



Self-adapting Goals Allow Transfer of Predictive Models to New Tasks

Kai Olav Ellefsen^{1(✉)} and Jim Torresen²

¹ Department of Informatics, University of Oslo, Oslo, Norway
kaiolae@ifi.uio.no

² Department of Informatics and RITMO, University of Oslo, Oslo, Norway

Abstract. A long-standing challenge in Reinforcement Learning is enabling agents to learn a model of their environment which can be transferred to solve other problems in a world with the same underlying rules. One reason this is difficult is the challenge of learning accurate models of an environment. If such a model is inaccurate, the agent's plans and actions will likely be sub-optimal, and likely lead to the wrong outcomes. Recent progress in model-based reinforcement learning has improved the ability for agents to learn and use predictive models. In this paper, we extend a recent deep learning architecture which learns a predictive model of the environment that aims to predict only the value of a few key measurements, which are indicative of an agent's performance. Predicting only a few measurements rather than the entire future state of an environment makes it more feasible to learn a valuable predictive model. We extend this predictive model with a small, evolving neural network that suggests the best goals to pursue in the current state. We demonstrate that this allows the predictive model to transfer to new scenarios where goals are different, and that the adaptive goals can even adjust agent behavior on-line, changing its strategy to fit the current context.

Keywords: Reinforcement Learning · Prediction · Neural networks · Neuroevolution

1 Introduction

Humans and animals rely on internal models (mental simulations of how objects respond to interaction) to predict the consequences of their actions and generate accurate motor commands in a wide range of situations [13, 17]. Recent advances in deep learning have enabled computers to learn such predictive models by gathering a large collection of observations from an environment [9, 16]. This opens up the possibility for guiding the actions of robots and computer agents by having them predict the consequences of each action and selecting the one leading to the best outcome.

When using predictions for guiding the actions of an agent, it is beneficial to limit the prediction to the most essential parts of the environment. For instance, if we would like to predict the effect of turning the steering wheel of a car, we should not try to predict the effect this has on birds we see in the sky,

trees far away from us, and so on. Rather, we should focus on the effect on some key observable measurements, such as the car speed, and the distance to other cars and pedestrians. If we focus on predicting the parts of the environment we currently care about, the prediction problem becomes much more manageable.

This intuition is the background for a recent, popular technique for learning to act by predicting the future [2]. The technique makes the assumption that we can analyze the outcome of an action by focusing on a *few measurements*. These measurements should be the observable quantities that are most related to success or failure in some scenario. The authors propose that a way to learn how to act in an environment is to learn to predict the effect one’s actions have on these measurements. Such predictions can readily be learned by gathering a large set of examples of observations, actions and resulting measurements from the environment (e.g. by simulating thousands of car trips).


Once one has learned to predict the measurements resulting from one’s actions, it is possible to select the best action for a given situation, by choosing the one giving the *most optimal predicted measurements*. This requires some way to map the predicted future measurements to a number representing the *utility* of this future. [2] solved this by defining a goal vector, which weights the different measurements according to user-defined rules. For instance, a rule could be that we give a very high weight to the measurement of distance to the nearest pedestrian, and a lower positive weight to the measurement of car speed, reflecting that we want to drive efficiently, but only if pedestrians around us are safe.

In this paper, we explore the potential for *automatically adapting* such goal vectors to a given scenario (Fig. 1). We show that this allows the reuse of a learned predictive model in new situations, where the best strategy has changed. The automatically adapting goal vectors, together with the already learned predictive model, can quickly generate behaviors for new scenarios where the underlying rules are the same (allowing reuse of learned predictions), but the goals are different. We further show that our adaptive goal vectors can adapt agent strategies on-line, responding to changes as they occur.


2 Background

2.1 Model-Based Reinforcement Learning

The problems we target in this paper are reinforcement learning (RL) problems, meaning an agent is tasked with learning how to solve some problem without any explicit guidance – relying instead on infrequent feedback in the form of rewards and punishment from the environment. Reinforcement learning can be divided into two high-level categories: Model-based and model-free RL. Model-free RL means we try to solve the problem without forming an explicit model of the environment, relying instead on learning a mapping from observations to values or actions. Many of the recent successes in deep RL have been in this category, including deep Q learning [10], which was the first example of power of



(a) The Goal-ANN produces goals adapted to the current situation



(b) The predictive ANN (developed by [2]) predicts the consequence of taking each action in the current situation.

