# Generalization of Agent Behavior through Explicit Representation of Context

Cem Tutum<sup>1</sup>, Suhaib Abdulquddos<sup>1</sup>, and Risto Miikkulainen<sup>1,2</sup>

<sup>1</sup>The University of Texas at Austin; <sup>2</sup>Cognizant Technology Solutions tutum@cs.utexas.edu, suhaib@utexas.edu, risto@cs.utexas.edu

## **ABSTRACT**

In order to deploy autonomous agents in digital interactive environments, they must be able to act robustly in unseen situations. The standard machine learning approach is to include as much variation as possible into training these agents. The agents can then interpolate within their training, but they cannot extrapolate much beyond it. This paper proposes a principled approach where a context module is coevolved with a skill module in the game. The context module recognizes the temporal variation in the game and modulates the outputs of the skill module so that the action decisions can be made robustly even in previously unseen situations. The approach is evaluated in the Flappy Bird and LunarLander video games, as well as in the CARLA autonomous driving simulation. The Context+Skill approach leads to significantly more robust behavior in environments that require extrapolation beyond training. Such a principled generalization ability is essential in deploying autonomous agents in real-world tasks, and can serve as a foundation for continual adaptation as well.

# 1 INTRODUCTION

To be effective, autonomous agents need to be able to perform robustly in previously unseen situations. Especially in real-world decision making and control applications such as games, simulations, robotics, health care, and finance, agents routinely encounter situations beyond their training, and need to adapt safely. A common practice is to train these models, mostly deep neural networks, with data collected from a number of hand-designed scenarios. However, the tasks are often too complex to anticipate every possible scenario, and this approach is not scalable.

One popular approach to address this problem is few-shot learning, in particular metalearning, either by utilizing gradients [3, 18, 22] or evolutionary procedures [2, 6]. In metalearning, systems are trained by exposing them to a large number of tasks, and then tested for their ability to learn relevant but previously unseen tasks. There are also a number of approaches mostly for supervised learning setting where new labels need to be predicted based on limited number of training data. However, applications of few-shot learning to control and decision making, including reinforcement learning problems, are limited so far [10].

The approach in this paper is motivated by work on opponent modeling in poker [12, 13]. In that domain, an effective approach was to evolve one neural network to decide what move to make, and another to modulate those decisions by taking the opponents playing style into account. When trained with only a small number of very simple but different opponents, the approach was able to generalize and play well against a wide array of opponents, include some that were much better than anything seen during training.

In poker, the opponent can be seen as the context for decision making. Each decision needs to take into account how the opponent is likely to respond, and select the right action accordingly. The player can thus adapt to many different game playing situations immediately, even those that have not been encountered before. In this paper, this approach is generalized and applied to control and decision making more broadly. A skill network reacts to the current situation, and a context network integrates observations over a longer time period in parallel. Together they learn to represent a wide variety of situations in a standardized manner so that a third, controller, network can make decisions robustly. Such a Context+Skill system can thus generalize to more situations than any of its components alone.

The Context+Skill approach is evaluated in this paper on several game domains: (1) Flappy Bird game extended to include more actions and physical effects (i.e. flap forward and drag in addition to flap upward and gravity); (2) LunarLander-v2 (from OpenAI Gym) extended with variations in main and side engine power as well as mass of the lander; and (3) CARLA autonomous driving environment where the steering curve, torque curve, and map can vary. Such extensions allow generating a range of unseen scenarios both by extending the range of effects of those actions as well as their combinations. The approach generalizes remarkably well to new situations, and does so much better than its components alone. Context+Skill approach is thus a promising approach for building robust autonomous agents in real-world domains.

## 2 METHODOLOGY

This section introduces the experimental domains, the Context+Skill approach, and the evolutionary training methodology.

