

600.408
Empirical Research Methods in Computer Science
Fall 2005

Noah Smith and David Smith
<http://nlp.cs.jhu.edu/~nasmith/erm>

1 Official Course Description

Computer programs are real-world processes that, like complex physical and biological systems, can be studied in controlled experiments and analyzed statistically. This course gives a rigorous grounding in empirical methods for students interested in any area of applied computer science. Topics: experimental design, probabilistic modeling, exploratory data analysis, hypothesis testing, and system tuning.

2 Logistics

The course meets 4-5:15 pm Wednesdays from October 12, 2005, through December 7, 2005, in Shaffer 300. There will be no class on November 23.

3 Instructors

Noah Smith (nasmith@cs.jhu.edu) and David Smith (dasmith@cs.jhu.edu), two PhD students from the Computer Science department. Our office is 332 New Engineering Building. Office hours by appointment.

4 Supplementary Text

Paul R. Cohen, *Empirical Methods for Artificial Intelligence*. MIT Press, 1995. (All readings will be provided.)

5 Grading Policy

Grades will be assigned as follows.

- 65% individual assignments (between four and six)
- 30% final exam
- 5% class participation

All work is to be completed **individually** unless otherwise noted. You are free and encouraged to discuss course content outside of class, but all work handed in must be your own. We implore you to stay on top of the course and turn in all assignments on time. Late assignments will be penalized according to a (ruthless) stochastic function, in direct correlation to lateness.

6 Topics

This outline is tentative. If there is a particular topic you'd like us to cover, let us know.

- Lecture 1: introduction, review of statistics, hypothesis testing, sampling
- Lecture 2: statistics of interest: means, quantiles, variance
- Lectures 3–4: experiments with runtime and “space”
- Lecture 5: exploratory data analysis
- Lecture 6: parametric modeling, regression, and classification
- Lecture 7: statistical debugging and profiling
- Lecture 8: summary and review