

Building a Better Course Recommender Tool

Using data to help students select courses



UNIVERSITY OF MICHIGAN

Abstract

Despite the vast amount of course registration and evaluation data available, there is relatively little being done to leverage that information to help current students figure out what courses to take and when best to take them. Our goal for this project was to design a student-facing course recommender tool to facilitate social navigation within this space.

The Google Doc

Aside from word-of-mouth, one of the only resources we had as graduate students in the School of Information for tips on what courses to take was a Google Doc with comments made by former students. Inspired by the surprising effectiveness of this simple document, we wanted to see if we could create a better, more permanent solution using past course registration history and other sources of relevant data.

SI 622 - Usab Eval & Needs *****

- Very important class for HCI. One of the fundamental courses. Introduces students to various methods in UX and User research. Usually is a client project.
- Good class, but be warned it is a lot of work, basically a 501 research report every two weeks or so. You do learn a lot and you get PEP credit.
- Yes, a lot of work, but definitely worth it ++
- This is a must-take course if you are a first-year HCI student. HCI students often say that this course was most important to getting them their internship. ---YES+++
- I would recommend this even for non-HCI people. The LIS specialization and I learned a lot in the course and really enjoyed it, even though the workload felt overwhelming at times. It counts for the methods requirement, so if you still need to fulfill that, definitely consider this class. Yes, it's 501-style, but much, much better!
- I really disliked this course, but mainly because I had group trouble. The skills could be helpful, though.

SI 500 - Info in Soc Systems - - - - -

- This too shall pass.
- Yes.
- Can we - this class or would it just be splinter? ++
- Honestly found the case study useful - did mine at work and turned it into a job opportunity (improving an existing system). Just saving, take the opportunities you have.
 - I did like that we could choose a case study around our own interests.
 - The case study was okay (not super enjoyable, but not the worst), as was most of the content. It was the disorganization of the course that really got to me ++
 - Agreed with this comment about making the best of an opportunity. I did mine on syllabi and it changed how I write mine for the better.
- Just as a note, the way 500 is taught has basically changed every single time.
- This! You probably won't have a case study. You probably won't have the same topic, hopefully, if they're mixing up who's teaching it, hopefully it will be more enjoyable.

The First Attempt

Our first attempt at creating a recommender system had mixed results. Using course registration data from the School of Information, we allowed users to input classes they had already taken and recommend additional courses based on what other students with similar schedules had previously taken. While there would occasionally be nuggets of interesting suggestions, the system would also often recommend unfulfilled required courses, which was not very useful. Another problem we did not expect was getting recommendations for courses that no longer existed because we had failed to consider that courses and requirements changed over time (or at least more than we expected).

Course Recommender System Beta
University of Michigan: School of Information

CourseRecommender uses the past three years of UMSI course registration data to recommend you classes based on what students have taken in the past. The system will filter and rank results using other students who have taken similar courses to your own. As this is still a beta version of the system, some new classes will not be in the system and some old classes that are no longer offered may be displayed.

Please enter course numbers to select classes you have taken

Add Classes you have taken:

Selected Classes

SI 500
(Removed)

Information in Social Systems: Collecting Flows and Processing --- The core properties of information, of people, and of technologies, create constraints and opportunities for analysis, design and management. This course introduces students to those core properties and their implications. In addition, the course introduces a perspective on the central responsibilities of professionals who bring information, people and technology together in more valuable ways.

SI 582
(Removed)

Introduction to Interaction Design --- Intro to Interaction Design will provide students with a hands-on introduction to interaction design. The course will focus on design methods and design thinking, and will allow students to develop their design sensibilities and practical skills through a series of

Recommend me some classes!

Recommended Courses



SI 612 - Pervasive Interact Des
(Removed)

Pervasive Interaction Design --- This course is advanced interaction design course that focuses on designing interactive applications for emerging mobile and context-aware technologies. It follows a similar format to 682 (Interaction Design) in that course is studio-based, consists largely of a seminar-long group project with multiple milestones, and is supplemented by readings and discussion relevant to the topic. Class meetings consist of brief lectures to introduce and frame course concepts, design methods, and technical topics; discussion of readings; in-class design exercises; and group work/lab time. Some programming will be required in order to complete prototyping activities, though students have flexibility in choosing the platforms and languages used. Programming will not be taught in the course.

SI 643 - Prof Prac Inf Centr
(Removed)

Exploring the Data


We wanted get a better idea of what kind of data we were working with, so we created several visualizations to help us with this. The School of Information registration data that we had spanned a total of 18 years from 1995-2013. During this time, enrollment has grown from tens of students at its inception to over 400 in the present day, with numerous specializations being added over time as well.



