

Gradient-free policy architecture search and adaptation

Sayna Ebrahimi

Department of Mechanical Engineering
Department of Electrical Engineering and Computer Sciences
University of California Berkeley United States
sayna@eecs.berkeley.edu

Anna Rohrbach

Max Plank Institute for Informatics
Saarland Informatics Campus, Germany
arohrbach@mpi-inf.mpg.de

Trevor Darrell

Department of Electrical Engineering and Computer Sciences
University of California Berkeley United States
trevor@eecs.berkeley.edu

Abstract: We develop a method for policy architecture search and adaptation via gradient-free optimization which can learn to perform autonomous driving tasks. By learning from both demonstration and environmental reward we develop a model that can learn with relatively few early catastrophic failures. We first learn an architecture of appropriate complexity to perceive aspects of world state relevant to the expert demonstration, and then mitigate the effect of domain-shift during deployment by adapting a policy demonstrated in a source domain to rewards obtained in a target environment. We show that our approach allows safer learning than baseline methods, offering a reduced cumulative crash metric over the agent’s lifetime as it learns to drive in a realistic simulated environment.

Keywords: gradient-free optimization, architecture search, domain adaptation, autonomous driving

1 Introduction

Deep architectures have become popular as function approximators to represent action-selection policies. Common approaches to learn the parameters of such models include reinforcement learning [20] and/or learning from demonstration [1]: both learn model parameters to maximize expected reward, mimic human behavior, and/or achieve implicit goals. However, the design of policy architectures, especially in a deep learning paradigm, remains relatively unexplored. Architectures are typically selected through a combination of intuition and/or trial and error.

Learning to learn, including the learning of learning architectures, is long-articulated goal of AI, and many meta-learning and lifelong learning schemes have been proposed (e.g., [21] offered seminal views; see [18] for a survey). Recently, renewed interest in this topic has focused on models which explicitly search over the structure of deep architectures, including models which fuse non-parametric Bayesian inference with deep learning to select the number of channels for visual recognition tasks [6], models which use reinforcement learning to directly optimize over deep architectures for recognition [24], and models which use a gradient-free optimization method (evolutionary search) to infer optimal network structure [12].

We investigate policy architecture search using gradient-free optimization, and learn optimal policy structure for autonomous driving tasks. We propose a model which learns jointly from demonstration and optimization, with the goal of safe training: minimizing the amount of damage a vehicle incurs to learn a threshold level of performance. We base our approach on exploration-based schemes due to their ability to optimize model weights and architecture hyperparameters, leverage expert demonstrations, and adapt to reward obtained in new domains. We believe that a model which can initialize from demonstration, and learn an optimal policy from that foundation, is likely to achieve higher performance while maintaining the constraint of safe training, compared to models which must randomly search through action space during initial learning, or which learn from a reasonably safe demonstration but cannot further optimize performance based on environmental reward.

Prior approaches to combine demonstration with reward-based learning have had mixed successes [15, 9, 1] mainly due to the poor generalization of the policy learned on demonstrations. We posit that effective behavior cloning requires learning a visual agent architecture that has sufficient structure to perceive the state of the world deemed relevant to the expert providing the demonstration. This may or may not be the case with existing, off-the-shelf visual models. We thus think it is wise to optimize over architectures and parameters when performing expert behavioral cloning.

Often, deep models which learn to perform in one domain fail to perform well when deployed in another setting, such as differing weather or lighting conditions. Models learned from demonstration are also well known to fail when the learned policy takes the agent away from the region of the state space where demonstration was provided [13]. We show that our method can effectively and safely adapt a model demonstrated in one environment but deployed in a visually different environment based on the the reward signal in the latter domain, even when the agent is initialized far from initial demonstrations. Our approach leverages only target domain reward, and makes no assumptions about domain alignment, explicit or implicit, nor assumes any demonstration supervision in the target domain.