## 2.1 Experimental Domains

Three different environments are used for experiments. The first one, Flappy Ball (FB; Fig. 1(a)), is an extension of the popular Flappy Bird computer game [24] The agent, controlled by a neural network, aims to navigate through the openings between pipes without hitting them for a certain length of time. The agent can flap forward and flap upward; gravity will pull it down and drag will slow it down. The agent gets a reward of +1 every time it passes a pipe successfully, and a penalty of -1 at each time step it crashes into the pipes and -5 when it crashes in the ceiling or ground. At every time step, the agent receives six-dimensional sensory information: vertical position, horizontal and vertical velocities, horizontal distance to the right edge of the closest pipe, and the height of the top and bottom pipes, normalized to [0,1]. The effects of flap upward and forward as well as gravity and drag can change

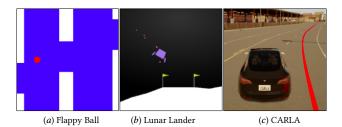


Figure 1: Scenes from Flappy Ball (FB; (a)), LunarLander-v2 (LL; (b)), and CARLA ((c). The red circle in the FB domain represents the agent and the white columns are pipes that move from right to left as the game progresses. The purple object in the LL domain represents the spacecraft that controls its main and side engines to land safely on the white surface designated with flags. In CARLA, the agent controls the throttle and steering of a car to reach destination staying as close as possible within the lane indicated by the red curve. The effects of the actions are varied across episodes to evaluate how well the controller adapts to previously unseen situations. For animated examples, see https://drive.google.com/drive/folders/1GBdJzD9tDHJkd59YbQUOIQua6nCiLjXa.

between episodes; Therefore, the agent has to infer such variations from its interactions with the environment over time.

The second domain is LunarLander-v2 (LL; Fig. 1(b)) from OpenAI Gym suite, modified to allow variations in mass of the spacecraft and the effect of the main and two side engine thrusters. As in the original LL-domain, the agent receives eight-dimensional numerical sensory information as its input (i.e., position and velocity of the lander in 2D, its angle and angular velocity, and whether each leg touches the ground). The purpose is to safely land on the lunar surface in a designated region indicated with the flags. The episode finishes if the lander crashes or comes to rest.

The third domain is the CARLA open-source autonomous driving simulation environment (Fig. 1(c)). The agent has both lateral (steering) and longitudinal (throttle) control of the car while driving between two points in a given certain amount of time. The steering and the torque curves are modified to evaluate generalization. Furthermore the agent is placed in completely unseen tracks as it tries to complete the same task albeit under more difficult driving conditions (as imposed by the modified steering and torque curves). As in the other environments, the agent receives numerical sensory information as five numerical values: Rangefinder coordinates which describe the distance between the agent and the lane boundaries along five axes. All control actions are given as continuous values varying within [-1,1].

These domains can be seen as proxy for control and decision making problems where the changes in the environment require immediate adaptation, such as operating a vehicle under different weather conditions, configuration changes, wear and tear, or sensor malfunctions. The challenge is to adapt the existing policies to the new conditions immediately without further training, i.e. to generalize the known behavior to unseen situations.

#### 2.2 Context+Skill Model

The core idea evaluated in this paper is to implement the agent as a Context+Skill network that takes advantage of an explicit representation of context. The Context+Skill Network consists of three components: the Skill and the Context modules and the Controller (Figure 2(a)). The first two modules receive sensory information from the environment as numerical values, as described in previous subsection. They send their output to the Controller, a fully connected feedforward neural network that makes the decisions on which actions to take.

The Skill module is also a fully connected feedforward network. Together with the Controller they form the Skill-only Network S (Fig. 2(c)). The Skill module used in this study has 10 hidden and five output nodes and the Controller has 20 hidden hidden nodes. S is used as the baseline model throughout the study. In principle it has all the information for navigating through the pipes, but does not have the benefit of explicit representation of context.

The other main component in the Context-Skill framework is the Context module. It is composed of a vanilla Long Short Term Memory (LSTM) cell ([8]). There are three gates in this recurrent memory cell: input, forget, and output. The gates are responsible for learning what to store, what to throw away, and what to read out from the long-term memory of the cell. Thus, the cell can learn to retain information from the past, update it, and output it at an appropriate time, thereby making it possible to learn sequential behavior [5, 7].

