

The effect of periodic changes in the fitness landscape on brain structure and function

Nitash C G^{1,2,6}, Barbara Lundrigan^{2,3,5,6}, Laura Smale^{2,3,4,6} and Arend Hintze^{1,2,3,6,*}

¹Department of Computer Science and Engineering ²BEACON Center for the Study of Evolution in Action

³Department of Integrative Biology ⁴Department of Psychology ⁵Michigan State University Museum

⁶Michigan State University, East Lansing, MI 48824, USA

*hintze@msu.edu

Abstract

Natural organisms have transitioned from one niche to another over the course of evolution and have adapted accordingly. In particular, if these transition go back and forth between two niches repeatedly, such as transitioning between diurnal and nocturnal lifestyles, this should over time result in adaptations that are beneficial to both environments. Furthermore, they should also adapt to the transitions themselves. Here we answer how Markov Brains, which are an analogue to natural brains, change structurally and functionally when experiencing periodic changes. We show that if environments change sufficiently fast, the structural components that form the brains become useful in both environments. However, brains evolve to perform different computations while using the same components, and thus have computational structures that are multifunctional.

Introduction

Changing fitness landscapes affect evolutionary dynamics profoundly (Wilke et al., 2001; Laakso et al., 2001). Slow changes can favour specialists, while generalists can have an advantage in environments that change often (Travis and Travis, 2002; Li and Wilke, 2004). At the genetic level, changing environments modify the effect of mutations, and can alter their rate of fixation. A mutation that is beneficial in one environment could be even more beneficial in the next, thereby accelerating fixation. Conversely, a deleterious mutation effect can be mitigated by an environmental change, thus allowing valley crossings to occur more often.

As the fossil record makes abundantly clear, the history of the planet is peppered with sudden, drastic environmental changes that impose vastly different constraints on surviving species, e.g. The Great Oxidation Event. Another example is the mass extinction of dinosaurs in the cretaceous, which allowed previously nocturnal rodents to fill the diurnal niche recently vacated by sauropods (Gerkema et al., 2013; Maor et al., 2017). Transitioning from a nocturnal to a diurnal lifestyle, or vice versa, is arguably one of the toughest transitions to make for a vertebrate. A transition of this magnitude requires that not only entire metabolic pathways adapt to the changed temperature, but also that the sensory modalities adapt to cope with different constraints on foraging. A

nocturnal lifestyle either requires extreme adaptations of the vision system to deal with the lack of light, or requires that organisms rely on other sensors, such as olfaction or echolocation in the case of bats. On the other hand, the light of day allows distant objects to be detected easily, emphasizing the importance of vision. On the downside, the higher temperature during the day causes odors to evaporate more quickly, reducing the importance of olfaction. Lastly, a new environment entails new predators (who may already have adapted to this environment), and hence the entire behavioral repertoire with respect to predation response must be adapted. As a consequence, we expect to see wide ranging structural and functional changes in the neural architecture of organisms that have to adapt to such drastic changes. Additionally, the changes between nocturnal and diurnal preference is also believed to drive speciation (Santini et al., 2015). We expect that periodic transitions not only force the organism to evolve to both environments, but may also cause the organism to adapt to the transitions themselves (as observed in the digital evolutionary model from (Li and Wilke, 2004)). Here, we are specifically interested in these repeated changes, since they should drastically affect the organization and topology of the evolved neural architectures.

There have been several computational approaches that have tried to address this question. For example, it has been shown that neural modularity can mitigate catastrophic forgetting (Ellefsen et al., 2015), i.e. modular structures allow organisms to retain cognitive functions when experiencing a different environment for a number of generations. At the same time, it is not clear that changing environments necessarily promote modular structures (Hüsken et al., 2002; Schlosser and Wagner, 2004; Kashtan and Alon, 2005; Hintze and Adami, 2008). At the genetic level, it has been shown that an organism's reproductive strategy affects the modularity and epistasis of the underlying genetic architecture (Misevic et al., 2006). At the moment, it is unclear if these results will translate to neural architectures.

