April 2022

The AWKWARD Real-Time Adjustment of Reactive Planning

Leila METHNANI [a,1], Andreas ANTONIADES [b] and Andreas THEODOROU [a]
[a] *The Department of Computing Science, Umeå University, Umeå, Sweden.*
[b] *Independent scholar, United Kingdom.*

Keywords. hybrid systems, reacting planning, social norms, games AI

Developing agents in highly dynamic environments involves employing techniques such as Behaviour-Based Artificial Intelligence (BBAI), which strictly avoids modelling the environment and focuses on the actions that an agent can take instead [1]. By performing local search within a pre-defined plan, BBAI allows for responsive, yet goal-oriented behaviour without the need to model the environment. This is made possible by the fixed relationship between sensory information and agent actions [2,3]. While BBAI can indeed increase search speed, it does not account for explicitly modelling expected social norms and interactions between agents.

In [4], we presented a hybrid architecture, Agents With KnoWledge About Real-time Duties (AWKWARD), as a solution to this problem. AWKWARD, depicted in Figure 1, combines BBAI approaches with formal reasoning. More specifically, we integrate the OperA framework [5] with Behaviour-Oriented Design (BOD) [2] to produce socially-aware self-adjusting and self-organising BBAI agents. AWKWARD consists of three modules, each with a distinct purpose: 1) the OperA module; 2) the reactive planner; and 3) the behaviour library.

The OperA module validates the social behaviour of the agent and provides direction to the reactive planner upon the completion of a behaviour—regardless of success. It contains normative descriptions of how the agent is expected to behave under specific social scenarios. The planner module allows the agent to act intuitively. Like all reactive planning paradigms, the planner conducts local search within a pre-defined plan based on environmental changes, as well as those internal to the agent. As changes are recorded, plan elements may be triggered or halted, allowing the agent to focus on what is considered most important in any given environment [2]. While each agent has its own plan, the behaviour library is a collection of primitives, i.e. senses and actions, shared across all AWKWARD agents deployed in the Multi-Agent System or Agent-Based Model. The behaviour library is accessible by both the planner and the OperA module of each agent, offering them the means to exercise control over the agent's environmental perception, internal status, and actuators available. By having a single behaviour library for all modules, we enable not only code re-usability, reducing the memory footprint of the agents, but also the ability to host the behaviour library remotely, if needed to use high-performing hardware to run complex actions and senses [6].

[1] Corresponding Author: The Department of Computing Science, Umeå University, Umeå, Sweden; E-mail: leila.methnani@cs.umu.se.

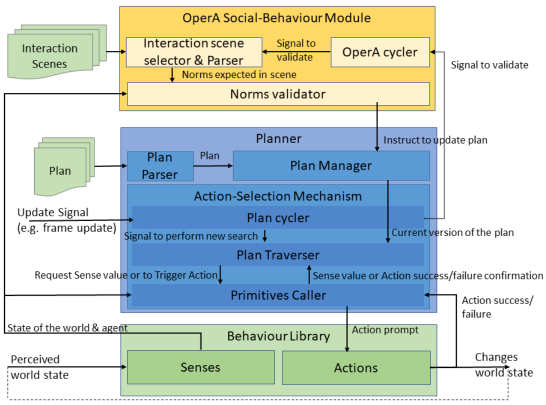

Figure 1. Conceptual diagram of the AWKWARD architecture with each of the 3 major components presented in a different colour [4].

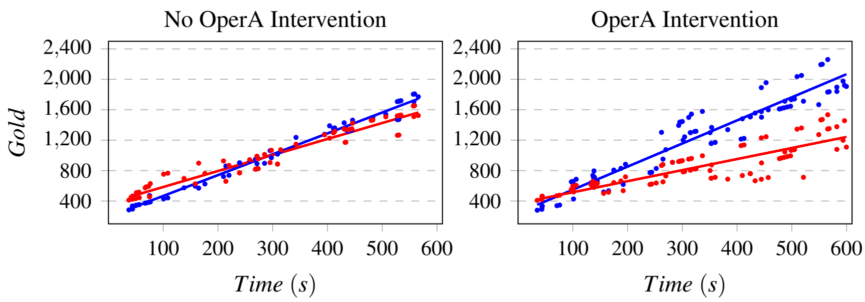

Figure 2. Gold acquisition over time for default position 1 (blue) and AWKWARD position 5 (red) over five trials. [4].

We developed a sample implementation in the popular video game DOTA2 to evaluate our system. It presents results from the first 10 minutes of the game, during which *priority farming* is an important strategy that encapsulates some expected social norms: one agent (position 1), must act selflessly and sacrifice wealth (i.e. the in-game currency, *gold*) for an allied agent (position 5), to benefit. We capture and define this expectation in what OperA calls an *interaction scene*, and use gold acquisition as a metric for performance. Upon each scene initiation, if a norm violation occurs, OperA prompts the agent to alter its plan. When a scene ends, the plan is returned to its initial structure. Figure 2 shows our preliminary results: a divergence in gold acquisition over time between the AWKWARD bot (position 5 in red) compared to the default DOTA2 bot (position 1 in blue) as shown by the right subplot. This divergence is a result of the AWKWARD bot's social behaviour change; OperA bans farming behaviour, resulting in the agent's own sacrifice of gold in favour of its ally's acquisition of gold instead. Contrastingly, the left subplot shows how each agent's gold acquisition is on par with one another without social norm enforcement by OperA.

April 2022

References

[1] Brooks RA. Intelligence without representation. Artificial Intelligence. 1991.

[2] Bryson JJ. Intelligence by design: Principles of modularity and coordination for engineering complex adaptive agents [Ph.D. thesis]. Massachusetts Institute of Technology; 2001.

[3] Brooks RA. New Approaches to Robotics. Science. 1991;253(5025):1227-32.

[4] Methnani L, Antoniades A, Theodorou A. Embracing AWKWARD! Real-time Adjustment of Reactive Planning Using Social Norms. arXiv preprint arXiv:220410740. 2022.

[5] Dignum V. A model for organizational interaction: based on agents, founded in logic [Ph.D. thesis]. SIKS; 2004.

[6] Bryson JJ, Theodorou A. How Society Can Maintain Human-Centric Artificial Intelligence. In: Toivonen-Noro M, Saari E, Melkas H, Hasu M, editors. Human-Centered Digitalization and Services. Springer; 2019. .