Inverse Cognition in Search and Rescue using Probabilistic Programming

Problem Statement:

Building effective and prescient recommendation/coaching systems using Artificial Intelligence is hard. Traditionally, such systems are more of a hinderance than an enhancement to user productivity (i.e. Microsoft Clippy)

We seek to build a coaching system that can effectively improve human performance in complex tasks. Specifically, we explore applications in Urban Search and Rescue using a simulated environment built in Minecraft.

Solution/Approach

A probabilistic model of human cognition that determines what the player is thinking, we call it the Cognitive Inverter (CI). Specifically the CI is a Dynamic Bayesian Network (DBN) that contains three key components:
(1) Victim Prioritization Strategy, The CI uses this information to infer which victims the player is attempting to save next.
(2) Room Search strategy which treats the map of room, doors, and halls within the building as a traversal graph of areas. The CI uses this to infer where the player is going to move next.
(3) World State. This includes a list of areas the player has and has not visited as well as which victims a player saw but did not save yet.

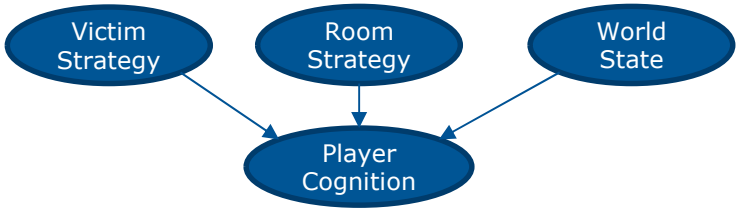

Experiments/Results

The figures below (read left to right and top to bottom) shows the CI changing its beliefs about the player's intended victims to save based on the observed player behavior

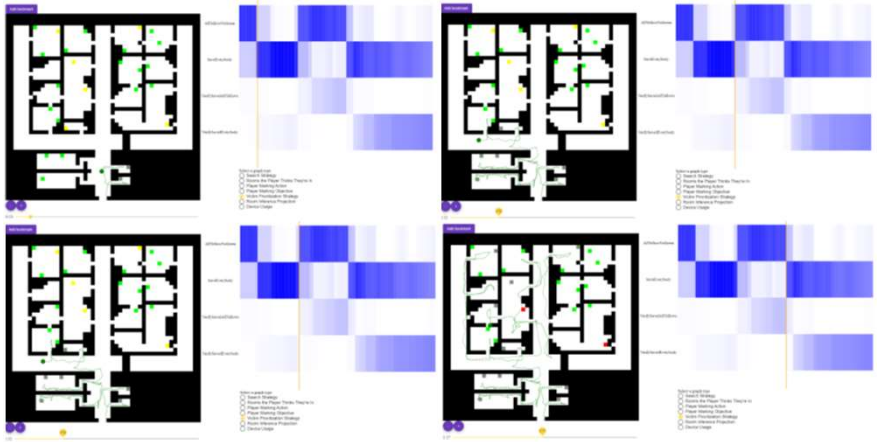

Conclusions and Future Work

Overall, using probabilistic programming as a means of inferring a person's cognitive state has proven promising on complex tasks such as urban search and rescue, especially when determining player intent. This will prove invaluable to both a human and artificial coach, as they can use this information to determine the most appropriate interventions that maximizes said person's productivity and effectiveness.

However, there remain several pieces of the CI that need iteration and improvement.

1. The CI needs to generalize better on more complex maps. We will need to revisit and possibly remove certain assumptions we made about the (relatively) small Minecraft map in our current experiments

2. The CI's inferences on user intent is all done offline, and the model will require performance improvements before it can be deployed in a live/online environment. One approach for this would be to memorize computations about path planning and world state in our inference.

3. The CI currently only infers cognition for one player. In the real world, Urban Search and Rescue is done in teams.

Contact: Kenneth Lu, (617) 491-3474 x571, klu@cra.com

charles river analytics