Recommender Systems

Course Syllabus for CS 538 (*Recommender Systems*), Spring 2023

OVERVIEW

This course covers the basic concepts of recommender systems, including personalization algorithms, evaluation tools, and user experiences. We will discuss how recommender systems and user models are deployed in e-commerce sites, social networks, and many other online systems, with readings from current and past research in the field.

COURSE LOGISTICS

Course Title	CS 538: Recommender systems
Credits	3
Schedule	MW 3:00–4:15 in CCP 259
Readings	*Practical Recommender Systems* (textbook)
	Additional research papers and online articles
Prerequisite	CS 533 (*Introduction to Data Science*)
Software	Python, PyData, and LensKit

INSTRUCTOR

Michael D. Ekstrand

E-mail	michaelekstrand@boisestate.edu
Office Hours	By appointment on Zoom or in CCP 354 (use link to schedule)
Home page	https://md.ekstrandom.net

BSU COVID-19 STATEMENT[1]

At Boise State, we're looking forward to another successful semester. This success depends on all of us taking small but necessary steps to keep students, faculty, staff and

[1] This section is required verbatim by the university. It is not how I would choose to write about COVID and our classroom, regardless of the first-person tone.

the local community healthy and safe during the pandemic. For its part, the university has put public health requirements measures in place to support our success. For example, Boise State Public Health staff monitor case rates and perform contract tracing in every class where there may have been a COVID-19 exposure.

As students and faculty, our part of this work is to understand and adhere to Boise State's rules and strategies to maintain public health. If you are experiencing any COVID-19 symptoms, you should stay home, test, and wait for the results before attending in-person campus activities. Layered COVID-19 prevention strategies are outlined on this University's COVID-19 web page.

I encourage you to learn about other precautions you can take to keep yourself and your peers safe; you can find the latest recommendations on the COVID-19 web page. Please note that failing to follow these rules and precautions is a violation of Boise State's Student Code of Conduct and will subject you to university sanctions and discipline.

While facial coverings are not currently required in classrooms and most other campus buildings, students who feel more comfortable wearing a mask are encouraged to do so. The Centers for Disease Control and Prevention continue to view masking as "a critical public health tool" alongside "vaccination, self-testing, and physical distancing," and Boise State recommends the use of facial coverings to provide an extra layer of protection from COVID-19, especially during periods of time or in places associated with high transmission. As the University allows students and faculty to choose whether or not to wear a facial covering in classrooms, we will respect one another's decisions in this regard.

Specific protocols may change during the semester, and additional measures may become necessary as the public health situation evolves. Guidance for such change will come from the Boise State University Office of the President. Additional details regarding guidelines, protocols, etc., may be found on the University's Public Health Response website.

We are taking these and other health precautions so that this course and others can continue to meet in person. Of course, even implementing these strategies can't keep us 100% safe from infection. Because we are breathing common air in the classroom and sitting close together for entire class meetings, all of us remain at risk of potential

exposure to the virus. To preserve this learning community for the duration of the semester, it is imperative that we all engage in behaviors that protect the overall public health.

This course is specifically designed for the in-person [or hybrid] learning environment. The face-to-face format of this course offers a number of benefits that appeal to many students. While I will accommodate the learning needs of students who need to isolate or quarantine, you should expect that this course will continue to meet in person all semester. Should you prefer a course that meets online or remotely, please talk with me, and we will see if another section of this course or a similar one is available.

LEARNING OUTCOMES

If we are all successful in this teaching and learning in this course, by the end you will be able to:

1. Identify and describe a recommender system in practical use.
2. Design, train and evaluate a recommendation algorithm.
3. Understand the work needed to go from a recommendation model to a live system with users.
4. Use data collected from a recommender system to understand user preferences and/or behavior.
5. Read current research on recommender systems, understand what it contributes to knowledge, and apply it to new settings.

RESOURCES AND READINGS

Textbook

Our primary (and required) textbook is *Practical Recommender Systems* by Kim Falk Jørgensen (Manning, 2018).

You may also find the video lectures in the *Recommender Systems* specialization on Coursera a useful addition to your study. The lectures are free; payment is only required if you wish to take the (outdated) Coursera assignments and get a certificate.

Supplemental Books

The following optional textbook may be useful if you wish to dive deeper into some of the topics we cover in a more traditional textbook:

> *Statistical Methods for Recommender Systems* by Agarwal and Chen (Cambridge, ISBN 978-1-13-956586-8; available electronically at the library)

You may also find this survey article useful to read:

> *Collaborative Filtering Recommender Systems* by Ekstrand, Riedl, and Konstan (now publishers; *Foundations and Trends in Human-Computer Interaction* **4**(2)). Available free from my web site.

Finally, the *Recommender Systems Handbook* (edited by Ricci, Rokach, and Shapira, published by Springer) is a good resource for overviews of many topics in recommender systems. The third edition was published in 2022.

Research Readings

Throughout the course, we will be reading and discussing a number of research papers. Links to each paper will be provided in Blackboard. For many class topics, I will be providing research papers to read. I encourage you do at least a first-pass reading of the paper **before** class, to prepare for our discussion. You may often find it useful to revisit the paper after we have discussed it to strengthen your knowledge.

There are three reasons for this. First, it takes time for knowledge to make its way from current papers to textbooks. Second, people building and deploying recommender systems in practice regularly consult recent papers to learn new techniques and apply them to their products. Third, as graduate students you need to be able to engage with (and eventually write) research literature, and this will give you practice.