Fig. 1. Combining adaptive goals with prediction-based action selection. Top: Our proposed Goal-ANN produces a weighting of the agent’s goals adapted to the current situation (indicated by current measurements). Bottom: The resulting weights are given to the predictive ANN, together with the current state and current measurements, resulting in a vector of predicted future values for the measurements.

deep RL for playing computer games, and the deep deterministic policy gradient algorithm (DDPG) which demonstrated that deep neural networks can also learn to solve continuous control problems [8].

Model-based RL attempts to solve two key challenges with model-free approaches: (1) They require enormous amounts of training data, and (2) there is no straightforward way to transfer a learned policy to a new task in the same environment [12]. To do this, model-based RL takes the approach of first learning a predictive model of the environment, before using this model to make a plan that solves the problem. An internal predictive model facilitates transfer of knowledge to new tasks in the same environment: Once the predictions are learned, they can be used for planning ahead to solve many different tasks.

Despite these advantages of model-based algorithms, model-free RL has so far been most successful for complex environments. A key reason for this is that model based RL is likely to produce very bad policies if the learned predictive model is imperfect, which it will be for most complex environments [12]. Recently, algorithms have been developed which address this problem, for instance by learning to interpret imperfect predictions [12], applying new video prediction techniques [3] and dynamics models [4], or periodically restarting prediction sequences, reducing the effect of accuracy degrading for long-term prediction [5].

Another way to mitigate the problem of having an imperfect predictive model is to keep the prediction task as simple as possible. While the methods above generally try to predict the entire future state (more specifically, entire frames of

input in the example of video games), one could predict more accurately by focusing on predicting exactly the values needed to guide action making. Dosovitskiy and Koltun [2] presented impressive results on the VizDoom RL environment by teaching agents to predict just a few key measurements, and combining this with a *goal vector* that defines a weighting among the measurements, effectively defining which type of future we most wish to observe. Here, we propose to combine this method with adaptive goal vectors, which has the potential to allow reuse of a predictive model in new scenarios where the goals are different.

2.2 Combining Deep Learning and Neuroevolution

Our proposed method combines two neural networks (Fig. 1), which are trained in different ways. The Prediction ANN is trained in a self-supervised manner, using the difference between predicted and actual future states as loss function. For the Goal-ANN, we have no target value, since we do not know the “correct” goal. We therefore optimize this networks by using neuroevolution, a technique employing a population of neural networks, having the ones performing their task best become further adapted and specialized to the problem [14].

A few other papers have recently combined deep learning with neuroevolution, aiming to get the benefits of both. Neuroevolution (NE) requires far more computation to solve problems than backpropagation-based deep learning. On the other hand, NE does not rely on a differentiable architecture, and works well in problems with sparse rewards, which are a challenge for most deep reinforcement learning algorithms [11]. A promising way to combine NE and DL is therefore to let the deep learning do the “heavy lifting”, for instance learning to make predictions or recognize objects based on a large number of examples, and train a small action-selection component using NE with the pre-trained deep neural network as a back-end. This approach was taken by [3], who trained a self-supervised deep neural network to predict the future, and then evolved a much smaller network for selecting actions based on internal states in the predictive network. A similar idea is to train a convolutional neural network to translate raw pixels to compact feature representations, before training a smaller evolving ANN to use the learned features to choose the right actions in an RL problem. This has been done successfully for simulated car racing [7], for learning to aim and shoot in a video game [11] and for a health-gathering VizDoom scenario [1].

Similarly to the work described above, our method lets the DL component do the data-intensive job, which is to learn to predict the consequences of actions from a large number of examples. Unlike the methods above, we only require our deep network to predict a few key future measurements, greatly simplifying its task, and potentially reducing the impact of errors due to faulty predictions.

