Dealing with Noise and Other Agents The Empowerment of Controller Selection

Christoph Salge
Adaptive Systems Research Group
University of Hertfordshire, Hatfield, United Kingdom
supported by H2020-641321 socSMCs
Email: c.salge@herts.ac.uk

Daniel Polani
Adaptive Systems Research Group
University of Hertfordshire, Hatfield, United Kingdom
supported by H2020-641321 socSMCs
Email: d.polani@herts.ac.uk

Empowerment [1] is an information-theoretic, intrinsic motivation that quantifies how much an agent can influence the world it perceives in the future. Empowerment-based behaviour generation produces a number of interesting behavioural phenomena, such as the avoidance of unpredictable agents [2] and the production of artefacts in the environment that reflect the agent's embodiment [3]. Furthermore, there is also the hypothesis that coupled empowerment maximisation, a model where an agent considers both its own and another agent's empowerment, can be used to formalize supportive and antagonistic behaviour to foster creativity [4]. Empowerment is formalised as the channel capacity from an agent's actions to the same agent's sensors at a later point in time. Importantly, this is traditionally computed as an open-loop channel capacity. Starting from the state of the empowerment computation, the agent considers all possible n-step action sequences, and determines what outcome they would produce after n steps, assuming those where acted out regardless of the sensor feedback produced by them in the meantime. The channel capacity is then computed between the action sequences and the resulting state distributions (for details, see [1]).

The inability to take sensor feedback into account produces some unintuitive results for the empowerment calculation for dynamic environments. Consider, for example, an agent that could enter a number of doors, but would first need to cross an open field, where it is subject to a degree of movement noise, such as wind. If the agent could consider its sensor feedback, then it would be no problem to dynamically correct a path that would lead it through any of the doors. But if it was to only consider its open-loop empowerment, then there would be no pre-selected action sequence that would get an agent to reliably pass a door. Open-loop empowerment would therefore underestimate what the agent could actually do. Similar effects can be seen when considering a multi-agent system, where the other agent is basically just a source of noise. In theory, it would be possible to design an agent to minimize the empowerment of another. For example, if a second agent would always act in a way to undo the actions of the first agent, then the empowerment of the first agent would be zero, as its actions would be cancelled out. This cannot be modelled with classical open-loop empowerment, as it is not possible to select a fixed action sequence to do this before the first agent acts. To compute empowerment that considers dynamically reacting to other agents, we need a different empowerment formalism.

In this paper, we want to introduce empowerment based on controller selection. Instead of selecting an action sequence, which is then acted out, we consider selecting a controller $p(a|s)$, a function that maps sensor inputs to actions. We can then ask:"what is the channel capacity between the variable encoding all the possible controllers, and the observed outcome after several steps?" We will demonstrate with some simple simulations how this produces different results in dynamic environments, and how it can be used much more elegantly to model the interaction between two empowerment-aware agents. We will also discuss some of the immediate formal implications, such as upper bounds, and the expected computational costs.

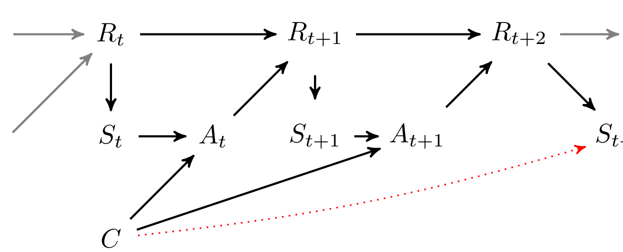

Figure 1. The time-unrolled perception-action-loop, modelling the sensors S and actions A of the agent and the environment R. C is a variable representing the selected controller, its states are strategies $p(a|s)$. The red arrows indicate the causal flow relevant for empowerment, in this case from the controller to the sensors.

References

[1] A. S. Klyubin, D. Polani, and C. L. Nehaniv, "Keep Your Options Open: An Information-Based Driving Principle for Sensorimotor Systems," *PloS one*, vol. 3, no. 12, pp. 1–14, 2008.

[2] C. Salge, C. Glackin, and D. Polani, "Empowerment and state-dependent noise-an intrinsic motivation for avoiding unpredictable agents," in *Advances in Artificial Life, ECAL*, vol. 12, 2013, pp. 118–125.

[3] ——, "Changing the Environment Based on Empowerment as Intrinsic Motivation," *Entropy*, vol. 16, no. 5, pp. 2789–2819, 2014.

[4] C. Guckelsberger, C. Salge, R. Saunders, and S. Colton, "Supportive and Antagonistic Behaviour in Distributed Computational Creativity via Coupled Empowerment Maximisation," in *Proc. 7th Conf. Computational Creativity (to appear)*, 2016.