FINAL PROJECT INSTRUCTIONS

MAS S.62: MIT Fall 2013

SCOPE & REQUIREMENTS

You will complete a final project on a topic related to interactive machine learning.

Requirements:

1) There should be significant attention to **supporting useful and effective interactions** between human(s) and machine learner(s) in the chosen context. That is, although you are free to focus your work on developing or improving upon a machine-learning algorithm, your project should be strongly motivated by an interactive application.

2) Each project should include a **substantially new contribution** to research and/or practice in interactive machine learning. **Please read** http://chi2012.acm.org/cfp-contribution-types.shtml **for descriptions of different types of research contributions**, and discussion of what makes a good paper for each contribution type.

3) Your project should involve a system that fits the strict definition of interactive machine learning used in this course. If you need refreshing, you'll find it in the second section of the Power to the People draft. If you are in doubt about whether a potential system would count as interactive machine learning, ask me.

A note on implementation projects: Don't just build something that has been built before. Don't just build something because you think it's "cool"—build it because it addresses a real need, shows that some technique is possible, enables you to learn something about whether and how some idea works in practice, etc.

Students are permitted to tie the project to their outside research activities.

You are encouraged, but not required, to consider submitting your project as a poster, paper, or demo to an upcoming conference that involves machine learning and/or human-computer interaction. Good conferences to consider are Intelligent User Interface (IUI), Computer-Human Interaction (CHI), Autonomous Agents and Multiagent Systems (AAMAS), Association for the Advancement of Artificial Intelligence (AAAI), and User Interface Software and Technology (UIST). If your project involves a robot, HRI, ICSR, Ro-Man, ICRA, and IROS are all also good candidates. I believe that all of these conferences have accompanying workshops with paper deadlines that are later than those for conference papers. I'm happy to provide advice on which conference might be the best fit for your work, and on how to tailor your submission to a particular venue.

Your final project grade will be based on a proposal, an oral presentation, and the research contribution as described in a written paper submitted on 12/18. No work submitted or completed after that date will be graded.

COLLABORATING

You may work alone or in a team of up to 3 people. The amount of work should be commensurate with the number of people in the group. I do encourage working with at least one other person to learn from each other and create a more significant research contribution. (At the end of the semester, group members will additionally submit a short statement describing how each member contributed to the project.)

TIMELINE & DELIVERABLES

0. By 10/27:
If you need guidance choosing a topic before writing a proposal, please contact me before 10/27 to discuss. I will be travelling from 10/25 to 10/30 but should have intermittent email access during that period.

1. Written proposal due 11/1 before class (10% of project grade)
Submit a one-page proposal to Brad via email. (If you are unsure of what project you'd like to do, you may submit two complete proposals, and we'll discuss them both.)

Include in your proposal:

- A detailed description of the proposed work
    - Specifically, what will be built, studied, etc.?
    - How will the work be evaluated?
- Motivation for the project: why is this interesting, important, new?
- Indication of the intended contribution type (from http://chi2012.acm.org/cfp-contribution-types.shtml)
- Preliminary bibliography of related work (at least 5 papers). Use a standard academic style for citations and references (e.g., IEEE or Chicago author-date style). Include information about why each cited reference is relevant.
- Names of group members (if any)

The proposal grade will be based on clarity and completeness of proposal text, and on the quality of the bibliography. We'll discuss your proposal the week of 11/5 and make any necessary modifications.

2. Proposal discussion: meet with Brad, 11/4, 11/5, or 11/6
Sign up for a half-hour slot at http://tinyurl.com/IMLProjectMeetings

3. Project check-in: meet with Brad, 11/26
Sign up for a fifteen-minute Skype slot at http://tinyurl.com/IMLProjectMeetings

4. Oral presentation: 12/6 and 12/11 (15% of project grade)

Each student/group will give an oral presentation describing the motivation, work, and findings. The presentation should include a live demo if possible. I will determine the talk durations once the groups have formed.

The presentation grade will be based on the clarity of the presentation, appropriateness of tone and slides to an academic setting, adherence to time limits, and ability to effectively field questions from the class.

Your grade will also reflect the quality of your presentation with respect to general guidelines and subjective metrics for a good academic presentation, including how engaging it is, how easily your point can be summarized in a sentence or two, and whether you are making eye contact, avoiding filler words (e.g., um, like, and right). I highly encourage you to read the following pieces on giving good academic talks and to follow the advice therein as much as possible:
http://www.cs.berkeley.edu/~jrs/speaking.html
http://pne.people.si.umich.edu/PDF/howtotalk.pdf

5. Written paper and any supplemental materials: due 12/18, 5:00pm (75% of grade)
Describe your research project in a 4-to-10 page conference-style paper. Adhere to the LaTeX or Word templates for the HCI Archive Format from CHI, available at http://chi2013.acm.org/authors/format/. (If you plan to submit your work to a different conference, you can use their template instead. Note the corresponding venue's page limit, though you are not bound to it for the final-project paper.)