Here we study how an agent controlled by a Markov Brain (Edlund et al., 2011; Marsteller et al., 2013; Hintze et al., 2017) evolves to forage in an environment that peri-

odically changes over the course of evolution. In one environment, the agent can “see” the food, while in the other the agent can only “smell” the food. We show that periodic transitions cause agents to rearrange their neural architecture in such a way that their structural components become useful in both environments - a phenomenon we term *structural amalgamation*. However, we also show that the structurally amalgamated brains do not create functionally overlapping modules. The brain states that an agent experiences are mostly unique to each environment, and are not shared between environments; the degree to which functional states are shared we term *functional amalgamation*.

Materials and Methods

Evolutionary Framework

All experiments were performed using the MABE (Bohm et al., 2017) framework. This allows us to evolve different neural substrates, in different environments, and using different selection methods, and allows the results of these experiments to be compared directly. The specific MABE modules used in these experiments are detailed below.

Markov Brains

Using natural organisms for an evolutionary experiment such as we are going to perform here, would be extremely time consuming, inconvenient, and costly. As an alternative, we use evolvable Markov Brains (MBs) (Edlund et al., 2011; Marstaller et al., 2013), which are networks that can be composed of probabilistic and deterministic logic gates (for a detailed technical description of Markov Brains see (Hintze et al., 2017)). Here we only use deterministic gates. Each MB is encoded by a genome, which contains genes, and each gene encodes one of the logic gates by specifying what function it performs and how it connects to other gates, the sensors, and the actuators of the agent the MB controls. Figure 1 shows the schematic of an MB, and Figure 2 shows the encoding that is used to construct MBs from the genome. Among other applications, MBs have been used to understand the evolution of animal behavior (Olson et al., 2013), decision making (Kvam et al., 2015), and rates of evolutionary adaptation in different contexts (Schossau et al., 2016). These MBs are used to control an agent that interacts with the environment using different kinds of sensors, actuators, and internal machinery. We study how these MBs respond to evolutionary pressures, as an abstract model of a natural brain. Clearly, neurons, tissues, sensory cells, and muscles are very different from logic gates and the binary sensors and actuators we use here. However, the analogy between these systems can be drawn on a different level; both systems have components that are involved in different processes, and we are interested in how evolution affects how these components are used, and in turn how their use affects their future evolution. In natural organisms we can ablate cells (say in

the visual cortex), and observe if the loss of these cells modulates behavior. Similarly, we can toggle logic gates on and off (like a cell ablation or genetic knockout), which allows us to draw analogous conclusions about the function of those logic gates. As a consequence, our results will not make specific implications about natural neural architecture, but rather how functionality is distributed over the components optimized by selection.

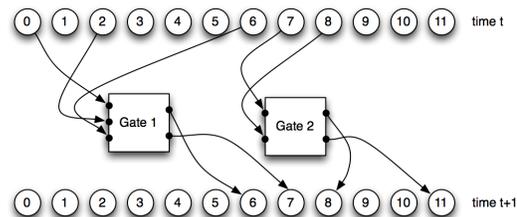


Figure 1: Schematic view of a Markov Brain Network. The sensory inputs, hidden states, and motor outputs at time point t at the top are read by two gates. The result of their computation propagates to new states at time point $t + 1$. In case two gates try to write into the same state, their outputs will be integrated by applying a logical OR operation.

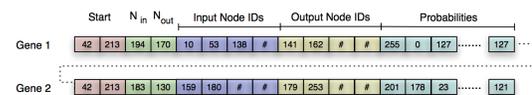


Figure 2: Encoding of a Markov Brain. A string of numbers is used as the genome. The genome is read from end to end, and specific pairs of numbers (42,213 for deterministic logic gates) indicate the start of a “gene”. Each gene encodes the connections of one logic gate: red sites are used for the number of inputs, green for the number of outputs, blue for the actual addresses of the states the gate reads from, and yellow for the actual addresses of the states the gate writes into. The remaining cyan colored sites are used to specify the logic of the gate.

