EvoJAX: Hardware-Accelerated Neuroevolution

Yujin Tang
yujintang@google.com
Google Brain

Yingtao Tian
alantian@google.com
Google Brain

David Ha
hadavid@google.com
Google Brain

ABSTRACT

Evolutionary computation has been shown to be a highly effective method for training neural networks, particularly when employed at scale on CPU clusters. Recent work have also showcased their effectiveness on hardware accelerators, such as GPUs, but so far such demonstrations are tailored for very specific tasks, limiting applicability to other domains. We present EvoJAX, a scalable, general purpose, hardware-accelerated neuroevolution toolkit. Building on top of the JAX library, our toolkit enables neuroevolution algorithms to work with neural networks running in parallel across multiple TPU/GPUs. EvoJAX achieves very high performance by implementing the evolution algorithm, neural network and task all in NumPy, which is compiled just-in-time to run on accelerators. We provide extensible examples of EvoJAX for a wide range of tasks, including supervised learning, reinforcement learning and generative art. Since EvoJAX can find solutions to most of these tasks within minutes on a single accelerator, compared to hours or days when using CPUs, our toolkit can significantly shorten the iteration cycle of evolutionary computation experiments. EvoJAX is available at **https://github.com/google/evojax**

1 INTRODUCTION

Hardware accelerators have played an important role in advancing the state-of-the-art for deep learning (DL), enabling rapid training of neural networks and shorter research iteration cycles for their development [12]. But much of this progress is restricted to systems that rely on gradient descent, a highly effective optimization method when we provide it with a well-defined objective function. But in areas such as artificial life, complex systems, computational biology, and even classical physics [18], much of the interesting behaviors we observe take place near the chaotic states, where a system is constantly transitioning between order and disorder. It can be argued that intelligent life and even civilization are all complex systems operating at the *edge of chaos* [3, 16]. If we wish to study these systems, we need efficient methods to simulate and find solutions in complex systems.
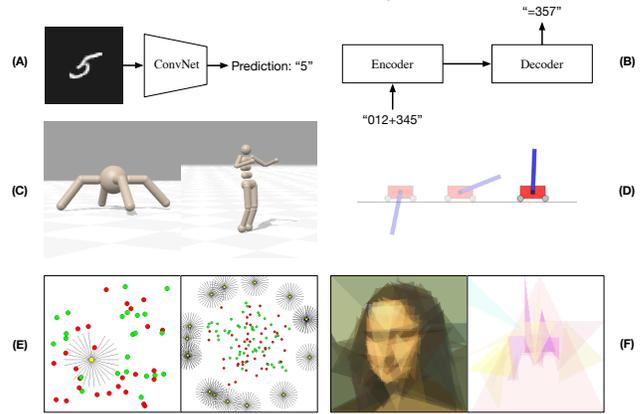
Figure 1: EvoJAX Examples. (A) MNIST classification. (B) Seq2Seq learning. (C) Robotic control. (D) Cart-pole swing up. (E) Left: WaterWorld wherein the agent (yellow) tries to get food (green) while avoiding poison (red). Right: A version of WaterWorld with multiple agents. (F) Abstract painting with only triangles. Left: Painting a concrete image. Right: Painting the concept "Walt Disney World".

Neural networks are a promising approach for modeling complex systems [9, 19], and neuroevolution has made great progress in developing methods for evolving neural networks to solve a wide range of problems. Evolution-based methods have been shown to find state-of-the-art solutions for reinforcement learning (RL) [8, 13, 22, 25, 29]. A policy with non-differentiable operations can solve many more tasks than one that is fully differentiable [20, 27, 28, 33]. More importantly, the removal of the requirement of a differentiable policy also liberates the researchers' mind, enabling higher levels of creativity for looking at problems and directions differently from the mainstream. In a sense, enabling researchers to use neural networks beyond gradient-based methods also enables the broader machine learning (ML) research community to explore in a way that is also less "grad student descent" [7]-based.

