

Augmentative Topology Agents For Open-ended Learning

Muhammad Umair Nasir, Michael Beukman, Steven James and Christopher Cleghorn

Abstract—In this work, we tackle the problem of open-ended learning by introducing a method that simultaneously evolves agents and increasingly challenging environments. Unlike previous open-ended approaches that optimize agents using a fixed neural network topology, we hypothesize that generalization can be improved by allowing agents’ controllers to become more complex as they encounter more difficult environments. Our method, Augmentative Topology EPOET (ATEP), extends the Enhanced Paired Open-Ended Trailblazer (EPOET) algorithm by allowing agents to evolve their own neural network structures over time, adding complexity and capacity as necessary. Empirical results demonstrate that ATEP results in general agents capable of solving more environments than a fixed-topology baseline. We also investigate mechanisms for transferring agents between environments and find that a species-based approach further improves the performance and generalization of agents.

I. INTRODUCTION

Machine learning has successfully been used to solve numerous problems, such as classifying images [1], writing news articles [2, 3] or solving games like Atari [4] or chess [5]. While impressive, these approaches still largely follow a traditional paradigm where a human specifies a task that is subsequently solved by the agent. In most cases, this is the end of the agent’s learning—once it can solve the required task, no further progression takes place. *Open-ended learning* is a research field that takes a different view: rather than converge to a specific goal, the aim is to obtain an increasingly growing set of diverse and interesting behaviors [6, 7]. One approach is to allow both the agents, as well as the environments, to change, evolve and improve over time [8, 9]. This has the potential to discover a large collection of useful and reusable skills [10], as well as interesting and novel environments [11]. Open-ended learning is also a much more promising way to obtain truly general agents than the traditional single task-oriented paradigm [12].

The concept of open-ended evolution has been a part of artificial life (ALife) research for decades now, spawning numerous artificial worlds [13, 14, 15, 16, 17]. These worlds consist of agents with various goals, such

as survival, predation, or reproduction. Recently, open-ended algorithms have received renewed interest [7], with Stanley et al. [6] proposing the paradigm as a path towards the goal of human-level artificial intelligence.

A major breakthrough in open-ended evolution was that of *NeuroEvolution of Augmenting Topologies* (NEAT) [18], which was capable of efficiently solving complex reinforcement learning tasks. Its key idea was to allow the structure of the network to evolve alongside the weights, starting with a simple network and adding complexity as the need arises. This inspired future research about open-endedly evolving networks indefinitely [17]. Specifically, novelty search [19], used the idea of *novelty* to drive evolution, instead of traditional objective-based techniques. This in turn led to the emergence of quality diversity (QD) algorithms [20, 21, 22, 23], which are based on combining novelty with an objective sense of progress, where the goal is to obtain a collection of diverse and high-performing individuals.

While QD has successfully been used in numerous domains, such as robotic locomotion [24, 21, 25], video game playing [22] and procedural content generation [26, 27], it still is not completely open-ended. One reason for this is that the search space for phenotypical behavior characteristics (or behavioral descriptors) remains fixed [21]. A second reason is that in many cases, the environment remains fixed, which limits the open-endedness of the algorithm [9]. A way to circumvent this is to co-evolve problems and solutions, as is done by Minimal Criterion Coevolution (MCC) [8]. This co-evolutionary pressure allowed more complex mazes to develop, and better agents to solve them emerged, giving rise to an open-ended process.

However, MCC had some limits; for instance, it only allows new problems if they are solvable by individuals in the current population. This leads to only slight increases in difficulty, and complexity which only arises randomly. Taking this into account, *Paired Open-ended Trailblazer* (POET) [9] builds upon MCC, but instead allows the existence of unsolvable environments, if it was likely that some individuals could quickly learn to solve these environments. POET further innovates by transferring agents between different environments, to increase the likelihood of solving hard problems. While

*University of the Witwatersrand, South Africa

POET obtained state of the art results, its diversity slows down as it evolves for longer. Enhanced POET [28] adds improved algorithmic components to the base POET method, resulting in superior performance and less stagnation. Enhanced POET, however, uses agents with fixed topology neural network controllers. While this approach works well for simple environments, it has an eventual limit on the complexity of tasks it can solve: at some point of complexity, the fixed topology agents may not have sufficient capacity to solve the environments.

To address this issue, we propose *Augmentative Topology Enhanced POET* (ATEP), which uses NEAT to evolve agents with variable, and potentially unbounded, network topologies. We argue that fixed-topology agents will cease to solve environments after a certain level of complexity and empirically show that ATEP outperforms Enhanced POET (EPOET) in a standard benchmark domain. Finally, we find that using NEAT results in improved exploration and better generalization compared to Enhanced POET.

II. RELATED WORK

POET [9] and EPOET [28] are the founding algorithms of the field of *open-ended reinforcement learning*, building upon prior approaches such as MCC [8]. This has led to an explosion of new use cases such as PINSKY [29, 30], which uses POET on 2D Atari games. This approach extends POET to generate 2D Atari video game levels alongside agents that solve these levels. Quesy and Richardson [10] uses unsupervised skill discovery [31, 32, 33] in the context of POET to discover a large repertoire of useful skills. Meier and Mujika [34] also investigate unsupervised skill discovery through reward functions learned by neural networks. Other uses of POET include the work by Zhou and Vanschoren [35], who obtain diverse skills in a 3D locomotion task. POET has also been shown to aid in evolving robot morphologies [36] and avoiding premature convergence which is often the result when using handcrafted curricula. Norstein et al. [37] use MAP-Elites [21] to open-endedly create a structured repertoire of various terrains and virtual creatures.