The C-module used in this study consists of an LSTM cell size of 10. The memory of the C-module ( $h_{t-1}$  and  $c_{t-1}$ ) is reset at the beginning of each new task, and accumulated (transferred) across episodes within each task. It can therefore form a representation of how actions affect the environment. The output of the LSTM ( $h_t$ ) is sent to Controller as the context. Together the C-module and the Controller form the Context-only network C (Fig. 2(b)). It serves as a second baseline, allowing integration of observations over time, but without a specific Skill network to map them directly to action recommendations.

The complete Context-Skill Network, CS (Fig. 2(a)) consists of both the Context and Skill modules as well as the Controller network of the same size as in C and S. The motivation behind the CS architecture, i.e. of integrating the Context module into S, is to make it possible for the system to learn to use an explicit context representation to modulate its actions appropriately. The method for discovering these behaviors is discussed next.

## 2.3 Evolutionary Learning Process

The goal in each domain is to find a safe solution that optimizes the performance objective. Although it is possible to formulate the optimization process based on that single objective, it turns out the diversity resulting from the multiobjective search speeds up training and helps discover well-performing solutions [11]. In the experiments in this paper, the solution that is perfectly safe or close to it is returned as the final result, and its generalization ability evaluated.

In each domain, the first objective,  $f_0$ , measures safety, and the second,  $f_1$ , measures performance, and is in conflict with safety. Accordingly in the FB domain, the number of any type of collisions

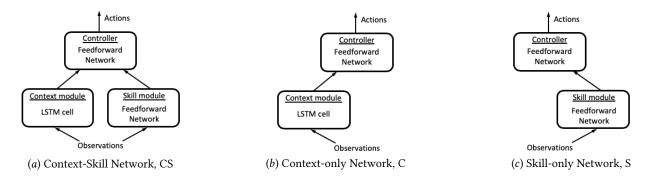


Figure 2: The architecture of the Context+Skill network and its ablations. (a) The full network consists of three components: a Skill module that processes the current situation, a Context module that integrates observations over the entire task, and a Controller that combines the outputs of both modules, thereby using context to modulate actions. This architecture is compared to (b) context-only ablation, and (c) skill only ablation in the experiments. Each component is found to play an important role, allowing the CS network to generalize better than its ablations.

Table 1: Parameter ranges during training

FB (±20%)	LL (±10%)	CARLA (±15%)
Flap <sub>base</sub> = -12.0	Main <sub>base</sub> = 20.0	$\alpha_{\rm base}$ =1.0
$Fwd_{base} = 5.0$	Side <sub>base</sub> =1.0	$\beta_{\rm base}$ =1.0
Gravity <sub>base</sub> = 1.0	Mass <sub>base</sub> =8.0	
Drag <sub>base</sub> = 1.0		

(or hits) is minimized, whereas the number of successfully passed pipes is maximized. In the LL-domain, the total rewards (indicating a successful landing) is maximized, whereas the landing time is minimized. In the CARLA domain, safety  $f_0$  is optimized by minimizing a cumulative weighted sum of a lane-distance penalty and a wobbliness penalty:

$$f_0 = \sum_{t=1}^{t=T} (d(t) + \lambda abs(s(t) - s(t-1))), \tag{1}$$

where d(t) is the distance from the center of the lane at time t (shown as a red line in Fig. 1(c)), s(t) is the steering output at time t, and  $\lambda$  is a proportionality coefficient (= 5.5 in the experiments below). Time is incremented in 1/30 sec intervals; each episode lasts T=20 seconds, or less if the agent reaches a specified target zone before that. Performance  $f_1$  is optimized by minimizing the Euclidean distance between the agent and the target zone at the end of the episode.

