Project Title: Socially compliant navigation for ground robots

Contact: Jean Oh (*jeanoh@cmu.edu*)

Project Description:

When navigating in crowds, people tend to change their high-level navigation strategies to blend in with social norms. People are also adaptive in making reasonable decisions even in new, unexpected situations. In this project, we aim to develop an approach where an autonomous ground robot can learn from human demonstrations how to adaptively switch its behavior according to social context. We have several social navigation algorithms including the baseline [Alahi16] and our approaches [Vemula18] that have been evaluated on pedestrian datasets. We plan to integrate these algorithms on a Clearpath Husky robot platform and evaluate their performances in real-life settings.

Preferred and/or Required Skills:

Knowledge of ROS and/or deep learning is preferred. Hands on experience with robot platforms and/or camera and 3D LIDAR sensors is a plus. The students must be able to work independently with minimum supervision, have strong time management skills, and clearly understand the responsibility of deliverables and time deadlines.

Project Goals:

This project expects four milestones:

1) Integration of social navigation planning algorithms on a Husky platform. The communication interface will be ROS. In this task, we develop a software interface to control/drive the robot to follow the waypoints generated by each planner. The task includes communicating through a low-level platform interface for the vehicle state information.
2) Data collection for navigation among human crowds. We will collect a new dataset that is similar to the public datasets, e.g., ETH pedestrian dataset, on CMU campus, with and without including the robot. All sensor data will be logged. We will also set up an overhead camera to establish ground truth trajectories of the agents. This task includes the design, setting up, tagging, and postprocessing the data. (IRB approved)
3) Training new models using the data collected in step 2) for each planning algorithm. Evaluate the results using the same evaluation method as before by measuring the displacement error in the recorded trajectories.
4) Running new experiments using the learned models and analyzing results. In the same environment as the data collection, we will run a set of experiments using each planning algorithm trained using the data from 2). We will create a mechanical turk task to collect human evaluation on this set of results.

The deliverables are the software and data from meeting the milestones above, and a final report in a technical report format, specifically detailing the experimental setup, parameters, results in each condition and for each algorithm. The students will co-author the paper that will be written based on the outcome of this project.

References:

[Alahi16] Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social LSTM: Human trajectory prediction in crowded spaces. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 961–971, 2016.

[Vemula18] A. Vemula, K. Muelling, J. Oh. Social Attention: Modeling Attention in Human Crowds. In Proc. of IEEE Conference on Robotics and Automation (ICRA), 2018