

Towards Automating the Generation of Human-Robot Interaction Scenarios

Matthew C. Fontaine

University of Southern California
mfontain@usc.edu

Abstract

My work studies the problem of generating scenarios to evaluate interaction between humans and robots. I expect these interactions to grow in complexity as robots become more intelligent and enter our daily lives. However, evaluating such interactions *only* through user studies, which are the de facto evaluation method in human-robot interaction, will quickly become infeasible as the number of possible scenarios grows exponentially with scenario complexity. Therefore, I propose automatically generating scenarios in simulation to explore the diverse possibility space of scenarios to better understand interaction and avoid costly failures in real world settings.

Introduction

The human robot interaction (HRI) community currently evaluates their algorithms via hand-authored user studies. When proposing a novel algorithm, each researcher designs an experimental setup to evaluate how their new algorithm performs with human subjects. While such studies are essential to evaluating how a real human will interact with a robot, robots deployed in the real world will encounter novel scenarios not evaluated in experimental settings. To discover scenarios outside of human subjects experiments, I propose simulating HRI scenarios and searching scenario space, where a scenario constitutes both an environment and simulated human agents. The environment consists of all object locations in a scene and the initial configuration of a robot, while the simulated agents produce actions real humans might take when interacting with the robot. The HRI field evaluates algorithms as closed loop systems, meaning the robot’s actions affect human behavior and vice versa. A goal of human robot interaction is eventually to deploy robots for complex real world scenarios. However, individual user studies can only evaluate a small fraction of these scenarios due to the cost of researcher time for running these studies. To complement user studies, I propose automatically generating scenarios in simulation to evaluate human robot interaction algorithms. By thoroughly evaluating such algorithms, researchers can develop more complex HRI algorithms and the industry will be able to better trust the capabilities of proposed methods from the HRI field.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Example failure scenario in shared autonomy.

Previous Work

Scenario Generation in Shared Autonomy I first studied the scenario generation problem in the area of shared autonomy (Fontaine and Nikolaidis 2021b), where a robotic system assists a human user in reaching their desired goal by inferring the user’s intent based on user actions taken. Framing the problem of discovering *failure scenarios* as optimization to maximize completion time yields trivial scenarios such as the human behaving erratically or hindering intent inference by placing two objects close together. Instead of optimization, we frame the problem as a quality diversity (QD) problem. In addition to an objective, the QD formulation permits the definition of measure functions whose output we would like to span. By specifying distance between objects and measured rationality as measures, the QD algorithm MAP-Elites searches the scenario space for scenarios with diverse measure function outputs, while still maximizing the task completion objective. Our approach discovered expected failures, such as the robot approaching the wrong goal given erratic inputs, but also surprising failures, such as the robot failing to reach the desired goal given a nearly optimal user if two objects are placed in a column (Fig. 1).

Efficient QD While the MAP-Elites algorithm is ideal for covering the measure space with high quality solutions, the algorithm is non-adaptive and perturbs existing solutions via a fixed Gaussian. To improve search efficiency, the CMA-ES algorithm adjusts a sampling Gaussian through complex adaptation mechanisms. I proposed combining these adaptation mechanisms with MAP-Elites to form a new QD algorithm called Covariance Matrix Adaptation MAP-Elites (CMA-ME) (Fontaine et al. 2020). CMA-ME finds better-

quality solutions and finds hard to reach solutions better than MAP-Elites in continuous benchmark domains and a reinforcement learning setting. Overall, CMA-ME more than doubles the performance of MAP-Elites across standard QD metrics. The next section explains how we incorporated CMA-ME into an efficient scenario generation system.

Generating Realistic and Valid Scenarios My goal is to create a scenario generation framework capable of generating complex scenarios. In the shared autonomy setting these scenarios were relatively simple, composed of only object locations and human joystick trajectories. Instead, imagine cooking with a robot. Kitchens are filled with many diverse objects and the human-robot team must fluently coordinate their actions. However, the space of kitchen environments is vast and many possible configurations are not likely to align with human designs. For example, it would be unusual to place the stove on top of the refrigerator.

