

Leveraging Roles in Robot Teams to Encourage Human Participation

Soheil Habibian

Virginia Tech

habibian@vt.edu

Abstract: Recent advances in multi-agent reinforcement learning enable teams of intelligent agents to collaborate and complete tasks. These approaches work well when every agent is a robot; but what happens when we introduce human partners into these robot teams? Prior work from behavior economics suggests that humans have standardized expectations when working in teams (e.g., we expect teams to be fair and interpretable). In this paper we explore how multiple centralized robots can leverage these expectations to create teams which encourage humans to join in and participate. Our hypothesis is that the way the robots divide the task --- i.e., how the robots split a larger task up into subtask allocations --- should be aligned on two axes: legibility and fairness. We first extend work on legible motion planning for individual robots to produce legible allocations for centralized robot teams. These legible allocations use a cognitive model of the human to identify robot actions that make the human's role clear. We then test our resulting algorithm for selecting subtask allocations across online and in-person user studies. Within these studies we show robot team behaviors to participants and ask them to (a) predict what their role is and (b) indicate which team they would prefer to work with. We find that our approach to generating legible teams does indeed make the human's role more clear, and that humans typically prefer to work with legible teams. Humans also prefer to work with fair teams (as opposed to teams where one or more robots do all the work) with one important caveat: humans are ambivalent to fairness when they are purely watching teams, but sensitive to fairness when they are actually participating within those teams.