To achieve these goals, we present a gradient-free optimization algorithm inspired by [16] with a modification in noise generation that results in estimating the gradients more efficiently and accurately (Sec. 3.1). We then apply this algorithm to search over variable length architectures which are rewarded for low errors (or high accuracies) and penalized for excessive growth of the network (Sec. 3.2). Next, we combine our gradient-free policy search with demonstrations to learn a better policy that adapts to the new environment by receiving rewards as feedback (Sec. 3.3). We experimentally show that our architecture search model finds a policy on the GTA game environment that outperforms previously published methods (e.g., [3]) in end-to-end steering prediction from demonstrations, and that it can be efficiently adapted to learn to drive in previously unseen scenarios (Sec. 4). Our model reduces the number of crashes incurred while learning to drive, compared to baselines based only on reward or demonstration but not both, or compared to previously proposed fixed architectures that were not optimized for the domain.

2 Related work

Architecture search has been investigated through different frameworks including reinforcement learning [24, 2, 22] and evolutionary techniques [12]. In [24], a recurrent neural network (RNN) was used to generate fixed-length architecture descriptions from a predefined search space and trained it with policy gradient methods. They were able to surpass the state of the art results on CIFAR-10 and Penn Treebank datasets. A meta-modeling algorithm was proposed in [2] which used Q -learning to sequentially search for convolutional layers for image classification tasks. They showed that their approach outperforms other existing meta-models and manually-designed architectures with similar types of layers. Recently, [22] introduced Budgeted Super Networks which are inspired by the REINFORCE algorithm with an objective function that maximizes prediction quality and computation cost simultaneously. Various versions of biologically-inspired methods, or neuroevolution strategies, have been proposed for architecture search ever since they were introduced by [5]. Most of them are based on biological genetics algorithms where there is a *fitness* function that gets re-evaluated at each “generation” to determine whether “genotypes” are perturbed in the correct direction to evolve appropriately [12, 11, 19, 17]. I.e., they initialize a model and evolve it based on its performance. This paradigm was recently re-visited as an alternative to reinforcement learning algorithms where optimization is performed in a gradient-free fashion and the algorithm was shown to be highly parallelizable resulting in significant speedups in playing MuJoCo and Atari games [16].



Figure 1: (a) Sample images used in the architecture search for behavioral cloning task; (b) Sample images from a target domain that were not seen during the architecture search for behavioral cloning.

Policy search in autonomous driving application has been largely focused on demonstration-based optimization approaches with [4] or without [3, 10] affordance measurements. It dates back to the classic ALVIN model [10] which was a shallow architecture that could map from pixels to simple driving actions. Several years after, researchers demonstrated end-to-end deep learning models for steering control of small-scale cars [8], and recently NVIDIA followed the same path and showed success in predicting steering angle on a full-size vehicle from raw pixels using a convolutional network [3]. A novel FCN-LSTM architecture was proposed on large scale crowd-sourced data to perform egomotion predictions conditioned on the previous temporal states [23]. They used dashcam camera videos to derive a generic driving model that predicted trajectory angle (not steering angles).

3 Approach

We propose a learning-to-learn model which includes architecture optimization, parameter learning, and representation adaptation over different time scales. Our approach can be summarized with the following two steps. (1) Given expert demonstration, search over architectures and parameters to find a policy that best mimics performance by monitoring the obtained accuracy and number of parameters. (2) Having learned from demonstration, adapt the model to the reward provided by the target environment. In both steps, it is essential to derive a function approximator that optimizes an objective function. We use a gradient-free optimization algorithm [16] that maximizes a parametrized reward function using gradient estimation to perform architecture search (Sec. 3.2) and policy learning (Sec. 3.3).

3.1 Gradient-free optimization algorithm

Let $F(\cdot)$ be our objective function parametrized by θ which is an n -dimensional vector. F can be the reward that an environment provides for an agent when it executes a policy with parameters θ ; our goal is to maximize the expected reward by perturbing the policy parameters, denoted as $\hat{\theta}$, by moving in particular directions. The parameter estimate update can be performed using a general stochastic form:

$$\hat{\theta}_{t+1} = \hat{\theta}_t + \alpha_t \nabla_{\hat{\theta}} y(\hat{\theta}) \quad (1)$$

where $y(\cdot)$ is an approximation of the objective function (i.e. $y(\cdot) = F(\cdot) + \text{noise}$) and $\nabla_{\hat{\theta}} y(\hat{\theta})$ is the gradient of objective estimate that can be approximated by any gradient estimator in the family of finite difference methods. The gradient is estimated in a randomly chosen direction by perturbing all the elements of $\hat{\theta}_t$ to obtain two measurements of $y(\cdot)$ as follows:

$$y_t^{(+)} = F(\hat{\theta}_t + c_t \Delta_t) + \epsilon_t^{(+)} \quad (2)$$

$$y_t^{(-)} = F(\hat{\theta}_t - c_t \Delta_t) + \epsilon_t^{(-)} \quad (3)$$

where $\Delta_t \in R^n$ is a vector of n mutually independent randomly perturbed variables taken from a zero-mean distribution. While there is no restriction for it to have a specific type of distribution, we

Table 1: Experimental search space defined for each layer type.

