

Overcoming Dynamicity with Plasticity: Neuromodulation for Lifelike Systems

Chloe M. Barnes¹, Anikó Ekárt¹, Kai Olav Ellefsen², Kyrre Glette^{2,3}, Peter R. Lewis⁴ and Jim Tørresen^{2,3}

¹Aston University, Birmingham, B4 7ET

²Department of Informatics, University of Oslo, Oslo, NO-0316

³RITMO, University of Oslo, Oslo, NO-0316

⁴Ontario Tech University, Oshawa, ON L1G 0C5

c.barnes1@aston.ac.uk

Introduction

Natural beings are often situated in dynamic and unpredictable environments, and have evolved to use mechanisms such as neuromodulation – the ability to change behaviour via changes to synaptic activity in the brain – to adapt their behaviour over time to survive. The ability to change behaviour in this way is referred to as ‘behavioural plasticity’. In this extended abstract, we summarise the findings from an exploration of how plasticity can affect how artificial agents evolve when solving tasks of different complexity (Barnes et al., 2020, 2021a), and when evolving in dynamic and unpredictable environments (Barnes et al., 2021b).

Plasticity in Natural and Artificial Life

Behavioural plasticity – specifically activational plasticity (Snell-Rood, 2013) – is the ability for an individual to change its behaviour immediately in response to new environments or stimuli, by making temporary phenotypic changes. Neuromodulation is a biological process found in animal brains that can facilitate this type of behavioural plasticity (Hamood and Marder, 2014), where synaptic activity between neurons is modified or regulated temporarily to produce reversible behavioural changes that do not affect learnt behaviour (Abbott and Nelson, 2000).

This type of plasticity has inspired research into designing adaptive, lifelike, artificial systems – especially those underpinned by neural networks since they themselves are inspired by connectionist models of the brain. In artificial systems, activational plasticity can be achieved by regulating or modulating connection or ‘synaptic’ activity locally in a neural network, or in a separate modulatory network (Beaulieu et al., 2020).

Operationalising Neuromodulation

In the studies discussed in this extended abstract (Barnes et al., 2020, 2021a,b), the effects of plasticity were explored using the River Crossing Dilemma (RCD) testbed, first proposed by Barnes et al. (2019) – used to explore how artificial agents evolve to solve tasks in shared 2D grid-world

environment. The RCD is characterised by a lethal, vertical river of water in the centre of a 19×19 grid, which agents must learn to cross to achieve their goal. In doing so, they are presented with a social dilemma, as a bridge for safe passage requires two stones – each with an increasing personal cost to place. Using the RCD testbed, we explored how artificial agents with neural controllers learnt to achieve goals when alone and when situated in an environment with another, unknown agent; the presence and actions of this other agent makes the environment unpredictable for all, to study how neuromodulation and behavioural plasticity affects goal-achievement in dynamic environments. These agents learnt using neuroevolution, whereby the weights of a population of neural networks are evolved or modified over time using an evolutionary algorithm.

Specifically, neuromodulation is operationalised in these studies *within* a single neural network, by temporarily gating/regulating the outgoing signals of neurons depending on the incoming signal. Hidden neurons in the network could evolve to be non-modulatory (standard) or modulatory (will gate or ‘turn off’ outgoing signals depending on the input). Effectively, modulatory neurons ‘fire’ when the incoming signal is negative, changing the outgoing signal from the neuron (i.e. the weights of the connections to the next layer of neurons) to be 0. In this way, an agent can temporarily change its behaviour to respond to its environmental stimuli without changing learnt or encoded knowledge, since neural network activity is regulated locally without modifying connection weights permanently. This is intended to help agents overcome dynamicity in their environments.

Task Complexity and Plasticity

Both natural and artificial agents are often presented with environments that change over time, are shared with others, and involve tasks that require multiple steps to complete (Dezfouli and Balleine, 2019). This environmental dynamicity and uncertainty can make it challenging to learn to complete tasks when the full state-space is not known – which is often the case. Consequently, the effect of behavioural plasticity via neuromodulation was ex-

plored in artificial agents to ascertain whether plastic behaviour is beneficial when learning to solve tasks with multiple stages (Barnes et al., 2020, 2021a).

The results show that the activity-gating neuromodulation described above has a significant effect on an agent's ability to solve tasks, when evolving in both single- and multi-agent (paired with one other, unknown agent) environments and when agents are presented with either a single- or multi-stage task. The expected fitness of agents evolving to solve these tasks of varying difficulties in variable conditions was also seen to increase when agents are capable of neuromodulation; this shows that behavioural plasticity can be beneficial for creating adaptive agent controllers that are able to overcome the dynamicity and uncertainty that often characterises realistic environments. Despite the significant benefit that behavioural plasticity has on these agents for goal-achievement and fitness, this does come at the cost of evolutionary volatility – that is, agent fitness is observed to fluctuate more often during evolution compared to agents without neuromodulation. This creates a trade-off between fitness and goal-achievement, and evolutionary volatility.