3 Methods

3.1 Learning to Act by Predicting the Future

The algorithm we propose in this paper is an extension of “Direct Future Prediction” (DFP) from the paper *Learning to act by predicting the future* [2]. DFP is

based on the idea of transforming an RL problem into a self-supervised learning problem of predicting parts of future states and selecting actions that produce those states that align best with an agent’s goals. A deep neural network is trained to predict the future value of a set of different measurements, \mathbf{m} , given the current state and action. \mathbf{m} should contain metrics that give an indication of how successfully an agent is solving its task. The scenario explored here and in [2] is a video game where an agent attempts to survive and kill monsters. \mathbf{m} therefore contains measurements of an agent’s ammunition, health and the number of enemies killed, which are all interesting indicators of an agent’s performance.

If an agent can learn a model of how its actions affect future measurements, selecting the best action is reduced to a problem of finding which of the potential future \mathbf{m} -vectors correspond best with the agent’s goals. [2] solve this by formulating a goal-vector \mathbf{g} , where each element indicates how much we care about the corresponding measurement in \mathbf{m} . For instance, given that we use ammunition, health and enemies killed as measurements, the goal-vector $[0.5, 1, -1]$ would indicate that we are interested in collecting ammunition, twice as interested in collecting health, and interested in *avoiding* attacking enemies (due to the negative value for that goal).

Training this network to predict the effect of one’s actions can now be done by collecting many episodes of gameplay, storing at each timestep t current measurements (\mathbf{m}_t), sensory input (\mathbf{s}_t – frames of images from gameplay in this case), the action taken (a_t) and the goal vector (\mathbf{g}_t – this is not changed during an episode while training). A sample of $(\mathbf{s}_t, \mathbf{m}_t, a_t, \mathbf{g}_t)$ would now be the training input (\mathbf{I}), while the target output value (\mathbf{O}) would be the change in the measurements at selected future timesteps ($\mathbf{m}_{t+\tau_1}, \dots, \mathbf{m}_{t+\tau_n}$). The network is trained using backpropagation with the loss function depending on the difference between its predicted change in future measurements ($\hat{\mathbf{O}}$) and the actual observed values (\mathbf{O}).

3.2 Adaptive Goals Guiding Action Selection

The contribution of this paper is a technique for adapting the goals (\mathbf{g}) of DFP-agents, which enables them to (1) transfer predictive models to new tasks and (2) adjust goals on-line according to current measurements. To do so, we take a pre-trained predictive model from the DFP-algorithm, and extend it with an additional neural network that suggests the best goal-vector given the current measurements (Fig. 1).

This Goal-ANN is trained using neuroevolution [14]: A population of neural networks compete for their ability to produce relevant goals. The best networks are randomly changed, by adding or removing nodes and connections, while worse networks are discarded. To select between networks, they are each assigned a *fitness score* reflecting how well they perform their task. In our setup, we calculate the fitness by inserting the agent governed by the evaluated Goal-ANN

and the (pre-trained) Predictive ANN into the game scenario, and measuring how well it performs its task (which varies slightly between our different experiments – see Sect. 4). To counter effects of randomness, each evaluation tests the agent 8 times and calculates the average performance.

The evolving networks are relatively small feed-forward ANNs where the connectivity and number of neurons are being optimized. The inputs are the 3 measurements (\mathbf{m}_t), and the output is the goal vector (\mathbf{g}_t) with the 3 corresponding goal weights. The inputs give the current value for the amount of ammunition, health and number of enemy kills (in that order), and the outputs represent the weight of the ammunition-objective, the health-objective, and the killing monsters-objective, all in the range $[-1, 1]$.

Our setup uses the popular neuroevolution algorithm NEAT [15], more specifically the NEAT-Python implementation¹, with the parameters shown in Table 1. With these parameters, each run of the algorithm results in 5,000 evaluations, corresponding to 5,000 1-min episodes of game play. The original DFP-algorithm was trained for around 95,000 episodes, demonstrating that our proposed goal-adaptation is an order of magnitude faster than training new strategies from scratch.

