

# Evolutionarily-Curated Curriculum Learning for Deep Reinforcement Learning Agents

Michael Cerny Green, Benjamin Sergent, Pushyami Shandilya, Vibhor Kumar  
Imbellus, Los Angeles, CA

## Abstract

In this paper we propose a new training loop for deep reinforcement learning agents with an evolutionary generator. Evolutionary procedural content generation has been used in the creation of maps and levels for games before. Our system incorporates an evolutionary map generator to construct a training curriculum that is evolved to maximize loss within the state-of-the-art Double Dueling Deep Q Network architecture with prioritized replay (Wang et al. 2016) (Schaul et al. 2015). We present a case-study in which we prove the efficacy of our new method on a game with a discrete, large action space we made called *Attackers and Defenders*. Our results demonstrate that training on an evolutionarily-curated curriculum (directed sampling) of maps both expedites training and improves generalization when compared to a network trained on an undirected sampling of maps.

## 1 Introduction

The use of games as benchmarks for AI progress has propagated to nearly the entire AI research community, including Chess, Atari Breakout, and more recently with Go as superhuman agents have been developed. Many recent papers document new AI methods being used within game environments. But the current state-of-the-art training method to ensure generalization in a neural network system of any kind remains brute force random sampling training, i.e. give a network enough unique states to train on and hope generalization naturally occurs. We propose a new method of directed sampling training called 'evolutionarily-curated curriculum learning' (ECCL), which we argue results in faster and better network generalization.

Past experiments have shown the potential in teaching simple concepts first, on which more complicated ones can then be taught, as a successful way to train networks (Elman 1993). To go one step further, we propose a method which specifically identifies weaknesses in the network and then generates content that force the network to face these weaknesses head-on. Our system dynamically evolves a curriculum by searching for content that maximizes the network's loss, which makes the network generalize faster and perform better.

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

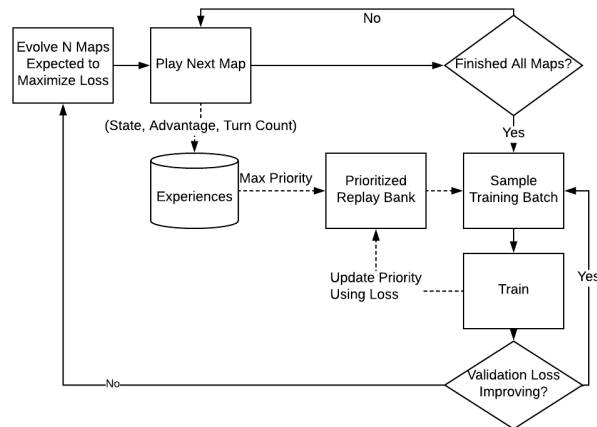


Figure 1: Evolutionarily-based curriculum learning in an agent's training loop

In this paper, we give a brief overview of research within reinforcement learning, the use of evolutionary algorithms in procedural content generation, and curriculum learning research for networks in Section 2. In Section 3 we discuss the theory of evolutionarily-based curriculum learning and how it could be applied to a reinforcement learning agent. We then use *Attackers and Defenders* as a case study in Section 4, with results and discussion from our experiment in Section 5, and conclude in Section 6.

## 2 Background

This section begins with a brief overview of deep reinforcement learning research, beginning with Minsky in 1954 (Minsky 1954) and finishing with the state-of-the-art DDDQN (Wang et al. 2016), AlphaZero (Silver et al. 2017b; 2017a), and ExIt (Anthony, Tian, and Barber 2017) agent architectures. It then discusses evolutionary algorithms and how they can be applied toward procedural content generation in games. The section concludes with the concept of curriculum learning for machines and the admittedly scant amount of research within this area.