For this reason, I started to investigate how to generate *realistic* scenarios, that is, scenarios that stylistically match human data. I first studied this problem by procedurally generating video game levels for Mario via generative adversarial networks (GANs) (Fontaine et al. 2021a). This work demonstrated that QD algorithms can efficiently search the latent space of a generative model to create playable, human-like levels that varied according to specified measures.

In the above work, the QD objective maximized playability assessed by observing the percentage of a level completed by an AI agent. However, often we know explicitly what constraints need to be satisfied by a solvable generated scenario. For example, in a kitchen environment we know that the human and robot must be able to reach the stove or sink from their starting locations. By ensuring these constraints, we guarantee that a scenario is solvable by a human-robot team. Our approach explored how to repair GAN-generated game levels via mixed integer programming (MIP) (Zhang et al. 2020). The repair guarantees a min-cost edit distance between the GAN level and the repaired level.

Next, I showed that these two methods could be combined into one framework to evaluate human-robot coordination (Fontaine et al. 2021b). Fig. 2 provides an overview of the system. We evaluate this framework on the Overcooked game domain, which is becoming increasingly popular as a test domain for AI/robot coordination algorithms. In this domain, a scenario is a kitchen environment, where a human-robot team must fulfill food orders.

By MIP repairing generated kitchens, the QD objective is free to optimize some other aspect of the scenario generation. In Overcooked, we maximize the gap between an optimal MDP controller that jointly plans the human-robot team and a human agent model teamed with a human-aware robot that must infer intent of its teammate through observed actions. Our work showed that more emphasis should be placed on studying the effect of the environment when studying human-robot coordination, rather than just studying how the human and robot perceive and account for one another. An online user study confirmed that environments that induce unbalanced workloads for simulated humans result in unbalance for real human teammates when evaluated on the same environment.




Figure 2: My proposed scenario generation framework

Ongoing Work

My current work explores how to incorporate differentiable quality diversity (DQD) algorithms (Fontaine and Nikolaidis 2021a) into my scenario generation framework to more efficiently search scenario space. Current QD algorithms treat the objective and measure functions as “black boxes”, ignoring gradient information. However, my QD framework integrates with deep generative models, which are differentiable functions. For this reason, I introduced the DQD problem, which I defined as a special case of QD where gradient information is available for the objective and measure functions. I have proposed a DQD variant of CMA-ME called Covariance Matrix Adaptation MAP-Elites via a Gradient Arborecence (CMA-MEGA). To evaluate CMA-MEGA, I have searched the latent space of StyleGAN guided by OpenAI’s CLIP model as differentiable objective and measure functions. My next goal is integrating this new class of QD algorithm into my scenario generation framework.

References

- Fontaine, M.; Liu, R.; Khalifa, A.; Modi, J.; Togelius, J.; Hoover, A.; and Nikolaidis, S. 2021a. Illuminating Mario Scenes in the Latent Space of a Generative Adversarial Network. In *35th AAAI Conference on Artificial Intelligence*.
- Fontaine, M.; Togelius, J.; Nikolaidis, S.; and Hoover, A. 2020. Covariance Matrix Adaptation for the Rapid Illumination of Behavior Space. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference (GECCO)*.
- Fontaine, M. C.; Hsu, Y.-C.; Zhang, Y.; Tjakana, B.; and Nikolaidis, S. 2021b. On the Importance of Environments in Human-Robot Coordination. In *Robotics Science and Systems (RSS)*.
- Fontaine, M. C.; and Nikolaidis, S. 2021a. Differentiable Quality Diversity. In *Advances in Neural Information Processing Systems*, volume 34. (oral presentation).
- Fontaine, M. C.; and Nikolaidis, S. 2021b. A Quality Diversity Approach to Automatically Generating Human-Robot Interaction Scenarios in Shared Autonomy. In *Robotics Science and Systems (RSS)*.
- Zhang, H.; Fontaine, M.; Hoover, A.; Togelius, J.; Dilkina, B.; and Nikolaidis, S. 2020. Video Game Level Repair via Mixed Integer Linear Programming. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*. (oral presentation).