Environment

Agents are tasked with solving a foraging problem and their performance is measured by how much food they can collect over a specified period of time. The environment is a square 2D grid with periodic boundary conditions (effectively turning the grid into a torus). There are 20 pieces of food randomly distributed over the grid, and the size of the grid is 10×10 . When an agent is directly on top of a resource, it can consume it. To keep the quantity of food constant, each time a food is consumed, a new piece of food is placed randomly into the environment on an empty space. An agent can take one of 4 different actions at every time

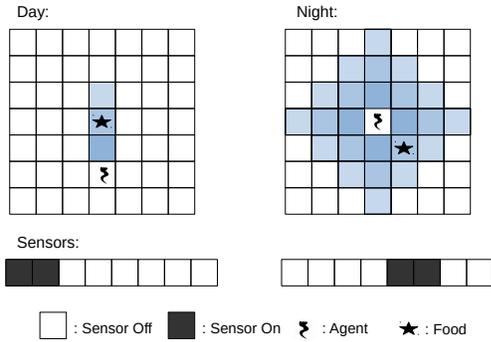


Figure 3: Environment types as perceived by an agent. An agent is represented by a squiggle and food by a star. The Day panel shows the agent’s line of sight with blue squares. The Night panel similarly shows the region that can be smelled by the agent with blue squares. The sensors that light up in the respective environments are shown below each panel.

step, namely, move forward, turn 45 degrees left or right, or attempt to consume food, which is only successful if food is directly under them. The environment can be of two different types, which we call *day* and *night*. During the *day*, agents can see food and not smell the food, while during the *night*, agents can only smell food and not see it. Sensors are implemented as four separate binary inputs for each type of environment. In the case of vision, the sensors light up if the agent is directly facing the food, and within 4 steps of the food. The number of sensors that light up is equal to 4 minus the number of steps from the closest food (see Figure 3 day). In the case of olfaction, the sensors light up if the agent is within a Manhattan distance of 4 from the food, in any direction. Similarly to the *day* environment sensors, the number of sensors that light up in the *night* environment is equal to 4 minus the Manhattan distance to the closest food (see Figure 3 night). In either environment, if all the corresponding sensors of an agent light up, this indicates that the agent is standing directly on top of the food. This allows agents to have similar sensory information that correspond to both the environmental types.

Agents experience an environment for 100 time-steps and can collect food during this time. The environment can technically change from *day* to *night* in two different ways: either during the lifetime of the agent, which allows them to experience both types of environments during their lifetime, or only over generational time-scales. Here we are interested in the changes to the neural architecture and topology resulting from generational changes to the environment, and as a result, individual agents will only ever encounter a single type of environment during their lifetimes.

Selection Method

Populations of agents were allowed to evolve for 10,000 generations in periodically cycling environments. At each generation, agents were selected to propagate offspring to the next generation, proportionally to their ability to collect food, which implements a Moran process (Moran, 1958).

Amalgamation

To understand how the architecture and function of the evolved MBs are affected by the periodic environmental changes, we perform two kinds of analyses on these agents.

Knockout Experiments and Structural Amalgamation

After agents were allowed to evolve, their performance along the line of decent (LOD) (Lenski et al., 2003) (the path from a random organism in the final population traced over its ancestry, back to the start population), was reconstructed. Each agent on the LOD was then tested 50 times over 500 time steps, both in a purely day environment and in a purely night environment in order to more accurately assess an individual’s ability to forage in these different types of environments. This also defined the baseline of performance for subsequent knockout experiments.

Each agent has a set of logic gates that define what we call *structural components*. Obviously each gate also conveys a function, but we assess the functional contribution differently (see functional amalgamation). After all, different functions could be implemented on the same computational structure, analogous to how computers are capable of running different programs using the same hardware. We test the contribution of each gene by sequentially removing one gate after another and testing how the agent performs in both the day and night environments. The loss (or gain) in performance, for each individually knocked out component, defines the extent to which it is involved in computations pertaining to the day or night type of environment, and to what degree this gene is involved in both computations.

Modularity measures the degree of connectivity of a component within a module as opposed to its connectivity to the remaining components. Since evolved MBs have to solve at most two tasks, we would a priori expect them to have either one or two modules. However, typical modularity measures do not differentiate between having only one module or no modules at all (for an elaboration on this topic see (Hintze and Adami, 2010)). We will therefore quantify the “Amalgamation” of the evolved brains and not the modularity. Specifically, we want to quantify to what degree the genes that determine the function and connectivity of logic gates in MBs, are involved in diurnal or nocturnal foraging, and if they are needed in both environments.