However, the progress of hardware-accelerated computational methods for evolution has not kept pace with ML, or even RL. Much of computational evolution is still conducted using CPU clusters, largely ignoring the recent breakthroughs in hardware accelerators such as GPUs/TPUs. Recent work started to demonstrate effectiveness of GPUs for neuroevolution [25], but so far such demonstrations are tailored for specific tasks [24], limiting their applicability to other domains. To enable greater access to hardware accelerators for neuroevolution researchers, we developed EvoJAX, a scalable, general purpose, neuroevolution toolkit. Building on the JAX library [1], our toolkit enables neuroevolution algorithms to work with neural networks running in parallel across multiple TPU/GPUs. EvoJAX achieves very high performance by implementing the evolution algorithm, neural network and task all in NumPy, which is compiled just-in-time to run on accelerators.
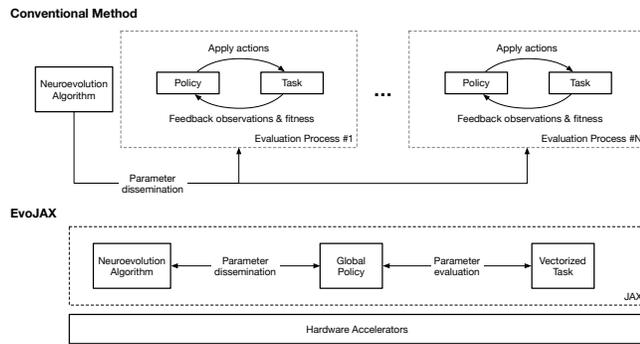
Figure 2: Architectural Overview of EvoJAX.

In this paper, we describe the design of EvoJAX and show how one can use and extend EvoJAX for neuroevolution research. We showcase several extensible examples of EvoJAX for a wide range of tasks, including supervised learning (image classification, seq-to-seq), RL (cart-pole swing-up [6], Brax locomotion [5], multi-agent water world), and generative art (image approximation with shapes, CLIP-guided abstract art [30]). We show that EvoJAX can find solutions to most of these tasks within minutes on GPU/TPUs, compared to hours or days when using CPUs. We believe our toolkit can significantly shorten the experimental iteration cycle for researchers working with evolutionary computation. We have also created several tutorials and notebooks as part of this open-source project to make adapting EvoJAX for novel use cases straightforward.

2 SYSTEM DESIGN

EvoJAX aims to improve the neuroevolution training efficiency by implementing the entire pipeline in modern ML frameworks that support hardware acceleration. We choose JAX[1] in our current implementation due to its wide variety of hardware support and its matured features of auto-vectorization, device-parallelism, just-in-time compilation, etc. As we will see in Section 4, as long as the component interfaces are properly implemented, EvoJAX also allows user extensions with other frameworks.

Figure 2 gives an overview of how EvoJAX works. There are three major components – the neuroevolution algorithm, the policy and the task. Although these components are common in conventional neuroevolution implementations, we highlight the key differences that make EvoJAX much more efficient:

Modern ML Optimizers Researchers and practitioners in the field of DL have been focusing on inventing optimization algorithms [21] and techniques [15, 32, 34] that are both fast and effective. Although these techniques were tailored for gradient-based optimizations, they can be directly applied to gradient estimation-based evolutionary algorithms [17, 23] too. By leveraging JAX-based libraries [1, 10, 11], EvoJAX not only achieves significant speed-up but also provides the users with the tools and the interfaces to develop their own implementations in a mature framework.

Global Policy In conventional neuroevolution implementations, it is a common practice to spawn multiple processes for parameters evaluation. To achieve hardware acceleration, the implementation adopts one of the DL frameworks and then each of the evaluation processes maintains a separate computational graph for the same policy. Unfortunately, most DL frameworks

are not designed for multi-process training scenarios and often cause difficulties. Moreover, when these processes are run on the same accelerator, maintaining identical copies of the computational graph is a waste of resource. Conforming to the "Single-Program, Multiple-Data" (SPMD) model [4], EvoJAX solves this by building a global policy and treat both the task observations and the policy parameters as data for the computational graph. This global policy design is easy to implement as it is consistent with DL frameworks, and in the experiments we observe high data-throughput.