## 2.1 Deep Reinforcement Learning

Reinforcement Learning (RL) concerns itself with the idea of learning through trial-and-error interactions with a dynamic environment and balancing the reward trade-off between long-term and short-term planning (Sutton and Barto 1998). RL has been studied since Minsky (Minsky 1954) in the 1950's. Since then, important improvements to the concept have been advanced including the temporal difference learning method (Sutton 1984; 1988), on which q-learning (Watkins and Dayan 1992) and actor-critic (Barto, Sutton, and Anderson 1983) techniques are built. Gullapalli (Gullapalli 1990) and Williams (Williams 1992) are early examples of the use of RL within artificial neural networks (ANNs). When Rumelhart et al discovered the *backpropagation* algorithm (Rumelhart, Hinton, and Williams 1986), deep learning took off in popularity. This has been bolstered recently by the rise in the capability and affordability of computer-processing units and graphical-processing units. Further readings of work in RL can be found in reviews by Schmidhuber (Schmidhuber 2015) and Szepesvári (Szepesvári 2010).

RL applied to deep learning has been only recently successful due to some key advancements. Mnih et al proposed Deep Q Networks (DQNs) (Mnih et al. 2015) using target networks and experience replay to improve the known divergence issues present in RL. van Hasselt et al built Double Deep Q Networks (Van Hasselt, Guez, and Silver 2016) which help reduce the overestimation errors that normal DQNs suffer from. Hessel et al wrote a paper surveying state-of-the-art improvements to DQN within the Atari framework, including Double Dueling Deep Q Networks, Distributional Deep Q Networks, and Noisy Deep Q Networks (Hessel et al. 2017). Prioritized experience replay (Schaul et al. 2015) allows DQNs to remember past experiences in a prioritized fashion, rather than at the same frequency that they were experienced.

Further DQN stability has been attained through the introduction of dueling networks (Wang et al. 2016), which are built on top of Double DQNs. Dueling networks use two separate estimators for the state value function and state-dependent action advantage function to generalize learning across actions without effecting the underlying RL algorithm. Double Dueling DQNs are currently considered state-of-the-art reinforcement learning algorithms.

Asynchronous Advantage Actor Critic (A3C) networks were created by Mnih et al in a successful attempt to apply neural networks to actor-critic reinforcement learning. The "asynchronous" part of A3C makes training parallelizable, allowing for massive computation speedups.

The AlphaGo algorithm combined Monte Carlo Tree Search (MCTS) with deep neural networks to play the game of Go and become the first AI to beat the human world champion of the game (Silver et al. 2016). The more advanced version, AlphaZero, was able to learn only by self-playing and outperformed AlphaGo (Silver et al. 2017b). The same AlphaZero architecture was then applied to both Chess and Shogi to convincingly beat world-champion programs (Silver et al. 2017a). At the same time, Anthony et al discovered the Expert Iteration algorithm (Anthony, Tian, and Bar-

ber 2017), which also uses a neural network policy to guide tree search. Since the advent of these state-of-the-art algorithms, much research has been done to either improve or apply them to different problems with varying degrees of success.

## 2.2 Evolution and Procedural Content Generation

Evolutionary algorithms (EA) fall within the area of optimization search inspired by Darwinian evolutionary concepts such as reproduction, fitness, and mutation (Togelius, Shaker, and Nelson 2016). EA has been used within games to procedurally generate levels, game elements within them, and sometimes even games themselves (Khalifa et al. 2017). Puzzle generation is a primary example of this kind of search-based generation (Ashlock 2010), which can be used to create puzzles with a desired solution difficulty. *Checkpoint based fitness* allows for fitness function parameterization (Ashlock, Lee, and McGuinness 2011), affording substantial control over generated properties. Stylistic generation is made possible by using fashion-based cellular automata (Ashlock 2015). An EA generator for a given game can also evolve many things at once by decomposing level generation into multiple parts. McGuinness et al did this by creating a micro evolutionary system which evolves individual tile sections of a level and an overall macro generation system which evolves placement patterns for the tiles (McGuinness and Ashlock 2011). Evolutionary search can be used for generalized level generation in multiple domains such as General Video Game AI (Khalifa et al. 2016) and PuzzleScript (Khalifa and Fayek 2015a). In later work by Khalifa et al. (Khalifa et al. 2018), they worked on generating levels for a specific game genre (Bullet Hell genre) using a new hybrid evolutionary search called Constrained Map-Elites. The levels were generated using automated playing agents with different parameters to mimic various human play-styles. Green et al. used EA to evolve *Super Mario Bros* (Nintendo, 1985) scenes which taught specific mechanics to the player (Green et al. 2018). We recommend Khalifa's review on searched-based level generation for further reading into EA for generation in games (Khalifa and Fayek 2015b).

