Toward Learning-Aided Interactive and Inclusive Robot Museum docent

Tribhi Kathuria, X. Jessie Yang, and Maani Ghaffari
University of Michigan, Email:{tribhi,xijyang,maanigj}@umich.edu

Abstract: This work reports on developing an integrated-two level architecture for planning the behavior of a Socially Assistive Robot (SAR), deployed in a partially known environment such as a museum. The two layers are designed to fulfil different objectives. The global planning layer is responsible for sharing domain expertise with interacting Humans (provide assistance), which in terms of a museum docent robot does intelligent tour planning. The global planning module optimizes a tour plan given a navigation graph, desired constraints and user preference. The local planning layer however, is responsible for executing the plan given by a global planner in a way that is sociable and interactive. A learning-based local navigation module integrates multi-modal feedback, e.g., vision and language, into the planning architecture and enables detour and replanning for a meaningful and interactive human-robot interaction. The tour design problem, i.e., global planning, is similar to the vehicle routing problem, and we formulate it using linear programming. This formulation allows for converting user preference into weights for the navigation graph and dropping nodes at the user's request or for constraint satisfaction. The proposed formulation of the problem is amenable to future developments in the type and number of constraints that can be specified to create customizable objectives for a tour guide robot.

We develop a learning-based agent for the local planner to navigate between points selected by the global planner. While the museum can be mapped beforehand and the robot can move using classical planning methods, it should be noted that the environment keeps changing with the furniture and artwork being constantly moved around during the day. Thus, we propose a learning based local planner that assumes that the traditional metric map of the environment is unavailable. The agent is a visual navigation agent that only takes the depth image as a sensor input. The learning based framework also allows for future extension into multi-modal navigation using cross modal models based on vision and language.

For the initial development phase, we learn to plan and give tours in a photo-realistic simulation environment called Habitat-AI [1]. We use the global planning path as the oracle and translate our navigation problem into a point goal navigation problem solved using a DD-PPO (Decentralized Distributed Proximal Policy Optimization) agent[2]. Learning to navigate in these photo-realistic environments is essential to a learning based navigation system to allow for good performance in real life for agents trained in simulation[3]. Finally, we develop a ROS bridge for the Habitat-AI environment for seamless deployment on the Fetch robot, which will give tours at the University of Michigan Museum of Arts in the future.

Supplemental Material

Figure (1) is the chosen scene from the Habitat-lab environment [1]. The simulation environment is photo-realistic (made from scans of actual houses) and physics-based, which allows for good sim2real behavior [3]

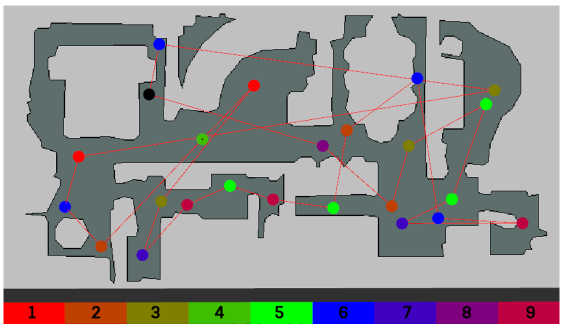

Figure (2) shows the unoptimized tour plan, the naive plan violates design constraints and visits all locations in the museum. The color coded points represent the user-attention demand (a constraint described in the study)

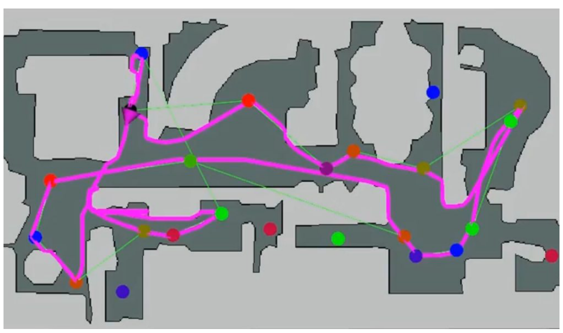

Figure (3) shows the optimized tour plan, the global planner (Green path) here decides to drop nodes in the navigation graph based on user constraints and also optimizes the total path length.

Finally a learning based agent is able to navigate between the different nodes of the graph only getting the ego-centric depth images as sensor input and assuming that the occupancy map of the environment is unknown. The actual path (Pink) taken by the robot avoids the obstacles in the scene.

References:

[1] Savva, M., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., Koltun, V., Malik, J. and Parikh, D., 2019. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 9339-9347).
[2] Wijmans, E., Kadian, A., Morcos, A., Lee, S., Essa, I., Parikh, D., Savva, M. and Batra, D., 2019. DD-PPO: Learning near-perfect pointgoal navigators from 2.5 billion frames. *arXiv preprint arXiv:1911.00357*.
[3] Kadian, A., Truong, J., Gokaslan, A., Clegg, A., Wijmans, E., Lee, S., Savva, M., Chernova, S. and Batra, D., 2019. Are we making real progress in simulated environments? measuring the sim2real gap in embodied visual navigation.