	Search space
Convolutional layer followed by Max-pool and/or Dropout	Filter height (FH) $\in [1, 3, 5, 7]$
	Filter width (FW) $\in [1, 3, 5, 7]$
	Stride height (SH) $\in [1, 2, 3]$
	Stride width (SW) $\in [1, 2, 3]$
	Number of filters (NF) $\in [16, 24, 32, 64, 128, 256]$
	Max-pool size (MP) $\in [1, 2, 3]$
	Dropout (DO1) $\in [0.3, 0.5, 0.7, 1.0]$
Fully-connected layer followed by Dropout	Number of units (NU) $\in [8, 16, 32, 64, 128, 256, 512]$
	Dropout (DO2) $\in [0.3, 0.5, 0.7, 1.0]$

use Laplace distribution, as it tends to choose orthogonal directions in the long run. Other recent efforts [24] utilized Gaussian noise to sample mirrored projections. Figure 2 shows a comparison between the two distributions. c_t is a small positive number and $\epsilon_t^{(+)}$ and $\epsilon_t^{(-)}$ are the noise associated with evaluating $F(\cdot)$ such that: $E(\epsilon_t^{(+)} - \epsilon_t^{(-)} | \{\theta_1, \theta_2, \dots, \theta_t\}, \Delta_t) = 0$. The gradient estimate can then be computed as:

$$\nabla_{\hat{\theta}_t} y(\hat{\theta}_t) = \left[\frac{y_t^{(+)} - y_t^{(-)}}{2 c_t \Delta_{t1}} \quad \dots \quad \frac{y_t^{(+)} - y_t^{(-)}}{2 c_t \Delta_{tn}} \right]^T \quad (4)$$

Parameter estimates can be updated by replacing the gradients in Eq. 1 with those found in Eq. 4.

3.2 Learning an optimal initial policy from demonstrations

Inspired by [24] we have used a recurrent neural network to sequentially generate the description of layers of an architecture from a given design space defined by the user. The RNN acts as a controller which generates the architecture description defined by its hyper-parameters chosen from a pre-defined search space. In [24], the authors used policy gradients to train the RNN which was able to produce fixed-length convolutional and recurrent architectures. Given demonstrations and having a *child* network defined by the RNN, they trained the child network using supervised learning and obtained an accuracy metric on the given task on a held-out validation set and used that accuracy as a reward signal to train the RNN.

Our model uses the demonstrations to provide the reward function, ((i.e., F above), to train the RNN. Unlike backpropagation which suffers from gradient vanishing while training RNNs, gradient-free algorithms do not have such an issue [16]. Our RNN controller specifies three types of layers: convolutional, fully connected, and max-pool which can have inter-layer dropouts. For the reward signal, we use the negative value of *total* loss function. At the last layer of the network we regress to three real-valued numbers, each having a mean-squared loss. The total loss is the sum of all three losses. We use a novel reward function (Eq. 5 that not only results in the minimum total loss, but also grows the architecture as long as the loss keeps decreasing. Note that in case of having a classification problem, an accuracy metric can replace loss value, hence the goal will be maximizing the accuracy while controlling number of parameters. Here we have a regression problem and our goal is to search for an architecture that is guaranteed to achieve a low loss on the given task and grow further with adding more layers (i.e. producing more parameters) to decrease this while being penalized for adding more parameters in turns of no gain in loss reduction. We propose to use a ReLU-based Lagrange-multiplier reward function as below:

$$R = R_1 - \lambda(R_1) R_2 \quad (5)$$

where R_1 is the negative of the minimum loss (or maximum accuracy in a classification problem) on the validation set for the last 5 epochs and R_2 is the total number of parameters in the child network. λ is the Lagrange parameter defined as a function of the first sub-reward in a ReLU-based fashion:

$$\lambda(R_1) = \begin{cases} 0 & \text{if } R_1 < A \\ \mu(R_1 - A) & \text{if } R_1 \geq A \end{cases} \quad (6)$$