## 2.3 Curriculum Learning in Machines

The concept of curriculum learning (CL) in machines can be traced back to Elman in 1993 (Elman 1993). The basic idea is to keep initial training data simple and slowly ramp up in difficulty as the model learned. Krueger and Dayan (Krueger and Dayan 2009) did a cognitive-based analysis with evidence that shaping data provided faster convergence. CL was further explored by Bengio et al. in 2009, in an attempt to define several machine learning guided training strategies (Bengio et al. 2009). Their experiments suggested that incorporating CL into training a model could both speed up training and significantly increase generalization. Recently, Curriculum Learning within adversarial network training (Cai et al. 2018) was explored by Cai et al in an attempt to mitigate "forgetfulness" and increase gener-




Figure 2: A visualization of a map in *Attackers and Defenders*

alization to reduce the effectiveness of adversarial network attacks.

## 2.4 Evolution within Networks

Using evolutionary strategy for neural networks is a well-researched topic. Genetic algorithms were used in the node weight balancing of a network by Ronald et al (Ronald and Schoenauer 1994) to evolve a controller for soft-landing a toy lunar module in a simulation. Simultaneously, Gruau evolved the structures and parameters of networks using cellular encoding (Gruau and others 1994). Cartesian Genetic Programming was first designed by Miller et al (Miller, Thomson, and Fogarty 1997) to design digital circuits using genetic programming. It is called ‘Cartesian’ because of the way it represents a program using a 2-dimensional set of nodes.

All of these methods and more may be housed under the umbrella of “neuroevolution” which is well-defined by Floreano et al (Floreano, Dürr, and Mattiussi 2008). A survey of neuroevolution within games is written by Risi and Togelius (Risi and Togelius 2017). We mention neuroevolution to highlight the major difference that whereas it is used to evolve the parameters or architecture of a network, our approach evolves training data as part of a curriculum for a constant architecture.

## 3 Evolved Curriculum in the Training Loop

Traditional training of a neural network involves a training schedule established by taking random batches of a fixed training set, which is assumed to be an unbiased random sample of the data space. For game-playing agents specifically, this training set is the set of levels or maps that the network is exposed to. The hope is that the sample is sufficient to train the neural network to generalize to the entire data space (i.e. all possible maps).

Instead, ECCL relies on producing a training curriculum composed of a biased sampling of the data space, specifically designed to improve generalization. ECCL involves

two parts: the agent to be trained and an evolutionary generator. Unlike neuroevolution methods, the use of evolution in evolved curriculum is *not* to evolve the weights or architecture of the network. The evolutionary generator’s sole purpose is to evolve *scenarios* which have the best potential to improve the agent’s generalization. To do this, the goal of the evolutionary generator is to produce maps which maximize the loss of the agent network.

Figure 1 shows the training loop that uses evolutionarily-based curriculum learning. In this figure, the agent is a Double Dueling Deep Q Network (DDDQN) with Prioritized Replay (Wang et al. 2016; Schaul et al. 2015) which we used in our case study, explained in Section 4. One reason for using this specific type of network architecture is the choice of our library- Tensorflow. The network used in experiments presents in this paper is near state-of-the-art with a few deletions that we deemed unnecessary to use.

When the network requests more maps to train on, the evolutionary generator is tasked to evolve maps which maximize the amount of loss in the network by querying the loss directly. By maximizing loss, the network is necessarily seeing a valid map in map space that it failed to generalize to properly. Concretely, this could be a edge-case map that where an uncommon move is optimal or requires another strategy altogether from maps the network has already seen. The agent plays this map, as well as any others in the batch.

After finishing a batch, the system divides the different game states caused by moves from all games into experience snapshots, and sorts the experience snapshots by priority to be stored in a prioritized replay bank. It then uses these experiences to train the weights in the network appropriately, and updates priorities in the bank using network loss. The network then asks the generator to produce more maps to repeat the process until training terminates.