Adversarial approaches are commonly adopted when developing open-ended algorithms. Dennis et al. [38] propose PAIRED, a learning algorithm where an adversary would produce an environment based on the difference between the performance of an antagonist and a protagonist agent. Domain randomization [39], prioritized level replay [40] and Adversarially Compounding Complexity by Editing Levels (ACCEL) [41] adopt

a similar adversarial approach, where teacher agents produce environments and student agents solve them.

Several domains and benchmarks have been proposed with the aim of encouraging research into open-ended, general agents. Team et al. [12] introduce the XLand environment, where a single agent is trained on 700k 3D games, including single and multi-agent games, resulting in zero-shot generalization on holdout test environments. Barthet et al. [42] introduced an autoencoder [43, 44] and CPPN-NEAT based open-ended evolutionary algorithm to evolve Minecraft [45, 46] buildings. They showed how differences in the training of the autoencoders can affect the evolution and generated structures. Fan et al. [47] create a Minecraft-based environment, MineDojo, which has numerous open-ended tasks. They also introduced MineCLIP as an effective language-conditioned reward function that plays the role of an automatic metric for generation tasks. Gan et al. [48] introduce the Open-ended Physics Environment (OPEN) to test learning representations, and tested many RL-based agents. Their results indicate that agents that make use of unsupervised contrastive representation learning, and impact-driven learning for exploration, achieve the best result.

III. ENHANCED POET

Since our method is heavily based on EPOET, we briefly describe this method, as well as the original POET algorithm. **POET** focuses on evolving pairs of agents and environments in an attempt to create specialist agents that solve particular environments. POET uses the 2D *Bipedal Walker Hardcore* environment from OpenAI Gym [49] as a benchmark. The first environment is a flat surface, and as evolution progresses, the environments become harder with the addition of more obstacles. POET also transfers agents across environments, which can prevent stagnation and leverage experience gained on one environment as a step towards solving another. An Environment-Agent (EA) pair is eligible to reproduce when the agent crosses a preset reward threshold on this environment. The next generation of environments is formed by mutating the current population and selecting only those environments that are neither too easy nor too hard. Finally, environments are ranked by novelty, and only the most novel children pass through to the next generation. More information about the hyperparameters of POET is listed in the supplementary material.

EPOET improves upon POET by adding in two algorithmic improvements: (1) a general method of evaluating the novelty of challenges and (2) an improved approach to deciding when agents should transfer to new

environments. In the original POET, the way to evaluate novelty was to compare the *environment characterization* (EC) of different environments. This is obtained by using some fixed, domain-specific static features, such as the roughness of the terrain. This inherently limits the exploration of the algorithm, as it is restricted to explore within these preset confines. Enhanced POET introduces an improved EC, Performance of All Transferred Agents EC (PATA-EC), which is based on the performance of different agents in the environment. Secondly, the original transfer mechanism in POET was generally inefficient, as it increased the required computation (as each agent needed to be fine-tuned), and resulted in subpar transfers as it was too easy to qualify for transfer. Enhanced POET makes this process more strict, only transferring very promising agents.

Enhanced POET also improves upon the environmental encoding used in the original algorithm, which was fixed and thus had a limited number of unique and diverse environments it could represent. The solution to this problem is to use a more expressive encoding in the form of compositional pattern producing networks (CPPNs) [50]. A CPPN is a specific neural network, which can take in x, y coordinates and produce a specific pattern when evaluated across an entire region. These CPPNs are evolved using the NEAT [18] algorithm, which increases the complexity of the environments as evolution progresses.

Lastly, the authors introduce *Accumulated Number of Novel Environments Created and Solved* (ANNECS), a metric for open-ended learning that, intuitively, describes the amount of interesting new content that is generated by the algorithm. ANNECS counts the number of environments that satisfy two constraints: (1) it must neither be too easy nor too hard and (2) it must be eventually solved by some agents in the future. Thus, if the ANNECS metric increases as time goes on, it indicates that the algorithm is continually producing novel and interesting environments.

IV. OPEN-ENDEDLY EVOLVING THE TOPOLOGY OF AGENTS

Many of the approaches introduced in prior work have been implemented using a fixed topology approach in conjunction with optimizers such as evolutionary strategies (ES) [51], V-MPO [52] (a modified version of maximum a posteriori optimization [53] which relies on value functions) and Proximal Policy Optimization [54], which motivates us to explore NEAT and the benefits it brings to the open-ended learning framework. We first describe the use of NEAT in Section IV-A and describe the overall approach in Section IV-B.