Dynamicity and Plasticity

A further study explored the effect of neuromodulation and behavioural plasticity on agents that are situated in environments with increasing variability (Barnes et al., 2021b); in the natural world, environmental dynamicity can arise from the unpredictable actions of others, which is becoming increasingly common in artificial systems as components may interact unintentionally (Hähner et al., 2015). To explore the effect of plasticity on agents in variable environments, this study observed agents evolving in an increasing number of environments with another, unknown agent that either stays consistent throughout evolution (less variable), or is random at each generation (more variable); the actions of a consistent partner would theoretically make the environment more predictable over time than a randomised partner, which would be inherently unpredictable.

The study found that modulatory agents achieved a significantly higher fitness than non-modulatory agents in all areas of the study. Further, a correlation was found between the variability in the environment and the strength of the effect that neuromodulation has on agent fitness, where neuromodulation has a stronger benefit on agents as variability increases. Evolving artificial agents to achieve goals in highly variable environments is challenging, but this study shows that a biologically-inspired mechanism like neuromodulation can increase agent fitness by enabling them to temporarily change their behaviour and phenotype – even when there are unknown entities in the environment.

Outlook

Artificial systems are growing in size and it is increasingly likely that the components they are comprised of will inter-

act in unintended ways (Hähner et al., 2015). By designing artificial and technical systems to be more lifelike in their behaviour – such as employing them with the ability to express behavioural plasticity – one would hope that these systems could combat the dynamicity and uncertainty that characterises the realistic environments inhabited by the natural beings Artificial Life researchers are inspired by. Neuromodulation is thus shown in these studies to be a viable option for such plastic behaviour, by enabling agents to adapt their behaviour temporarily in response to environmental changes, without affecting learnt knowledge, and without requiring knowledge of others in the environment.

References

- Abbott, L. F. and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nature Neuroscience*, 3(11):1178–1183.
- Barnes, C. M., Ekárt, A., Ellefsen, K. O., Glette, K., Lewis, P. R., and Tørresen, J. (2020). Coevolutionary learning of neuromodulated controllers for multi-stage and gamified tasks. In *Proceedings of the IEEE 1st International Conference on Autonomous Computing and Self-Organizing Systems (ACSOS)*, pages 129–138. IEEE.
- Barnes, C. M., Ekárt, A., Ellefsen, K. O., Glette, K., Lewis, P. R., and Tørresen, J. (2021a). Behavioural plasticity can help evolving agents in dynamic environments but at the cost of volatility. *ACM Transactions on Autonomous Adaptive Systems*, 15(4).
- Barnes, C. M., Ekárt, A., Ellefsen, K. O., Glette, K., Lewis, P. R., and Tørresen, J. (2021b). Evolving neuromodulated controllers in variable environments. In *Proceedings of the IEEE 2nd International Conference on Autonomous Computing and Self-Organizing Systems (ACSOS)*, pages 164–169. IEEE.
- Barnes, C. M., Ekárt, A., and Lewis, P. R. (2019). Social action in socially situated agents. In *Proceedings of the IEEE 13th International Conference on Self-Adaptive and Self-Organizing Systems*, pages 97–106.
- Beaulieu, S., Frati, L., Miconi, T., Lehman, J., Stanley, K. O., Clune, J., and Cheney, N. (2020). Learning to continually learn. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI)*, pages 992–1001. IOS Press.
- Dezfouli, A. and Balleine, B. W. (2019). Learning the structure of the world: The adaptive nature of state-space and action representations in multi-stage decision-making. *PLOS Computational Biology*, 15(9):1–22.
- Hähner, J., Brinkschulte, U., Lukowicz, P., Mostaghim, S., Sick, B., and Tomforde, S. (2015). Runtime self-integration as key challenge for mastering interwoven systems. In *Proc. of the 28th Intl. Conf. on Architecture of Computing Systems (ARCS)*, pages 1–8. VDE.
- Hamood, A. W. and Marder, E. (2014). Animal-to-animal variability in neuromodulation and circuit function. In *Cold Spring Harbor Symposia on Quantitative Biology*, volume 79, pages 21–28. Cold Spring Harbor Laboratory Press.
- Snell-Rood, E. C. (2013). An overview of the evolutionary causes and consequences of behavioural plasticity. *Animal Behaviour*.