In the evolutionary learning process, the weights of all three neural networks described in Section 2.2 are evolved while the network architecture remains fixed. The goal is to maximize the average fitness across multiple tasks, where each task is based on different physical parameters in the FB, LL and CARLA domains. The FB-domain has four parameters (Flap, Fwd, Gravity, Drag), the LL-domain has three parameters (Main, Side, Mass) and the CARLA-domain has two parameters ( $\alpha$  and  $\beta$ , which control the steering angle and torque curves of the car, respectively). In each task during evolution, only one parameter is subject to change, while the rest are fixed at their base values (given in Table 1). Thus, the number of tasks is equal to the number of parameters. The

Algorithm 1 Evolutionary process for training networks

```
1: procedure EVOLVE
        stop := False
        parents := random_init_individuals(μ)
        task\_params = prepare\_task\_params(n_{episodes}, n_{tasks})
4:
5:
        fitness = distribute(eval_fitness(), (parents, task_params))
6:
        for gen from 1 to ngen do
            offspring = tournament_sel_DCD(parents, \mu)
7:
            for i from 1 to \lambda do
8:
9
                 if random() \leq p<sub>crossover</sub> then
                     SBX(offspring[i], offspring[i+1])
10:
                 Polynomial_Mutation(offspring[i])
11:
12:
                 Polynomial_Mutation(offspring[i+1])
            params = prepare\_task\_params(n_{episodes}, n_{tasks})
13:
             fitness = distribute(eval fitness(), (parents, task params))
14:
            for j from 1 to \lambda do
15:
                 if fitness[j][0] \ge pipes_{max} then
16:
                     if fitness[j][1] \le hits_{max} then
17:
                         stop := True
18
             parents := tournament_sel_DCD(parents + offspring, \mu)
19:
            if stop == True then
20:
                 return parents
21:
                 break
22:
```

parameters in each domain are varied during training within  $\pm 20\%$ ,  $\pm 10\%$  and  $\pm 15\%$  in the FB, LL and CARLA-domains, respectively.

Each task, and therefore each parameter, is uniformly sampled  $n_{\rm episodes}$ =5 (except in the CARLA domain where each parameter is sampled six times) times within the limits specified above. Thus, there are 20 fitness evaluations per individual in each generation in the FB domain, 15 evaluations in the LL-domain and 12 evaluations in the CARLA-domain. The fitness of every individual in the population, i.e., average score of all episodes, is evaluated in parallel on the same task distribution for a fair comparison. The memory of C-module in CS and C is reset at the beginning of each task, and transferred from episode to episode otherwise.

Non-dominated sorting genetic algorithm (NSGA-II) [1] was implemented in DEAP [4] as the optimization method. The overall procedure is shown in Algorithm 1. It receives  $n_{tasks}$ ,  $n_{episodes}$ =5, perturb=0.2 in FB, 0.1 (in LL) or  $n_{episodes}$ =12, perturb=0.15 (in CARLA), base parameter values,  $\mu$  = 96 (48 in CARLA),  $p_{crossover}$  = 0.9,  $p_{gen}$  = 2,500 as input. The population size ( $\mu$ ) is chosen as a multiple of 24 since the fitness evaluations are distributed among 24 threads on a cluster (i.e., Dell PowerEdge M710, 2× Xeon X5675, six cores @ 3.06GHz). Default NSGA-II settings, including SBX (Simulated Binary Crossover), Polynomial Mutation, tournament\_sel\_DCD (Tournament Selection Based on Dominance), and ( $\mu$  +  $\lambda$ ) elitist selection strategy were used [1]. The process is robust to minor variations of these choices.

#### 3 RESULTS

The evolutionary learning results are first described, followed by evaluation of the generalization ability of the best solutions. For animated examples of behaviors in each domain, see https://drive.google.com/drive/folders/1GBdJzD9tDHJkd59YbQUOIQua6nCiLjXa.

# 3.1 Learning

CS, C, and S architectures were evolved with different random seeds six times in FB, five times in LL, and once in CARLA (due to the computational complexity of this domain). To avoid overfitting, specific stopping criteria were selected for each domain after some experimentation. In FB, training was run until an individual in the population achieved a fitness scores of at least  $f_0$ =0.01 (hits) and  $f_1$ =22.0 (pipes). In LL, training was run until an individual reached  $f_0$ =200 (total rewards). In CARLA, evolution was run until an individual reached  $f_0$ =2043 (safety penalty) and  $f_1$ =55 meters (distance from the target). In FB and CARLA a minimum performance criteria was included to make sure the agent would not simply optimize safety by staying still. The final Pareto-optimal set in each run contained multiple individuals; in FB and LL a safe network (i.e. satisfying the stopping criteria) with the highest performance, and in CARLA the network with the least Euclidean distance to the origin in the Pareto front, was selected as the final network, and its generalization ability evaluated.