A. NeuroEvolution of augmenting topologies

We leverage NeuroEvolution of Augmenting Topologies (NEAT) to evolve the structure of an agent’s controller. NEAT starts with a population of simple neural networks (NNs), where the input neurons are directly connected to the output neurons without any hidden layers. Crossover is performed between two parents and the resulting children are mutated by adding connections and nodes, or perturbing weights. In this way, the NN will gradually be complexified. One of the major problems to overcome is the *Permutations or Competing Convention Problem* [55, 56]. Competing conventions describes the case in which the crossover of networks that represent the same solution but are encoded differently (e.g. a different ordering of neurons) can lead to a loss of information and a significantly worse child. NEAT addresses this by introducing a method to keep track of the historic origin of a gene by using the *innovation number*. Using this innovation number, identical genes from two parents can be aligned, while genes that only occur in one (denoted excess or disjoint genes depending on their position) can be inherited from the fitter parent. Finally, NEAT introduces speciation [57], where individuals with similar topologies are grouped together, and share a fitness. This protects innovation and ensures diversity. This speciation calculation is shown in Equation 1. In this equation, c_1, c_2 , and c_3 are coefficients that indicate the importance of each factor while N is the number of genes in the larger genome. E and D denote the number of excess and disjoint genes respectively. W is the average weight difference of similar genes. δ , then, indicates how close two genomes are; if δ is less than some threshold, then the two genomes belong to the same species.

$$\delta = \frac{c_1 E}{N} + \frac{c_2 D}{N} + c_3 \cdot W \quad (1)$$

NEAT has demonstrated superior performance when compared to fixed topology approaches, and has been used in numerous subsequent research works to great success [50, 19, 58, 59, 60].

B. Augmentative Topology Enhanced POET (ATEP)

In this section, we discuss the basic building blocks of our algorithm and the different variants we experimented with. ATEP combines EPOET with NEAT to allow the agents’ network topologies to evolve. This means that the algorithmic steps are very similar to EPOET, and the main differences are (1) the optimizer used: we use NEAT to optimize the variable-topology agents whereas EPOET used Evolution Strategies to optimize fixed-topology agents; and (2) the transfer mechanism, which

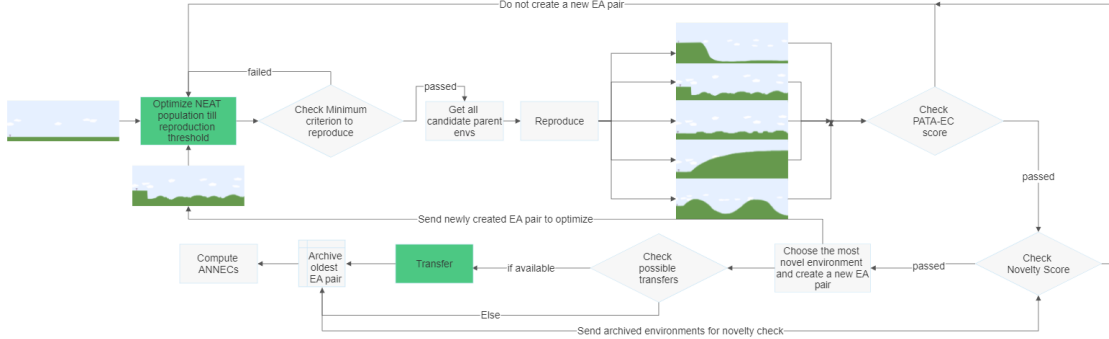


Fig. 1: A flowchart demonstrating the flow of the ATEP framework, with blocks in green being where ATEP differs from POET. For both EPOET and ATEP, each environment is associated with an agent, represented by an ES population for EPOET and a NEAT population for ATEP. The environment images used in the chart were created by ATEP. Please refer to the supplementary material for pseudocode describing the transfer mechanisms used in ATEP.

will be discussed later in this section. The detailed flow of ATEP is described in Figure 1.

We first use NEAT to evolve a population for each environment. The valid environments (those that pass the minimal criterion) then reproduce to create a new generation of (slightly harder) environments. We then take the environment that is the most novel (as measured by the Euclidean distance between the *PATA-EC* scores), and create a new environment-agent pair. The transfer eligibility of these environments is then evaluated, and if there are valid transfers available, we can move agents between environments. In EPOET, transfer is performed as follows: we compare the fitness of the candidate agent to the fitness of the target agent, over the previous 5 generations. If the candidate’s fitness is greater than all previous 5 fitness scores, we fine-tune it on the target environment and again compare it against the best fitness from the previous 5 generations. If both of these checks are passed we transfer the candidate and replace the target. For ATEP, we experiment with two different transfer mechanisms, the first being inspired by the approach used by EPOET, denoted as *Fitness-Based Transfer ATEP* (FBT-ATEP). In this case, we compare the best genome in the candidate population to the best genome from the target population. We then perform the same checks as EPOET, and if both are passed, we replace the entire target population with the candidate.

For the second transfer mechanism, we use the speciation inherent in NEAT to influence transfer. Specifically, we check if the best genome in the candidate population is within a δ threshold (using the speciation calculation in Equation 1) of any target environment’s best genome. If this is the case, we transfer the candidate species

and replace the target species with it. This approach, called *Species-Based Transfer ATEP* (SBT-ATEP), skips the step of comparing fitness scores and has its own advantages which we discuss in the next section. Finally, we also consider random transfer (RT-ATEP) and no transfer (NT-ATEP) to investigate whether the transfer mechanisms have a large impact on the results.

C. Experimental setup

Now we describe the experimental setup for ATEP, its variants, and our baselines. In ATEP, we use NEAT as the algorithm to evolve the topology and weights.¹ To reduce the computational load, we change one aspect of the original EPOET paper, reducing the number of active environments from 40 to 20. We make this change for both EPOET and ATEP, so the results are still comparable.

