This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Evolving inborn knowledge for fast adaptation in dynamic POMDP problems

PLEASE CITE THE PUBLISHED VERSION

<https://doi.org/10.1145/3377930.3390214>

PUBLISHER

Association of Computing Machinery (ACM)

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

© ACM 2020. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in GECCO '20: Proceedings of the 2020 Genetic and Evolutionary Computation Conference, <http://dx.doi.org/10.1145/3377930.3390214>

LICENCE

In Copyright

REPOSITORY RECORD

Ben-lwhiwhu, Eseoghene, Pawel Ladosz, Jeff Dick, Wen-Hua Chen, Praveen Pilly, and Andrea Soltoggio. 2020. "Evolving Inborn Knowledge for Fast Adaptation in Dynamic POMDP Problems". Loughborough University. <https://hdl.handle.net/2134/12278087.v1>.

Evolving Inborn Knowledge For Fast Adaptation in Dynamic POMDP Problems

Eseoghene Ben-Iwhiwhu
Department of Computer Science,
Loughborough University,
Loughborough, UK
e.ben-iwhiwhu@lboro.ac.uk

Pawel Ladosz
Department of Computer Science,
Loughborough University,
Loughborough, UK
p.ladosz2@lboro.ac.uk

Jeffery Dick
Department of Computer Science,
Loughborough University,
Loughborough, UK
j.dick@lboro.ac.uk

Wen-Hua Chen
Department of Aeronautical and
Automotive Engineering,
Loughborough University,
Loughborough, UK
w.chen@lboro.ac.uk

Praveen Pilly
HRL Laboratories, California, USA
pkpilly@hrl.com

Andrea Soltoggio
Department of Computer Science,
Loughborough University,
Loughborough, UK
a.soltoggio@lboro.ac.uk

ABSTRACT

Rapid online adaptation to changing tasks is an important problem in machine learning and, recently, a focus of meta-reinforcement learning. However, reinforcement learning (RL) algorithms struggle in POMDP environments because the state of the system, essential in a RL framework, is not always visible. Additionally, hand-designed meta-RL architectures may not include suitable computational structures for specific learning problems. The evolution of online learning mechanisms, on the contrary, has the ability to incorporate learning strategies into an agent that can (i) evolve memory when required and (ii) optimize adaptation speed to specific online learning problems. In this paper, we exploit the highly adaptive nature of neuromodulated neural networks to evolve a controller that uses the latent space of an autoencoder in a POMDP. The analysis of the evolved networks reveals the ability of the proposed algorithm to acquire inborn knowledge in a variety of aspects such as the detection of cues that reveal implicit rewards, and the ability to evolve location neurons that help with navigation. The integration of inborn knowledge and online plasticity enabled fast adaptation and better performance in comparison to some non-evolutionary meta-reinforcement learning algorithms. The algorithm proved also to succeed in the 3D gaming environment Malmo Minecraft.

CCS CONCEPTS

• **Computing methodologies** → **Lifelong machine learning**; **Neural networks**; *Online learning settings*; *Partially-observable Markov decision processes*; *Genetic algorithms*;

KEYWORDS

lifelong learning, adaptive agent, self modifying network, neuroevolution, neuromodulation, few-shots learning, Hebbian learning

ACM Reference Format:

Eseoghene Ben-Iwhiwhu, Pawel Ladosz, Jeffery Dick, Wen-Hua Chen, Praveen Pilly, and Andrea Soltoggio. 2020. Evolving Inborn Knowledge For Fast Adaptation in Dynamic POMDP Problems. In *Genetic and Evolutionary Computation Conference (GECCO '20)*, July 8–12, 2020, Cancun, Mexico. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3377930.3390214>

1 INTRODUCTION

The field of deep reinforcement learning (RL) has showcased amazing results in recent time, solving tasks in robotic control [4, 12], games [15] and other complex environments. Despite such successes, deep RL algorithms are sample inefficient and sometimes unstable. Furthermore, they usually perform sub-optimally when dealing with sparse reward and partially observable environments. One further limitation of deep RL is when rapid adaptation to changing tasks (dynamic goals) is required. Established methods only work well in fixed task environments. In an attempt to solve this problem, deep meta-reinforcement learning (meta-RL) methods [5, 6, 20, 31, 34] were specifically devised. However, these methods are largely evaluated on dense reward, fully observable MDP environments, and perform sub-optimally in sparse reward, partially observable environments.