## 4 Case Study: Attackers and Defenders

The following section concerns a case study in which we compare our method’s impact against other state-of-the-art algorithms. We hypothesized that the evolutionary generator would create higher quality training data which would more effectively improve network performance, and our experiment attempts to prove this.

Section 4.1 explains the game of *Attackers and Defenders* a simple tower defense game which we created as a testbed. Section 4.3 explains the generator and how it produces maps. Section 4.4 describes the training/testing methods used to validate our claims.

### 4.1 Attackers and Defenders

To prove the concept of ECCL, we created a discrete, large action space, tower-defense game called *Attackers and Defenders* as a test-bed. Figure 2 displays a visualization of a game map. The objective of the player in *Attackers and Defenders* is to prevent enemies from reaching their *home* tile for as long as possible. This game is a model of sequential decision making that applies broadly to other game and non-game domains, which makes it an appropriate testbed for our algorithm.

Game Entity	Description
Neutral	an empty tile with no penalties
Slow	a tile making attackers 2 turns to move
Block	a tile preventing attackers from moving
Home	the tile attackers are trying to move onto
Source	the tiles from which attackers spawn
Attacker	automatous entities which are moving toward the home tile
Defender	entities which the player can place; these do damage to all attackers within range

Table 1: A table with all entities in the game

**Game Entities** *Attacker* entities have hit-points (HP), which may vary in number and generally increase over the course of a single play session. To facilitate this survival goal, the player is given *Defender* entities which do damage to *attacker* HP, *slow* tiles which penalize *attacker* movement when traveling through these tiles, and *block* tiles which prohibit *attacker* movement. Table 1 displays all tiles/entities within the game and how they work.

**Game Loop** Each turn, the player is prompted to place a *defender*, *slow*, or *block* tile on the game map. After the player places an entity, the game advances forward one turn. During this period, a *source* tile may spawn an *attacker*, which will then slowly advance toward the *home* tile. If an *attacker* moves into a space within a *defender*'s attack range (which may overlap with other *defenders*), the *attacker* will suffer damage equivalent to the sum of all in-range *defender* damage. If an *attacker* runs out of HP, it will be destroyed. If an *attacker* manages to move onto the *home* tile, the game will end.

## 4.2 Constructive Generator

In order to appropriately measure the effect of an evolved curriculum versus the impact of increased access to additional training points, we design the constructive generator. With access to a set of underlying parameters of the data, each with a set of acceptable values, and a global set of constraints, there exists a constructive generator that produces unbiased random samples of this data by simply permuting over possible combinations of parameter values, and simply throwing away those combinations that do not satisfy the constraints. Training with such a generator (which we refer to as our "constructive" network) is analogous to the undirected sampling case where the training data is fixed.

In *Attackers and Defenders*, the constructive generator is given the available tile types and where they can be placed (i.e. parameter values and set of possible values), along with the constraints presented in Table 2, and simply selects random combinations of tile values, outputting only those combinations that satisfy the constraints and discarding the rest.

## 4.3 Evolutionary Generator

Our system uses the Feasible Infeasible 2-Population (FI-2Pop) genetic algorithm (Kimbrough et al. 2008) to evolve boards. FI-2Pop is an evolutionary algorithm which uses two

populations: a feasible population and an infeasible population. The infeasible population aims at improving infeasible solutions to "legally-playable" threshold, when they become feasible and are transferred to the feasible population. The feasible population, on the other hand, aims at improving the quality of feasible chromosomes. If one becomes infeasible, it is then relocated to the infeasible population. After evolving solutions for several generations, the system outputs the board with the highest fitness.

## Chromosomal Representation, Crossover, and Mutation

A board chromosome is represented as a 2-dimensional array of tile types. Crossover (Figure 3) is done using 2-d array crossover, by picking a sub-array within one parent and swapping it with the other, creating a new board as a result. Mutation is done by selecting a random tile and changing its type. Mutation may be performed multiple times on a single board after crossover is completed.




Figure 3: The 2D array representation of an *Attackers and Defenders* board. Crossover is shown, using Parent 1 as a template and Parent 2 as a replacement sub-array (black outline). A tile from resulting board is then mutated (red outline), with the fitness of the new child to be calculated later.