This reward function acts as follows. The RNN keeps generating new layers by being rewarded only based on the obtained total loss, until it produces a child network that achieves a desired value of loss A (or accuracy A) on the validation set. Once it reaches this threshold, it will be penalized for further growing the architecture if loss does not decrease consequently. The parameter μ in Eq. 5 defines the respective threshold. E.g. choosing $\mu = 0.01/x$ allows architecture growth by x

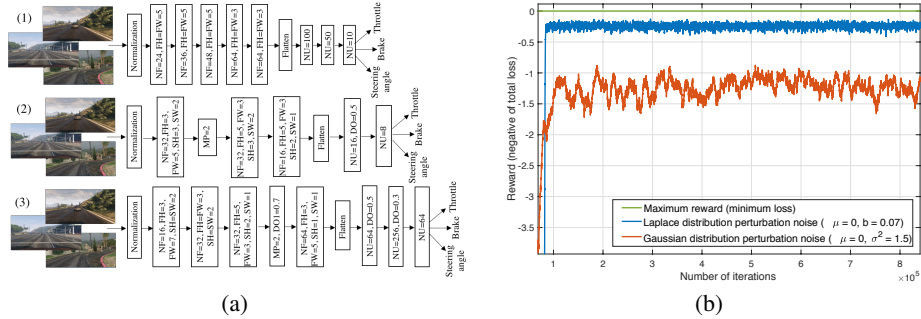


Figure 2: (a) Illustration of the baseline and learned architectures: (1) prior work [3], (2) our small network, (3) our large network; (b) Comparison of random Gaussian vs. random Laplace distribution in terms of reward vs. number of iterations.

number of parameters in the new layer if it causes the overall loss decrease (or accuracy increase in classification) by 1%. The thresholds are adjustable based on the problem at hand and the desired trade-off between computational cost and loss minimization.

Our RNN controller is a three-layer LSTM network followed by a softmax layer. The inputs to the RNN are the hyper-parameters that describe a layer (see Tab. 1 for our search space.). Training the RNN starts with randomly initializing the hyper-parameters of the child network which initially has only one layer. RNN uses the reward function to update its parameter weights such that those which contributed more in the obtained reward, receive a higher weighting factor during the update and hence, we move in a hill-climbing direction which eventually maximizes the reward. We use the algorithm described in Sec. 3.1 to generate an architecture that yields the minimum loss. The process of generating a new layer terminates when we achieve convergence in the received reward.

3.3 Adapting a demonstrated policy to a new driving domain

As confirmed experimentally below, it is well known that a policy learned from behavioral cloning can perform poorly when evaluated on inputs with a domain shift relative to the demonstration supervision. To overcome this, we further use the gradient-free search algorithm described in Sec. 3.1 to adapt a driving policy learned from demonstration in a source domain based on rewards in a target domain. We experiment with the setting where the initial agent state (e.g., location), and/or weather and lighting conditions, are substantially different than provided as demonstration. We compare to baselines where we perform reward-based optimization using initial demonstrations instead of a randomly initialized policy, which makes our reward function converge faster, and more safely.

In the driving scenario, we wish to learn to drive with optimal or near-optimal performance, defined by the reward in the target domain. Specifically, the reward function used in our experiment is composed of two factors (we receive +1 if obeyed and 0 if violated): 1) No crashes with other objects 2) Staying within the lane lines if they are available in the driving scene. We have used the lane reward function and accident detection function defined in [14] source code. The necessary information is provided by a *paths.xml* file in [14].

In our model, an *episode* is the time interval that the agent has successfully driven without having a car crash. Note that not all deviations from the middle of the road necessarily result in an accident. In case of a minor deviation, while the car receives 0 as its reward, it continues driving until it makes a mistake that causes it to crash and the game restarts. There are distinct thresholds for middle-lane deviation defined in [14] for different roads (highway, urban, etc.) and different vehicle types.

4 Experimental evaluation

We implemented the method described above and ran comprehensive experiments to show the efficiency and applicability of our approach in searching for an optimal driving policy that has the minimum number of catastrophic failures. Details of the experiments along with the results are provided in the following subsections. All the experiments are executed in the GTA game environment

using a publicly available plugin [14] that allowed us to have control over driving conditions such as lighting, weather, car model, and reward function.

