R. By examining how different scientific disciplines leverage AI, students gain perspective on the breadth of applications and can identify opportunities for their own research. The emphasis on scientific context—rather than methods in isolation—ensures that students understand not just how to apply AI tools, but when and why they are appropriate.

The integration of ethics throughout the course prepares students to be responsible practitioners of AI-driven science, aware of both the potential and the pitfalls of these powerful methods.

---