**Evaluating Feasibility and Fitness** Each board chromosome contains two fitness functions which determine where they fail in terms of feasibility and fitness. The *constrained* fitness dictates whether they are within the infeasible population, and the *feasible* fitness ascertains how optimal of a board it is.

*Constrained fitness* is calculated by averaging the constraint factors listed in Table 2. If the constrained average is 1, then this chromosome is feasible. *Feasible fitness* is measured by calculating the loss from the agent's loss network on the given map. The larger the loss, the higher the feasible fitness for the chromosome.

Factor	Description
Separate Quads	% of <i>sources</i> in different board quadrants
Home Paths	% of <i>sources</i> that initialize with a path to <i>home</i> tile
Home Center	1 if <i>home</i> is near center of board, else 0
Home Blocks	1 if no <i>blocks</i> near <i>home</i> , else 0

Table 2: The constraint factors present in the generator

#### 4.4 Procedure

To evaluate the effectiveness of an evolutionarily-based curriculum inside a reinforcement learning training loop, we created several training schedules. A Double Dueling Deep Q Network (Wang et al. 2016) (DDDQN) using prioritized replay and a separate loss network was trained from initialization on each schedule. Figure 4 displays the architecture of this network, and Figure 6 displays the separate loss network.

The replay bank of the DDDQN holds 20,000 training experiences using hyperparameters  $\alpha = 0.6$ ,  $\beta_0 = 0.4$  annealed to  $\beta = 1.0$  over the first 1000 games. To create experiences for the replay bank, the network plays a map from *Attackers and Defenders*, after which it stores all encountered initial states, actions, next states, and rewards as experiences in the bank. The network updates its weights every 5 maps it plays. A training cycle consists of 250 batches which contain 32 samples selected according to prioritized replay. The Q-value update uses a future discount factor  $\gamma = 0.99$ . The loss of each individual experience is used to update the priority. Afterwards, the loss network is trained using the initial state as input and the loss as the target.

All schedules begin with 50 maps created using a constructive generator to train the loss network such that its output is usable to assess a maps loss potential. This constructive generator provides an undirected random sampling of the game space. The 50 starting maps are used are the identical for every schedule. After these initial maps, the schedules differ. The first of these schedules continues to contain only maps constructed by the constructive generator. The second of these schedules contains only evolutionarily-curated maps. The third network contained equal mix of randomly generated maps and evolutionarily-curated maps. Table 3 defines each network’s training curriculum.

For each schedule, the network was tested and scored on a fixed set of 1000 randomly generated maps after every 200 training maps. The network was optimizing to slay the maximum amount of attackers over the course of play and was scored based on how many of these attackers were slain before an attacker reached the home tile. Training continued until the network failed to improve for two consecutive testing cycles.

## 5 Results & Discussion

Here, we present the results of the previously described case study. We compare the results of the fully-evolved-curriculum-trained network (full network), the mixed-evolved-curriculum-trained network (mixed network), and

Network	Curriculum
DQN 1	50 + 100% randomly constructed maps
DQN 2	50 + 100% evolutionarily-curated maps
DQN 3	50 + 50% randomly constructed maps & 50% evolutionarily-curated maps

Table 3: The networks and their training curriculum ratios

the randomly-generated-curriculum-trained (constructive network) trained on randomly sampled maps.

As Figure 7 demonstrates, the full network generalizes well within 400 maps of training with a score of 21.47, peaking after just 1,600 maps at 22.14. In comparison, the constructive network never reaches this score. Even after training on 6,800 maps, it only peaks at 20.83 at 1,600 maps. The mixed network peaked at 20.22 after 2,000 maps which is slightly below the constructive network’s peak.

Figure 5 displays the loss of each network. Each tick displays the average loss of from a training cycle containing 250 batches, out to a total of 3,500 maps. The full network starts out at 14.44, higher than the constructive network at 12.67. This suggests that the full network learning was presented with maps which high learning potential, as expected. Very quickly the two networks converge and hover between 7.5 and 9 which corresponds the full network’s peak performance within the first several hundred maps. The constructive network gradually increases over time to 10 then stabilizes after it reached peak performance. Contrary to expectations, the mixed network on the other hand shows a much higher loss than either of the other two networks starting at 40.95 and remaining substantially higher ranging from 10 to 14.