If all gates are used in both environments we argue that the brain is fully amalgamated. To that end, performance was measured in both environments as W_{day} and W_{night} . After that, each component was knocked out one after an-

other, and the performance after each knockout in each environment was assessed as $K_{day,i}$ and $K_{night,i}$ (where i indicates the index of each component). We can now quantify the contribution of each component i to the performance of day and night as:

$$\Delta W_{day,i} = W_{day} - K_{day,i} \quad (1)$$

$$\Delta W_{night,i} = W_{night} - K_{night,i} \quad (2)$$

A gate (i) that contributes to only one type of environment, say the day, will have a high $\Delta W_{day,i}$ and a low $\Delta W_{night,i}$. A gene that contributes to both types of environments, will however result in both delta values being high, while an obsolete or redundant gene will show no effect on the deltas. Additionally, we observed cases where, when genes are deleted, the performance of the agent in one or both environments increased, effectively creating negative values for $\Delta W_{night,i}$ or $\Delta W_{day,i}$. Since we are only concerned with the degree of the effect the deleted gate has, we used the absolute of these delta values. The overlap or degree of amalgamation for each gene A_i can now be quantified as the lower of the values $|\Delta W_{night,i}|$ and $|\Delta W_{day,i}|$ for each gene i . The total degree of structural amalgamation (A_S), normalized to the number of genes N is thus:

$$A_S = \frac{1}{N} \sum_{i=0}^N \min(|\Delta W_{day,i}|, |\Delta W_{night,i}|) \quad (3)$$

A brain that has high A_S is a highly amalgamated brain, in that more of its structural components are involved in the function of both types of environments.

State Transition Experiments and Functional Amalgamation We also analyzed the “experience” that brains have in both types of environments. The experience of a brain can be captured by recording the internal states of a brain. This is analogous to experiments done on natural organisms’ brain states (Werner, 2009). In a nutshell, the state a brain is in at every moment defines the experience the brain has. Since we are working with digital organisms, we recorded the internal states of a brain at every single update as well as the states of the environment. We then reconstructed the brain state transition graph. This graph shows how each brain state transitions to the next one, based both on the sensory inputs the agent receives, and what the agent does. In this graph, each node represents a brain state (experience) and edges show the transitions from state to state given the inputs. By running the agent in different environments, we color coded which node was visited in what environment. We expect some states (experiences) to be visited only in one environment, while other states might be visited in both environments. This allows us to quantify how much each state of the brain is involved in experiencing and

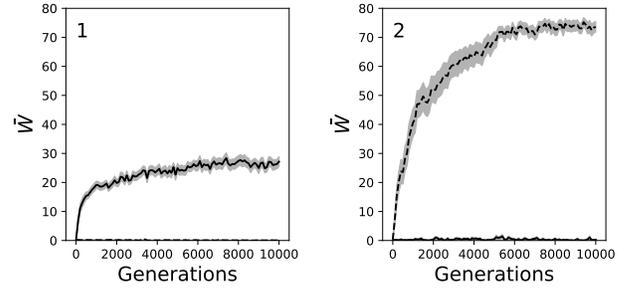


Figure 4: Average fitness of agents from the LOD evolved in only the night environment (panel 1) or the day environment (panel 2) without any periodic changes. The fitness measured in the day environment is shown as a dashed line. The fitness in the night environment is shown as a solid line. Gray shadows indicate the standard error over all 40 replicates.

processing each environment (e.g. Fig 9). Functional amalgamation (A_F) is therefore V (the number of states visited in both environments) over N (the total number of states experienced):

$$A_F = \frac{V}{N} \quad (4)$$

Results

We evolved populations of 100 agents starting with the day or night environment and switched the environments in periodic conditions (no switches, every 1 to 30 generations, 40, 50, 60, 70, 80, 90, 100, 150, 200, 250, 500, 750, 1000, 2500, and 5000 generations) over 10,000 generations. For each experimental condition, 40 replicates were performed. For each replicate the line of decent (LOD) was reconstructed (Lenski et al., 2003) and each agent on the LOD was tested in the day and night environment again.

