RESEARCH STATEMENT

SRAVYA KONDRAKUNTA
EMAIL: KONDRAKUNTA.SRAVYA@GMAIL.COM

My research focus is to develop an autonomous agent that can operate robustly in real-world environments. The quest to build autonomous agents that perform as efficiently as humans in the real world is on the bucket list of computer scientists for decades. Although some intelligent autonomy developed in specific domains like chess or AlphaGo that outperform some human experts, an agent to exceed a human on many tasks is hard to achieve. Two traditional approaches to achieving such autonomy are the classical approaches that contain goals, plans, and actions to achieve the goals. The other is learning approaches that use large amounts of data to build a model and make the autonomous agent behave in a certain way. However, both approaches do not provide a set answer to a critical challenge that stands in the way of achieving an efficient and robust autonomy in its response to anomalies(unexpected events in the real world). Hence, I follow an approach that is a subcategory of the symbolic AI called Goal-Driven Autonomy (GDA) to tackle the issue.

GDA follows a goal-centered approach to achieve and manage an agents' goals. The main steps involved in goal achievement in the GDA process include detecting an anomaly; Explaining and reasoning about the anomaly to detect a problem, and finally, Generating a Goal to handle the anomaly if the anomaly is a potential problem. An agent detects anomalies when certain expectations on the agent or the world do not meet the observation. After anomaly detection, the agent then checks if the anomaly poses a threat or potential threat to either the agent or its goals (a problem anomaly) by using causal explanation patterns. Finally, the agent formulates a goal such that the agent is no longer in the anomalous state. Apart from anomaly detection and goal formulation, a GDA agent also concerns itself with goal management. GDA agent tries to achieve goal management using several different goal operations named selection, change, formulation, monitoring, achievement, delegation.

My current research focus is goal management using goal operations. In particular, I am interested in three different goal operations selection, change, and formulation. Goal selection operation tries to select the current goal to work on from all the goals provided to the agent. Goal change transforms the agents' current goal into an achievable goal when its current goal is not achievable due to various reasons in the dynamic environment. Goal formulation comes with a goal when the agent detects a problem anomaly to respond to the problem or gain more knowledge. I implemented the three-goal operations in several domains separately. I am currently investigating the interaction of the three-goal operations in a single domain. I have developed an algorithm such that the agent can decide upon a single goal operation when there is a possibility of multiple. The algorithm uses several different factors and analyzes the adverse effects of the anomaly on the agent's health; negative effects of the anomaly on the agent's goals; anomaly repetition; resources vs. benefit of achieving a goal, and finally, the importance of a particular goal to its mission.

After getting all the factors in a multiple-goal operation scenario, the agent then tries to choose a single goal operation based on outputs obtained from the five factors listed above. The algorithm can easily extend to include several of the other goal operations, and it is also generalize-able across different domains. I implemented my algorithm in two domains: the underwater domain to survey fish in water and study their habitat. The other construction domain is an extension to the simple blocks world. I implemented the GDA agent in both simulated and real-world settings. I used MOOS-IVP as a simulator to simulate underwater vehicles and our simulator to simulated the construction domain. I also used underwater gliders to embody such an agent. The agent presented promising results in both settings. My research now offers a path to manage the goals of autonomy on its own, thus making it more robust.