One key aspect in achieving fast adaptation in dynamic partially observable environments is the presence of appropriate learning structures and memory units that fits the specific class of learning problems. Therefore, standard model-free RL algorithms do not perform well in dynamic environments because they are tabula-rasa systems. They hold no knowledge in their architectures to allow a fast and targeted learning when a change in the environment occurs. Upon a task change, these algorithms will try to randomly explore the action space to relearn from scratch a different, new policy. On the other hand, model-based RL, holds knowledge of the structure of the environment, which in turn allows for rapid adaptation to changes in the environment, but such a knowledge needs to be built manually into the system.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

GECCO '20, July 8–12, 2020, Cancun, Mexico

© 2020 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-7128-5/20/07...\$15.00

<https://doi.org/10.1145/3377930.3390214>

In this paper, we investigate the use of neuroevolution to autonomously evolve inborn knowledge [26] in the form of neural structures and plasticity rules with a specific focus on dynamic POMDPs that have posed challenges to current RL approaches. The neuroevolutionary approach that we propose is designed to solve rapid adaptation to changing tasks [26] in complex high dimensional partially observable environments. The idea is to test the ability of evolution to build an unconstrained neuromodulated network architecture with problem-specific learning skills that can exploit the latent space provided by an autoencoder. Thus, in the proposed system, an autoencoder serves as a feature extractor that produces low dimensional latent features from high dimensional environment observations. A neuromodulated network [25] receives the low dimensional latent features as input and produces the output of the system, effectively acting as high level controller. Evolved neuromodulated networks have shown computational advantages in various dynamic task scenarios [25, 26].

The proposed approach is similar to that proposed in [1]. One key novelty is that our approach seeks to evolve selective plasticity with the use of modulatory neurons, and therefore, to evolve problem-specific neuromodulated adaptive systems. The relationships among image-pixel inputs and control actions in POMDPs is highly nonlinear and history dependent, therefore, an open question is whether neuroevolution can exploit latent features to evolve learning systems with inborn knowledge. Thus, we test the hypothesis that a neuromodulated evolved network can discover neural structures and their related plasticity rules to encode required memory and fast adaptation mechanisms to compete with current deep meta-RL approaches.

We call the proposed system a Plastic Evolved Neuromodulated Network with Autoencoder (PENN-A), denoting the combination of the two neural components. We evaluate our proposed method in a POMDP environment where we show better performance in comparison to some non-evolutionary deep meta-reinforcement learning methods. Also, we evaluated the proposed method in the Malmo Minecraft environment to test its general applicability.

Two interesting findings from our experiments are that (i) the networks acquire through evolution the ability to recognise reward cues (i.e. environment cues that are associated with survival even when reward signals are not given) and (ii) the networks can evolve *location* neurons that help solving the problem by detecting, and becoming active at, specific location of the partially observable MDP. The evolved network topology allows for richer dynamics in comparison to fixed architectures such as hand-designed feed-forward or recurrent networks.

The next section reviews the related work. Following that, a formal task definition is presented. Next is the description of the proposed method employed in this work, followed by the evaluation of results. The PENN-A source code is made available at: <https://github.com/dlpbc/penn-a>.

2 RELATED WORK

In reinforcement learning (RL) literature, meta-RL methods seek to develop agents that adapt to changing tasks in an environment or a set of related environments. Meta-RL [23, 24] is based on the general idea of meta-learning [2, 10, 30] applied to the RL domain.

Recently, deep meta-RL has been used to tackle the problem of rapid adaptation in dynamic environments. Methods such as [5, 6, 14, 17, 20, 31, 34] use deep RL methods to train a meta-learner agent that adapts to changing tasks. These methods are mostly evaluated in dense reward, fully observable MDP environments. Furthermore, most methods are either memory based [5, 14, 31] or optimization based [6, 34]. Optimization based methods seek to find an optimal initial set of parameters (e.g. for an agent network) across tasks, which can be fine-tuned with a few gradient steps for each specific task presented to it. Therefore, a small amount of re-training is required to enable adaptation to every change in task. Memory based methods (implemented using a recurrent network or temporal convolution attention network) do not necessarily require fine tuning after initial training to enable adaptation. This is because memory-based agents learn to build a memory of past sequence of tasks and interactions, thus enabling them to identify change in task and adapt accordingly.

In the past, neuroevolution methods have been employed to solve RL tasks [13, 28], including adapting to changing tasks [3, 25] in partially observable environments. These methods were evaluated in environments with high level feature observations. Recently, several approaches have been introduced that combine deep neural networks and neuroevolution to tackle high dimensional deep RL tasks [1, 9, 16, 22, 29]. These approaches can be divided into two major categories. The first category uses neuroevolution to optimize the entire deep network end to end [18, 19, 22, 29]. The second category splits the network into parts (for example, a body and controller) where some part(s) (e.g. body) are optimized using gradient based methods and other part(s) (e.g. controller) are evolved using neuroevolution methods [1, 9, 16]. Current deep neuroevolution methods are usually evaluated in fully observable MDP environments, where the task is fixed. Furthermore, after the training phase is completed, the weights of a trained network are fixed (the same is true for standard deep RL). The recent attention to neuroevolution for deep RL aims to present such approaches as a competitive alternative to standard gradient based deep RL methods for fixed task problems.