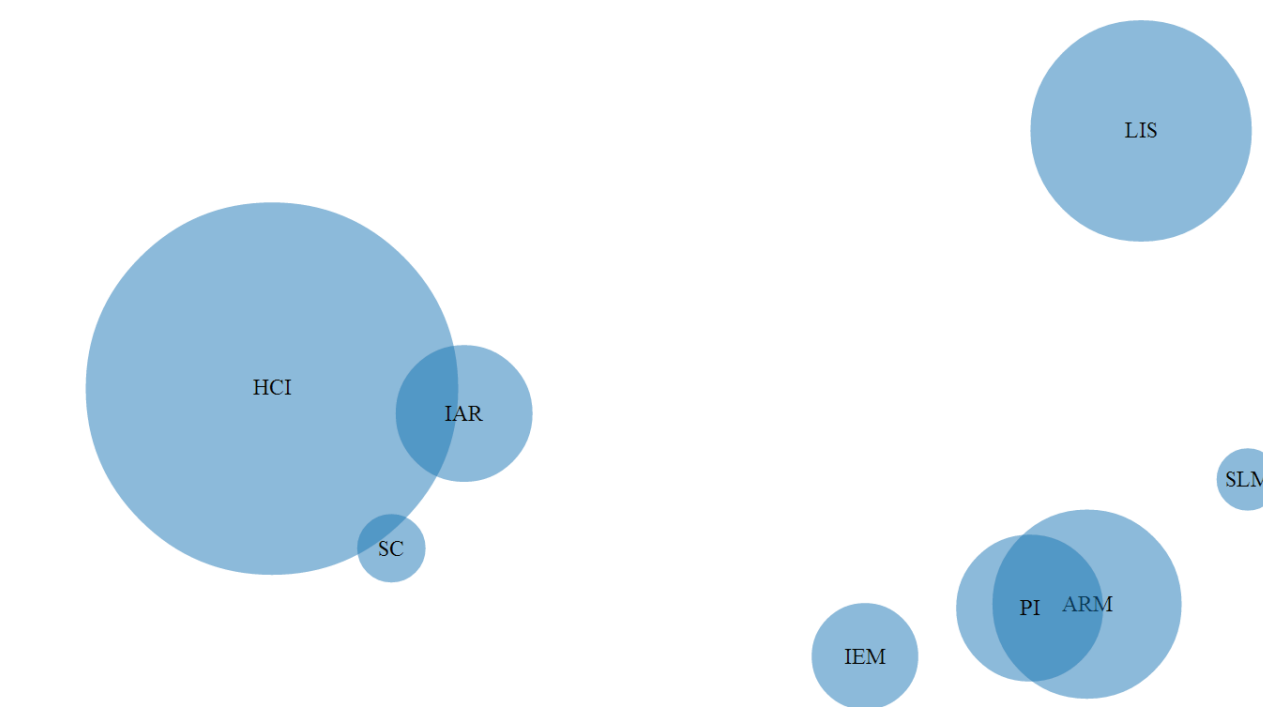
Bubbles representing individual courses sized by total number of enrolled students.

- | | |
|---|---|
| INFO Information | LIS Library and Information Services |
| HCI Human Computer Interaction | SLM School Library Media |
| IAR Information Analysis and Retrieval | ARM Archives and Records Management |
| SC Social Computing | PI Preservation of Information |
| ICD Incentive Centered Design | CI Community Informatics |
| IEM Information Economics for Management | Tailored Tailored |
| IPOL Information Policy | |

Special thanks to Kevin Gao, Tiffany Liu, Chloe Ng, & Krishna Vadrevu for their work in helping to build this visualization



To further investigate how course enrollment changes over time, we measured differences in the number of students per course from year to year (normalized by the total number of students enrolled). Aside from the first few years when the school was rapidly expanding, we found that things change a little bit, but not much, from one year to the next. However, over many years, these differences do add up. In addition, there are occasionally large spikes when major changes are made (e.g. 2007-2008 when a new mandatory course was introduced).



To visualize differences between the various specializations within the School of Information, we compared common degree requirements and used multidimensional scaling to display them in two-dimensional space. Specializations that are closer together are more similar (i.e. share more requirements with each other). Nodes are sized by the current number of students belonging to that specialization.

Conclusion and Next Steps

There are many sources of data that could potentially be incorporated into this system, however, many of these sources are scattered and in a form not easily accessible or usable to students. One of the challenges of this project was gathering and consolidating lots of data from many different sources (e.g. course registration data from the Data Warehouse, course catalog data using the UM Developer API, etc.).

In addition to course registration and requirements data, we are also interested in incorporating course evaluation data to provide better recommendations. This would give the system much more power to take into account particular student interests and individual preferences. We realize that this information can be quite sensitive, so we are still figuring out how best to utilize and implement this.

The prototype we are developing will only use School of Information data, which is a decision we made due to time constraints and because we are most familiar with this data. We hope that in the future, all students will be able to use a system like ours find out everything they would want to know all in one place - how many people took a class, when they took it, and how much they liked it - in order to painlessly figure out the best courses to take.