

CS 584: Data Mining

George Mason University

Department of Computer Science

Course title and number	CS 584: Data Mining
Term	Spring 2024
Class times and location	Mon 7:20pm-10:00pm, Innovation Hall Room 105
Piazza	https://piazza.com/gmu/spring2024/cs584002

1. Course Description and Prerequisites

The amount of data available for analysis continues to increase exponentially across a broad range of areas. This leads to the need for development of techniques to discover useful and interesting information from these large collections of data. This course aims to provide an overview of key data mining methods and techniques, including classification, regression, dimension reduction, clustering, association rule mining, recommendation, and text mining. The emphasis will be on developing basic skills for data processing, modeling, prediction, and performance evaluation. Besides, topics about dark side of data mining techniques and applications will also be introduced, like various types of unfairness and bias in different data mining applications.

Formally, you must have received a grade of C or better in CS 310 and STAT 344. Programming experience in Python is preferred, although Java or C will work as well (assignments will use the Python framework). Students should be familiar with probability and statistics concepts, as well as linear algebra. Please expect lots of programming in all the assignments and class projects.

2. Learning Outcomes

- 1) The ability to apply computing principles, probability and statistics relevant to the data mining discipline to analyze data.
- 2) A thorough understanding of model programming with data mining tools, algorithms for estimation, prediction, and pattern discovery.
- 3) The ability to analyze a problem, identifying and defining the computing requirements appropriate to its solution: data collection and preparation, functional requirements, selection of models and prediction algorithms, software, and performance evaluation.
- 4) The ability to understand performance metrics used in the data mining field to interpret the results of applying an algorithm or model, to compare methods and to reach conclusions about data.
- 5) The ability to communicate effectively to an audience the steps and results followed in solving a data mining problem.

The learning outcomes will be assessed based on a combination of homework assignments, exams, projects, and presentations.

3. Instructor Information

Instructor: Ziwei Zhu

Email: zzhu20 at gmu dot edu

Office: Engineering 4609

Office hours: 4:30 pm to 5:30 pm every Monday

Attention: write '[CS 584]' in the subject if the email is related to this course

4. TA Information

TA: Archange Giscard Destine

Email: adestine@gmu.edu

Office hours: 10 am to 11 am every Wednesday (online)

Zoom link:

<https://gmu.zoom.us/j/97206558941?pwd=Mnl1RG1tNHpucTRma21tNXJLVWp1Zz09>

5. Textbook and/or Resource Material

- Pang-Ning Tan, Michael Steinbach, Anuj Karpatne and Vipin Kumar *Introduction to Data Mining (Second Edition)*. Website: <https://www-users.cse.umn.edu/~kumar001/dmbook/index.php>
- Jure Leskovec, Anand Rajaraman, and Jeff Ullman *Mining of Massive Datasets*. Website: <http://www.mmids.org/>

6. Grading Policies

Your final letter grade will be given based on:

Letter grade	Points (out of 100)
A+	98-100
A	94-97
A-	90-93
B+	86-89
B	83-85
B-	80-82
C+	76-79
C	73-75
C-	70-72
D	60-69
F	0-59

The overall course points will be determined using the following weights.

- 1) Six homework assignments: 45%
- 2) In-class quiz: 20%
- 3) Midterm exam: 15%
- 4) Final exam: 20%

Details are as follows:

Homework assignments: the first one is for students to get familiar with Python and submission systems, which accounts for 5%. And each of the rest five assignments accounts for 8%. ***** Late assignments will not be accepted *****

In-class Quiz: there will be 4 in-class quizzes, each will be 20 min, and each accounts for 5% towards the final grade.

Midterm Exam: will be a two-hour closed book exam. (One cheatsheet of standard US letter size paper is allowed).

Final Exam: will be a two-hour closed book exam. (One cheatsheet of standard US letter size paper is allowed)

Attention: A missed exam/quiz cannot be made up. No late days for all assignments are allowed. So, manage your time wisely. All assignments must be performed individually.

7. Course Topics, Calendar of Activities, Project Milestones (subject to change)

Week	Topic	Assignments	Quiz
1 (01/22)	Introduction to Data Mining	HW0 out	
2 (01/29)	Classification Basics and K Nearest Neighbor	HW1 out	
3 (02/05)	Principle Component Analysis	HW0 due	Quiz 1
4 (02/12)	Linear Regression		
5 (02/19)	Logistic Regression and Neural Networks	HW1 due, HW2 out	Quiz 2
6 (02/26)	Deep Learning and Fairness in Machine Learning		Explain Quiz 1&2
7 (03/04)	Spring Break no class		
8 (03/11)	Midterm Exam	HW2 due, HW3 out	
9 (03/18)	Explain midterm exam, Clustering		
10 (03/25)	Clustering		
11 (04/01)	Associate Rule Mining	HW3 due, HW4&5 out	Quiz 3
12 (04/08)	Recommender Systems: Basics		
13 (04/15)	Recommender Systems: Advance		
14 (04/22)	Fairness and Bias in Recommender Systems		Quiz 4
15 (04/29)	Final review	HW4&5 due	Explain Quiz 3&4
16 (TBA)	Final Exam		

8. Academic Integrity and GMU Honor Code

Collaboration in thinking through problems can be highly beneficial and is allowed in this class. However, you may not share or look at any written material (code, answers to problems) that will be part of your or another student's submission. Please make sure you are cognizant of the GMU Honor Code: <https://oai.gmu.edu/mason-honor-code/full-honor-code-document/>. In addition, the CS department has its own Honor Code policies (<https://cs.gmu.edu/resources/honor-code/>) regarding programming assignments. Any deviation from the GMU or the CS department Honor Code is considered a Honor Code violation.

9. Accommodations and Resources for Disabilities

If you have a documented learning disability or other condition that may affect academic performance you should: 1) make sure this documentation is on file with the Office of Disability Services (SUB I, Rm. 222; 993-2474; <http://www.gmu.edu/student/drc> to determine the accommodations you need; and 2) talk with the instructor to discuss your accommodation needs.

10. Campus Closure or Emergency Class Cancellation/Adjustment Policy

If the campus closes, or if a class meeting needs to be canceled or adjusted due to weather or other concern, students should check Piazza for updates on how to continue learning and for information about any changes to events or assignments.