In the past, neural network based agents employing Hebbian-based local synaptic plasticity have been used to achieve behavioural adaptation with changing tasks [3, 8, 25]. Such methods use a neuroevolution algorithm to optimize the parameters of the network when producing a new generation of agents. As an agent interacts with an environment during its lifetime in training or testing, the weights are adjusted in an online fashion (via a local plasticity rule), enabling adaptation to changing tasks. In [3, 8] this technique was employed, and further extended to include a mechanism of gating plasticity via neuromodulation in [25]. These methods were evaluated in environments with low dimensional observations (with high level features) and not compared with deep (meta-)RL algorithms.

3 TASK DEFINITION

A POMDP environment E , defined by a sextuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \mathcal{O}, \Omega)$ is employed in this work. $\mathcal{S}$ defines the state set, $\mathcal{A}$ the action set, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ the environment dynamics, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ the reward function, $\mathcal{O}$ the observation set, and Ω the function that maps observations to states.

Figure 1: Illustration of a dynamic environment and required behavior of a learning agent. An agent is required to learn to perform optimally and then exploit the learned policy until a change in the environment occurs, at which point the agent needs to learn again before exploiting.

The environment E contains a number of related tasks. A task $\mathcal{T}_i$ is sampled from a distribution of tasks $\mathcal{T}$. The task distribution $\mathcal{T}$ can either be discrete or continuous. A sampled task is an instance of the partially observable environment E . The configuration of the environment (for example, the goal or reward function) varies across each task instance. An optimal agent is required to adapt its behaviour to task changes in the environment (and maximize accumulated reward), only from few interactions in the environment. When presented with a task $\mathcal{T}_i$, an optimal agent should initially explore, and subsequently exploit when the task is understood. When the task is changed (a new task $\mathcal{T}_j$ sampled from $\mathcal{T}$), the agent needs to re-explore the environment in few-shots, and then to start exploiting again when the new task has been understood.

In each task, an episode is defined as the trajectory τ of an agent interactions in the environment, terminating at a terminal state. A trial consist of two or more tasks sampled from $\mathcal{T}$. The total number of episodes in a trial is kept fixed. A trial starts with an initial task $\mathcal{T}_i$ that runs for a number of episodes, and then the task is changed to other tasks (one after another) at different points within the trial (see Figure 1). The points at which a task change occurs are stochastically generated, and the task is changed before the start of the next episode. For example, when the number of tasks is set as 2 (i.e. $\mathcal{T}_i$ and $\mathcal{T}_j$), the trial starts with task $\mathcal{T}_i$ which runs for a number of episodes, and it is replaced by task $\mathcal{T}_j$ for the remaining episodes in the trial. An agent is iteratively trained, with each iteration consisting of a fixed number of trials. The subsections below describes two environments where the proposed system is evaluated.

3.1 The Configurable Tree Graph Environment

The configurable tree graph (CT-graph) environment is a graph abstraction of a decision making process. The complexity of the environment is specified via configuration parameters; branching factor b and depth d , controlling the width and height of the graph. Additionally, it can be configured to be fully or partially observable. It contains the following types of state; start, wait, decision, end (leaf node of graph) and crash. Each observation $o \in \mathcal{O}$ is a 12x12 grey-scale image. The total number of end states grows exponentially as the depth d of the graph increases (see Figure 2A and B).

In the experiments in this study, partial observability is configured by mapping all wait states to the same observation, and all decision states to the same observation. Also, b is set to 2. Therefore, each decision state has two choices, splitting into two sub-graphs. The discrete action space is defined as; *choice 1*, *choice 2*, *wait action*, thus discrete. The *wait action* is the correct action in a wait state. In a decision state, *choice 1* or *choice 2* is the correct subset from which to select. All incorrect actions lead to the crash state and episode termination.

An agent starts an episode in the start state, and the episode is completed when the agent traverses the graph to an end state or takes a wrong action in a state. Once an agent transitions from one state to the next, it cannot go back. In a task instance, one of the end states is set as the goal location. An agent receives a positive reward when it traverses to the goal location, and reward of 0 at other non-goal states. The agent may receive a negative reward in a crash state.

