EE 106B Project Proposal and Guidelines

EE 106B Course Staff

March 18, 2020

1 Overall Project Guidelines

The overall goal of this final project is to give you the opportunity to tackle a research problem of your own choosing and gain experience writing a full-length academic paper. While we don't expect your final product to be at all worthy of publication, we want you to pick a problem at approximately the scale of a conference paper and perform a preliminary analysis. If all your initial analysis or experiments end up working (which we don't necessarily expect to happen), you should be far enough along that you could submit to a conference after a couple months more work (the ICRA deadline is in September).

The usual requirements for this project are quite open-ended:

- The project must be a research project (not an implementation/engineering project like most 106A projects)
- The project must relate to robotics in some way
- You must do at least some work on hardware

Due to the exenuating circumstances this semester, i.e. the Covid19 pandemic and the city lockdown, we are relaxing the project requirements even further. Your project need not use hardware, and it can lie anywhere between a full research contribution and a literature review. Your project should be some convex combination of

- A theoretical contribution: Extend the mathematical theory of robotics in some way, shape, or form. Prof. Sastry's lab does a lot of this sort of work.
- An experimental contribution: Devise or implement an algorithm or sensor and test or characterize it in the real world (or simulation if appropriate). Prof. Bajcsy's lab does a lot of this sort of work.
- A design contribution: Design, analyze, and perhaps prototype a novel robotic system. [1] Prof. Fearing's lab does a lot of this sort of work.
- A literature review: Collect, organize, and summarize recent papers in the field. All good research includes some amount of literature review, and literature reviews are important to find in the current body of work (new problems to solve) as well as to educate researchers new to the field.

Examples of all four types of papers (and many combinations of those types) can be found in the list of paper presentations. These are an excellent starting point to think about new problems in the field.

In the sections below we discuss some aspects of the project in greater detail. We are deliberately vague in defining the project requirements because we don't want to constrain you from doing interesting research. However, we expect you to put a significant amount of work into this project (well over 50 hours per person) and produce something that you are truly proud of. This is not a project you can start at the last minute.

[1] Some papers of this type instead detail the design of novel software (tools, simulators, solvers, etc) for researchers to use. Examples include Trac-IK, IPOPT, MoveIT, OMPL, Tensorflow, etc. We *strongly discourage* you from trying to design novel software. It's not easy to make high-performance software and takes a long time, and you usually need to invent a new algorithm before you make the software package

Project groups should be composed of 2-4 students. This is the ideal size for a good research project; it's good to have more than one person so you can catch each others' mistakes, but research projects are rarely dividable into many independent tasks and thus aren't suitable for large groups. Exceptions can be made on an ad hoc basis for larger groups or solo projects.

2 Final Project Deliverables

The deliverables for your final project will be

Deliverable	Percent of Grade	Due Date
Project Proposal	5%	Friday March 27
Intermediate presentation	15%	Late April (TBA)
Final Presentation	30%	Friday May 8 (Dead Week)
Final Report	40%	Friday May 15 (Finals Week)
Website with Videos	10%	Friday May 15 (Finals Week)

A brief description of each deliverable is listed here. More information will be forthcoming.

Proposal See below.

Intermediate Presentation Each group will be giving a 10-15 minute presentation on their progress as of late April (approximately halfway through) 10% of the grade is based on the presentation, and 5% is based on the feedback you give to other groups. These presentations were originally meant to occur during lab section, but we may decide to move them to lecture or make them fully recorded.

Final Presentation Each group will give a 20-30 minute final presentation on their work during dead week. This presentation should include video of any hardware demos, simulations, etc. Once again, we'll need to figure out how we want to do these remotely.

Final Report Each group will write a full-length research paper on their work. The minimum paper length is 5+n pages (5 pages not including the bibliography), barring ad-hoc exceptions. Ideally papers won't exceed 6+n pages (the limit at most conferences), but certain papers, like literature reviews, might need more space.

Website with Videos Each group will make a website to showcase their project. The main purpose of this website is to provide videos to supplement your paper and to market your work, so it will be much shorter than your website for 106A.

