Research Statement

Evana Gizzi, Tufts University

Introduction

One of my all time favorite movies is *The Shawshank Redemption*. In this movie, a man named Andy Dufranes is sentenced to life in prison after being wrongly convicted of double murder. Unbeknownst to the viewer, the man spends his first 20 years in jail planning and preparing for his escape. But how does one escape from a maximum security prison? This resourceful man used a rock hammer, that he kept hidden in a book, to chisel a tunnel through his cell wall, and then a 6ft rope to tie his belongings to his foot while he crawled through a sewage pipe, and escaping to freedom. How cool is that?!

My research investigates how to solve these 'seemingly unsolvable problems' through creative problem solving (CPS) in intelligent agents. Although there is no formally accepted definition of CPS, I define it as *problem solving which occurs when the existing capabilities and knowledge of an agent about itself and its environment are insufficient to accomplish the goal, and thus, the agent expands its own capabilities and knowledge about the environment in a way that allows it to accomplish the goal* [5]. But how, exactly, does this expansion happen? I aim to tackle the question of knowledge base expansion by developing methods for CPS on a theoretical basis, and then implementing and testing those methods in a proof-of-concept model in a cross-domain manner. Although CPS can be used to solve any problem, it is most useful in the case of anomalous situations, where traditional problem solving methods are less useful.

How can we define and implement CPS, and can CPS be used to understand/react to anomalous situations? In order to answer this question, I investigate three *key areas*. (1) What exactly is CPS? (2) What kinds of problems require CPS? (3) How can CPS be implemented in intelligent agents, specifically in the context of anomalies?

Being able to implement CPS in intelligent agents is extremely important in our digital world, where many processes becoming automated. These processes work by predetermining possible contexts and situations, and developing automatic reactions to those known, predetermined situations. What many of these systems lack is the ability to react to off-nominal situations (called *outliers*), especially those which require actionable outcomes (called *anomalies*). In most modern applications of automation, like digital credit card transactions or computer networking, anomalies which go unaddressed can have catastrophic consequences.

Current Major Limitations The limitations of this problem in the research community start with detecting outliers and distinguishing anomalies from outliers. This can be particularly challenging because there is no way to identify these circumstances a priori. Moreover, anomalies look very different depending on the domain, context, and/or specific problem. This not only affects the identification problem, but also the method for resolving or solving the anomaly.

Defining Creative Problem Solving

In order to address key areas (1) and (2) in my research, I have cultivated a collaboration with Lakshmi Velayudhan Nair (PhD candidate at Georgia Institute of Technology) in developing a comprehensive survey in CPS, including a definition of CPS developed using the breath of research in the field. In this work, we define CPS problems as *problems in which the existing capabilities and knowledge of the agent about itself and its environment, are insufficient to accomplish the goal.* From this, we define CPS as *the process by which the agent discovers new concepts that were not in the initial conceptual space of the agent, allowing it to accomplish a previously impossible goal* [4].

Anomaly Detection

In order to address the key areas (2) in my research, I developed a research project devoted to the disambiguation of anomalies from outliers, and providing a method for finding the saliency responsible for the anomaly. I address the identification issue by developing a generalized framework for anomaly detection, which scrutinizes 3 *key elements* to make its anomaly/outlier distinction: *Context*, *Prediction*, and *Observation* [2]. I then tested this generalized framework in an aviation

domain, on a real-time, human-in-the-loop rapid prototyping flight simulator at NASA Langley Research Center, and found that the system was able to detect anomalies in all trials, invariant of the subject flying the aircraft, timing of anomaly, and nature of anomaly [3].

Robotic Action Discovery

In my project *Research in Action Primitive Variation in Robotics (RAPDR)*, I address key area (3), by investigating methods for expanding the knowledge base of an agent as a means for CPS. I focus specifically on ways that the agent can learn new actions to expand its knowledge base, and provide it with additional, formerly unknown capabilities. One method for discovery is by taking the known actions, and attempting to segment them into smaller actions, such that the segmentations provide additional effects that were not formerly possible with the initial known actions [1]. Currently, I am developing an action discovery method which takes the known actions as input, and attempts to find new actions by varying the parameters on those known actions. For example, the 'speed' parameter could be increased in the 'push' action to yield a new action, 'strike' [4].

Spacecraft Fault Reasoning

In my project *Research in Artificial Intelligence for Spacecraft Resilience (RAISR)*, I address key area (1) (2) and (3), investigating computational methods for identifying types of safe-mode[1] status' in spacecrafts, along with diagnostic techniques using data driven paradigms for explaining the faults. In order to identify the type of safe-mode, I used a Dempster-Shafer Theoretic framework combined with classical logic to deduce the level of severity of the safe-mode. I then use associative algorithms, neural networks, and graph theory to investigate the diagnostic of the safe-mode event. The ensemble of techniques allows me to diagnose both nominal and outlier/anomaly cases. This work is currently in process as an Internal Research and Development (IRAD) project at NASA Goddard Space Flight Center, in which I serve as Principle Investigator.

Proposed Future Work

In future work, I would like to further investigate the *relevancy* question in the context of generalized anomalies [6, 7]. I plan on using the RAISR project to help crystallize the direction for the this work for the Fall of 2020. I believe that furthering research into non-trivial relevancy relationships will help contribute to handling generalized anomalies using CPS.

References

[1] E. Gizzi, M. G. Castro, and J. Sinapov. Creative problem solving by robots using action primitive discovery. In *2019 Joint IEEE 9th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, pages 228–233. IEEE, 2019.

[2] E. Gizzi, L. Le Vie, M. Scheutz, V. Sarathy, and J. Sinapov. A generalized framework for detecting anomalies in real-time using contextual information. In *MRC@ IJCAI*, pages 1–7, 2018.

[3] E. Gizzi, L. Le Vie, M. Scheutz, V. Sarathy, and J. Sinapov. Knowledge acquisition in the cockpit using one-shot learning. In *NAECON 2018-IEEE National Aerospace and Electronics Conference*, pages 208–213. IEEE, 2018.

[4] E. Gizzi, L. Nair, J. Sinapov, and S. Chernova. From computational creativity to creative problem solving agents. In *11th International Conference on Computational Creativity*, pages 370–373, 2020.

[5] E. Gizzi, L. Nair, J. Sinapov, and S. Chernova. Creative problem solving in intelligent agents: A survey. *Journal of Artificial Intelligence Research (JAIR)*, [**To Be Submitted:** Fall 2020].

[6] V. E. Houston, W. J. Manuel, E. Gizzi, and B. Barrows. Advancing aircraft operations in a net-centric environment with the incorporation of increasingly autonomous systems and human teaming. In *AIAA Aviation 2019 Forum*, page 3517, 2019.

[1] A safe-mode status is when a spacecraft shuts down all of its non-essential functions in order to protect itself from an event which caused a detected on-board fault

[7] A. Mashiku, C. Frueh, N. Memarsadeghi, E. Gizzi, and A. Burton. Predicting satellite close approaches using statistical parameters in the context of artificial intelligence. 2019.