Vectorized Tasks Same as the policies, conventional methods also create copies of the tasks in the spawned processes for independent parameters evaluations. To be compliant with EvoJAX's global policy design, we propose to group these tasks in a vectorized form. In terms of implementation, this can be achieved by either creating the task in auto-vectorizaton supported frameworks or by creating a task observations collector on top of all the evaluation processes. EvoJAX adopts the first method.

Device Parallelism Thanks to the device-parallelism support in JAX, EvoJAX is capable of scaling its training procedure almost linearly to the available hardware accelerators. Utilizing EvoJAX's training pipeline, this device parallelism is automatically managed and is transparent to the users. As we will see in Section 3, together with the previously mentioned features, EvoJAX significantly shortens the training time for novel and non-trivial tasks.

EvoJAX defines simple yet functionally complete interfaces for the three components, any implementations that are compliant with the interfaces can be seamlessly integrated (see Section 4).

Finally, in addition to the mentioned major components, EvoJAX also comes with a trainer and a simulation manager that help orchestrate and manage the training process. They contain detailed implementations of task roll-out seeds generation, efficient training loops, time profiling and logistics operations such as logging, testing and periodic model saving. Convenient as they are, we point out that EvoJAX is a flexible toolkit, where it is possible to use any component independently (e.g., using a custom training loop).

3 EVOJAX EXAMPLES

We provide a total of six examples (see Figure 1) to showcase the capacity, efficiency and the usage of EvoJAX online in the format of Python scripts and notebooks. The examples are designed to feature different aspects of EvoJAX and are in three categories: Supervised Learning Tasks, Control Tasks and Novel Tasks. As the experimental setups, "Robotic Control" was trained with TPUs, "Concrete and Abstract Painting" was trained with 8 NVIDIA V100 GPUs, and the rest were trained with 1 NVIDIA V100 GPU.

Supervised Learning Tasks They provide both the data and the ground-truth labels to train the policy. In EvoJAX, supervised learning tasks are modelled as single-step tasks, the examples in this category are thus isolated from other factors to prove the correctness and efficiency of our algorithms' implementation.

- *MNIST Classification*. Here, we train a convolutional neural network (ConvNet) with 10K parameters with EvoJAX. Although MNIST is a solved problem in DL, it is non-trivial for neuroevolution in terms of achieving high test accuracy within a short time (e.g., in minutes). We show that EvoJAX can train the ConvNet to reach $> 98\%$ test accuracy within 5 minutes.

- *Seq2Seq Learning*. It has recently been shown that genetic algorithms (GA) can train large models [20]. Here, we show that EvoJAX is also capable of training a large network with hundreds of thousands of parameters. We adopt a seq-to-seq task where the policy is required to output a sequence after observing a query sequence. Concretely, the query is a sequence that represents the addition of two randomly generated integers (e.g., "012+345=", we pad the numbers with leading 0's so that they have equal lengths) and the result is a sequence representing the answer. Using an LSTM-based seq2seq [26] model, EvoJAX achieves $> 99\%$ test accuracy within tens of minutes.

While one would obviously use gradient-descent for such tasks in practice, the point is to show that neuroevolution can also solve them to some degree of accuracy within a short amount of time, which will be useful when these models are adapted within a more complicated task where gradient-based approaches may not work.

Control Tasks The purpose of including control tasks are two-fold: 1) Unlike supervised learning tasks, control tasks in EvoJAX have undetermined number of steps, we thus use these examples to demonstrate the efficiency of our task roll-out loops. 2) We wish to show the speed-up benefit of implementing tasks in JAX and illustrate how to implement one from scratch.

- *Robotic Control*. Brax [5] is a differentiable physics engine implemented in JAX that simulates environments made up of rigid bodies, joints, and actuators. We show that it is easy to wrap Brax tasks in EvoJAX, and it takes EvoJAX tens of minutes to solve a robotic locomotion task on Colab TPUs.
- *Cart-Pole Swing Up*. Through this classic control task, we illustrate how a task is implemented from scratch in JAX and integrated into EvoJAX's training pipeline. In our implementation, a user can command the initial states to be randomly sampled from a narrow (easy version) or a wide (hard version) range of possible settings, with the latter being much harder to solve. EvoJAX solves both versions within minutes.

