

COACH

Course of Action arCHitecture

Project Resource:

GitHub repository

Introduction:

COACH (Course of Action arCHitecture) is a strategic planning tool designed to enhance the decision-making capabilities of AI systems. It integrates with AI frameworks such as [MarcoPolo](#) to help operators plan and understand complex sequences of actions for AI agents, which is vital in environments where uncertainty is high.

Challenges in planning have traditionally been approached by the fields of operational research (OR) and logistical planning, where statistical models are employed to represent the systems of interest.[1] For example, OR tools can be used to optimize the routes and schedules of delivery trucks, taking into account the constraints and objectives of the human operator. They work best on specific problems where all of the relevant details are known and the system of interest is static, rather than dynamic. This is an issue owing to the dynamic nature of most problems in life, where static, unchanging solutions are not suitable. Further, although they are rooted in solid statistical and mathematical theory, they fail to capture the full risk profiles of plans of action. These are major shortcomings that are present in current iterations of planning software required for critical decision-making scenarios, such as contested logistics.

Another branch of research that focuses on solutions that adapt to changing situations is that of reinforcement learning (RL), in which artificial intelligence (AI) agents are trained to behave a certain way based on their given reward structure. RL's strength lies in how it enables AI to address problem solving with scenarios both old and new contested logistics and supply chain resiliency planning.

Both RL and OR offer their own unique benefits to the challenges in planning, but seldom have these subjects been united. When they have, they tend to address simple tasks with discrete actions such as "make a step forward".[2] COACH combines the best of both of these disciplines with our course of action architecture (COACH), a hierarchical AI planning tool which we built to clarify the reasoning of AI in planning courses of action (COAs).

COACH counters OR's inflexibility in addressing uncertainty by giving human operators a more nuanced look at risk profiles. This lets operators exercise their own intuition along with guidance provided by algorithmic estimates in the form of automatically generated COAs. These COAs are made by having RL agents solve the given scenario, with each solution having some variability in how the agents were able to solve the problem, such as in fuel efficiency or timeliness. When our solution is paired with COA adjustment capabilities, this enables highly efficient counterfactual analysis for rapid, AI-backed decision-making. This type of framework, in which humans are monitoring and having a direct hand in the activity of AI, is referred to as human-on-the-loop.

The part of the solution that requires the most effort in implementation is making simulation models match reality. We have already built a robust simulation space that works very well for space environments. We have also proven this capability in the ground, air, and sea domains in logistical planning. Additionally, this framework is functional on the scale of tracking thousands of agents.

Lastly, our team has done extensive research in implementing RL for training AI agents, and our latest development in this endeavor

is skill-agents. These agents specialize in certain tasks given the provided action space, such as efficient routing for vehicles or object tracking for satellites. By using a skill-based framework, we are able to both reduce redundancies in training agents and represent resulting COAs in terminology that is readily understandable by human operators.

The GitHub repository for this product can be visited [here](#).

References:

- [1] F. Petropoulos et al., "Operational Research: methods and applications," Journal of the Operational Research Society, vol. 75, no. 3, pp. 423–617, Mar. 2024, doi: 10.1080/01605682.2023.2253852.
- [2] C. F. Hayes et al., "A practical guide to multi-objective reinforcement learning and planning," Auton Agent Multi-Agent Syst, vol. 36, no. 1, p. 26, Apr. 2022, doi: 10.1007/s10458-022-09552-y.

Mobius Logic

Mobius Logic has been helping healthcare and financial services companies with their data analysis and modeling methods since 2004. The company has built data analysis platforms that help Fortune 500 companies obtain reliable and validated insights from their data. Mobius Logic has also been successful in modifying one of its commercially available products, MAKANA, for automated Machine Learning Ops. The modified MAKANA is a specialized tool for a USAF ISR end user that has been implemented in the program called Airman Role-based Intelligence Search Engine (A-RISE), a successful Phase I, II and now Phase III SBIR project (Contract # FA865022C1037). Several other modifications to the MAKANA product have been made, such as in the IDIA and CLAWS projects (Contract #s: FA864922P0686, FA864921P1020, FA864920C0130). CLAWS is now a commercial off the shelf (COTS) product marketed on the Microsoft app store. All of the projects described here have started as SBIR Phase I or Direct to Phase II (D2P2) efforts, which validates our ability to mature SBIR-funded tech development to the commercial level.

Contact Us

1775 Tysons Blvd, 5th Floor
Tysons, VA 22102
www.mobiuslogic.com

Tel: 855.965.6442
info@mobiuslogic.com