The evolution of S in general took the shortest amount of generations since it has the least number of parameters (287 compared with 982 in C and 1207 in CS). The same architecture is used for all three domains. To make sure the number of parameters was not a factor, another S with a larger Skill module, with the same number of parameters as CS, was also evolved with the same stopping criterion. However, it performed poorly compared to the smaller S in the generalization studies, apparently because it was easier to overfit. Thus, it was excluded from the comparisons.

## 3.2 Generalization

To evaluate the generalization performance of the best performing networks, the task parameters in each domain were changed in the following ways:

- The range of variation in the task parameters was increased significantly beyond the training limits (Table 2);
- All parameters were varied simultaneously instead of one at a time; and

Table 2: Parameter ranges during generalization

FB (±75%)	LL (±50%)	CARLA (±35%)
Flap <sub>min</sub> = -21.0	$Main_{min} = 10.0$	$\alpha_{\min} = 0.65$
$Flap_{max} = -3.0$	$Main_{max} = 30.0$	$\alpha_{\text{max}} = 1.35$
$Fwd_{min} = 1.25$	$Side_{min} = 0.5$	$ \bar{\beta}_{\min} = 0.65 $
$Fwd_{max} = 8.75$	$Side_{max} = 1.5$	$\beta_{\text{max}} = 1.35$
Gravity <sub>min</sub> = $0.25$	$Mass_{min} = 4.0$	
Gravity <sub>max</sub> = 1.75	$Mass_{max} = 12.0$	
Drag <sub>base</sub> = 0.25		
Drag <sub>base</sub> = 1.75		

 In the CARLA domain, the agents were evaluated on an entirely new track that had never been encountered during training. This track included curves with significantly higher curvature than those encountered during training.

More specifically, each parameter axis was divided into equal steps (i.e. 10, 20 and 35 in FB, LL and CARLA, respectively). In FB and LL, each set of task parameters were sampled three times and averaged; in CARLA, only one sample was used. Therefore, all three networks were tested for  $3\times 10^4$ ,  $24\times 10^3$  and  $1.225\times 10^3$  episodes in these domains. Figures 3, 4, and 5 show the difference in generalization between CS and S, CS and C and C and S. They are histograms indicating how often each difference was observed in the episodes. The first network in the subtraction performs better wrt. a minimization objective if the histogram is skewed to the left, and wrt. a maximization objective if the histogram is skewed to the right.

Thus, the plots show that in both FB and LL, CS generalizes much better than S and C both in terms of safety and performance. In CARLA, CS generalizes better than S and C in terms of performance; in terms of safety, CS generalizes much better than S in terms of lane distance and about the same in wobbliness, and it generalizes much better than C in terms of wobbliness and about the same or slightly less in terms of lane distance. Thus, CS consistently achieves superior performance and mostly superior safety across multiple domains.

## 4 BEHAVIOR ANALYSIS

To understand how the CS architecture outperforms its individual components C and S, an FB task with parameters [Flap=-7.0, Gravity=0.58, Fwd=8.75, Drag=0.58] was evaluated further. This setting has a previously unseen exaggerated effect for forward flap, and a previously unseen diminished effects for upward flap, gravity and drag. Thus, actions tend to push up and speed up the agents more than expected, and it is difficult for it to slow down and come down.