We set up two baselines: the first, denoted as EPOET40x40, is EPOET with the original controller consisting of two hidden layers with 40 nodes each. The second baseline, EPOET20x20, is a controller with two layers of 20 nodes each. This allows us to evaluate the effect of having a small fixed topology, a larger fixed topology, and a variable topology. Furthermore, this allows us to confirm our hypothesis that fixed topology agents will stagnate after a certain level of complexity. For further details on the controllers, please refer to the supplementary material.

V. RESULTS AND DISCUSSION

In this section, we discuss and analyze our results. We break the results into 3 different categories: Open-

¹Hyperparameter settings for the various methods are listed in the supplementary material.

Endedness, nodes complexity exploration and generalization ability. All results are gathered based on 2 seeds due to the expensive computational load, with each algorithm requiring approximately 50,000 to 200,000 CPU hours for a single run. EPOET20x20 required the least amount of computation, while SBT-ATEP required the most. Each algorithm was run in parallel on a cluster consisting of 264 Intel Xeon cores, with the runtime ranging between 10 and 30 days.

A. Open-endedness

As mentioned, Wang et al. [28] introduce the ANNECS metric to capture the open-endedness of an algorithm; we take it as our most important score to judge which algorithm performs better on complex environments.

Figure 2 shows the ANNECS score as a function of training time. We see that there is a significant difference between EPOET20x20 and FBT- and SBT-ATEP, indicating that the small network results in solving fewer environments. EPOET40x40 performs substantially better than EPOET20x20, and is competitive with ATEP early on during training. The rate of increase in ANNECS, however, does decrease after about 13k iterations, whereas ATEP increases at a consistent rate. This substantiates our hypothesis that fixed topology agents will start stagnating at some level of environment complexity, due to capacity issues. While we can improve the results by increasing the size of the network, that will merely delay the onset of stagnation.

FBT-ATEP outperforms EPOET40x40, although it also slows down slightly as time progresses. This is due to replacing the entire target population with the transferred population, which may eliminate all useful skills learned by the target population. SBT-ATEP, on the other hand, only replaces a single species that is close to the candidate species, leaving the rest of the population intact. We also find that SBT-ATEP has negligible delays in solving environments and, even though it performed similarly to FBT-ATEP and EPOET40x40 early on during training, it starts to outperform these in the second half of the experiment. This, as we will show later, is partly due to SBT-ATEP exploring more actions. We further note that the variations using no transfer (NT-ATEP) or random transfer (RT-ATEP) perform poorly, indicating that intelligent transfer mechanisms are necessary.

Although ATEP outperforms EPOET, it is more computationally expensive, as measured by the number of function evaluations. One function evaluation means one individual being evaluated on an environment. SBT-ATEP has the most function evaluations since once

a species transfers from one population to another, it becomes highly probable that it can transfer in the opposite direction because they may now be within the $\delta_{threshold}$ range. This increases the population size, resulting in more function evaluations. The tradeoff here is of function evaluations to performance, which is justified as the performance confirms our hypothesis. Figure 3 displays the total number of function evaluations.

B. Nodes complexity exploration

We have now shown that SBT-ATEP outperforms all of the other tested methods based on the ANNECS score. We also find that it generally uses a smaller neural network with fewer nodes than the other algorithms. Figure 4 shows the number of nodes and corresponding fitness value for each algorithm. We can see that SBT-ATEP generally has a high fitness, but fewer nodes than the other approaches. This is echoed in Figure 5c, where SBT-ATEP has the least number of nodes for most of the experiment, although it gradually adds nodes and complexity. FBT-ATEP, on the other hand, adds nodes very rapidly. This again indicates that the transfer mechanism in EPOET is critical.

Inspired by this, we further look into a simple *Fitness to Nodes* ratio (FNR) metric, shown in Figure 5a, and find that SBT-ATEP outperforms all other algorithms on this metric for the majority of the run. This indicates that SBT-ATEP outperforms all other algorithms on a per-environment basis, while using fewer nodes. This leads us to believe that a better-curated transfer mechanism, based on SBT-ATEP, will sustain the FNR for longer runs.

Furthermore, in Figure 5b, we calculate an *ANNECS to Nodes* ratio (ANR) metric with the intent to observe the role of nodes in the Open-Endedness of the agents, i.e. to have the ability to complexify over time. We observe that SBT-ATEP performs significantly better than the other models. FBT-ATEP has the lowest ANR, as it adds nodes much faster than the rate of increase in ANNECS.

C. Generalization Ability

We next evaluate the generalization ability of our open-ended agents, as prior work [12] has shown that these agents have the potential to generalize to new unseen environments. To concretely test this, we first take the 20 latest environments from each method. For each environment, we take the latest agent that could solve this environment from the method under consideration. Each of these agents is now evaluated on the selected environments from the other methods




Fig. 2: Accumulated Number of Novel Environments Created and Solved (ANNECS). To save compute, we stop the NT- and RT-ATEP experiments early, as it is clear that they perform poorly.




Fig. 3: Cumulative sum of the number of function evaluations, with the Y-axis converted to a log-scale. While ATEP requires significantly more function evaluations, we find that its total wall-time is only 3 times more than EPOET, as the neural networks are generally smaller and each evaluation does not take as long.