When evolving agents in only the day or night environment we find that performance over evolutionary time converges to a maximum that is different for the day and night environment (see Figure 4). We conclude that in this computational model vision is the superior sense, since it does not require agents to forage for food as much as olfaction requires them to. This also shows that evolving from the night to the day environment is easier than the other way around, which creates an asymmetry within the experiment. However, this is also true for natural evolution, where the transitions between environments are often easier in one direction than the other.

Generally, agents experiencing environmental changes very often (up to about changes every 10 generations) keep improving their performance regardless of the transitions over the entire course of the experiment (see Figure 5 panel 1 and 2). Agents whose environment changes less often (every

20 to 50 generations) improve their performance less well. They encounter an occasional drop in fitness once in a new environment (see Figure 5 panel 3). However, they still manage to optimize their performance overall, just not to the degree they would have in a constant environment (compare to Figure 4) or if switches occurred more frequently (see Figure 5 panel 1 and 2). Lastly, if environmental switches happen rarely (above 100 generations) agents keep experiencing a drop in fitness once the environment switches, and they are not capable of improving their general performance further over time (see Figure 5 panel 4).

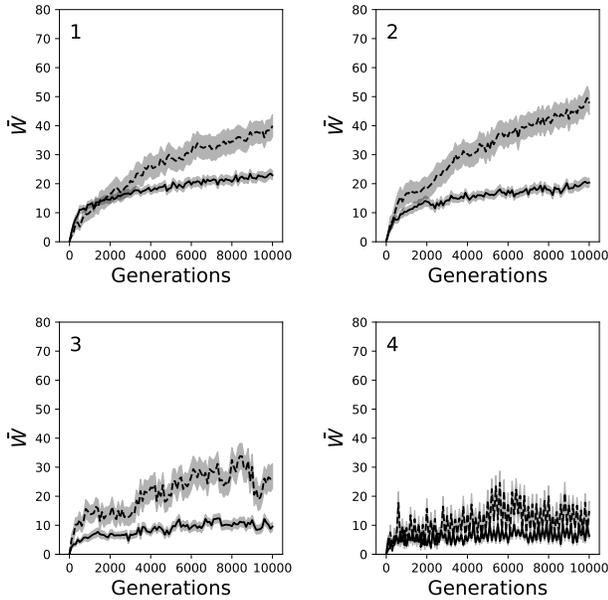


Figure 5: Average fitness of agents from the LOD evolved with periodic environmental switches 1) every 4 generations, 2) every 10, 3) every 50, and 4) every 100 generations. The fitness measured in the day environment is shown as a dashed line. The fitness in the night environment is shown as a solid line. Gray shadows indicate the standard error over all 40 replicates.

For an overview of the attained fitness in the last generation, i.e. the performance after 10,000 generations, for different periods between switches, see Figure 6.

These results confirm that when environments change rapidly enough agents can not only adapt to one environment at a time, but also adapt to the transition itself (Li and Wilke, 2004). We did not measure if the loss of function when transitioning to a new environment is in itself adaptive i.e. if the rate of function loss decreases over time. However, how this adaptation works in agents controlled by Markov Brains is likely to be different than in (Li and Wilke, 2004). This previous work used AVIDA (Ofria and Wilke, 2004) which evolves sequences of computer commands that control each individual. Rapidly changing environments cause

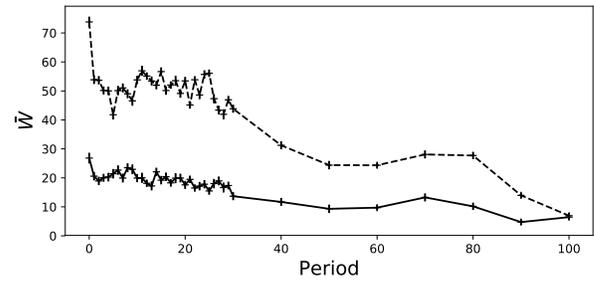


Figure 6: The average fitness of organisms starting their evolution in the day environment. The solid line shows their performance in the night environments, and the dashed line their performance in the day environment for different periods between switches (x axis) after 10,000 generations. Each data point is the average of 40 replicate experiments, error bars show the standard error.

those commands beneficial in one environment to be interspersed with commands that are beneficial in the other environment. Markov Brains do not have this sequential structure but instead form networks, which in our opinion much more closely resemble actual neural structures. The question now is how Markov Brains change structurally and functionally under these conditions.