3 Project Proposal

Project proposals will be due on Friday March 27rd at 11:59pm. Your proposal should include:

- **About You**. Include your names, years, majors, relevant courses/research experience, and a one-to-two sentence description about your research interests (or the aspects of robotics that interest you).
- **Research Question**. What problem are you trying to solve? What contribution are you trying to add to the field? This should be phrased as a question – if you already know the answer it's not interesting research.

  You should also discuss where on the spectrum you'll be aiming for between a literature review or a novel contribution. You're free to change this target over the course of the project.
- **Motivation**. Why is the project of interest? Are there motivating applications?
- **Proposed Methodology**. How do you plan to approach the problem? What literature will you be surveying, or how will you undertake your research project? What assumptions will you be making?

- **Related Work**. Include at least three[2] references to papers on your topic. Summarize them in a couple lines and justify why you picked them. You don't need to have performed a full second pass on them yet, but you should have gone deeper than a first pass. Note that you may need to first-pass more than three papers in order to find three papers to start with, so please remember your your literature review skills and be efficient with your time. Some tips:
  - S Keshav's How to Read a Paper is a great resource. Since you guys are still pretty inexperienced, each of Keshav's three passes will likely take you several read-throughs. Since you are all still relatively inexperienced, a first pass will probably take you 10-30 minutes per paper.
  - Remember that your very first read of any paper should just encompass the abstract, section titles, and figures. Your next read should include the introduction and conclusion. Don't read the proofs until after you've determined that the paper is worthwhile.
  - If you find a (high quality) recent paper, look at its references to find seminal older work. Similarly, if you find a useful older paper, you can use google scholar to search through the papers that cite it.
  - Keep the academic family tree in mind. If you find a researcher doing relevant work, it's likely that their advisor or students do similar things.
- **Experimental Plan**. If you'll be performing experimental or design work, elaborate on the methodology section above. What hardware or simulator do you plan to use? Provide a system diagram to characterize the inputs and outputs to your system. Describe any software modules you will need and provide references to third party software (e.g. ROS pacakages) you plan to use. Describe what experiments you plan to run. The purpose of this section is to give the course staff more information so we can help you out.

4 Potential Topics

Here we list a couple potential topics for projects, which were drawn from Prof Sastry, Prof Bajcsy, and Valmik. However, the projects to do not have to be limited to these areas, as long as they relate to robotics as a field. Example projects from previous years can be found on the website.

You are free to use your own research (masters project, work for a lab, etc) as long as you're clear about what work you did previously and what you're using to satisfy this project.

4.1 Grasping

1. Grasping or segmentation of objects in clutter:

   One of the major goals in grasping is the ability to pick up objects in cluttered spaces (like product bins in Amazon warehouses). What current solutions exist? How might you redefine grasp metrics or grasping algorithms to work well in clutter?

   A similar problem is the difficulty of segmenting or identifying distinct objects in cluttered environments (like a box of odds and ends at a garage sale), as they may be partially occluded or difficult to distinguish from nearby objects. One way to deal with this is by using a robot to push, grasp, separate, or otherwise manipulate the objects into positions in which they can be identified. This is an example of *Active Perception* a field which Prof Bajcsy started in the 80's, or the integration of perception and control to learn new information about a robot's environment.

2. Part-based grasp analysis:

   One hypothesis about how humans quickly grasp objects is by identifying affordances, or parts that are easy to grasp, such as handles, cylindrical, or spherical parts. What solutions already exist?

   A possible direction to take this is to segment 3D models and analyze grasp quality for each segment. What's the proper way to segment a model? If you have two objects with similarly shaped affordances, how can you transfer grasps between the objects?

[2] Potentially you might find fewer than three papers that are actually useful for your project (if your field is very new or esoteric, for example). If this is the case, feel free to include fewer papers and justify why you did so.

4.2 Soft Robotics

1. Soft robot design and sensor integration:

   Researchers and hobbyists have been designing soft robotic actuators for a number of years, but relatively little work has been done in soft robotic sensing and control. Often these systems are evaluated qualitatively or via external cameras, but in order to deploy soft robots in the real world, we need reliable onboard sensing.