Novel Tasks In this last category, we go beyond simple illustrations and show examples of novel tasks that are more practical and attractive to researchers in the genetic and evolutionary computation area, with the goal of helping them try out ideas in EvoJAX.

- *WaterWorld*. In this task [14], an agent tries to get as much food as possible while avoiding poisons. EvoJAX is able to train the agent in tens of minutes. Furthermore, we demonstrate that multi-agents training in EvoJAX is possible. Here, we spawn the entire population in the same task roll-out and directly measure each agent's performance in a multi-agent world. This training scheme automatically generates task complexity beyond human design, and is beneficial for learning policies that can deal with interactions between agents and environmental uncertainties.
- *Concrete* and *Abstract Painting*. We reproduce the results from a computational creativity work [30]. The original work, whose implementation requires multiple CPUs and GPUs, could be accelerated on a single GPU efficiently using EvoJAX, which was not possible before. Moreover, with multiple GPUs/TPUs, EvoJAX can further speed up the mentioned work almost linearly. We also show that the modular design of EvoJAX allows its components be used independently – in this case it is possible to use only the

Table 1: Time Comparisons. We report the training time for both methods to achieve widely accepted test scores.

	Baseline	EvoJAX
MNIST	36 min	3 min
Cart-Pole Swing Up (Hard Version)	37 min	2 min
Locomotion (Ant)[1]	201 min	9 min

neuroevolution algorithms from EvoJAX while leveraging one's own training loops and environment implantation.

We summarize EvoJAX's benefit via these examples. First of all, EvoJAX brings significant training speed up. In Table 1 we show the time costs of training some popular tasks with both a conventional setup and EvoJAX.[1] On modest hardware accelerators, EvoJAX trains $10 \sim 20$ times faster which leads to quicker idea iterations. Secondly, the capability of training multi-agents in a complex setting that is beyond human design supplies training environmental richness. And finally, EvoJAX puts the entire pipeline on unified hardware setups and that allows the practitioners to simplify complex hardware arrangements. As an example, for the substantial load of computation in our Abstract Painting example, the baseline needs to use both GPUs and CPUs, while EvoJAX only uses GPUs.

4 EXTENDING EVOJAX

A goal of EvoJAX is to provide researchers with an infrastructure that allows fast idea iterations. With EvoJAX it is possible to devise more effective neuroevolution algorithms, to explore novel policy architectures, and to experiment with new tasks. EvoJAX has carefully defined interfaces, as long as these interfaces are properly implemented, a user extended module can be integrated into the pipeline seamlessly.

```python
import jax.numpy as jnp

class TaskState: obs: jnp.ndarray
class PolicyState: keys: jnp.ndarray
class NEAlgorithm:
    def ask(self) -> jnp.ndarray: pass
    def tell(self, fitness: jnp.ndarray) -> None: pass
class PolicyNetwork:
    def reset(self, states: TaskState) -> PolicyState: pass
    def get_actions(self, t_states: TaskState, params: jnp.ndarray,
                    p_states: PolicyState) \
                        -> Tuple[jnp.ndarray, PolicyState]: pass
class VectorizedTask:
    def reset(self, key: jnp.ndarray) -> TaskState: pass
    def step(self, state: TaskState, action: jnp.ndarray) \
                -> Tuple[TaskState, jnp.ndarray, jnp.ndarray]: pass
```

Figure 3: Major Component Interfaces in EvoJAX.

Devising New Algorithms Users interested in inventing new neuroevolution algorithms should implement *NEAlgorithm* in Figure 3, which serves as the base class for all neuroevolution algorithms in EvoJAX. Being consistent with most conventional implementations, *NEAlgorithm* adopts the "ask" and "tell" interfaces, where the former requests the algorithm to generate a population of parameters and the latter reports the parameters evaluation results back to the algorithm for internal states update. Taking on the conventional interfaces for the neuroevolution algorithms not only brings familiarity to the developers and thus reducing the required

---

[1]We use the code from [27] as the baseline. For the Locomotion task, we use PyBullet Ant in the baseline and Brax Ant in EvoJAX. The baseline is trained with 96 CPUs.

learning effort, but also allows the practitioners to quickly plug in existing algorithms for sanity checks by writing a simple wrapper.

