CISC889 Empirical Methods for Computer Science, Fall 2015

Time:	MW 3:35–4:50pm
Location:	307 ISE Lab
Web page:	Sakai (`http://sakai.udel.edu`)
Professor:	Ben Carterette (carteret@udel.edu)
Office:	440 Smith Hall
Office hours:	TR 3:30–4:30 or by appointment

The course web page will be the primary source of information about the course. There you will find readings, assignments, a course schedule, slides, and additional resources. Please check it often.

Course Description
Computer science is increasingly an empirical discipline, with solutions to problems motivated by many factors that can be measured but are difficult to reason about. Thus like any other empirical discipline, doing good computer science requires an understanding of experimental design, statistical analysis, data analysis, hypothesis generation, modeling and simulating phenomenon, and synthesizing results.

Unlike many other empirical disciplines, computer scientists can generate vast amounts of data at very low cost; the challenge is not in obtaining observations but in performing the proper analysis and ensuring that conclusions do not overstep the limitations of either the analysis or the experiment design.

Textbooks: (not required)
Empirical Methods for Artificial Intelligence, Paul Cohen.
Introductory Statistics with R, Peter Dalgaard
Experimental and Quasi-Experimental Designs, Shadish, Cook & Campbell.

Tentative schedule: Subject to change!

Week starting	Topics		Assignments
Aug. 31	introduction	– short week –	
Sep. 7	what is science?	– short week (holiday) –	
Sep. 14	selecting projects		report: project ideas
Sep. 21	project lab		report: project proposal
Sep. 28	–	– no class (prof. travel) –	review: project proposal
Oct. 5	causality/data analysis		
Oct. 12	measurement/variability	– short week (prof. travel) –	report: data analysis
Oct. 19	hypotheses		review: data analysis
Oct. 26	experimental design		
Nov. 2	experiment design lab		report: experiment
Nov. 9	stat hypothesis testing		review: experiment
Nov. 16	experimental validity 1	– short week (prof. travel) –	
Nov. 23	experimental validity 2	– short week (holiday)–	
Nov. 30	advanced modeling		report: threats to validity
Dec. 7	last week of class		review: threats to validity
			report: final project report

Total number of class sessions: 23 over 15 weeks; readings will be assigned every week except the 1st, 4th, 5th, and 10th.

Prerequisites: None, but it is good to be doing research with an advisor already. If you are not, that is OK.

Project: The project will involve selecting a research question and seeing it through to a final report including data analysis, formulation of a precise (interesting) hypothesis, an experiment to test it, results from the experiment, and possible threats to the validity of the experiment.

There will be six written project reports to turn in over the course of the semester:

1. project ideas: 1–2 pages
2. project proposal: 1–2 pages
3. data analysis: about 5 pages
4. hypothesis & experiment design: 2–4 pages
5. threats to validity: 2–4 pages
6. final report: 5–10 pages, summarizing material from previous reports and including final experimental results

In addition, you will peer-review submissions from classmates. Reports numbered 2, 3, 4, and 5 will be reviewed by two classmates each.

Grading: You will graded on responses to assigned readings, class participation, research project reports, and your reviews of classmates' reports.

- 20%—responses to readings.
- 20%—discussion/participation.
- 40%—research project reports.
- 20%—reviews of submitted reports.

All work must be done as an individual; group work is not part of this class.

Reading responses will be given a letter grade based on your critique, with an A awarded for a good critique containing multiple original thoughts, to a C awarded for a submission that contains virtually no critique at all. Three readings can be skipped with no penalty to your final grade.

Project reports will be given a letter grade based on how well you have satisfied the report requirements. Peer reviews *will* influence the grades awarded on reports, but *will not* be the final arbiter of report grades.

Peer reviews will be given a letter grade based on the quality of your critique.

Lateness Policy: Reading responses must be turned in by the due date given on Sakai (usually 1pm on the day of class). Reading responses will *not* be accepted late under any circumstances.

Reports and responses will be penalized 10% (one letter grade) for each day late.

Academic Honesty Policy: All work you turn in must be your own. Students caught submitting someone else's work (either another student's or answers found on the web) will receive zero credit and may be prosecuted according to the University guidelines.