Structural Amalgamation

We defined the degree of structural amalgamation as the degree to which the computational components (logic gates) are used in both environments. We find that at the end of evolution, the degree of structural amalgamation depends strongly on the speed at which environments switch periodically (see Figure 7).

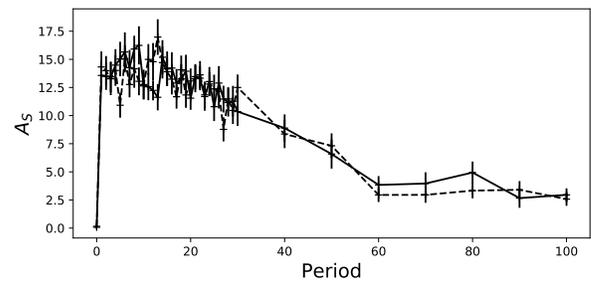


Figure 7: Structural amalgamation (A_S) for agents evolved under different periodic conditions (x axes). Agents that started to evolve in the night environment are shown as a solid line, and agents that started to evolve in the day environment are shown as dashed lines. Error bars indicate the standard error. Observe that the left most x-value (0) represents agents that never experienced any environmental change, and thus don't exhibit any structural amalgamation.

This result shows that fast changing environments indeed cause structural components of the brain to be used in both environments, while no changes or rarely changing environments do not cause structural amalgamation. This is not too surprising once the selection pressures these agents experience are considered. After a switch in environment, the selection pressure for components involved in the previous environment is lost, and thus genes (gates) start to drift. The only way this drift could be halted is by ensuring that these components are now also useful in the new environment.

Functional Amalgamation

As explained before, amalgamating the structural components of the brain does not imply that functions have to overlap or become intertwined as well. At the same time, structural amalgamation also does not preclude this from happening. We find that the degree of functional amalgamation does not depend on the frequency of periodic transitions (see Figure 8) and is generally low (less than 30% of states are experienced in both environment).

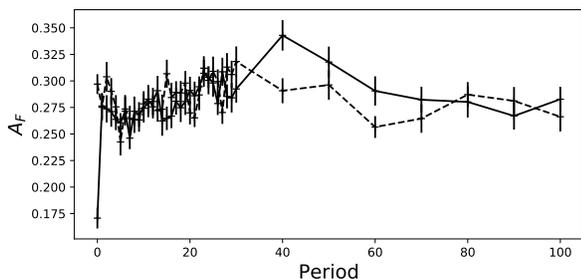


Figure 8: Functional amalgamation (A_F) for agents evolved under different periodic conditions (x axes). Agents that started to evolve in the night environment are shown as a solid line, and agents that started to evolve in the day environment are shown as dashed lines. Error bars indicate the standard error.

Discussion

We showed that MBs evolved to control an agent in periodically switching environments, adapt not only to each environment anew after a switch, but also experience structural and functional reorganization due to the switches. As long as environments change frequently enough (less than 50 generations in each environment) agents do not experience a significant loss in fitness after a switch, and can improve fitness over many switches. Natural organisms probably experience spontaneous or periodic switches on much larger timescales, however natural organisms certainly evolve much slower than the digital ones we used here. We argue that these two effects balance each other and that our results apply to natural organisms qualitatively if not necessarily quantitatively.