Exploring Novel Policy Architectures *PolicyNetwork* in Figure 3 defines the policy interface, all policies in EvoJAX implement the *get_actions* method. The method puts no restrictions on what the policy network should be or how it should behave, giving full freedom for neural architecture search (NAS). Because EvoJAX conforms to the SPMD model, *get_actions* accepts three parameters: the vectorized task states, the population parameters and the policy's internal states. At the beginning of a roll-out, each individual in the population sees identical observations, they will then diverge due to the population's different behaviors. Because JAX requires pure functions, the policy's states (e.g., random seeds, LSTM cell states, etc) are passed to *get_actions* via a Flax [10] dataclass *p_states*, which is initialized by *PolicyNetwork.reset*. The method returns the actions and the updated policy states. At runtime, calling *get_actions* is equivalent to passing a batch of data through the model.

Experimenting with More Tasks In Figure 3, *VectorizedTask* forms the base for all EvoJAX tasks. Similar to OpenAI's Gym environments [2], the interface defines the *reset* and the *step* methods. Following the pure-function principle of JAX, one major difference between EvoJAX tasks and Gym environments is that EvoJAX's tasks do not keep internal states. Instead, these states are encapsulated in a *TaskState* instance and carried over the roll-out steps. Similar to *PolicyState*, users can inherit *TaskState* and create one's own task specific state to encapsulate arbitrary information besides the environment observations. In most tasks, the initial states are generated via a procedure of randomness. The *reset* method thus accepts *key*'s that act as seeds for the random process.

5 LIMITATIONS AND FUTURE WORKS

EvoJAX is based on the JAX framework, which is based on the familiar NumPy and is thus friendly to researchers accustomed to such tools. However, practitioners may have to take effort to understand the subtleties of JAX in order to maximize its performance. The time spent on learning the JAX framework may translate to a delayed adoption of EvoJAX, hence much of our focus so far has been on creating examples and tutorials that others can use as templates to build upon. Another limitation of EvoJAX is the compatibility with existing non-parallelizable tasks. Although it is possible to create an observation collector on top of the evaluation processes to mimic the behavior of *VectorizedTask*, the operation involves inter-process communications that becomes a bottleneck, preventing such tasks from the benefit of hardware-acceleration.

In the future, we plan to release more neuroevolution algorithm implementations to EvoJAX in addition to PGPE [23, 31] in the current release. We will add more policies and tasks to both demonstrate a wider variety of examples in order to encourage greater adoption of EvoJAX, with the goal of further enhancing the computation tools available in evolutionary computation research.