3.2 Malmo Minecraft Environment

Malmo [11] is an AI research platform built on top of Minecraft video game. The platform is configurable, and it enables the construction of various worlds in which AI agents can be evaluated. In this work, a double T-maze was constructed, with discrete action space *left turn*, *right turn* and *forward action*. A task is defined based on the maze ends, requiring the agent to navigate to a specific maze end (goal location). The maze end that is set as the goal location varies across tasks. The agent only receives a positive reward when they navigate to the maze end that is the goal location. It receives reward of 0 in every other time step. If the agent runs into a wall, the episode is terminated and it receives a negative reward. The agent receives a visual observation of its current view at each time step (hence it does not fully observe the entire environment). Each observation is a 32x32 RGB image based on a first-person view of the agent at each time step.

4 METHODS

We seek to develop an agent that is capable of continual adaptation through its life time (across episodes) - exploring, exploiting, re-exploring when the task changes and exploiting again. The system (specifically the controller or decision maker) is evolved to acquire knowledge about both the invariant and variant aspects of an environment (e.g. changing tasks).

The agent is modelled using two neural components with separate parameters and objectives; a deep network F_θ (used as a feature extractor and parameterized by θ) and a neuromodulated network G_ϕ (serving as a controller and parameterized by ϕ). Both components make up the overall system model $\mathcal{M}_{\theta, \phi}$. See Figure 3 for a general system overview. The presented architectural style is similar to a standard deep RL setup. However, it differs on two fronts; (i) the controller is a neuromodulated network (described in Section 4.2) rather than a standard neural network, (ii) the training setup combines gradient based optimization method [21, 32], gradient free optimization method (neuroevolution [27, 33]), and Hebbian-based synaptic plasticity to train the system. Using this setup, each neural component therefore contains its own objective function. An autoencoder network was employed as the feature

Figure 2: Environments (note, during execution, goal location is dynamic across episodes). (A) CT-graph instance, $b = 2$ and $d = 2$. (B) CT-graph instance, $b = 2$ and $d = 3$. (C) Malmo Minecraft instance (a double T-Maze), bird's eye view on top, with some sample observations at the bottom. The maze-end with the teal colour is the goal location.

Figure 3: System overview, showcasing the feature extractor and controller components. In the controller, white and blue nodes are standard and modulatory neurons respectively. Modulatory connections facilitates selective plasticity in the network.

extractor, thus enabling the use of Mean Squared Error (MSE) or Binary Cross Entropy (BCE) objective function:

$$\arg \min_{\theta} \frac{1}{n} \sum_{i=1}^n (F_{\theta}(o_i) - o_i)^2$$

$$\arg \min_{\theta} -\frac{1}{n} \sum_{i=1}^n o_i \cdot \log(F_{\theta}(o_i)) + (1 - o_i) \cdot \log(1 - F_{\theta}(o_i))$$

where n is the number of training observations and $F_{\theta}(o_i)$ is the output of the autoencoder for observation i (reconstructed observation). Each agent in the population uses the same feature extractor. The fitness function of the evolutionary algorithm is given by:

$$\arg \max_{\phi} \sum_{\mathcal{T}_i \sim \mathcal{T}} \sum_{ep=1}^z R(\tau_{ep})$$

$\mathcal{T}_i$ represents a task sampled from the task distribution $\mathcal{T}$, and a single trial consist of two tasks as defined in Section 3. Also, z is the number of episodes in which a task is kept fixed within a trial. It is stochastically generated and may differ between tasks in a trial within an interval. $R(\tau_{ep})$ is the accumulated reward of a trajectory of an episode ep , defined as:

$$R(\tau_{ep}) = \sum_{t=0}^k \mathcal{R}(s_t, G_{\phi}(F_{\theta}^{enc}(o_t))) \quad (1)$$

where $\mathcal{R}(s, a)$ is the reward function that takes state and action as arguments and produces a scalar reward value. F_{θ}^{enc} is the same autoencoder feature extractor network earlier described, but denoting that we only want the output from the encoder (the latent features). Also, t represents discrete time steps and k is the length of the trajectory of an episode.

4.1 Feature Extractor

This neural component of the system is tasked with learning a good latent representation of the observations from the environment, which can be fed to the controller as input. In the CT-graph experiments, a fully connected autoencoder was employed (two layers encoder and decoder respectively). In the Malmo Minecraft experiments, a convolutional autoencoder was employed (four layers encoder and decoder respectively).

4.2 Control Network (Decision Maker)

This neural component takes the latent features of the feature extractor as its input, and produces an output which serves as the final output of the system (the action or behaviour of the system). It is a neuromodulated network (see Section 4.2.1), that reproduces the model introduced in [25]. The network can evolve two neuron

types - a standard and a modulatory neuron. The output neuron(s) always belong to the standard neuron type.