You are encouraged to also submit supplemental electronic material (especially video, software for download, etc.). If your system/code is a significant component of your research contribution, you should host it in an online repository and provide the URL in your paper.

Your grade will be based on both the quality of the project and the quality of the writing.

- Project quality:
    - Does the project provide a novel contribution to research and/or practice in interactive machine learning?
    - Is the motivation clear?
    - Is the implementation/execution appropriate?
    - Is the evaluation appropriate and technically sound?
    - Are the conclusions warranted?
- Writing quality:
    - Is the paper well-structured?
        - Should include Abstract, Introduction, Background or Related Work, Method or Implementation, Evaluation, Discussion, Conclusion, References, or some minor variation on that theme.
        - Should include figures, tables, and charts as appropriate.
    - Is the writing clear?
    - Is the paper's tone appropriate for academic writing?
    - Is the writing free of grammatical and spelling errors?

- Is the bibliography content appropriate?
- Are citations in a uniform format (e.g., IEEE or Chicago author-date style)?
- Is the writing (including motivation, language, level of detail, bibliography, etc.) tailored to the intended audience?

POSSIBLE PROJECTS:

Possible project types:

- Implementation and evaluation of a new IML system; novelty might be found in the interface and/or the learning algorithm.
- Evaluation of an existing IML system; note that such evaluation will be much more extensive or unique than that of a project that involves implementation.
- Evaluation of concept(s) across multiple IML systems. For an example, see suggested topic below on comparing ML metrics with HCI metrics.
- Other projects that contribute to the IML research community.

Potential specific topics:

- Interactive training of a personalized activity recognition system.
- Interactive training of person recognition using hybrid sensing (video, vocal audio, audio from movement like footsteps).
- A trainable system for classifying sound in an environment and taking action based on the class.
    - For example, one might want to train a classifier to recognize certain categories of sound that a pet or a baby makes and email a recording of the sound (and possibly audio of what precedes it) if the classification is one of fear or need.
    - Alternately, a sound classification could send a control signal to some device such as a food dispenser or a cooling fan.
    - If you wish to use Java, a UROP I mentor has created a Java module that will extract prosodic features such as vocal pitch and intensity. Using this module would allow you to focus more on the interface and learning algorithm.
- A classifier for a human's internal state that makes recommendations or takes action when in certain states.
    - For example, one could use various human-sensing (e.g., galvanic skin response and facial expression) to create time-series data while a person works on her computer. Later, the system could replay screen captures of previous computer activity and ask the person to label whether she was being productive, taking an intentional break from productivity, or unproductively working with mental fatigue. These labels would be linked to human-sensing data. In later sessions, the system could classify activity; when the person is fatigued it could suggest or even force a break from computer usage. If desired, the person could provide feedback on these actions, providing another source of data.

- Comparing ML metrics with HCI metrics.
    - Possibly to find a decent ML-metric proxy (e.g., cross-validated accuracy) for HCI metrics (e.g. user satisfaction) for lower-cost evaluation and model selection than when asking a human. A search over algorithm space could use many low-cost ML-proxy metrics and few HCI metrics.
- Interactive training of a recommender system for what to wear from one's own wardrobe.
- Work related to human interaction with a learning from demonstration or reinforcement learning algorithm, applied either in simulation or on a robot. We have not covered these techniques yet, but if you already have some experience or only want to focus on the interface side, I have deeper expertise in these topics and can guide you in your project.
    - With learning from human feedback, some potential projects include:
        - Examining the effect of different task narratives (e.g., is the trainer teaching a frog robot to go to water vs. a regular robot to go to a location) on how people give feedback. People may provide less feedback when they believe a state is naturally rewarding or punishing (such as getting to water, but not going to a location), which has implications for whether that state will be able to be identified as a goal by the learning algorithm.
        - Mapping natural language to feedback that can be generically applied to behavior within any Markov Decision Process. For instance, "good" might simply correspond to the previous action being correct, whereas "there you go" might mean the previous action was correct *and* should be chosen again at least a few more times.
        - Adding knobs to a training interface, driving level of exploration, willingness to revisit recent states, and other high-level characteristics of behavior.