Table 1. Parameters for Python-NEAT

Parameter	Value	Parameter	Value
Add Connections	0.15	Population Size	50
Delete Connection	0.1	Generations	100
Add Node	0.15	ANN Inputs / Outputs	3 / 3
Delete Node	0.1	Weights Range	$[-30, 30]$
Weight Mutation	0.8	Activation Function	Clamped linear response in range $[-1, 1]$
Weight Replacement	0.02		

(a) Mutation Rates

(b) Other parameters

3.3 Training Scenario

The scenario we use for training and testing the Goal-ANN is from the original paper suggesting DFP [2]. The scenario is based on the VizDoom platform [6], a popular platform for developing and testing reinforcement learning algorithms. VizDoom is based on the first-person shooter (FPS) Doom, and offers the potential to train agents to handle complex 3D environments directly from pixel inputs. VizDoom scenarios can be used to train and test skills such as understanding one’s surroundings, navigation, exploration and dealing with opponents/enemies. The violent nature of the game is a concern, and we are eager to test our technique on more peaceful scenarios soon. However, to test extending the exact model trained in [2], we needed to reuse one of their scenarios.

¹ <https://neat-python.readthedocs.io/>.

The specific VizDoom scenario we use here is the one titled “D3-Battle” in [2]. This is a challenging scenario where the agent is under attack by alien monsters inside a maze, and has to try to kill as many of the monsters as possible. To aid the agent, ammunition and health kits are scattered around the maze. The agent is provided with (and learns to predict) three measurements: Its current amount of ammunition and health, as well as how many enemies it has killed. We use a trained predictive model from the original paper, which was trained in a “goal-agnostic” manner, that is, the goal vector \mathbf{g} was randomized between episodes (each value uniformly sampled from the interval $[-1, 1]$). Such goal-agnostic training was found to generalize better to new tasks than fixed-goal predictors [2], and we therefore focus on this model for our study.

4 Results

We want to explore two possible advantages of evolved goal weights: (1) Their ability to adapt an existing predictive model to new scenarios where strategies need to be different, and (2) Their ability to produce context-dependent goals, that is, goals that vary depending on current measurements.

To do so, we compare three main techniques: (1) Acting by following the same static goal-vector applied in [2], (2) Acting by following a simple rule for adapting the goal to the current situation and (3) Acting by following the evolved Goal-ANN. The static goal vector from [2] is $[0.5, 0.5, 1.0]$ where the three numbers represent the importance of the ammunition-, health- and enemy kills-objectives, respectively. This goal captures the intuition that the aim of the game is to kill as many monsters as possible, but collecting some ammunition and health is a good side-goal to help the agent in its primary mission. The simple rule to improve this static goal (called “Hardcoded” in plots below) is to switch to the goal $[0, 1.0, -1.0]$ whenever the agent’s health is below 50%, which adds the intuition that we should stop attacking and focus on gathering health when injured.

In all plots below, fitness values are averaged over 20 independent evaluations of each agent.

4.1 The Original Scenario

As mentioned above, we test our method on a VizDoom scenario developed by [2], using their pre-trained agent together with our evolved adaptive goals. The scenario consists of a maze populated by monsters and the agent. The goal of the trained agent is to kill as many monsters as possible, and to do so, it can benefit from collecting ammunition and health packages along the way. In this original scenario, the final reward (also referred to as Fitness below, as is common in Evolutionary Algorithms) is based only on the number of monsters killed.

Figure 2 shows the average reward (number of monsters killed in one minute) of the three compared techniques on the original scenario. There are no significant differences between the compared techniques. In other words, there is no advantage to adapting the goals in this scenario.

Analyzing the agents’ gameplay in this scenario reveals why this is the case: Agents are very fast, killing monsters without taking much damage, which makes the default goal of aggressively attacking monsters while picking up any health or ammo ahead a very good choice. In other words, there is no conflict among the goals and no reason why a different choice of goal weights should improve performance.




Fig. 2. The original scenario. No significant differences in fitness values between using the original goal (Static), a simple adaptive goal (Hardcoded) and the evolved Goal-ANN.

4.2 A Hard Scenario

To test the ability of the adaptive goals to solve different scenarios, and the potential for adjusting goals during gameplay, we set up a much harder scenario. We introduce the following difficulties that force the agent to sometimes change its strategy between defensive and aggressive modes:

- Monsters are tougher (they have twice as much health as before).
- Player is weaker (it begins the game with only 10% health)
- A fitness penalty of 100 is given for dying (versus 0 before).
- Player starts with 0 initial ammunition (versus 20 bullets before).





Fig. 3. The hard scenario. Left: The evolved adaptive Goal-ANN significantly outperforms both the standard (static) and hardcoded (dynamic) goals. Right: The mean population value for each goal through evolution. The evolved strategies tend to focus more on health and less on attacking.

