SYLLABUS FOR Recommender Systems

What: CSci 5123, Recommender Systems
Instructor: Joseph Konstan
When: [TBD]
Where: [TBD]
Office hours: [TBD]

Catalog Description:
An overview of recommender systems, including content-based and collaborative algorithm for recommendation, programming of recommender systems, and evaluation and metrics for recommender systems.

(3 cr; prereq: Java programming and 2033 and 3081, or instructor consent; Student option; Offered Fall Odd Year)

This course is an introduction to recommender systems. The goal for the course is for you to learn:

- The basic concepts behind recommender systems--software tools to recommend or evaluate products, information, or other options for users
- The use of recommender systems in practice
- Techniques for making recommendations, including non-personalized, content-based, and collaborative recommendation techniques
- The details of algorithms for product-association, keyword-based information filtering profiles, user-user collaborative filtering, item-item collaborative filtering, and dimensionality-reduction collaborative filtering
- How to evaluate recommender systems, including a variety of metrics and the strengths, weaknesses, and use-cases for each
- How to program recommender systems using the LensKit open source toolkit
- Advanced concepts and current research in recommender systems

Course Materials
Option 1: This course may be offered using a traditional textbook on recommender systems such as Jannach and Zanker's *Recommender Systems: An Introduction*, or Aggarwal's *Recommender Systems: The Textbook* . These would be supplemented by class readings, software documentation, and other materials.

Option 2: This course may also be offered using material recorded for the University of Minnesota's online 4-course sequence on "Recommender Systems" offered through Coursera. This online sequence has approximately 4 hours of lecture material for each week of the course. When this option is used, contact time will be split between flipped classroom project exercise time and supplemental lecture and discussion time.

In either option, there will be seven non-programming assignments and five programming assignments completed during the course (and graded by course TAs).

Prerequisites and Expected Knowledge
This course requires Java programming and the use of matrix techniques common in linear algebra. If you do not have these skills, it will be very difficult to pass this course. You should have completed CSci 2033 and CSci 3081W or have equivalent skills.

Course Assignments, Exams, and Grading
The course has three types of deliverables from students:
1. Non-Programming Assignments. These seven assignments are designed to help teach you the fundamentals behind algorithms and metrics. They will involve submitting spreadsheets or specific answers.

2. Programming Assignments. Programming assignments will use the open source LensKit toolkit and the Java programming language. They are designed to help you learn how to implement and experiment with recommender systems. Assignments may be completed individually or in pairs; if completed in pairs, please make sure both students' name and ID number is included in the submission.

3. Exams. The course has four short exams (approx. 45 minutes each) to assess your knowledge of the course materials. Exams will be administered in class.

Late Assignments
Late assignments are due before the start of class on the due date. Students may late on assignments up to twice in the course without penalty. The late assignments must be completed within one week of the original due date. All assignments due on a particular date count as part of a single excused lateness. Assignments submitted more than one week late will receive no credit. The third (and subsequent) late submissions will be penalized by 30%.

Course Grades
Course grades will be determined based 40% on exams (10% each), 30% on programming assignments (6% each), and (30% on written assignments 4-5% each). The course is not graded on a curve; students who earn 93.5% or better will earn an A; 90-93.49 will earn an A-; 86.5 - 89.99 will earn a B+; and so forth. If this results in unfairly low grades, I will consider reducing the thresholds, but not increasing them. I hope all students earn an A in the course.

There is no final exam in this course.

Incompletes
Incompletes are only awarded in very rare circumstances when an unforeseeable event causes a student who has completed all coursework to date to be unable to complete a small portion of work (typically an assignment).Incompletes will not be awarded for foreseeable events including a heavy courseload and poorer-than-expected performance on assignments.

Scholastic Conduct
All work submitted for this course is expected to be your original work. Work copied or derived from other sources should credit these sources appropriately. Your assignments may be jointly submitted to another course with the approval of the other instructor (they will still be judged on CSci 5123 criteria for this course).

Special Circumstances
Students with special needs or circumstances should contact me as soon as possible to make any necessary arrangements. As with incompletes, extensions are only granted for unforeseeable events, but arrangements may be made to obtain materials and submit work in advance if needed. Other accommodations may be arranged in cooperation with disability services.

Everything else that is required
I'm required to tell you a lot of generic things. I do so by providing a link to a set of syllabus statements you should review: https://policy.umn.edu/education/syllabusrequirements-appa.

Course topics and tentative due dates:
Week 1: Intro to recommender systems and to course.
Weeks 2-3: Non-Personalized Recommenders (assignments due Week 3).
Weeks 4-5: Content-Based Recommenders (assignments due Week 5).
Weeks 6-7: User-Based Collaborative Filtering Recommenders (assignments due week 7)
Weeks 8-9: Item-Based Collaborative Filtering Recommenders (assignments due week 8)
Weeks 10-11: Metrics and Evaluation (assignments due week 11)
Weeks 12-13: Matrix Factorization and Advanced Techniques (assignments due week 14)
Week 14: Advanced topics, course conclusion