(60 in total). We perform 30 runs per environment-agent pair and calculate the mean and maximum of the rewards. We split the results into three categories: environments with fitness scores above 300, between 200 and 300, and below 200. Scores below 200 indicate that the environment has not been solved by the agent. Figure 6d shows the performance of each method when evaluated on the 60 other environments. We observe that SBT-ATEP outperforms all other models, with only 10% of the environments remaining unsolved.

Secondly, we test the generalization capabilities of agents on all of the environments created by their own algorithm. We exclude EPOET20x20 as it fails to solve

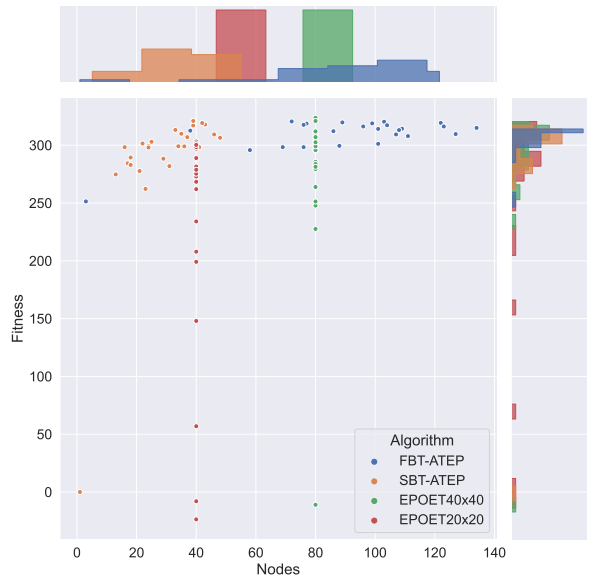


Fig. 4: Mapping of fitness to nodes. We plot values every 1000 iterations, starting at iteration number 150. Each dot represents a specific iteration, as well as the mean fitness over all environments and the mean number of nodes in the population.

80 environments in the whole run. We take into account 80 environments that were solved by the model itself and observe how each agent performs on all of them. Figures 6a, 6b and 6c show the results. Here, early-stage agents perform worse and late-stage agents are shown to have generalization abilities on previously




Fig. 5: Analysis with respect to the number of nodes. Figure (a) shows FNR along iterations, (b) shows ANR along iterations, (c) shows the addition of nodes along iterations

unseen landscapes. The transfer mechanism plays a key role in this generalization, as it exposes agents to more environments. Despite not having seen all environments, late-stage agents generalize much better. SBT-ATEP generalizes the best, with the lowest proportion of unsolved environments, in contrast to the lower-performing EPOET40x40 and FBT-ATEP.

Finally, we briefly investigate potential reasons why SBT-ATEP outperforms FBT-ATEP. We find that SBT-ATEP explores more actions, as it only transfers a single species instead of replacing the target population as is done by FBT-ATEP. This allows the new species to complement the actions that were already explored by the existing population. The action distributions of each action for SBT-ATEP, FBT-ATEP and EPOET40x40 are shown in the supplementary material.

VI. CONCLUSION AND FUTURE WORK

This work investigated the effect of having an Augmentative topology agent on an open-ended learning algorithm’s performance. We hypothesized that using a fixed topology would result in agents that exhibit delays in solving an environment after a certain point in environment complexity. We showed that this is indeed the case, and addressed this limitation by introducing ATEP, which allows the network topology of the agents to change and add complexity as necessary. We demonstrated that this approach outperforms existing methods in terms of the ANNECS score and generalization ability, while using fewer nodes in the neural networks. Our approach, however, does require more function evaluations than competing approaches. Thus, a promising future direction would be to use NEAT with Novelty Search [19] or Surprise Search which tends to converge faster than simple NEAT [61]. QD algorithms may also be worthwhile to explore in the context of open-ended learning as they have the ability to produce a population of high-performing and diverse individuals [62]. Exploring Neurogenesis [63, 64],

where neurons are added to a single neural network based on various external triggers, could also be a promising direction. To reduce computational load, it would also be promising to look into developing single-population open-ended learning methods without losing the exploration abilities of EPOET.

Furthermore, we have opened up possible future research into transfer mechanisms. We compared simple approaches such as FBT and SBT, but more advanced approaches could yield further performance improvements. For instance, we could combine both FBT and SBT in a weighted manner, or transfer only a certain percentage of a species or population. Finally, this work provides a starting point, like EPOET itself, into open-ended learning with augmentative topology agents. We therefore used the simple 2D BipedalWalker as our benchmark. Future work should compare ATEP with standard EPOET on different and more complex environments. Ultimately, we hope that this new approach furthers research into open-ended algorithms that do not slow down over time, and can keep up with an ever-changing environment.

REPRODUCIBILITY AND ETHICAL STATEMENT

For reproducibility, we have provided a GitHub repo² where users can follow instructions to reproduce the experiments. We also provide pseudocode and hyperparameter settings in the supplementary material. ATEP is an Open-Ended Learning algorithm that has stochastic elements, similar to many other machine learning algorithms. It is critical for users to perform standard evaluations as the user would do for other machine learning algorithms. A full run of ATEP may be computationally expensive and will take approximately 50,000 to 200,000 CPU hours.