This suggests that the mixed network failed to generalize when presented with an equal mix of evolved maps and generated maps. The network appears to have learned how to discriminate evolve maps from randomly generated maps in a manner that harmed performance on the test set. Specifically, the network learned two classes of strategies: evolved and constructive. As a result, the learning from evolved examples would not generalize correctly to constructive maps that were used in the testing set. By contrast, the full network only saw evolved maps after the first 50 maps and was able to generalize the strategies learned from evolved maps to randomly generated maps in the test set.

Lastly, the sole requirement of the evolutionary generator is to specify the parameters of the game itself and is generalizable. The network informs the generator of its loss function which makes it not specific to any domain. It is sufficient for the network set up to have the ability to interact with the game as the architecture is not dependent on it. Since both halves of the systems (generator and network) are generalizable themselves, slight modifications to the system can make it adapt very well to a new scenario.

## 6 Conclusion

In this paper, we have introduced evolutionarily-curated curriculum learning as a new methodology to train reinforcement agents. We performed a case study using a game we

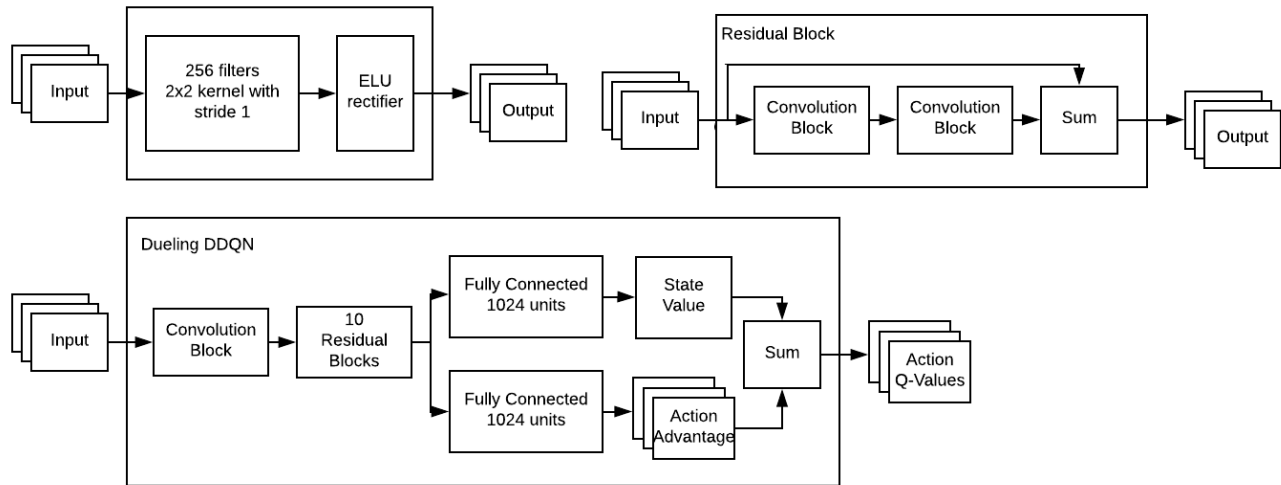


Figure 4: The architecture of our DDDQN consists of a convolutional layer followed by a tower of 10 residual blocks. The final residual block is fed to two separate fully-connected layers to produce the current states Q-value and the predicted advantage of each possible action. The streams are combined to produce predicted action Q-values.




Figure 5: The loss over time for the training of each network. Loss was collected every 5 maps of training by averaging the loss across the 5 maps.

created called *Attackers and Defenders* to prove the validity and effectiveness of this new method. Specifically, we tested a Double Dueling Deep Q-network (DDDQN) with prioritized replay and a separate loss network using this method.

Based on our results, our initial hypothesis that evolutionarily-curated curriculum learning helps networks generalize better and faster than undirected sampling, has been proven true in this environment. Even after nearly three times the amount of training time, the constructive trained network never approaches the performance of the full evolutionarily-curated network. Therefore, it appears that this new training methodology, ECCL, can be used to both expedite training and increase generalization or max performance.

