Project Descriptions
Decide by: Wed
Groups of Pairs or Threes (Rare exceptions for 4s or 1s)

* The End of Rush Hour

The goal of this project is to generate a solution to end rush hour. We assume that there is a city with streets with traffic lights and lanes. The streets are one-way or two-way. We further assume that there are people who are indoors in their homes in the morning or in their offices in the afternoon. Each person has a window of time that they need to be at work (morning) or at home (evening). The goal is to create an app that signals people when they should leave to get on the road. It will be ok if the solution includes changing the timing of the traffic lights. If needed, the solution could also change the one-way, 2-way directions of the streets, but that's less desirable. Start with a small simulated city (grid of streets) and an increasingly larger number of people to move. You could then move to a map of a real city or a section of it. You may also want to consider to start with very flexible calendars and then increasingly constrain them.

* Learning from Data about the Physical Space

Lots of data is collected about the activities of people in the cyber space, i.e., purchases, searches, etc. There is also a lot of data being captured about our physical space, using sensors. There is a lot of data that is publically available about the physical space, in nature, in space, or in human environments. An example is the traffic data captured by cameras in NYC (http://nyctmc.org). You are to pick on a real data source of your interest and analyze with learning tools.

* Human-Robot and Robot-Robot Interaction

We will consider tasks in which humans and robots may need to collaborate, and study such interaction. The goal is to work on multi-agent planning problem in which the algorithms plans for the different elements of the collaborative team (humans and robots) taking into account each other's limitations and capabilities. The task could be a rescue task (RoboCupRescue simulation), or two robots delivering items, in which one robot has arms but cannot move and the other has no arms but has a basket and can move, or a robot helping a person to put on some item of clothing. The goal is to plan for such distributed and coordinated execution.

* Solutions for a $100 Billion Problem: Healthcare appointment scheduling

Doctors and dentists appointments often are frustrating to both doctors and patients. Patients frequently experience long waiting periods for their scheduled appointments, and doctors often have some patients fail to show for scheduled opportunities. There are many interesting questions about how to do this better, since the problem involves scheduling, uncertain durations, hard constraints, and often the need for realtime rescheduling. Mobile phones might be an avenue for increasing efficiency, such as if nearby patients could be texted if an appointment opened up. Making some assumptions (based on prior models, such as those in related paper or based on data) create and test an algorithm for patient appointment scheduling.

http://economix.blogs.nytimes.com/2009/02/09/a-hidden-cost-of-health-care-patient-time/?_r=0
http://www.post-gazette.com/business/businessnews/2013/02/24/No-shows-cost-health-care-system-billions/stories/201302240381
http://www.palgrave-journals.com/jors/journal/v63/n6/full/jors201183a.html
http://ibisuva.nl/assets/publicaties/artikelen/2014-kuiper-optimized-appointment-scheduling.pdf
http://link.springer.com/article/10.1007/s00291-003-0122-x

* Learning and Teaching Agents

Adversarial and stochastic are two common settings considered in machine learning and reinforcement learning: the first considers when an adversary may select the worst possible outcome or sample (potentially in response to an agent's decision), and stochastic assumes the outcome/sample is sampled from a (typically stationary) distribution. But in some settings the world may be far more generous, like when a teacher is helping a student learn. In this case the samples or results may be chosen specifically to help the agent learn. Initial results suggest that this setting may lead to far faster learning (less experience) than other settings. Many open questions remain, including how this translates to other settings, what happens when the learner can interact, what if both the learner and the teacher are modeling each other's behavior (e.g. if the learner knows that the teacher is trying to help, should the learner use a different strategy), and are there cases where meta reasoning can cause problems (the learner thinking about the teacher thinking about the learner…)? Relevant literature includes:

http://www.research.rutgers.edu/~thomaswa/pub/uai12Dynamic.pdf
https://flowers.inria.fr/mlopes/myrefs/12-aaai-teach.pdf
http://pages.cs.wisc.edu/~jerryzhu/pub/MachineTeachingAAAI15.pdf
http://www.cs.ubc.ca/~kevinlb/pub.php?u=2012-AAMAS-BayesianParameterAnalysis.pdf

* Predicting student intent / goals

MOOCs and other online educational resources are used by a wide variety of learners with different goals. It would be helpful to automatically infer the intent of a learner, both at the macro level (what does he/she want out of this course) and/or at the micro level (what is he/she trying to do in the current assignment/code/etc.). To investigate this, there exist datasets on MOOC and other student learning on CMU Datashops. This problem can be viewed as goal recognition, which could involve automatically extracting hierarchical goals from data, and clustering, classifying and filtering methods to identify goals and a student state.