4.1 Dataset

For our evaluation we have collected a dataset of an expert policy by playing GTA collecting 2,267,662 images of size 66×200 similar to [3]. Labels include steering angle, brake, and throttle values. In order to learn from diverse driving scenarios, we have used data from different locations (highways, rural roads, urban streets) where weather and lighting conditions were adjusted using [14]. Our goal was to expose the learning algorithm to a comprehensive demonstration set yet to set aside some specific scenes for further testing the performance of behavioral cloning task. Sample images from demonstration are shown in Fig. 1(a). They include rainy (daytime), overcast (day and night), foggy (daytime), sunny scenes and thunderstorms (daytime). Some particular scenes such as rain and thunderstorm during nighttime as well as snow at daytime have been kept for our test set (see Fig. 1(b)). Our test set is composed of 100,000 images.

4.2 Learning a policy architecture from GTA demonstrations

The search space for the hyperparameters that describe a fully-convolutional architecture is presented in Tab. 1. The activation function is fixed to be a rectified linear unit. The RNN-controller has three LSTM layers, with 10 hidden units in each, and a softmax layer at the end to choose from the given search space. The RNN weights are initialized with a random Laplace distribution ($\mu = 0, b = 0.07$). Once the RNN predicts a new layer’s description, the child network is built and trained with batch size of 64 and Adam optimizer [7] with learning rate 10^{-3} . We train the child network for different number of epochs starting from 20 epochs (depending on which layer we are at) and compute the reward function as described in Sec. 3.2. In order to finish optimizing one layer we track the loss reduction of both validation and training sets between the first and the last epoch to avoid overfitting.

Our model is capable of generating architectures at low or high costs of architecture growth. In order to compare our designed architecture with [3],¹ we control the reward function such that it never produces an architecture with more than 252,241 parameters while minimizing the loss (maximizing the reward). We present our architectures and comparison to prior work in Fig. 2(a)(2). The corresponding performance comparison is shown in Tab. 2. The smallest architecture (in terms of number of parameters) that we have built is shown in Fig. 2(a), which obtains a smaller total loss than [3]. The network of [3] (Fig. 2(a)(1)) appears to suffer from overfitting which might be explained by the absence of Dropout or a pooling layer. Not restricting our architecture search algorithm to be bounded by a number of parameters, we learn a larger network (2(a)(3)) that has over $2M$ parameters and obtains a minimum total loss of 0.085 on the training set and 0.088 on the validation set. As discussed above, we test all models on substantially different driving scenes which were never seen during training. On this challenging test set our large network obtains the minimum loss of 0.195. We use this network as our initial driving policy in the next section and improve it further in the reward-providing GTA environment.

Table 2: Total MSE obtained using architecture proposed in prior work [3] and our models obtained by architecture search on demonstrations.

Model	# of parameters	Training loss	Validation loss	Test loss
Bojarski et al. [3]	252,241	0.098	0.11	0.212
Our small network	228,227	0.093	0.096	0.197
Our large network	2,198,723	0.085	0.088	0.185

We further empirically compare two different noise distributions for perturbing the parameters of the network: random Gaussian and random Laplace. We perform the architecture search over our

¹As no reference implementation of [3] is openly available we had to use our own, which may be suboptimal w.r.t. the authors’ as we did not have full access to their model parameters. Also their model was only used to predict steering angle, and overall it is not clear whether their goal was to maximize performance, find a model with relatively few parameters, or both, so they may not have explored the full design space with their model. Nonetheless, [3] was closest model in the literature for end-to-end steering angle prediction and thus the best available baseline.

Table 3: Comparison of two policies (our large network and [3]) learned based on target domain reward, with and without source-domain demonstrations.

Model	Wall-clock convergence	Total # of car crashes	Total # of middle-lane keeping violations
(1) Bojarski et al. [3], w.o demo	154 hours	15,565	18,662
(2) Bojarski et al. [3], w. demo	74 hours	1,387	3,243
(3) Our large network, w.o. demo	114 hours	6,877	8781
(4) Our large network, w. demo	53 hours	832	982

small model using both distributions with mean zero; variance is chosen using grid search. Fig. 2(b) shows the results for reward convergence versus number of iterations. Both distributions result in convergence to high reward values (minimizing loss), however, the Laplace distribution tends to be less noisy and reaches slightly higher reward values.