The control network is parameterized by ϕ . Unlike θ (which represents only the weights of the feature extractor network), ϕ consists of the weights, architecture and the co-efficients of Hebbian-based plasticity rule (described in 4.2.2) of the network, and it is evolved. Therefore, evolution is tasked with finding the architecture and plasticity rules, including selective plasticity enabled by modulatory neurons to target neurons. The large search space that is granted to evolution allows for rich dynamics that include memory in the form of both recurrent connections and temporary values of rapidly changing modulated weights.

The agent is never fed the reward signal explicitly. The reward signal is only used by the evolutionary process for the fitness evaluation, which in turn drives the selection process. Therefore, the network is tasked to learn the discovery of reward cues implicitly from the visual observations in the environment.

4.2.1 Neuromodulated Network Dynamics. Though processing is distributed across neurons, a standard neural network usually contains one type of neuron - where the dynamics of each neuron is homogeneous across the network. In a neuromodulated network, there can be two types of neurons, each type having different dynamics - thus heterogeneous. The two types of neurons are standard neurons and modulatory neurons [25]. The standard neurons have the same dynamics as the ones in standard neural network. The modulatory neurons are used to dynamically regulate plasticity in the network.

Each neuron i has one standard and one modulatory activation value that represent the weighted amount of standard and modulatory activity they receive from other neurons (see Equations 2 and 3). $a_{std,i}$ is the output signal of neuron i that is propagated to other neurons in its outgoing connections (this is true for both standard and modulatory neurons). $a_{mod,i}$ is used internally by the neuron itself to regulate the Hebbian-based plasticity of the incoming connections from other standard neurons, as described in Section 4.2.2. The framework allows for selective plasticity in the network, as parts of the network may become plastic or not plastic depending on the change of the modulatory activation signals over time. In turn, the final action of the network is affected in the current and future time steps - thus enabling adaptation.

$$a_{std,i} = \tanh \frac{\sum_{j \in \text{std}} w_{ji} a_{std,j}}{2} \quad (2)$$

$$a_{mod,i} = \tanh \frac{\sum_{j \in \text{mod}} w_{ji} a_{std,j}}{2} \quad (3)$$

4.2.2 Neuromodulated Hebbian Plasticity. The Hebbian synaptic plasticity of the control network is governed by the Equations 4, 5 and 6. A, B, C, D, α are the coefficients of the plasticity rule. The update of a weight is dependent pre-synaptic and post-synaptic standard activations, the plasticity co-efficients, and the post-synaptic modulatory activation. This is true for all weights in the neuromodulated network.

$$w_{ij} = w_{ij} + \Delta w_{ij} \quad (4)$$

$$\Delta w_{ij} = a_{mod,j} \cdot \delta w_{ij} \quad (5)$$

$$\delta w_{ij} = \alpha \cdot (A \cdot a_{std,i} \cdot a_{std,j} + B \cdot a_{std,i} + C \cdot a_{std,j} + D) \quad (6)$$

5 RESULTS AND ANALYSIS

Figures 4 and 6 show the results of the experiments in the CT-graph environment. Figure 7 shows the results of the experiment in the Malmo Minecraft environment. In addition, we present results obtained in the Malmo Minecraft environment (Figure 7), evaluating the general applicability of PENN-A.

5.1 Performance in CT-graph Environments

The proposed method (PENN-A) was evaluated on depth 2 and 3 CT-graph environments, with branching factor of 2. The controller was evolved for 200 generations, with population of 600 and 800 for depth 2 and 3 experiments respectively. Tournament selection with segment size of 5 was employed. Each controller was evaluated for 4 trials, with 100 episodes and 2 tasks per trial. The initial task is changed between episodes 35 and 65, determined stochastically for each trial. The depth 2 CT-graph experiment was employed as a baseline, and we compared PENN-A against some recent deep meta-RL methods (each with its own experimental setup). The depth 3 CT-graph experiment was employed to evaluate the PENN-A in a more complex configuration of the environment.

In order to ensure compatibility in the result presented across all methods, the number of evaluations (horizontal axis) were scaled to the approximate number of episodes equivalent. Additionally, the vertical axis is the average accumulated reward across all trials and episodes. In the depth 2 CT-graph result (Figure 4), we see that PENN-A performs optimally when compared to deep meta-RL methods; optimization-based (MAML [6] and CAVIA [34]) and memory-based (RL² [5] without extra input). Only the observations were fed as input to the neural network for all methods including PENN-A. We hypothesize the deep meta-RL methods perform sub-optimally due to the partial observability of the environment. When extra input (the reward, previous time step action and done state) are concatenated to the observation and fed to the RL² method (which is vanilla setup), then it is able to perform optimally (see Figure 5). We hypothesize that RL² exploits the actions fed as input to the network, ignoring the observations and other parts of the input. This reduces the problem complexity in comparison to conditions where only the observations are fed as input.