Neither the C nor the S network performed well in this task: The C network collided with six pipes and S with five. Remarkably, CS managed to pass all 21 pipes. Both C and S used all four actions (flap upward, forward, simultaneously upward and forward, and glide, i.e. do nothing), but CS interestingly never uses flap upward alone. That action simply lifts the agent up, which is rarely optimal action in this environment where it takes such a long time to come down. If it is necessary to go up it is because the opening is high,

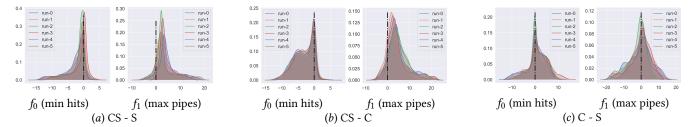


Figure 3: Comparing generalization of Context+Skill and its ablations in the FlappyBall domain. The x-axis shows the differences in generalization performance across the  $3 \times 10^4$  test episodes for the three pairs of architectures. With minimization objectives, a distribution that is skewed to the left of the 0 line (black dashed line) is better, and with maximization, a distribution that is skewed to the right is better. CS generalizes much better than S (a) and C (b), which are about equal (c).

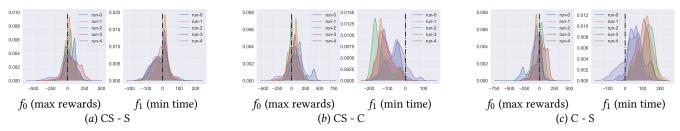


Figure 4: Comparing generalization of Context+Skill and its ablations in the LunarLander domain. As in FlappyBall, CS generalizes much better than S (a) and C (b), which are about equal (c).

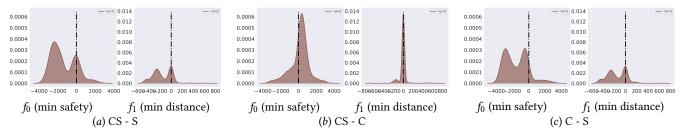


Figure 5: Comparing generalization of Context+Skill and its ablations in the CARLA domain. As in the other two domains, CS generalizes much better than S(a) and is slightly better in performance and about equal in safety than S(b). C in this domain performs better than S(c), suggesting that context dynamics are more important than reactivity in this domain.

and in that case it is more efficient to move forward at the same time.

As an illustration, second row in Fig. 6 shows a situation at the fourth and fifth pipe. Both C and S make a similar mistake by flapping up and forward. They end up too high too fast, do not have enough time to come back down, and crash into the fifth pipe. In contrast, even before the fifth pipe becomes visible, CS refrains from both actions while there is enough time for weaker gravity and drag to slow and pull down the agent, and it reaches the opening in the fifth pipe just fine.

To understand how CS manages to implement this behavior whereas C and S do not, the outputs of the C and S modules are compared between this generalization task (where C and S hit the fifth pipe) and a task where all parameters are at their base values (where C and S do not hit the fifth pipe). Their 10 and 5-dimensional outputs are first reduced to two dimensions through

principal component analysis and then subtracted. The top row of Fig. 6 shows these differences for the C- and S-modules of CS, for the C-module of C, and for the S-module of S, at the locations in the image below.

One might intuitively expect that C-module in C and S-module in S would not change their behavior much in the generalization task, but the C-module in CS would vary significantly to modulate the output of the S-module. Surprisingly, the opposite is true: Both the C and the S-module in CS vary very little compared to those in C and S (see also the quantitative comparison in Table 3). The generalization task presents novel inputs that results in novel outputs in C and S, and the controller does not know how to map them to correct actions. In contrast, the C and S-modules in CS have learned to standardize their output despite the change in context; their outputs are what the controller expects as its input, and is able to output

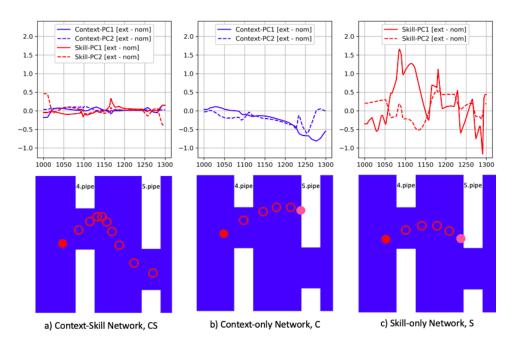


Figure 6: Contrasting the generalization ability of (a) CS, (b) C, and (c) S networks. Bottom row: At the fourth pipe in the generalization task, S and C flap upward and then forward, end up too high too fast without enough time to come back down, crashing into the fifth pipe. In contrast, the Context-Skill Network avoids the collision by correctly estimating the effects of its actions, giving itself enough time to come down. Top row: Differences between the principal components of the module outputs between the nominal and generalization tasks. They differ little in CS, making it easier for the controller to output the correct actions in new contexts. For an animation of these episodes, see https://drive.google.com/drive/folders/1GBdJzD9tDHJkd59YbQUOIQua6nCiLjXa.