4.3 Safe policy adaptation

Next, we want to learn the driving policy in a target game domain. We start with an initial model, either using a behaviorally cloned or randomly initialized policy and gradually improve it by receiving rewards from the environment. As stated in 3.3, an episode of the game starts at a random location and weather condition in the game. To initialize the policy, we use the larger architecture learned in the first step (with the model of [3] as the baseline). We evaluate both models with and without being adapted to demonstration forming four cases: (1) the baseline network of without demonstration (i.e., with randomly initialized weights) and (2) with behaviorally cloned initial weights, (3) our larger architecture without demonstration and (4) with behaviorally cloned initial weights. We run all models in the GTA environment to receive the reward described in Sec. 3.3. Once the reward is received, the weights are perturbed by a Laplace random noise ($\mu = 0, b = 0.07$) and the same procedure is repeated until the average reward in each episode of the game converges to its maximum value. Results averaged across several runs are presented in Tab. 3 where our model optimized with demonstration outperforms all other cases. In particular, our model has the least number of cumulative crash occurrence prior to converging to 100% of averaged reward (details below).

In Table 4 we have listed the results for test accuracies on a dataset taken from a target domain that has not been seen in the demonstrations. On the left we see behavioral cloning alone has poor performance when significant domain shifts occur. Our adapted model has performance over loss minimization of 10^{-5} . On the right we see that adapted performance is strong even without reward in the target domain, indicating that visual domain shift is a lesser issue than being off-demonstration; our model can adapt in the source domain and still be accurate on the target. Best performance is obtained with adaptation to reward in both source and target. This also shows that there is a improvement on loss minimization when we learn from rewards. It is worth noting that we can not judge the driving behavior only by looking at the total MSE loss as it is not a comprehensive representative of the driving task. Each one of the angle, brake, and throttle converges to a separate MSE loss among which steering angle has the least and brake has the largest loss values. This shows that learning steering angle is easier with demonstrations compared to brake and throttle which at each time step depend on multi previous frames. Table 5 shows all model predictions for a relatively complex image chosen from the target domain (Fig. 1 at top corner) where a pedestrian crossing the street when the signal light is green. The behaviorally cloned models tend to predict that the agent should keep going whereas the adapted models to the target domain with rewards are able to predict the correct decision despite the green light presence in the image. Our adapted large network rewarded on both source and target is able to make the best prediction for throttle and brake (steering angle is perfect across all models).

Fig. 3 illustrates the percentage of averaged reward per episode for the aforementioned four models until convergence. Our designed architecture which is adapted to the demonstrations on a source domain, starts with less than 10% reward in its first episode which lasts for 37 seconds. This is reasonable considering the fact that each episode of the game, is intentionally set up to start in a completely random new environment which is highly possible to be a significantly different domain that what the policy has seen up to that point. This again highlights the fact that a behaviorally cloned model is at high risks of failure when it is tested in a different domain. The models keep learning from the rewards until convergence. It can be seen in Fig. 3(a) that our designed architecture adapted

Table 4: Comparison of performance (average total loss) of two policies (our large network and [3]) at test time on target domain (T) when they are trained with rewards from target domain, source domain (S), and both (T+S).

	Behavioral cloning	Demo on S Rew. on T+S Test on T		Behavioral cloning	Demo on S Rew. on S Test on T	Demo on S Rew. on S+T Test on T	Demo on S Rew. on T Test on T
Bojarski et. al.[3]	0.212	10^{-4}					
Our large network	0.185	10^{-5}	Our large network	0.185	7×10^{-5}	10^{-5}	3×10^{-5}




Figure 3: Averaged reward per episode vs. number of episodes for our large network and [3] using (a) policy adapted to demonstrations; (b) a non-adapted policy. (Note different x-axis scales.)

to the demonstration reaches 100% of averaged reward after 53 hours in its last episode (episode number #832) which lasts for 90 minutes and is then terminated by user (no crashing happens). It is also shown that a suboptimal, yet adapted to the demonstrations policy [3], also converges but only to 97.3% of the maximum reward and then plateaus for more than 12 hours. Unadapted policies are also shown in Fig. 3(b) converging to an averaged reward of 87% in a drastically different time-scale confirming the positive effect of using demonstrations in policy learning. Supplementary video of results can be found in <https://saynaebrahimi.github.io/corl.html>