²https://github.com/umair-nasir14/ATEP_LLRL

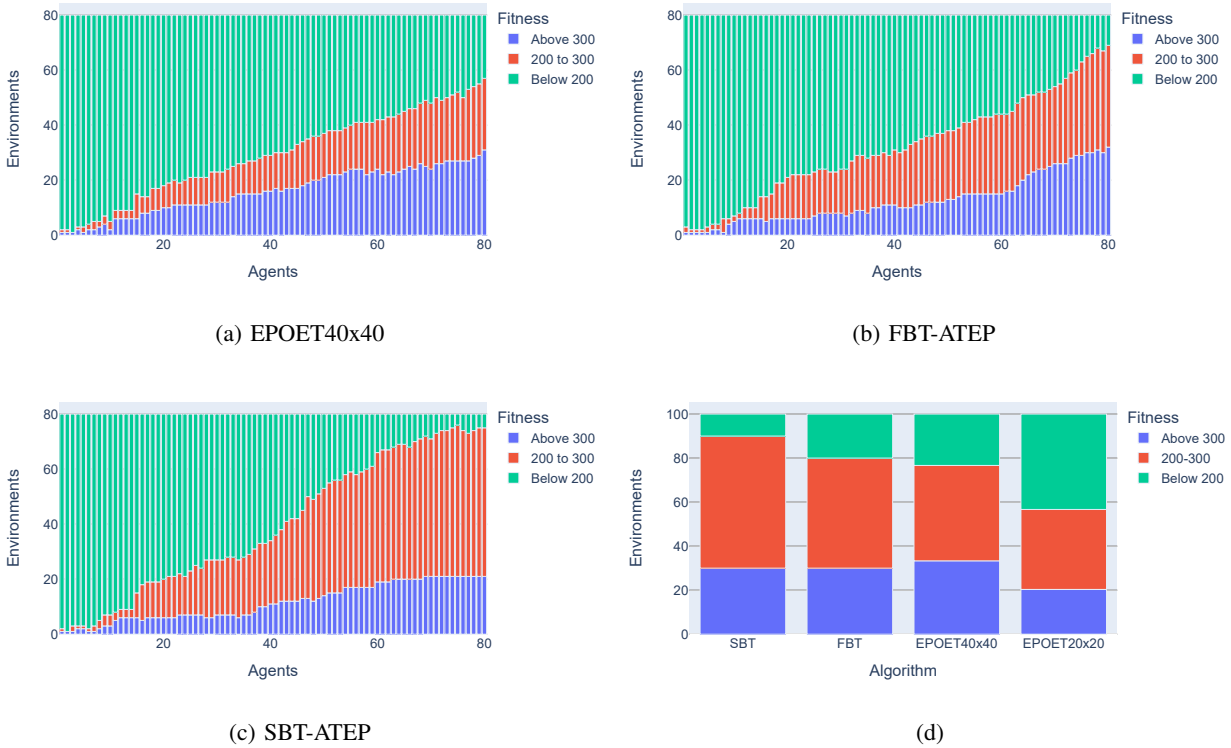


Fig. 6: Figures showing generalization capabilities. Figures (a), (b) and (c) show agents of 80 solved environments being tested on all 80 environments, for EPOET40x40, FBT-ATEP and SBT-ATEP respectively. Note that EPOET20x20 does not take part in this test as it failed to produce 80 environments in the run. Figure (d) shows each algorithm being tested on the 20 latest environments created by all other algorithms, i.e., each algorithm is evaluated on 60 environments. The Y-axis shows the percentage of environments in each category. Each test is conducted for 30 runs and the mean scores are taken.

REFERENCES

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems*, 2012, pp. 1106–1114.
- [2] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [3] T. Schick and H. Schütze, “It’s not just size that matters: Small language models are also few-shot learners,” in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021, pp. 2339–2352.
- [4] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, “Human-level control through deep reinforcement learning,” *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [5] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel *et al.*, “A general reinforcement learning algorithm that masters chess, shogi, and go through self-play,” *Science*, vol. 362, no. 6419, pp. 1140–1144, 2018.
- [6] K. O. Stanley, J. Lehman, and L. Soros, “Open-endedness: The last grand challenge you’ve never heard of,” *O’Reilly Online*, 2017. [Online]. Available: <https://www.oreilly.com/ideas/open-endedness-the-last-grand-challenge-youve-never-heard-of>
- [7] K. O. Stanley, “Why open-endedness matters,” *Artificial life*, vol. 25, no. 3, pp. 232–235, 2019.
- [8] J. C. Brant and K. O. Stanley, “Minimal crite-

