The AAAI-17 Workshop on
What's Next for AI in Games?
WS-17-15

Against Forward Models

Joseph C. Osborn, Michael Mateas
Expressive Intelligence Studio
University of California, Santa Cruz

We propose an agenda shift in academic video-game-playing AI and modeling. We call for less *simulation* and more *speculation*: for imprecise but analyzable *closed-form* abstractions instead of accurate but opaque forward models; for giving agents the game's instruction manual, or at least knowledge of game domain concepts; and for techniques that scale to real examples without requiring gigawatt-hours.

It is well-known that domain knowledge assists search, but current discourse in game-playing AI focuses too much on generic search and not enough on representing general game domains. Current approaches emphasize rapid sampling of thousands of alternatives rather than efficient extraction of design information using justifiable inductive biases. Black-box forward model search is the key enabler of this mindset, abetted by the low-level semantics of most game modeling tools. Video games have domain similarities beyond merely receiving button inputs at 60 frames per second, such as gravity and projectiles functioning similarly across games that feature them, similarities in enemy behavior, *et cetera*. Successful general video game players already leverage such shared schemata (Pérez-Liébana et al. 2016), as do *ad hoc* players created for commercial games.

Why are we pursuing general video game playing using tools meant for analyzing arbitrary black boxes? To recenter the importance of video game domain knowledge, we advocate 1.) reviving the GVG-AI learning track and 2.) introducing a rule-understanding track that gives agents time to read games' rules before playing. These build towards a substantial challenge: Can we develop AI that can beat, in reasonable time, an (arbitrary) 1980s-era role-playing game like *Dragon Warrior* or *Final Fantasy*? Humans readily abstract these systems into coarse hierarchical and parallel goals, but they seem far beyond current game playing techniques.

While human players primarily *speculate* what might happen if high-level actions are taken, machine players mainly *simulate* taking low level actions and go with the best observed outcome. People use less accurate models but extract more useful information from them. Moreover, even perfectly accurate models can't avoid re-planning in adversarial or stochastic games.

Separately, game programming and modeling languages often *only* have a frame-by-frame semantics and are effectively unabstractable. Even high-level languages like PuzzleScript and Machinations are difficult to meaningfully abstract; VGDL might be an exception, though it has no formal semantics. By analogy to mathematics, we want to work with *closed-form* models. Rather than defining Mario's jump as a transition function between physics states, we would prefer a piecewise-quadratic function we can evaluate anywhere; path planning becomes quadratic programming and requires much less search. Zelda-like adventure games are naturally phrased as partial order planning over locks and keys embedded in a topological map. We must develop modeling tools that are expressive at these multiple levels of abstraction.

In summary, forward models are useful and easy to obtain, but low-level forward search should be the last resort of automated game players and model checkers—not the first.

References

Pérez-Liébana, D.; Samothrakis, S.; Togelius, J.; Schaul, T.; and Lucas, S. M. 2016. Analyzing the robustness of general video game playing agents. In *Proceedings of the IEEE Conference on Computational Intelligence and Games*.