The difference between structural and functional amalgamation is an important observation. While agents amalgamate their components (genes/gates) to be used in both environments, they evolve to have state to state transitions organized so that states used in one context are separated from states used in another. In hindsight this might not be too surprising, since this form of organization separates mental states from each other as much as possible; environments are experienced differently instead of similarly. Also, at every time point, MBs can only be in one particular state, and a new input will change the state of the brain to a new state, given the current one. As a consequence, specific states can only be reached by a specific sequence of inputs. When the agent finds itself in a new environment that works differently than the one it evolved in, it is unlikely that a similar sequence of inputs will ever be observed. Instead, a different stream of inputs will necessarily lead to a different sequence of brain states.

What is particularly surprising, is the finding that even brains that have never experienced a different environment, also do not have significant functional amalgamation (see Figure 7 no transitions). Brains that have only evolved in one environment have no direct selection pressure against functional amalgamation. As a consequence, there should be no particular preference or aversion to functional amalgamation, however we still observe that only a few states are shared between both environments. This suggests that agents experience the environment they are evolved in properly, but when they encounter a new environment, they are experiencing distinctively different states. Again, figuratively speaking, when you see something unknown, instead of trying to explain how this new perception fits to something you already know you instead conclude that you haven't seen this before. The state to state transitions for evolved MBs resemble this principle. Observe that this is at best a loose analogy, albeit an interesting one.

We predict that sufficiently short periodic environmental changes should have a similar amalgamating effect on the structure of natural brains while preserving functional separation. Analyzing the anatomy of natural brains across differently evolved species should reveal this principle to be applicable to natural organisms as well.

Further investigations will consider the effect of periodic and aperiodic switching of environments on the evolvability and robustness of the computational systems. Similarly, we are curious about the effect of periodic transitions on agents that experience both the day and night environment during their lifetime. These studies will shed more light on the effect of evolutionary transitions on neural architectures.

References

- Bohm, C., CG, N., and Hintze, A. (2017). MABE (modular agent based evolver): A framework for digital evolution research. *Proceedings of the European Conference of Artificial Life*.

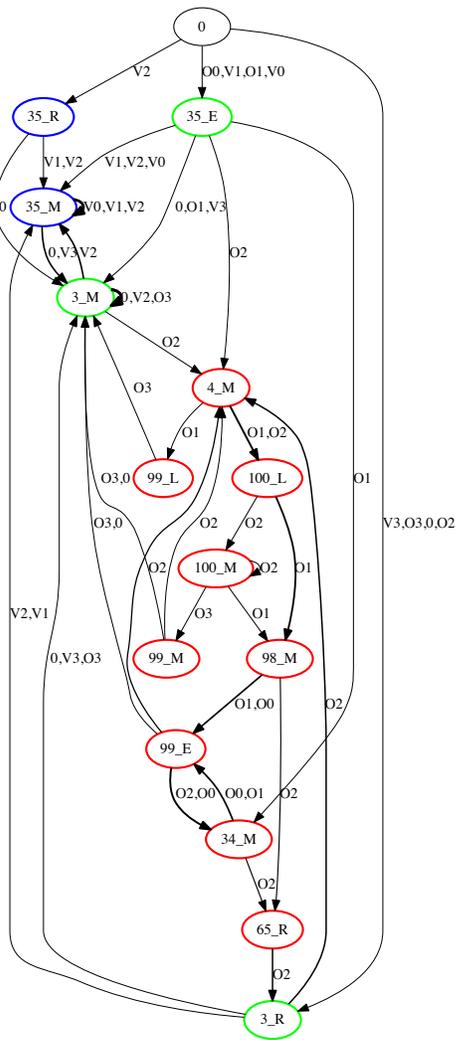


Figure 9: A state to state transition graph for an agent evolved only in the night environment. The nodes represent individual brain states and their corresponding actions (E eating, M for moving forward, R for turning right, and L for turning left). Arrows are labeled according to the sensory inputs experienced during the lifetime of the agent. A label starting with O indicates an olfactory input, while V indicates a vision signal. The number trailing the O or V specifies the distance of the food. After testing the agent multiple times in both environments, the states experienced during the day environment are colored in blue, and the night environment colored in red. The states visited in both environments are colored in green. The width of the edges indicates how often each transition was observed relatively to each other. Agents start in state 0.

Edlund, J. A., Chaumont, N., Hintze, A., Koch, C., Tononi, G., and Adami, C. (2011). Integrated information increases with fitness in the evolution of animats. *PLoS computational biology*, 7(10):e1002236.