- tion coevolution: a new approach to open-ended search,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2017, pp. 67–74.
- [9] R. Wang, J. Lehman, J. Clune, and K. O. Stanley, “Paired open-ended trailblazer (POET): Endlessly generating increasingly complex and diverse learning environments and their solutions,” *arXiv preprint arXiv:1901.01753*, 2019.
- [10] A. Quessy and T. S. Richardson, “Rewardless open-ended learning (ROEL),” 2021.
- [11] L. Gisslén, A. Eakins, C. Gordillo, J. Bergdahl, and K. Tollmar, “Adversarial reinforcement learning for procedural content generation,” in *2021 IEEE Conference on Games*. IEEE, 2021, pp. 1–8.
- [12] O. E. L. Team, A. Stooke, A. Mahajan, C. Barros, C. Deck, J. Bauer, J. Sygnowski, M. Trebacz, M. Jaderberg, M. Mathieu *et al.*, “Open-ended learning leads to generally capable agents,” *arXiv preprint arXiv:2107.12808*, 2021.
- [13] T. S. Ray, “An approach to the synthesis of life,” *Artificial life II*, pp. 371–408, 1991.
- [14] C. Ofria and C. O. Wilke, “Avida: A software platform for research in computational evolutionary biology,” *Artificial life*, vol. 10, no. 2, pp. 191–229, 2004.
- [15] L. Spector, J. Klein, and M. Feinstein, “Division blocks and the open-ended evolution of development, form, and behavior,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2007, pp. 316–323.
- [16] L. S. Yaeger and O. Sporns, “Evolution of neural structure and complexity in a computational ecology,” in *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, 2006, pp. 330–336.
- [17] L. Soros and K. Stanley, “Identifying necessary conditions for open-ended evolution through the artificial life world of chromaria,” in *ALIFE 14: The Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, 2014, pp. 793–800.
- [18] K. O. Stanley and R. Miikkulainen, “Evolving neural networks through augmenting topologies,” *Evolutionary computation*, vol. 10, no. 2, pp. 99–127, 2002.
- [19] J. Lehman, K. O. Stanley *et al.*, “Exploiting open-endedness to solve problems through the search for novelty,” in *ALIFE*, 2008, pp. 329–336.
- [20] J. Lehman and K. O. Stanley, “Evolving a diversity of virtual creatures through novelty search and local competition,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2011, pp. 211–218.
- [21] J.-B. Mouret and J. Clune, “Illuminating search spaces by mapping elites,” *arXiv preprint arXiv:1504.04909*, 2015.
- [22] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune, “Go-explore: a new approach for hard-exploration problems,” *arXiv preprint arXiv:1901.10995*, 2019.
- [23] O. Nilsson and A. Cully, “Policy gradient assisted map-elites,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2021, pp. 866–875.
- [24] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, “Robots that can adapt like animals,” *Nature*, vol. 521, no. 7553, pp. 503–507, 2015.
- [25] D. Tarapore, J. Clune, A. Cully, and J.-B. Mouret, “How do different encodings influence the performance of the map-elites algorithm?” in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2016, pp. 173–180.
- [26] A. Khalifa, S. Lee, A. Nealen, and J. Togelius, “Talakat: Bullet hell generation through constrained map-elites,” in *Proceedings of The Genetic and Evolutionary Computation Conference*, 2018, pp. 1047–1054.
- [27] S. Earle, J. Snider, M. C. Fontaine, S. Nikolaidis, and J. Togelius, “Illuminating diverse neural cellular automata for level generation,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2022, pp. 68–76.
- [28] R. Wang, J. Lehman, A. Rawal, J. Zhi, Y. Li, J. Clune, and K. Stanley, “Enhanced POET: Open-ended reinforcement learning through unbounded invention of learning challenges and their solutions,” in *International Conference on Machine Learning*, 2020, pp. 9940–9951.
- [29] A. Dharna, J. Togelius, and L. B. Soros, “Co-generation of game levels and game-playing agents,” in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 16, no. 1, 2020, pp. 203–209.
- [30] A. Dharna, A. K. Hoover, J. Togelius, and L. Soros, “Transfer dynamics in emergent evolutionary curricula,” *IEEE Transactions on Games*, 2022.
- [31] V. Campos, A. Trott, C. Xiong, R. Socher, X. Giró-i-Nieto, and J. Torres, “Explore, discover and learn: Unsupervised discovery of state-covering skills,” in *Proceedings of the 37th International Conference on Machine Learning*, vol. 119, 2020, pp. 1317–