Table 3: Change in C and S modules during generalization.

Network	PC	MSD	STD
CS	Context-PC1	0.004	±0.058
	Context-PC2	0.003	±0.039
	Skill-PC1	0.007	±0.082
	Skill-PC2	0.018	±0.133
С	Context-PC1	0.139	±0.282
	Context-PC2	0.061	±0.147
S	Skill-PC1	0.414	±0.600
	Skill-PC2	0.088	±0.286

the correct actions. Interestingly, this effect is similar to standardizing context in sentence processing, which makes it possible to generalize to novel sentence structures [14]. Remarkably, whereas in sentence processing the standardization was implemented by a hand-designed architecture, in Context+Skill it is automatically discovered by evolution.

#### 5 DISCUSSION AND FUTURE WORK

The Context+Skill approach represents context explicitly, and has a remarkable ability to generalize to unseen situations. In the experiments so far, the neural networks have a fixed topology; it may be possible to customize their architecture further through evolution [20, 21], and thereby delineate and optimize their roles

further. Besides the architecture, the choice of training tasks plays an important role; methods that automatically design a curriculum, i.e., a sequence of new training tasks [9, 15, 17, 19, 23], could lead to further improvements. Third, instead of using handcrafted features, convolutional layers added in front of the Context and Skill modules could be used to discover features while training, extending the approach to more general visual tasks.

Lifelong machine learning tries to mimic how humans and animals learn by accumulating the knowledge gained from past experience and using it to adapt to new situations incrementally [16]. The generalization ability of Context+Skill can serve as a foundation for continual learning. It provides an initial rapid adaptation to new situations upon which further learning can be based. How to convert generalization into a permanent ability in this manner is an interesting direction of future research.

## 6 CONCLUSION

A major challenge in deploying artificial agents in the real world is that they are brittle—they can only perform well in situations for which they were trained. This paper demonstrates a potential solution based on separating contexts from the actual skills. Context can then be used to modulate the actions in a systematic manner, significantly extending the unseen situations that can be handled. This principle was successfully evaluated in three domains: challenging versions of the Flappy Bird and LunarLander-v2 games as

well as the CARLA autonomous driving simulation. The results suggest that the Context+Skill approach should be useful in many control and decision making tasks in the real world.

## 7 ACKNOWLEDGMENTS

This research was supported in part by DARPA L2M Award DBI-0939454.

## REFERENCES

- Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197.
- [2] Fernando, C., Sygnowski, J., Osindero, S., Wang, J., Schaul, T., Teplyashin, D., Sprechmann, P., Pritzel, A., and Rusu, A. 2018. Meta-learning by the Baldwin effect. In Proceedings of the Genetic and Evolutionary Computation Conference Companion. ACM Press.
- [3] Finn, C., Abbeel, P., and Levine, S. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. (2017).
- [4] Fortin, F.-A., De Rainville, F.-M., Gardner, M.-A., Parizeau, M., and Gagné, C. 2012. DEAP: Evolutionary Algorithms Made Easy. *Journal of Machine Learning Research* 13 (jul 2012), 2171–2175.
- [5] Géron, A. 2017. Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems (1st ed.). O'Reilly Media. Inc.
- [6] Grbic, D. and Risi, S. 2019. Towards Continual Reinforcement Learning through Evolutionary Meta-Learning. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '19). Association for Computing Machinery, 119–120.
- [7] Greff, K., Srivastava, R. K., Koutnik, J., Steunebrink, B. R., and Schmidhuber, J. 2017. LSTM: A Search Space Odyssey. IEEE Transactions on Neural Networks and Learning Systems 28, 10 (2017), 2222–2232.
- [8] Hochreiter, S. and Schmidhuber, J. 1997. Long Short-Term Memory. Neural Comput. 9, 8 (1997), 1735–1780.
- [9] Justesen, N. and Risi, S. 2018. Automated Curriculum Learning by Rewarding Temporally Rare Events. (2018). arXiv:cs.AI/1803.07131
- [10] Kansky, K., Silver, T., Mély, D., Eldawy, M., Lázaro-Gredilla, M., Lou, X., Dorfman, N., Sidor, S., Phoenix, S., and George, D. 2017. Schema Networks: Zeroshot Transfer with a Generative Causal Model of Intuitive Physics. (2017).