Ellefsen, K. O., Mouret, J.-B., and Clune, J. (2015). Neural modularity helps organisms evolve to learn new skills without forgetting old skills. *PLoS Comput Biol*, 11(4):e1004128.

Gerkema, M. P., Davies, W. I., Foster, R. G., Menaker, M., and Hut, R. A. (2013). The nocturnal bottleneck and the evolution of activity patterns in mammals. *Proceedings of the Royal Society of London B: Biological Sciences*, 280(1765):20130508.

Hintze, A. and Adami, C. (2008). Evolution of complex modular biological networks. *PLoS Comput Biol*, 4(2):e23.

Hintze, A. and Adami, C. (2010). Modularity and anti-modularity in networks with arbitrary degree distribution. *Biology direct*, 5(1):32.

Hintze, A., Edlund, J. A., Olson, R. S., Knoester, D. B., Schossau, J., Albantakis, L., Tehrani-Saleh, A., Kvam, P., Sheneman, L., Goldsby, H., et al. (2017). Markov brains: A technical introduction. *arXiv preprint arXiv:1709.05601*.

Hüsken, M., Igel, C., and Toussaint, M. (2002). Task-dependent evolution of modularity in neural networks. *Connection Science*, 14(3):219–229.

Kashtan, N. and Alon, U. (2005). Spontaneous evolution of modularity and network motifs. *Proceedings of the National Academy of Sciences of the United States of America*, 102(39):13773–13778.

Kvam, P., Cesario, J., Schossau, J., Eisthen, H., and Hintze, A. (2015). Computational evolution of decision-making strategies. *arXiv preprint arXiv:1509.05646*.

Laakso, J., Kaitala, V., and Ranta, E. (2001). How does environmental variation translate into biological processes? *Oikos*, 92(1):119–122.

Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936):139–144.

Li, Y. and Wilke, C. O. (2004). Digital evolution in time-dependent fitness landscapes. *Artificial Life*, 10(2):123–134.

Maor, R., Dayan, T., Ferguson-Gow, H., and Jones, K. (2017). Temporal niche expansion in mammals from a nocturnal ancestor after dinosaur extinction. *bioRxiv*, page 123273.

Marstaller, L., Hintze, A., and Adami, C. (2013). The Evolution of Representation in Simple Cognitive Networks. *Neural Computation*, 25(8):2079–2107.

Misevic, D., Ofria, C., and Lenski, R. E. (2006). Sexual reproduction reshapes the genetic architecture of digital organisms. *Proceedings of the Royal Society of London B: Biological Sciences*, 273(1585):457–464.

Moran, P. A. P. (1958). Random processes in genetics. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 54, pages 60–71. Cambridge Univ. Press.

- Ofria, C. and Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial life*, 10(2):191–229.
- Olson, R. S., Hintze, A., Dyer, F. C., Knoester, D. B., and Adami, C. (2013). Predator confusion is sufficient to evolve swarming behaviour. *Journal of The Royal Society Interface*, 10(85):20130305–20130305.
- Santini, L., Rojas, D., and Donati, G. (2015). Evolving through day and night: origin and diversification of activity pattern in modern primates. *Behavioral Ecology*, 26(3):789–796.
- Schlosser, G. and Wagner, G. P. (2004). *Modularity in development and evolution*. University of Chicago Press.
- Schossau, J., Adami, C., and Hintze, A. (2016). Information-Theoretic Neuro-Correlates Boost Evolution of Cognitive Systems. *Entropy*, 18(1):6–22.
- Travis, J. and Travis, E. (2002). Mutator dynamics in fluctuating environments. *Proceedings of the Royal Society of London B: Biological Sciences*, 269(1491):591–597.
- Werner, G. (2009). Viewing brain processes as critical state transitions across levels of organization: neural events in cognition and consciousness, and general principles. *BioSystems*, 96(1):114–119.
- Wilke, C. O., Ronnewinkel, C., and Martinetz, T. (2001). Dynamic fitness landscapes in molecular evolution. *Physics Reports*, 349(5):395–446.