   How have roboticists designed soft robots, and how have they integrated sensing, if they have at all? Design a multi-DOF soft robot manipulator and specify the actuators you'd need to properly observe its state. If you have access to materials (or the lockdown ends) you can try prototyping.

   Valmik mostly wants this project done so he can try to outdo these guys, but an easy-to-make Open-Source soft robot with proper sensing would allow many other labs to start working with soft robots. This project could potentially serve as a contribution to the soft robotics toolkit or perhaps an entry to the soft robotics competition, if it happens again.

2. Soft robot grasp metrics:

   Soft robotic grippers are used to pick up delicate items like fruits, etc because they can wrap around the object and conform to its profile. They can better control the forces they apply and are robust to small errors in grasp pose. However, no one (that Valmik knows of) has designed grasp metrics for soft robots or mechanical models for how the grasps work.

   What are the current approaches being tried, if any? Try to develop a model for how a soft finger would grasp an object. One thing you can try is to imagine that the finger is composed of N rigid links in an open chain that share the same torque input (where N is on the order of 100). How do the model predictions change with N? How do you account for multiple fingers or strangely-shaped objects? If the lockdown ends (unlikely), you could try testing your model with the soft fingers in the lab.

3. Dynamics of hyper-redundant manipulators:

   Because soft robots deform smoothly, they can be expressed as an infinite chain of rigid joints, each with infinitely small length. A lot of work has been done on the analysis of hyper-redundant robot manipulators: those manipulators with arbitrarily many rigid joints, and this analysis could be very helpful in modeling soft robotics.

   This is an area into which Valmik has been planning on looking for his research, but he hasn't done so yet, so this would be just a literature review unless you find a gap in the work or have an idea for an extension.

4.3 Autonomous Vehicles

1. Decentralized motion planning for swarms:

   Shankar recently read this article and is a bit suspicious of its claims. Provably correct decentralized control for groups of mobile robots is a major research problem and Shankar is curious about how far this group has come. An interesting project could be designing a decentralized motion planning strategy to control a number of mobile robots (and hopefully prove things about it). You could potentially extend Wang and Rubenstein's work, or go in a different direction.

   Other professors whose work you could look into are Hadas Kress-Gazit at Cornell, and perhaps Magnus Egerstedt at Georgia Tech. This project would likely benefit from the use of the Robotarium.

2. Active Safety in Vehicles:

   There's been a lot of interest recently in *active safety* in vehicles, where the vehicle will take over from the human driver in order to prevent collisions. Examples include automatic lane keeping and autonomous emergency braking. There are a couple projects you could do in this area

- Identify a traffic scenerio and try to design an active safety controller to prevent collisions. Try to set guarantees on collision avoidance (for example, "If I stay 4 seconds behind the car in front of me I'm guaranteed to be able to avoid it if it suddenly brakes"). In which situations do your guarantees fail?
- The vehicle's definition of safety is based on its internal model of the world, and thus its sensors. This model might not match the model of the vehicles driver in some situations. For example, a noisy sensor might detect an obstacle that isn't there (or fail to detect one that is), or the human might need to break "safety" in order to perform some maneuver (like driving on the wrong side of the road to avoid a pothole). How might the vehicle and its driver communicate in these situations? Is there a natural way for the human to counteract the vehicle's input?

5 Resources

Here are some resources that might be helpful to you

- Gazebo and MuJoCo are two simulators you can use. Gazebo tends to be used more for ROS-based things, while MuJoCo is usually used in machine learning, but both can probably be used for either purpose. OpenAI Gym might also have some useful environments.
- The Robotarium is an excellent platform for mobile robot or swarm research.
- Harvard's Soft Robotics Toolkit is a great resource for soft robotics projects.
- Adafruit and similar sites sell relatively cheap hardware and sensors you can use for experiments at home.
- Most vision-based projects only require a (cheap) webcam. I personally use Logitech C920 cameras, and the lab uses Logitech C922's (the enterprise version of the C920). Intel sells RealSense cameras for only $200 on their website. Microsoft Kinects are also pretty cheap, and if you have an XBox you might have one lying around already.
- If you want to do some hardware work at home, but money is a major obstacle for you, let us know and we might be able to help out.