Software Resources

This course contains several programming assignments in Python using the LensKit library and PyData stack (as taught in CS 533). For more information on programming with PyData, I recommend *Python for Data Analysis*.

The current LensKit documentation can be found at https://lkpy.lenskit.org.

COURSE STRUCTURE

Class Sessions and Attendance

This is an in-person course, but you Zoom into class if necessary. You should attend all class sessions by one means or another if possible. If you need to be absent for some reason, please let me know as soon as you can. I will be recording the class sessions for you to reference when studying for exams, etc., but they will be interactive and watching after the fact is not a substitute for participating in class. Recordings will only be accessible to students enrolled in this class.

I will be wearing an N95 or KN95 mask in class in order to reduce the risk of spreading or contracting COVID. I would prefer if you wear masks as well. If you are feeling ill, please stay home, attending class through Zoom if you are able and viewing the recording if you are not.

This is a relatively small class, so if you are on Zoom, I ask that you turn on video if you are at all comfortable with that. It is easier to have an engaging discussion when we can see each other. I will not penalize keeping your camera off, I just strongly encourage video participation.

Course Components

The work in this course falls into several categories:

- **Assignments** to give you experience building, testing, and analyzing recommender systems.
- A **project** to building a recommender or carrying out a new experiment.
- **Midterm** and **final** exams.

Your final grade will be computed from these components as follows:

Category	*%*
Assignments	35
Midterm	15
Final	15
Project	35

The standard 70/80/90 scale determines the minimum grade you will receive (that is, if you have 80 total course points, you will receive at least a B-).

Assignments

There will be several assignments throughout the semester. Many of these assignments will require you to implement and/or test one or more recommender algorithms using the LensKit toolkit and Python data science / machine learning tools (the PyData stack). Some assignments will require a Jupyter notebook as part of the assignment to present your results. See the CS 533 Notebook Checklist for my formatting expectations.

Each assignment is due at **midnight on Sunday** the end of the week in which it is due.

You may collaborate with classmates on the **concepts** of the assignments and on debugging your code, but must submit your own work. In each assignment submission, list all classmates you collaborated with on that assignment. You may also borrow pieces of code, so long as (1) you clearly describe the origin of any code you did not write yourself in comments, and (2) borrowed code does not comprise more than 45% of the total code in your submission.

Project

The final component of this class is the project, in groups of up to 3. Each student or group will present their work at the end of the semester. More details will be available in the project description released in early February.

Exams

This course has midterm and final exams. Since Zoom participation is available, both exams will be written take-home exams. You will have 48 hours to complete the midterm and 72 hours to complete the final exam.

COURSE POLICIES

Late Work

You have a budget of 6 **late days** to use at your discretion throughout the semester. Each day extends a project or assignment deadline by 24 hours, no questions asked.

Conduct

I expect you to behave in a civil, respectful manner in all class interactions, both in official meetings such as lectures and out-of-classroom activities such as project group meetings and study sessions, and to contribute to a constructive learning environment.

The Recurse Center Social Rules are a good source of guidance on how to maintain a constructive and educational environment.

Disability Accommodations

If you need particular accommodations to be able to fully participate in this course, please talk with me as soon as possible. If you have documentation from Office of Disability Services for particular accommodations, please bring it, but I am happy to discuss with you anything needed for you to fully participate in the class.

COVID Accommodations

Any of us may encounter disruptions this semester either directly or indirectly due to the current state of public health and the challenges of working and learning remotely. This of course may include contracting COVID (although I hope that does not happen to any of us), but there are also many indirect effects such as health problems with roommates or family members, unexpected changes in childcare situations, etc.

My goal is to ensure you have a path to successfully completing the semester and achieving the course learning outcomes. If something arises, whether directly related to COVID or not, that affects your ability complete the assigned work on schedule, please talk with me as soon as practical to work out a plan. Because this is a smaller class, I am not making particular COVID policies (aside from the overall course structure), but want to work with you to make this semester a valuable remote learning experience.

Office Hours

My office hours this term are by appointment, although most days I am usually also available for a little while after class to answer questions in person. The logistics section of the syllabus and Canvas have a link to my Calendly page, where you can schedule an office hours appointment either in-person or via Zoom. It will create an appointment on my calendar (with a Zoom link if needed) and invite you to it.

S C H E D U L E

Following is a tentative schedule, that may be adjusted. **Bold** items are key dates for the project and exams.

Week	Date	Topic	Assignment
1	1/9	Fundamental Elements of Recommendation	A0
2	1/16	Recommendation Data *No class Monday 1/16 - MLK*	
3	1/23	Introductory Approaches	A1
4	1/30	Vectors and Neighborhoods	
5	2/6	Evaluating Recommendations	
6	2/13	Matrices, Embeddings, and Optimization	A2
7	2/20	*No class Mon – Presidents' Day*; **Midterm Exam**	
8	2/27	Building an Algorithm	
9	3/6		
10	3/13	Recommendation Tasks	
B	3/20	*Spring Break*	A3
11	3/27	Presenting and Explaining Recommendations	
12	4/3	User Response to Recommendations	
13	4/10	Social and Ecosystem Concerns	
14	4/17	Open Directions	*Draft Report*
15	4/24	**Project Presentations**	
F	5/1	**Final Exam**	*Final Report*

Assignments (A*n*) are due **on Sunday at midnight** the end of the week they are listed. Due dates may be adjusted later depending on how the semester goes.