Since the monsters are now stronger and the player weaker, the default static goals result in a strategy that is too aggressive and frequently results in the player

dying. The hardcoded strategy performs better, since it balances aggression and defense. The evolved strategy significantly outperforms both the others ($p < 0.05$), finding an even better balance between the three objectives (Fig. 3a).

We can see what the evolved strategy has learned by plotting the output of the Goal-ANN. In Fig. 3b we plot the output of the Goal-ANN per generation of evolution, averaged both over the entire population and all timesteps in each individual’s life. We see that the evolved strategies gradually move towards a focus on the health-objective, while reducing focus on attacking enemies.

Figure 3b summarizes strategies over complete game episodes. However, the Goal-ANN has the potential to produce strategies that vary through a game, depending on the current measurements. To investigate if this is taking place, we performed a sweep over sensible values for all three measurements, passed these values through the best evolved Goal-ANN and measured the resulting output goals. The results show that changing the current health or the number of kills has no effect on the produced goals. However, changing the amount of ammunition does affect the output goals, as plotted in Fig. 4. We see that evolution has discovered the strategy of focusing less on attacking and more on gathering ammunition when the current ammunition level is low².




Fig. 4. The hard scenario. The strategy discovered by evolution is to focus on gathering ammunition until some has been collected, after which the agent switches to attacking. Simultaneously, the agent focuses on collecting health.

4.3 The No-Ammunition Scenario

To test the potential for the Goal-ANN to adapt strategies to scenarios where rules are very different, we set up a scenario with no ammunition available to the agent. The scenario is otherwise identical to the hard scenario above. This change turns the game into a defensive exercise, where staying away from enemies and gathering health is the best strategy. Since the aggressive default goal now results in a very bad strategy, we here test an additional *defensive* goal, which has the maximum negative weight on the attack-objective, and maximum positive weights on the two others ($\mathbf{g} = [1, 1, -1]$).

As expected, the default aggressive strategy performs very badly in this scenario, while the hardcoded and the new defensive strategy do better (Fig. 5a). The evolved strategy significantly outperforms all others, with $p < 0.01$ according to the Mann-Whitney U test, for all pairwise comparisons. We were initially surprised to see the evolved strategy outperform the defensive strategy, since

² <https://youtu.be/NCzrO5KHMxQ> shows an agent playing according to this strategy.

we expected its full focus on avoiding confrontation to be optimal. We therefore took a closer look at the evolved strategy³.




Fig. 5. The scenario with no ammunition. Left: Evolved adaptive goals generated by the Goal-ANN significantly outperform both the static (Aggressive and Defensive) and Hardcoded goals. Right: The mean population value for each goal through evolution. The evolved strategies quickly learn to not focus on attacking.

As expected, we found the evolving strategy to give a strong negative weight to the attack-objective, and strong positive weights to the two others (Fig. 5b). Repeating the sweep across measurements, we found that the current health measure is the only value that leads to different goals when changed. The evolved strategy is to focus exclusively on collecting health when the measure is critically low, and otherwise to also include a drive for collecting ammunition, in both cases avoiding attacking as strongly as possible (Fig. 6). It may seem surprising that the ammunition objective is valuable at all here, since there is no ammunition in the environment. We hypothesize that a high value for this objective can be valuable, since it can encourage the agent to keep moving, heading towards objects with a small resemblance of ammunition, thus staying away from danger.




Fig. 6. The scenario with no ammunition. The strategy discovered by evolution is to focus on only on health if it is low, but add focus on gathering ammunition otherwise.

³ <https://youtu.be/6pTnkCGV6NI> shows an agent playing according to this strategy.

5 Conclusion

We proposed extending an existing technique for selecting actions based on predictions of the consequences of those actions with an adaptive goal-producing neural network. This Goal-ANN can modify the behavior resulting from predictions of the future, by changing the agent’s preference among different future outcomes.

We demonstrated that the Goal-ANN allows the transfer of a predictive model to scenarios where the optimal behavior is different, due to different amounts of resources in the environment. Such knowledge transfer has the potential to greatly improve the efficiency of Reinforcement Learning methods, since they allow training a model on one task, and adapting it rapidly to other, related problems.