- arXiv:cs.AI/1706.04317
- [11] Knowles, J. D., Watson, R. A., and Corne, D. W. 2001. Reducing Local Optima in Single-Objective Problems by Multi-objectivization. In Evolutionary Multi-Criterion Optimization, Eckart Zitzler, Lothar Thiele, Kalyanmoy Deb, Carlos Artemio Coello Coello, and David Corne (Eds.). Springer Berlin Heidelberg, 269–283.
- [12] Li, X. and Miikkulainen, R. 2017. Evolving Adaptive Poker Players for Effective Opponent Exploitation, In AAAI-17 Workshop on Computer Poker and Imperfect Information Games. Technical Reports of the Thirty-first AAAI Conference of Artificial Intelligence (AAAI-17).
- [13] Li, X. and Miikkulainen, R. 2018. Opponent Modeling and Exploitation in Poker Using Evolved Recurrent Neural Networks. In Proceedings of The Genetic and Evolutionary Computation Conference (GECCO 2018). ACM, Kyoto, Japan.
- [14] Miikkulainen, R. 1996. Subsymbolic Case-Role Analysis of Sentences with Embedded Clauses. Cognitive Science 20 (1996), 47–73.
- [15] Narvekar, S. and Stone, P. 2018. Learning Curriculum Policies for Reinforcement Learning. (2018). arXiv:cs.LG/1812.00285
- [16] Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., and Wermter, S. 2018. Continual Lifelong Learning with Neural Networks: A Review. (2018). arXiv:cs.LG/1802.07569
- [17] Risi, S. and Togelius, J. 2019. Procedural Content Generation: From Automatically Generating Game Levels to Increasing Generality in Machine Learning. (2019). arXiv:cs.AI/1911.13071
- [18] Schmidhuber, J. 1987. Evolutionary principles in self-referential learning, or on learning how to learn: The meta-meta-... hook. Ph.D. Dissertation. Institut für Informatik, Technische Universität München.
- [19] Schmidhuber, J. 2011. POWERPLAY: Training an Increasingly General Problem Solver by Continually Searching for the Simplest Still Unsolvable Problem. (2011). arXiv:cs.AI/1112.5309
- [20] Schrum, J. and Miikkulainen, R. 2014. Evolving Multimodal Behavior with Modular Neural Networks in Ms. Pac-Man. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '14). Association for Computing Machinery, 325–332.
- [21] Stanley, K. O. and Miikkulainen, R. 2002. Evolving Neural Networks through Augmenting Topologies. Evol. Comput. 10, 2 (2002), 99–127.
- [22] Thrun, S. and Pratt, L. 1998. Learning to Learn: Introduction and Overview. Kluwer Academic Publishers, USA, 3–17.
   [23] Wang, R., Lehman, J., Clune, J., and Stanley, K. O. 2019. Paired Open-Ended
- [23] Wang, R., Lehman, J., Clune, J., and Stanley, K. O. 2019. Paired Open-Ended Trailblazer (POET): Endlessly Generating Increasingly Complex and Diverse Learning Environments and Their Solutions. (2019). arXiv:cs.NE/1901.01753
- [24] Wikipedia, 2020. Flappy Bird. https://en.wikipedia.org/wiki/Flappy\_Bird. (2020). Accessed 3-February-2020).