Table 5: Predictions for a complex driving scene shown in Figure 1 (top right corner) using all our learned models

		Behavioral cloning	Demo on S Rew on S Test on T	Demo on S Rew on S+T Test on T	Demo on S Rew on T Test on T
Large network	Steering angle	-0.006	0.003	0.005	0.002
	Brake	0.191	0.889	0.931	0.956
	Throttle	0.665	0.083	0.010	0.052
Bojarski et. al. [3]	Steering angle	-0.005	-0.002	0.002	-0.001
	Brake	0.183	0.567	0.677	0.778
	Throttle	0.775	0.223	0.121	0.156

5 Conclusion

The goal of this work is to learn a policy for an autonomous driving task minimizing crashes and other safety violations while training. To this end we propose an algorithm which learns to generate an optimal network architecture from demonstration using a new reward function that optimizes accuracy and model size simultaneously. We confirm behavioral cloning alone can perform poorly when the target domain differs from source demonstrations. We show that our method can adapt the model learned by demonstration to a new domain relying on target environmental rewards. Experimental evaluation shows that our model achieves higher accuracy, fewer cumulative crashes, and higher target domain reward. We believe these results are encouraging and important steps towards the ultimate goal of learning complex driving policies with zero cumulative crashes or serious accidents either in simulation or the real world.

References

- [1] B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- [2] B. Baker, O. Gupta, N. Naik, and R. Raskar. Designing neural network architectures using reinforcement learning. *arXiv preprint arXiv:1611.02167*, 2016.
- [3] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [4] C. Chen, A. Seff, A. Kornhauer, and J. Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2722–2730, 2015.
- [5] M. Eigen. *Ingo Rechenberg Evolutionsstrategie Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. mit einem Nachwort von Manfred Eigen, Friedrich Frommann Verlag, Struttgart-Bad Cannstatt, 1973.
- [6] J. Feng and T. Darrell. Learning the structure of deep convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2749–2757, 2015.
- [7] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [8] Y. LeCun, U. Muller, J. Ben, E. Cosatto, and B. Flepp. Off-road obstacle avoidance through end-to-end learning.
- [9] M. N. Nicolescu and M. J. Mataric. Learning and interacting in human-robot domains. *IEEE Transactions on Systems, man, and Cybernetics-part A: Systems and Humans*, 31(5):419–430, 2001.
- [10] D. A. Pomerleau. Alvin, an autonomous land vehicle in a neural network. Technical report, Carnegie Mellon University, Computer Science Department, 1989.
- [11] J. K. Pugh and K. O. Stanley. Evolving multimodal controllers with hyperneat. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 735–742. ACM, 2013.
- [12] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, Q. Le, and A. Kurakin. Large-scale evolution of image classifiers. *arXiv preprint arXiv:1703.01041*, 2017.
- [13] S. Ross, G. J. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics*, pages 627–635, 2011.
- [14] A. Ruano. Deepgtav. <https://github.com/ai-tor/DeepGTAV>, 2017.
- [15] P. E. Rybski and R. M. Voyles. Interactive task training of a mobile robot through human gesture recognition. In *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, volume 1, pages 664–669. IEEE, 1999.
- [16] T. Salimans, J. Ho, X. Chen, and I. Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.
- [17] J. D. Schaffer, D. Whitley, and L. J. Eshelman. Combinations of genetic algorithms and neural networks: A survey of the state of the art. In *Combinations of Genetic Algorithms and Neural Networks, 1992., COGANN-92. International Workshop on*, pages 1–37. IEEE, 1992.
- [18] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- [19] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.

- [20] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [21] S. Thrun and L. Pratt. *Learning to learn*. Springer Science & Business Media, 2012.
- [22] T. Veniat and L. Denoyer. Learning time-efficient deep architectures with budgeted super networks. *arXiv preprint arXiv:1706.00046*, 2017.
- [23] H. Xu, Y. Gao, F. Yu, and T. Darrell. End-to-end learning of driving models from large-scale video datasets. *arXiv preprint arXiv:1612.01079*, 2016.
- [24] B. Zoph and Q. V. Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.