However, the mixed network does not appear to perform

as well despite the fact that its loss values are much higher. This was a result we did not initially expect, allowing for any amount of evolutionarily-based curriculum learning would improve network training. Upon inspection, the mixed network spent considerable effort in differentiating between maps coming from the constructive generator vs. the evolutionary generator. This suggests that a discriminator network trained to predict whether a map was constructed or evolved could be added to the evolutionary generator’s fitness functions. Another possibility would be using a similarity metric in fitness to ensure evolved maps are sufficiently different from previously evolved maps to prevent the network from learning to recognize evolved maps.

Given the generalizability of ECCL as in the discussion from the previous section, as it only requires a data generator and a game-playing agent architecture, we also expect it to work well with AlphaGo Zero-based agents as well, and leave that open for future work.

## References

- Anthony, T.; Tian, Z.; and Barber, D. 2017. Thinking fast and slow with deep learning and tree search. In *Advances in Neural Information Processing Systems*, 5360–5370.
- Ashlock, D.; Lee, C.; and McGuinness, C. 2011. Search-based procedural generation of maze-like levels. *IEEE Transactions on Computational Intelligence and AI in Games* 3(3):260–273.
- Ashlock, D. 2010. Automatic generation of game elements via evolution. In *Computational Intelligence and Games (CIG), 2010 IEEE Symposium on*, 289–296. IEEE.
- Ashlock, D. 2015. Evolvable fashion-based cellular automata for generating cavern systems. In *Computational*

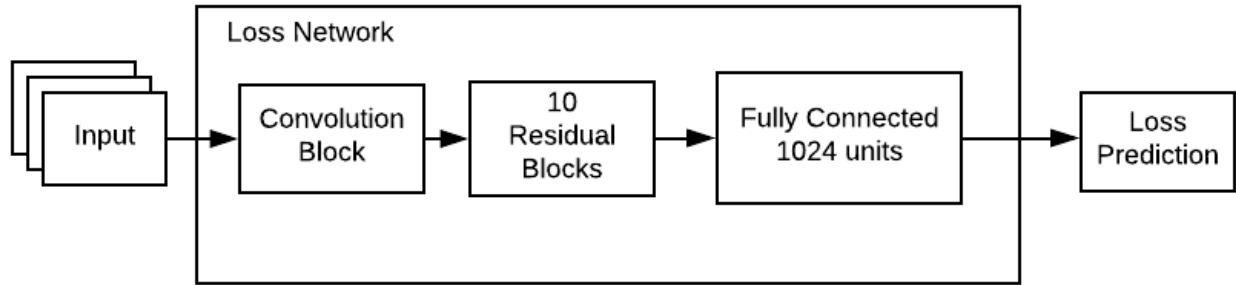


Figure 6: The architecture of our loss network consists of a convolutional layer followed by a tower of 10 residual blocks which feeds a fully connected layer that outputs the loss prediction.




Figure 7: The testing score results averaged in 1,000 map batches over the course of training. The y-intercept marks network performance at network initialization with randomized weights.

*Intelligence and Games (CIG), 2015 IEEE Conference on, 306–313. IEEE.*

Barto, A. G.; Sutton, R. S.; and Anderson, C. W. 1983. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE transactions on systems, man, and cybernetics* (5):834–846.

Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, 41–48. ACM.

Cai, Q.; Du, M.; Liu, C.; and Song, D. 2018. Curriculum adversarial training. *CoRR* abs/1805.04807.

Elman, J. L. 1993. Learning and development in neural networks: The importance of starting small. *Cognition* 48(1):71–99.

Floreano, D.; Dürer, P.; and Mattiussi, C. 2008. Neuroevo-

nual international conference on machine learning, 41–48. ACM.