REFERENCES

[1] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. *JAX: composable transformations of Python+NumPy programs.* http://github.com/google/jax
[2] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. 2016. Openai gym. *arXiv:1606.01540* (2016).
[3] Leon Chua, Valery Sbitnev, and Hyongsuk Kim. 2012. Neurons are poised near the edge of chaos. *International Journal of Bifurcation and Chaos* 22, 04 (2012).
[4] Frederica Darema. 2001. The spmd model: Past, present and future. In *European Parallel Virtual Machine/Message Passing Interface Users' Group Meeting.* Springer.
[5] C. Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. 2021. *Brax - A Differentiable Physics Engine for Large Scale Rigid Body Simulation.* http://github.com/google/brax
[6] Daniel Freeman, David Ha, and Luke Metz. 2019. Learning to Predict Without Looking Ahead: World Models Without Forward Prediction. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc.
[7] Oguzhan Gencoglu, Mark van Gils, Esin Guldogan, Chamin Morikawa, Mehmet Süzen, Mathias Gruber, Jussi Leinonen, and Heikki Huttunen. 2019. HARK Side of Deep Learning–From Grad Student Descent to Automated Machine Learning. *arXiv:1904.07633* (2019).
[8] David Ha. 2020. Slime Volleyball Gym Environment.
[9] David Ha and Yujin Tang. 2021. Collective Intelligence for Deep Learning: A Survey of Recent Developments. *arXiv:2111.14377* (2021).
[10] Jonathan Heek, Anselm Levskaya, Avital Oliver, Marvin Ritter, Bertrand Rondepierre, Andreas Steiner, and Marc van Zee. 2020. *Flax: A neural network library and ecosystem for JAX.* http://github.com/google/flax
[11] Matteo Hessel, David Budden, Fabio Viola, Mihaela Rosca, Eren Sezener, and Tom Hennigan. 2020. *Optax: composable gradient transformation and optimisation, in JAX!* http://github.com/deepmind/optax
[12] Sara Hooker. 2021. The hardware lottery. *Commun. ACM* 64, 12 (2021), 58–65.
[13] Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, et al. 2017. Population based training of neural networks. *arXiv:1711.09846* (2017).
[14] Andrej Karpathy. 2015. *REINFORCEjs.* https://cs.stanford.edu/people/karpathy/reinforcejs/waterworld.html
[15] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2016. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv:1609.04836* (2016).
[16] Roger Lewin. 1999. *Complexity: Life at the edge of chaos.* University of Chicago.
[17] Horia Mania, Aurelia Guy, and Benjamin Recht. 2018. Simple random search of static linear policies is competitive for reinforcement learning. In *The 32nd Conference on Neural Information Processing Systems.* 1805–1814.
[18] Luke Metz, C Daniel Freeman, Samuel S Schoenholz, and Tal Kachman. 2021. Gradients are Not All You Need. *arXiv:2111.05803* (2021).
[19] Sebastian Risi. 2021. The Future of Artificial Intelligence is Self-Organizing and Self-Assembling. https://sebastianrisi.com/self_assembling_ai.
[20] Sebastian Risi and Kenneth O Stanley. 2019. Deep neuroevolution of recurrent and discrete world models. In *Proceedings of GECCO.* 456–462.
[21] Sebastian Ruder. 2016. An overview of gradient descent optimization algorithms. *arXiv:1609.04747* (2016).
[22] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv:1703.03864* (2017).
[23] Frank Sehnke, Christian Osendorfer, Thomas Rückstieß, Alex Graves, Jan Peters, and Jürgen Schmidhuber. 2010. Parameter-exploring policy gradients. *Neural Networks* 23, 4 (2010), 551–559.
[24] Felipe Such. 2018. *Accelerating Deep Neuroevolution: Train Atari in Hours on a Single Personal Computer.* https://eng.uber.com/accelerated-neuroevolution/
[25] 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:1712.06567* (2017).
[26] I. Sutskever, O. Vinyals, and Q. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in NIPS.* 3104–3112.
[27] Yujin Tang and David Ha. 2021. The Sensory Neuron as a Transformer: Permutation-Invariant Neural Networks for Reinforcement Learning. In *The 35th Conference on Neural Information Processing Systems.*
[28] Yujin Tang, Duong Nguyen, and David Ha. 2020. Neuroevolution of Self-Interpretable Agents. In *Genetic and Evolutionary Computation Conference.*
[29] Yujin Tang, Jie Tan, and Tatsuya Harada. 2020. Learning agile locomotion via adversarial training. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).* IEEE, 6098–6105.
[30] Yingtao Tian and David Ha. 2021. Modern Evolution Strategies for Creativity: Fitting Concrete Images and Abstract Concepts. *arXiv:2109.08857* (2021).
[31] Nihat Engin Toklu, Paweł Liskowski, and Rupesh Kumar Srivastava. 2020. ClipUp: A Simple and Powerful Optimizer for Distribution-Based Policy Evolution. In *International Conference on Parallel Problem Solving from Nature.* 515–527.
[32] Twan Van Laarhoven. 2017. L2 regularization versus batch and weight normalization. *arXiv:1706.05350* (2017).
[33] Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. 2019. Paired open-ended trailblazer (poet): Endlessly generating increasingly complex and diverse learning environments and their solutions. *arXiv:1901.01753* (2019).
[34] Kaichao You, Mingsheng Long, Jianmin Wang, and Michael I Jordan. 2019. How does learning rate decay help modern neural networks? *arXiv:1908.01878* (2019).