Recommender Systems

Course Syllabus for CS 538 (*Recommender Systems*), Spring 2021

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 10:30–11:15 on Zoom
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
Home page	https://md.ekstrandom.net

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). If you are a member of the Association for Computing Machinery (a student membership costs $19/year), you can access the book for free via the ACM's subscription to the O'Reilly Learning Platform.

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 The lectures are free; payment is only required if you wish to take the (outdated) Coursera assignments and get a certificate.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. I expect the third edition to come out in mid- to late-2021.

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* or *Python Data Science Essentials*.

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

COURSE STRUCTURE

Class Sessions and Attendance

Class will be delivered through live Zoom lectures. You should attend all class sessions 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 is not a substitute for participating in class. Recordings will only be accessible to students enrolled in this class.

This is a relatively small class, so 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 notebook formatting expectations.

Each assignment is due at **midnight on Sunday** at 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, but clearly describe the origin of any code you did not write yourself in comments.

Project

The final component of this class is the project; I encourage you to work with a partner on the project. Each student or group will present their work at the end of the semester. More details will be available in the project description, which I will release in early February.

Exams

This course has midterm and final exams. Since the course is delivered remotely, 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 may submit **assignments** late with a 10% penalty for each day late. Project deliverables are due their stated due dates.

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

Due to the remote nature of the class and our work in general this semester, office hours are by appointment only. The logistics section of the syllabus and Blackboard have a link to my Calendly page, where you can schedule an office hours appointment. It will create an appointment on my calendar with a Zoom link, and invite you to it.

SCHEDULE

Following is an approximate schedule. I have kept the last part of the semester open for discussing a variety of topics from current literature, and I may adjust other parts of the schedule as we progress through the semester. **Bold** items are key dates for the project and exams.

Week	Date	Topic	Assignment
1	1/11	Fundamental Elements of Recommendation	A0
2	1/18	Recommendation Data and Evaluation	
3	1/25	Neighborhood-Based Recommendation	A1
4	2/1	Learning about Users	
5	2/8	Matrix Factorization and Embeddings	
6	2/15	Learning about Items	A2
7	2/22	Composing Recommendation Algorithms	
		Midterm Exam	
8	3/1	Learning to Rank	
9	3/8	Presenting and Explaining Recommendations	
10	3/15	User Response to Recommendations	A4
11	3/22		
12	3/29		
13	4/5		
B	4/12	*Spring Break*	
14	4/19		*Draft Report*
15	4/26	**Project Presentations**	
F	5/3	**Final Exam** (Wed. May 6 2:30-4:30 PM)	*Final Report*

Assignments (A*n*) are due **on Sunday at midnight** the end of the week they are listed.