Figure 6 presents result for a depth 3 CT-graph. We present result for only PENN-A in depth 3 CT-graph (a more difficult problem than depth 2 CT-graph) since the other methods performed sub-optimally in depth 2 CT-graph. We again observe PENN-A performing optimally in the more difficult CT-graph setting.

5.1.1 Network Analysis. To better understand the evolved solution and how the network implements policies, we analyzed the best performing networks after evolution in a depth 2 CT-graph environment. While different evolutionary runs produced highly different networks, we observed interesting patterns in the neural activations. For one network of 11 neurons (including the output neuron), the absolute activation value distribution (across trials and episodes per time step) is plotted for each neuron in Figure 8. We see that the absolute activation distribution of some neurons

Figure 4: Results for a CT-graph with depth 2. PENN-A is compared against non-evolutionary meta-RL methods.

Figure 5: RL^2 in the CT-graph with depth 2. The method is run with extra input to the network (reward, done state, and previous time step action concatenated with current observation to form input).

Figure 6: PENN-A performance in a CT-graph with depth 3.

Figure 7: Malmo Minecraft result.

Figure 8: Absolute activation values distribution (across trials and episodes) per time step of a sample evolved controller. (A) This neuron is active specifically at decision states (steps 3 and 5), while it remains low at wait states. (B) This neuron clearly identifies wait states (steps 2, 4 and 6) and remains inactive otherwise.

Figure 9: Distributions of the activation values for each neuron in a sample network when the goal location (reward) is found and vice versa. The neurons highlighted in green bounding box react differently to the presence or absence of reward cues from observation. (A) heat maps of grey-scale CT-graph observations. The top image is the observation presented when the goal location is found, with a bright square reward cue. The bottom image is the observation when the goal location is not found, with the reward cue absent. (B) Neurons 11 and 13 show complementary firing patterns based on reward cues. (C) Neurons 1, 2, 5 and 6 are active when a reward cue is observed, and have little or no activity when the reward cue is not observed.

are high at specific time steps, i.e., at specific points within the graph environment (see Figure 8A and B) - and therefore function as *location neurons*. Such kind of location neurons had been previously discovered in an evolutionary setting in [7]. In the current experiments, it is worth noting that location neurons are designed by evolution to exploit latent features and possibly help action-selection in a high-dimensional dynamic POMDP. In particular, the neuron in Figure 8A is active at decision states, while the neuron in Figure 8B is active at wait states.

One aspect of our experimental setting is that the reward signal is not fed to the network, but the environment provides reward cues embedded in the observations as it is shown in Figure 9A where a bright square represents a reward. The actual reward value is only accumulated in the fitness function, and is therefore not explicitly visible to the network. The surprising results that networks evolved to explore the environment and find the reward even if no reward signal was given suggests that the reward cue was recognised. In fact, in the example shown in Figure 9B, some neurons fire positively when a reward cue is observed and negatively when not observed or vice versa. Other neurons fire when a reward cue is observed and have little or no firing when not observed (see Figure 9C). Not all evolved networks appeared to have *reward neurons*. Nevertheless, the examples that evolved such reward cues detectors demonstrate that evolution is able to incorporate invariant knowledge of the environment to optimize the policy, in this case, reward seeking behaviour and fast adaptation speed to changing task.

5.2 Performance in Malmo Minecraft

To further assess the validity of our method, it is important to use a different benchmark environment with a larger input and RGB observations that offered a different feature space, hence the Malmo Minecraft environment. The controller was evolved with population size of 800, in 400 generations. The same selection strategy as used in the CT-graph was employed. Each controller was evaluated for 8 trials, with 50 episodes and 3 tasks per trial. The task is changed at two stochastically generated points within the trial. The result is presented in Figure 7, keeping the same axes format as with the results presented for the CT-graph environment. Again, the proposed method was able to perform optimally with a high average reward score, demonstrating its capability to scale to other high dimensional, less abstract environments.