- 1327.
- [32] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine, "Diversity is all you need: Learning skills without a reward function," in *International Conference on Learning Representations*, 2019.
- [33] A. Sharma, S. Gu, S. Levine, V. Kumar, and K. Hausman, "Dynamics-aware unsupervised discovery of skills," in *International Conference on Learning Representations*, 2020.
- [34] R. Meier and A. Mujika, "Open-ended reinforcement learning with neural reward functions," *arXiv preprint arXiv:2202.08266*, 2022.
- [35] F. Zhou and J. Vanschoren, "Open-ended learning strategies for learning complex locomotion skills," *arXiv preprint arXiv:2206.06796*, 2022.
- [36] E. H. Stensby, K. O. Ellefsen, and K. Glette, "Co-optimising robot morphology and controller in a simulated open-ended environment," in *International Conference on the Applications of Evolutionary Computation*, 2021, pp. 34–49.
- [37] E. S. Norstein, K. O. Ellefsen, and K. Glette, "Open-ended search for environments and adapted agents using map-elites," in *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*. Springer, 2022, pp. 651–666.
- [38] M. Dennis, N. Jaques, E. Vinitzky, A. Bayen, S. Russell, A. Critch, and S. Levine, "Emergent complexity and zero-shot transfer via unsupervised environment design," *Advances in Neural Information Processing Systems*, vol. 33, pp. 13 049–13 061, 2020.
- [39] F. Sadeghi and S. Levine, "Cad2rl: Real single-image flight without a single real image," *arXiv preprint arXiv:1611.04201*, 2016.
- [40] M. Jiang, E. Grefenstette, and T. Rocktäschel, "Prioritized level replay," in *International Conference on Machine Learning*, 2021, pp. 4940–4950.
- [41] J. Parker-Holder, M. Jiang, M. Dennis, M. Samvelyan, J. Foerster, E. Grefenstette, and T. Rocktäschel, "Evolving curricula with regret-based environment design," *arXiv preprint arXiv:2203.01302*, 2022.
- [42] M. Barthet, A. Liapis, and G. N. Yannakakis, "Open-ended evolution for Minecraft building generation," *IEEE Transactions on Games*, 2022.
- [43] M. J. Kusner, B. Paige, and J. M. Hernández-Lobato, "Grammar variational autoencoder," in *International Conference on Machine Learning*, 2017, pp. 1945–1954.
- [44] X. Chen, D. P. Kingma, T. Salimans, Y. Duan, P. Dhariwal, J. Schulman, I. Sutskever, and P. Abbeel, "Variational lossy autoencoder," *arXiv preprint arXiv:1611.02731*, 2016.
- [45] S. C. Duncan, "Minecraft, beyond construction and survival," 2011.
- [46] M. Cipollone, C. C. Schifter, and R. A. Mof-fat, "Minecraft as a creative tool: A case study," *International Journal of Game-Based Learning (IJGBL)*, vol. 4, no. 2, pp. 1–14, 2014.
- [47] L. Fan, G. Wang, Y. Jiang, A. Mandlekar, Y. Yang, H. Zhu, A. Tang, D.-A. Huang, Y. Zhu, and A. Anandkumar, "Minedojo: Building open-ended embodied agents with internet-scale knowledge," *arXiv preprint arXiv:2206.08853*, 2022.
- [48] C. Gan, A. Bhandwaldar, A. Torralba, J. B. Tenenbaum, and P. Isola, "Open: An open-ended physics environment for learning without a task," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5878–5885.
- [49] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "OpenAI gym," *arXiv preprint arXiv:1606.01540*, 2016.
- [50] K. O. Stanley, "Compositional pattern producing networks: A novel abstraction of development," *Genetic programming and evolvable machines*, vol. 8, no. 2, pp. 131–162, 2007.
- [51] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, "Evolution strategies as a scalable alternative to reinforcement learning," *arXiv preprint arXiv:1703.03864*, 2017.
- [52] H. F. Song, A. Abdolmaleki, J. T. Springenberg, A. Clark, H. Soyer, J. W. Rae, S. Noury, A. Ahuja, S. Liu, D. Tirumala *et al.*, "V-MPO: On-policy maximum a posteriori policy optimization for discrete and continuous control," *arXiv preprint arXiv:1909.12238*, 2019.
- [53] A. Abdolmaleki, J. T. Springenberg, J. Degraeve, S. Bohez, Y. Tassa, D. Belov, N. Heess, and M. Riedmiller, "Relative entropy regularized policy iteration," *arXiv preprint arXiv:1812.02256*, 2018.
- [54] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [55] N. J. Radcliffe, "Genetic set recombination and its application to neural network topology optimisation," *Neural Computing & Applications*, vol. 1, no. 1, pp. 67–90, 1993.
- [56] D. J. Montana, L. Davis *et al.*, "Training feedforward neural networks using genetic algorithms." in *International Joint Conference on Artificial Intel-*

ligence, vol. 89, 1989, pp. 762–767.

- [57] S. W. Mahfoud, “Niching methods for genetic algorithms,” Ph.D. dissertation, University of Illinois at Urbana-Champaign, 1995.
- [58] K. O. Stanley, D. B. D’Ambrosio, and J. Gauci, “A hypercube-based encoding for evolving large-scale neural networks,” *Artificial life*, vol. 15, no. 2, pp. 185–212, 2009.
- [59] J. Schrum, V. Volz, and S. Risi, “CPPN2GAN: Combining compositional pattern producing networks and gans for large-scale pattern generation,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2020, p. 139–147.
- [60] M. L. Clei and P. Bellec, “Neuroevolution of recurrent architectures on control tasks,” in *Genetic and Evolutionary Computation Conference, Companion Volume*, J. E. Fieldsend and M. Wagner, Eds., 2022, pp. 651–654.
- [61] D. Gravina, A. Liapis, and G. Yannakakis, “Surprise search: Beyond objectives and novelty,” in *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, 2016, pp. 677–684.
- [62] V. Bhatt, B. Tjanaka, M. C. Fontaine, and S. Nikolaidis, “Deep surrogate assisted generation of environments,” *CoRR*, vol. abs/2206.04199, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2206.04199>
- [63] K. Maile, E. Rachelson, H. Luga, and D. G. Wilson, “When, where, and how to add new neurons to ANNs,” *arXiv preprint arXiv:2202.08539*, 2022.
- [64] T. J. Draelos, N. E. Miner, C. C. Lamb, J. A. Cox, C. M. Vineyard, K. D. Carlson, W. M. Severa, C. D. James, and J. B. Aimone, “Neurogenesis deep learning: Extending deep networks to accommodate new classes,” in *2017 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2017, pp. 526–533.