We also showed that the Goal-ANN is capable of learning *adaptive* strategies, where goals change on-line, depending on the current state of the environment. This is a very valuable property, since most complex problems require one to modify one’s strategy depending on the current context.

A final valuable feature of the Goal-ANN is its *interpretability*: We can easily probe the rules it has learned by sweeping across input measurements and observing the resulting output goals. We demonstrated that in our target scenarios, such an analysis allows us to verify that the learned goal-switching behavior is sensible.

This has been an initial study of the potential for combining self-adapting goals with a predictive deep neural network. There are many future questions here that we aim to address, including tests on more scenarios, comparison with other techniques for transfer learning and testing if the Goal-ANN could perform even better if given more information about the current state of the world.

Acknowledgments. This work is supported by The Research Council of Norway as part of the Engineering Predictability with Embodied Cognition (EPEC) project #240862, and the Centres of Excellence scheme, project #262762.

References

1. Alvernaz, S., Togelius, J.: Autoencoder-augmented neuroevolution for visual doom playing. In: 2017 IEEE Conference on Computational Intelligence and Games (CIG) (2017)
2. Dosovitskiy, A., Koltun, V.: Learning to act by predicting the future. In: ICLR 2017, pp. 1–14 (2017)
3. Ha, D., Schmidhuber, J.: Recurrent world models facilitate policy evolution. In: Advances in Neural Information Processing Systems 31, pp. 2451–2463. Curran Associates, Inc. (2018)
4. Hafner, D., et al.: Learning latent dynamics for planning from pixels. arXiv preprint [arXiv:1811.04551](https://arxiv.org/abs/1811.04551), November 2018
5. Kaiser, L., et al.: Model-based reinforcement learning for atari. arXiv preprint [arXiv:1903.00374](https://arxiv.org/abs/1903.00374) (2019)

6. Kempka, M., Wydmuch, M., Runc, G., Toczek, J., Jaskowski, W.: ViZDoom: a doom-based AI research platform for visual reinforcement learning. In: IEEE Conference on Computational Intelligence and Games, CIG (2017)
7. Koutnik, J., Schmidhuber, J., Gomez, F.: Evolving deep unsupervised convolutional networks for vision-based reinforcement learning. In: Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation, pp. 541–548. ACM (2014)
8. Lillicrap, T.P., et al.: Continuous control with deep reinforcement learning. arXiv preprint [arXiv:1509.02971](https://arxiv.org/abs/1509.02971) (2015)
9. Luc, P., Neverova, N., Couprie, C., Verbeek, J., Lecun, Y.: Predicting deeper into the future of semantic segmentation. In: Proceedings of the IEEE International Conference on Computer Vision (2017)
10. Mnih, V., et al.: Human-level control through deep reinforcement learning. *Nature* **518**(7540), 529–533 (2015)
11. Poulsen, A.P., Thorhauge, M., Funch, M.H., Risi, S.: DLNE: a hybridization of deep learning and neuroevolution for visual control. In: 2017 IEEE Conference on Computational Intelligence and Games, CIG 2017 (2017)
12. Racanière, S., et al.: Imagination-augmented agents for deep reinforcement learning. In: Advances in Neural Information Processing Systems 30, pp. 5690–5701. Curran Associates, Inc. (2017)
13. Schillaci, G., Hafner, V.V., Lara, B.: Exploration behaviours, body representations and simulations processes for the development of cognition in artificial agents. *Front. Robot. AI* **3**, 39 (2016)
14. Stanley, K.O., Clune, J., Lehman, J., Miikkulainen, R.: Designing neural networks through neuroevolution. *Nat. Mach. Intell.* **1**(1), 24–35 (2019)
15. Stanley, K.O., Miikkulainen, R.: Evolving neural network through augmenting topologies. *Evol. Comput.* **10**(2), 99–127 (2002)
16. Villegas, R., Yang, J., Zou, Y., Sohn, S., Lin, X., Lee, H.: Learning to generate long-term future via hierarchical prediction. In: ICML, April 2017
17. Wolpert, D.M., Doya, K., Kawato, M.: A unifying computational framework for motor control and social interaction. *Philos. Trans. R. Soc. B Biol. Sci.* **358**(1431), 593–602 (2003)