6 CONCLUSION

This paper introduced an evolutionary design method for fast adaptation in POMDP environments. The system combines a feature extractor network and an evolved neuromodulated network with the aim of acquiring specific inborn knowledge and structure via evolution. While the suitability of evolved neuromodulated networks to solve environments with changing task was known [25, 26], we demonstrated that such advantages are scalable to high dimensional input spaces, and can be used in combination with an autoencoder. The results showed performance that compare or surpass some deep meta-RL algorithms. Interestingly, the evolved networks were capable of learning to recognise implicit reward cues, and therefore

Table 1: Network architecture for CT-graph experiments

Layer	Activation	Units
Input	N/A	144
FC	ReLU	64
FC	ReLU	16
FC	ReLU	64
FC	ReLU	144

Table 2: Network architecture for Malmo Minecraft experiments

Layer	Activation	Kernel	Stride	Channels
Input	N/A	N/A	N/A	N/A
Conv2D	ReLU	3x3	2	16
Conv2D	ReLU	3x3	2	32
Conv2D	ReLU	3x3	2	32
Conv2D	ReLU	3x3	2	8
ConvTranspose2D	ReLU	3x3	2	32
ConvTranspose2D	ReLU	3x3	2	32
ConvTranspose2D	ReLU	3x3	2	16
ConvTranspose2D	Sigmoid	4x4	2	3

could explore the environment in search for the goal location without an explicit reward signal. This ability that was acquired by the networks through evolution is an example of inborn knowledge that allow networks to be born with the knowledge of what are reward cues. Subsequently, this information can be used to direct fast adaptation when the optimal policy changes (e.g. the task change). The networks also evolved location neurons to help the deployment of a policy by distinguishing different states in the underlying MDP. We speculate that this approach might be promising when a combination of inborn knowledge and online learning are required to perform optimally in rapidly changing environments.

A EXPERIMENTAL SETTINGS

The PENN-A source code containing the experimental setup is made available at: <https://github.com/dlphc/penn-a>.

A.1 Feature Extractor

Mean Squared Error (MSE) was employed as the loss function across all CT-graph experiments, with a vanilla SGD (learning rate of 0.001) optimizer. Likewise, Binary Cross Entropy (BCE) was employed as the loss function in the Malmo Minecraft experiments, using RM-Sprop (learning rate of 0.0005) optimizer. The network architectures for the CT-graph and Malmo Minecraft experiments are presented in Table 1 and 2. The double horizontal line in both tables highlight the split between the encoder and decoder (i.e. specifications on top of the double line are for the encoder and likewise bottom for the decoder) of each autoencoder. In the CT-graph experiments where a Fully Connected (FC) autoencoder was employed, each input observation is flattened into a vector before feeding it to the network.

A.2 Control Network

Excluding population size and number of generations, the evolutionary parameters from [25] were followed.

The latent features from the feature extractor network were scaled between 0 and 1. To further restrict the latent features, a transformation operation was applied to the scaled latent features v before it was fed to the control network as shown below.

$$w = \begin{cases} 1 & \text{if } \sigma'(v) > 1 \\ 0 & \text{if } \sigma'(v) < 0 \\ \sigma'(v) & \text{otherwise} \end{cases}$$

$$\sigma'(v) = \log \frac{v}{1-v}$$

where $\sigma'(v)$ is an inverse sigmoid operation on v , and w is the transformed feature space. The scaling and transformation operations were performed independently of the feature extractor optimization (i.e. the operations were applied on copies of the latent features), and were applied across all experiments.

In this work, both evaluation environments were designed to work with discrete action space (3 actions each). Therefore, a single output neuron was employed across all experiments. The tanh activation value of the neuron was discretized to produce the actions of an agent. An activation value within the interval $[-1.0, -0.33]$ mapped to one action, the interval $[-0.33, +0.33]$ mapped to another action, and the interval $(+0.33, 1.0]$ mapped to the last action.

ACKNOWLEDGMENTS

This material is based upon work supported by the United States Air Force Research Laboratory (AFRL) and Defense Advanced Research Projects Agency (DARPA) under Contract No. FA8750-18-C-0103. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force Research Laboratory (AFRL) and Defense Advanced Research Projects Agency (DARPA).

REFERENCES

- [1] Samuel Alvernaz and Julian Togelius. 2017. Autoencoder-augmented neuroevolution for visual doom playing. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 1–8.
- [2] Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. 1992. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*. Univ. of Texas, 6–8.
- [3] Jesper Blynel and Dario Floreano. 2002. Levels of dynamics and adaptive behavior in evolutionary neural controllers. In *Proceedings of the seventh international conference on simulation of adaptive behavior on From animals to animats*. MIT Press, 272–281.
- [4] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. 2016. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*. 1329–1338.
- [5] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. 2016. RL²: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779* (2016).
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 1126–1135.
- [7] Dario Floreano and Laurent Keller. 2010. Evolution of adaptive behaviour in robots by means of Darwinian selection. *PLoS Biol* 8, 1 (2010), e1000292.
- [8] Dario Floreano and Francesco Mondada. 1996. Evolution of plastic neurocontrollers for situated agents. In *Proc. of The Fourth International Conference on Simulation of Adaptive Behavior (SAB), From Animals to Animats*. ETH Zürich.
- [9] David Ha and Jürgen Schmidhuber. 2018. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems*. 2450–2462.