- lution: from architectures to learning. *Evolutionary Intelligence* 1(1):47–62.
- Green, M. C.; Khalifa, A.; Barros, G. A.; Nealen, A.; and Togelius, J. 2018. Generating levels that teach mechanics. In *Proceedings of the Foundation of Digital Games*.
- Gruau, F., et al. 1994. Neural network synthesis using cellular encoding and the genetic algorithm.
- Gullapalli, V. 1990. A stochastic reinforcement learning algorithm for learning real-valued functions. *Neural networks* 3(6):671–692.
- Hessel, M.; Modayil, J.; Van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2017. Rainbow: Combining improvements in deep reinforcement learning. *arXiv preprint arXiv:1710.02298*.
- Khalifa, A., and Fayek, M. 2015a. Automatic puzzle level generation: A general approach using a description language. In *Computational Creativity and Games Workshop*.
- Khalifa, A., and Fayek, M. 2015b. Literature review of procedural content generation in puzzle games. <http://www.akhalifa.com/documents/LiteratureReviewPCG.pdf>.
- Khalifa, A.; Perez-Liebana, D.; Lucas, S. M.; and Togelius, J. 2016. General video game level generation. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, 253–259. ACM.
- Khalifa, A.; Green, M. C.; Perez-Liebana, D.; and Togelius, J. 2017. General video game rule generation. In *Computational Intelligence and Games (CIG), 2017 IEEE Conference on*, 170–177. IEEE.
- Khalifa, A.; Lee, S.; Nealen, A.; and Togelius, J. 2018. Talakat: Bullet hell generation through constrained map-elites. In *Proceedings of The Genetic and Evolutionary Computation Conference*. ACM.
- Kimbrough, S. O.; Koehler, G. J.; Lu, M.; and Wood, D. H. 2008. On a feasible–infeasible two-population (fi-2pop) genetic algorithm for constrained optimization: Distance tracing and no free lunch. *European Journal of Operational Research* 190(2):310–327.
- Krueger, K. A., and Dayan, P. 2009. Flexible shaping: How learning in small steps helps. *Cognition* 110(3):380–394.
- McGuinness, C., and Ashlock, D. 2011. Decomposing the level generation problem with tiles. In *Evolutionary Computation (CEC), 2011 IEEE Congress on*, 849–856. IEEE.
- Miller, J. F.; Thomson, P.; and Fogarty, T. 1997. Designing electronic circuits using evolutionary algorithms. arithmetic circuits: A case study.
- Minsky, M. L. 1954. *Theory of neural-analog reinforcement systems and its application to the brain model problem*. Princeton University.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529.
- Risi, S., and Togelius, J. 2017. Neuroevolution in games: State of the art and open challenges. *IEEE Transactions on Computational Intelligence and AI in Games* 9(1):25–41.
- Ronald, E., and Schoenauer, M. 1994. Genetic lander: An experiment in accurate neuro-genetic control. In *International Conference on Parallel Problem Solving from Nature*, 452–461. Springer.
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1986. Learning representations by back-propagating errors. *nature* 323(6088):533.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Schmidhuber, J. 2015. Deep learning in neural networks: An overview. *Neural networks* 61:85–117.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of go with deep neural networks and tree search. *nature* 529(7587):484.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2017a. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017b. Mastering the game of go without human knowledge. *Nature* 550(7676):354.
- Sutton, R. S., and Barto, A. G. 1998. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge.
- Sutton, R. S. 1984. Temporal credit assignment in reinforcement learning.
- Sutton, R. S. 1988. Learning to predict by the methods of temporal differences. *Machine learning* 3(1):9–44.
- Szepesvári, C. 2010. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning* 4(1):1–103.
- Togelius, J.; Shaker, N.; and Nelson, M. J. 2016. The search-based approach. In Shaker, N.; Togelius, J.; and Nelson, M. J., eds., *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*. Springer. 17–30.
- Van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep reinforcement learning with double q-learning. In *AAAI*, volume 2, 5. Phoenix, AZ.
- Wang, Z.; Schaul, T.; Hessel, M.; Hasselt, H.; Lanctot, M.; and Freitas, N. 2016. Dueling network architectures for deep reinforcement learning. In *International Conference on Machine Learning*, 1995–2003.
- Watkins, C. J., and Dayan, P. 1992. Q-learning. *Machine learning* 8(3-4):279–292.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8(3-4):229–256.