6.883: Program Analysis

Syllabus

Course website:
`http://pag.csail.mit.edu/6.883/`

Class meetings:
TR 11-12:30, 24-407

Lecturer:
Michael Ernst
32-G718
(617) 253-0945
`mernst@csail.mit.edu`
`http://pag.csail.mit.edu/~mernst/`
Office hours: by appointment. While there are no set office hours, I welcome discussions with students and encourage you to contact me if you wish to meet.

1 Structure of the course

6.883 is a graduate paper-reading seminar. Each class session will begin with a brief discussion and presentation of material, after (or during) which the floor will be open to rebuttals, discussion of related work, criticism, brainstorming about follow-on research, etc. At this level in your career, you should no longer be a passive listener to lectures but an active participant in the discussion.

To help you prepare, you will write a one-paragraph commentary on each paper no later than 9am on the day *before* the class meets to discuss the paper. You will post your commentary to the course webpage for viewing by the instructor and by other students. The commentary should reflect your understanding and analysis of the issues raised by the paper, and should also help direct (both your and others') preparation for in-class discussion.

You may write the commentary in whatever style you prefer that meets the goals listed above. One good format for the commentary is to critique the paper, listing the following three points: its biggest contribution (and, briefly, why that result was not already obvious), its biggest mistake (in motivation, methodology, algorithm, data analysis, conclusions, or some other area), and the biggest question that it raises (or the most interesting and important follow-on work that it suggests). Another acceptable format is to summarize the paper, describing its thesis, approach, and conclusions, and stating why it is significant. The commentary should also list questions that you have about the paper, such as about technical points or connections to related work.

If you read the paper and there are issues you do not understand, then by all means ask questions about them and we will all gain by the discussion. Remember, if you have a question, it is likely that many other people have the same question but are too shy (or vain, or insecure) to ask it; they will appreciate your raising the point. However, do come to class prepared: carefully read the paper and get as much as you can out of it on your own. Doing so will make the class time that much more productive.

Each student will present one paper during the semester. (The presentations will be done by pairs of students if there are enough students registered.) The lecturer will present the other papers. Presenters will meet with the lecturer a week before the class meeting to receive feedback and improve their presentation.

2 Projects

The centerpiece of the course is a semester-long project, done in groups of 2–4 students. The project should make a research contribution in the area of tool support for programming tasks. As a research contribution, it should answer a question whose answer no one previously knew.

A non-exhaustive list of types of acceptable projects are:

- proposing and evaluating a fundamental new technique
- developing and assessing new algorithms to replace currently-used ones
- translating a methodology to a new problem domain
- applying known techniques to new problem domains, such as operating systems, networks, embedded systems, security, biology, aerospace, etc.
- evaluation of existing techniques or tools, via case studies or controlled experiments
- port an existing tool to a new domain (e.g., a new programming language or a new version of one)

A list of specific potential projects will be distributed at the second class meeting. (It is intentionally not being distributed until after you have completed Assignment 1.) You may select (or modify) a project from among the list of suggested ones, from your answers to Assignment 1, from other students' answers to Assignment 1, or from other sources. The lecturer will help you choose and refine a project.

For instance, you might start from a programming problem that has vexed you and devise a solution (or demonstrate that it is a broader problem than previously recognized). You might take a paper of particular interest to you and attack one of the problems listed in its future work section. You might take a promising paper that lacks experimental evaluation, or whose methodology is flawed, and perform a proper experimental evaluation.

You should also consider building on your strengths and finding ways to leverage other work that you are doing. For instance, if you do work in AI, you might consider how machine learning could be applied to problems of program comprehension or to filtering output from program analysis tools. You also might consider using a project that you are working on, or that you worked on in the past, as a test case for a tool. Synergies with your other work are welcomed!

Your final project report should be written in the style of a conference paper, and you will present it in a 6.883 conference at the end of the semester. (You are not required to submit your work to a conference or workshop, but you might choose to do so, and this is the standard of work expected.) You need not close the book on the area of research you choose, but you should do a good job of learning and clearly communicating new knowledge in that area. Do not feel overwhelmed by the final project or final report. The instructor will meet with you repeatedly for discussions, feedback, and assistance on all aspects of your project. For tips about writing a technical paper, you should read (among other resources) `http://pag.csail.mit.edu/~mernst/advice/write-technical-paper.html`.

2.1 Timetable

- By September 29: select a group, select a project, write a short (1–2 pages) project proposal that includes its thesis and technical approach, and meet with the instructor for feedback on and approval of your project proposal.
- By October 27: submit a version of your report that includes the introduction, description of the technical approach (e.g., algorithmic details), related work, and experimental methodology (i.e., evaluation plan). The evaluation plan should indicate what subject programs you will use and what metrics you will measure. It should explain why those are the right metrics, and what they reveal about your work. (For example, do they address the novel part of your technique, or are they end-to-end measures? Can they be directly compared to previous work, or not?) It should also indicate your criteria for success. (You should lay those out now, so that you have an objective measure when the research is complete.) Include the division of labor among members of your group. Do not stint on this component of the project: a prototype implementation and draft report is due in just three weeks.
- By November 17: submit a preliminary version of the report, including at least partial results. In particular, you should have a prototype implementation that can run end-to-end on a small example program. (For example, your compiler should be able to compile the factorial program.) Having a working implementation now will enable you to enhance it and to run experiments during the next several weeks; if you do not even yet have working code yet, you are unlikely to be able to complete a quality project in time.
- Beginning of December: present your project in class. You should plan to talk for 20–25 minutes, followed by 5–10 minutes for questions though the questions may come during rather than after the talk. As with any talk, be sure to practice it before you present in class. For more tips, see `http://pag.csail.mit.edu/~mernst/advice/giving-talk.html`.

- By December 14 (last day of MIT classes): submit your final report.

3 Grades

Grades will be assigned based on the project (this is the largest factor), class participation (during class meeting) and occasional homework. As this is a graduate class, the class is likely to be A-centered, but students are not guaranteed a grade of A (or even a passing grade).