CAREER: Robust Online Decision Procedures for Societal Scale CPS

Abhishek Dubey (Vanderbilt University)

In collaboration with Ayan Mukhopadhyay (Senior Research Scientist, Vanderbilt University), Baiting Luo (PhD Candidate) and Yunuo Zhang (PhD Candidate), Jacob Buckelew (PhD Candidate)

CPS applications affect entire cities and societies, which we refer to as **Societal-Scale Cyber-Physical Systems (SCPS),** where sequential decisions must be taken under uncertainty in a manner that balances safety, system objectives, and societal impact.

The problems described above can be modeled as variants of Markov decision processes (MDP), which can be solved offline (e.g., by reinforcement learning (RL)) or online (e.g., Monte Carlo search). However, these approaches have limitations for SCPS.

The key challenges are:

1. Offline planning can use vast volume of data, but how can we make it adaptable?
2. Online planning is adaptable, but how can we make it scalable?
3. How can we detect when the *world* changes?

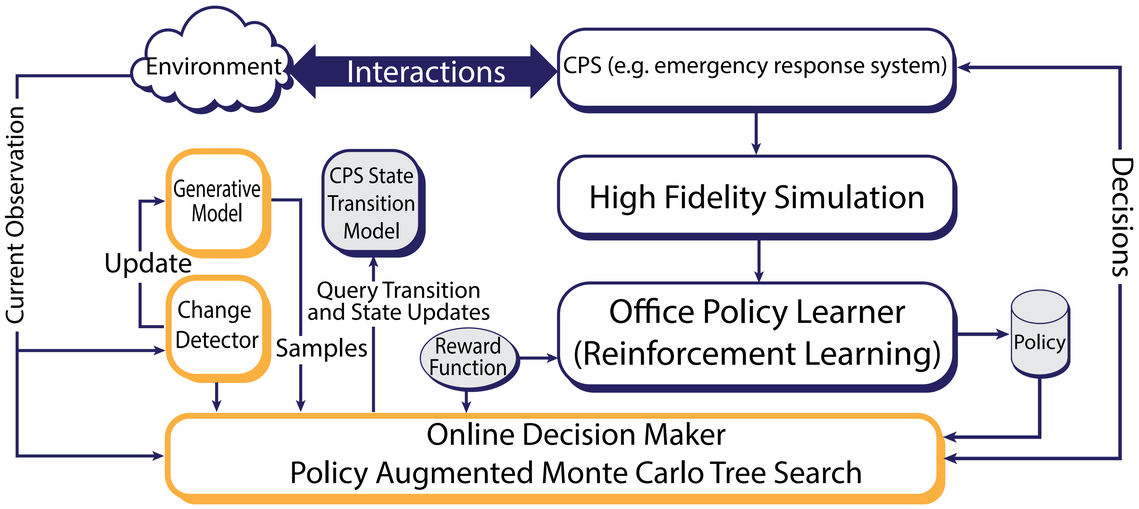

Our Approach

We have made **three fundamental algorithmic advances:**

1. We have designed a novel online planning algorithm (Policy-Augmented Monte Carlo Tree Search) that can extract useful information from a policy trained offline and use during online planning. This approach makes offline planning adaptable and online planning scalable **(AAMAS 2024a).**
2. If the decision-maker is unaware of *how the world has changed,* we have developed an approach where the agent can explore the world safely, and slowly adapt to the updated environment **(AAMAS 2024b).**
3. We have developed a tractable density estimator that can be used to detect if the environment has change (under review at KDD 2024).

Broader Impact

1. Our work is bridging the gap between data-driven decision-theoretic algorithms and CPS, by ensuring that such algorithms can be applied robustly and safely to SCPS such as transportation, emergency management, and traffic operations.
2. Impact in education: We are working on designing planning gym, which is a training environment and can be used for testing and training on large scale societal systems.

(AAMAS 2024a): Decision Making in Non-Stationary Environments with Policy-Augmented Search, International Conference on Autonomous Agents and MultiAgent Systems 2024.
(AAMAS 2024b): Act as You Learn: Adaptive Decision-Making in Non-Stationary Markov Decision Processes, International Conference on Autonomous Agents and MultiAgent Systems 2024.

Award ID#: 2238815