- [10] Sepp Hochreiter, A Steven Younger, and Peter R Conwell. 2001. Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks*. Springer, 87–94.
- [11] Matthew Johnson, Katja Hofmann, Tim Hutton, and David Bignell. 2016. The Malmo Platform for Artificial Intelligence Experimentation.. In *IJCAI*. 4246–4247.
- [12] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* (2015).
- [13] Gary McHale and Phil Husbands. 2004. Gasnets and other evolvable neural networks applied to bipedal locomotion. *From Animals to Animats* 8 (2004), 163–172.
- [14] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. A Simple Neural Attentive Meta-Learner. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=B1DmUzWAW>
- [15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fiedjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529.
- [16] Andreas Precht Poulsen, Mark Thorhauge, Mikkel Hvilshj Funch, and Sebastian Risi. 2017. DLNE: A hybridization of deep learning and neuroevolution for visual control. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 256–263.
- [17] Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. 2019. Efficient Off-Policy Meta-Reinforcement Learning via Probabilistic Context Variables. In *International Conference on Machine Learning*. 5331–5340.
- [18] Sebastian Risi and Kenneth O Stanley. 2019. Deep neuroevolution of recurrent and discrete world models. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 456–462.
- [19] Sebastian Risi and Kenneth O Stanley. 2019. Improving Deep Neuroevolution via Deep Innovation Protection. *arXiv preprint arXiv:2001.01683* (2019).
- [20] Jonas Rothfuss, Dennis Lee, Ignasi Clavera, Tamim Asfour, and Pieter Abbeel. 2019. ProMP: Proximal Meta-Policy Search. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SkxXCi0qFX>
- [21] David E Rumelhart, Geoffrey E Hinton, Ronald J Williams, et al. 1988. Learning representations by back-propagating errors. *Cognitive modeling* 5, 3 (1988), 1.
- [22] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864* (2017).
- [23] Juergen Schmidhuber, Jieyu Zhao, and MA Wiering. 1996. Simple principles of metalearning. *Technical report IDSLA 69* (1996), 1–23.
- [24] Nicolas Schweighofer and Kenji Doya. 2003. Meta-learning in reinforcement learning. *Neural Networks* 16, 1 (2003), 5–9.
- [25] Andrea Soltoggio, John A Bullinaria, Claudio Mattiussi, Peter Dürri, and Dario Floreano. 2008. Evolutionary advantages of neuromodulated plasticity in dynamic, reward-based scenarios. In *Proceedings of the 11th international conference on artificial life (Alife XI)*. MIT Press, 569–576.
- [26] Andrea Soltoggio, Kenneth O Stanley, and Sebastian Risi. 2018. Born to learn: the inspiration, progress, and future of evolved plastic artificial neural networks. *Neural Networks* (2018).
- [27] Kenneth O Stanley, Jeff Clune, Joel Lehman, and Risto Miikkulainen. 2019. Designing neural networks through neuroevolution. *Nature Machine Intelligence* 1, 1 (2019), 24–35.
- [28] Kenneth O Stanley and Risto Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary computation* 10, 2 (2002), 99–127.
- [29] Felipe Petroski Such, Vashisht Madhavan, Edoardo Conti, Joel Lehman, Kenneth O Stanley, and Jeff Clune. 2017. Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *arXiv preprint arXiv:1712.06567* (2017).
- [30] Sebastian Thrun and Lorien Pratt. 1998. Learning to learn: Introduction and overview. In *Learning to learn*. Springer, 3–17.
- [31] Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Rémi Munos, Charles Blundell, Dhharshan Kumaran, and Matt Botvinick. 2016. Learning to reinforcement learn, 2016. *arXiv preprint arXiv:1611.05763* (2016).
- [32] Paul J Werbos. 1982. Applications of advances in nonlinear sensitivity analysis. In *System modeling and optimization*. Springer, 762–770.
- [33] Xin Yao. 1999. Evolving artificial neural networks. *Proc. IEEE* 87, 9 (1999), 1423–1447.
- [34] Luisa Zintgraf, Kyriacos Shiarli, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. 2019. Fast Context Adaptation via Meta-Learning. In *International Conference on Machine Learning*. 7693–7702.