

Query-Efficient Imitation Learning for End-to-End Simulated Driving

Jiakai Zhang

Department of Computer Science
New York University
zhjk@nyu.edu

Kyunghyun Cho

Department of Computer Science
Center for Data Science
New York University
kyunghyun.cho@nyu.edu

Abstract

One way to approach end-to-end autonomous driving is to learn a policy that maps from a sensory input, such as an image frame from a front-facing camera, to a driving action, by imitating an expert driver, or a reference policy. This can be done by supervised learning, where a policy is tuned to minimize the difference between the predicted and ground-truth actions. A policy trained in this way however is known to suffer from unexpected behaviours due to the mismatch between the states reachable by the reference policy and trained policy. More advanced algorithms for imitation learning, such as DAgger, addresses this issue by iteratively collecting training examples from both reference and trained policies. These algorithms often require a large number of queries to a reference policy, which is undesirable as the reference policy is often expensive. In this paper, we propose an extension of the DAgger, called SafeDAgger, that is query-efficient and more suitable for end-to-end autonomous driving. We evaluate the proposed SafeDAgger in a car racing simulator and show that it indeed requires less queries to a reference policy. We observe a significant speed up in convergence, which we conjecture to be due to the effect of automated curriculum learning.

Introduction

We define end-to-end autonomous driving as driving by a single, self-contained system that maps from a sensory input, such as a front-facing camera, to actions necessary for driving, such as the angle of steering wheel and braking. In this approach, the autonomous driving system is often *learned* from data rather than manually designed, mainly due to sheer complexity of manually developing a such system.

This end-to-end approach to autonomous driving dates back to late 80's. ALVINN by Pomerleau (Pomerleau 1989) was a neural network with a single hidden layer that takes as inputs a frame from a front-facing camera and a response map from a range finder sensor and returns a quantized steering wheel angle. The ALVINN was trained using a set of training tuples (image, sensor map, steering angle) collected from simulation. A similar approach was taken later in 2005 to train, this time, a convolutional neural network to drive an off-road mobile robot (Muller et al. 2005). More recently, Bojarski et al. (Bojarski et al. 2016) used a similar, but deeper,

convolutional neural network for lane following based solely on a front-facing camera. In all these cases, a deep neural network has been found to be surprisingly effective at learning a complex mapping from a raw image to control.

A major learning paradigm behind all these previous attempts has been supervised learning. A human driver or a rule-based AI driver in a simulator, to which we refer as a reference policy drives a car equipped with a front-facing camera and other types of sensors while collecting image-action pairs. These collected pairs are used as training examples to train a neural network controller, called a primary policy. It is however well known that a purely supervised learning based approach to imitation learning (where a learner tries to imitate a human driver) is sub-optimal (see, e.g., (Daumé Iii, Langford, and Marcu 2009; Ross, Gordon, and Bagnell 2010) and references therein.)

We therefore investigate a more advanced approach to imitation learning for training a neural network controller for autonomous driving. More specifically, we focus on DAgger (Ross, Gordon, and Bagnell 2010) which works in a setting where the reward is given only implicitly. DAgger improves upon supervised learning by letting a primary policy collect training examples while running a reference policy simultaneously. This dramatically improves the performance of a primary policy. However, DAgger needs to constantly query a reference policy, which is expensive especially when a reference policy may be a human driver. It also uses an exploration policy that may not be safe to itself or its environment. To reduce the cost of querying the reference policy, the active imitation learning approach (Judah et al. 2014) allows to choose most useful states where advice from the teacher can be useful. Kim and Pineau (Kim and Pineau 2013) tried to address the DAgger problems by proposing to learn multiple policies, each specializing in a particular region of state space. They used a metric based on state distance to decide if the reference policy needs to be queried. Instead of learning multiple policies, Laskey et al. (Laskey et al. 2016) recently introduced the SHIV approach (Svm-based reduction in Human InterVention) which converged to a single policy.

In this paper, we propose a query-efficient extension of the DAgger, called *SafeDAgger*. We first introduce a safety classifier that learns to predict the error made by a primary policy *without* querying a reference policy. This safety classifier is incorporated into the DAgger's iterations in order to select

only a small subset of training examples that are collected by a primary policy. This subset selection significantly reduces the number of queries to a reference policy.

We empirically evaluate the proposed SafeDagger using TORCS (tor accessed May 12 2016), a racing car simulator, which has been used for vision-based autonomous driving research in recent years (Koutník et al. 2013; Chen et al. 2015). In this paper, our goal is to learn a primary policy that can drive a car indefinitely without any crash or going out of a road. The experiments show that the SafeDagger requires much fewer queries to a reference policy than the original DAgger does and achieves a superior performance in terms of the average number of laps without crash and the amount of damage. We conjecture that this is due to the effect of automated curriculum learning created by the subset selection based on the safety classifier.

Imitation Learning for Autonomous Driving

In this section, we describe imitation learning in the context of learning an automatic policy for driving a car.

State Transition and Reward

A surrounding environment, or a world, is defined as a set of states S . Each state is accompanied by a set of possible actions $A(S)$. Any given state $s \in S$ transitions to another state $s' \in S$ when an action $a \in A(S)$ is performed, according to a state transition function $\delta : S \times A(S) \rightarrow S$. This transition function may be either deterministic or stochastic.

For each sequence of state-action pairs, there is an associated (accumulated) reward r :

$$r(\Omega = ((s_0, a_0), (s_1, a_1), (s_2, a_2), \dots)),$$

where $s_t = \delta(s_{t-1}, a_{t-1})$.

A reward may be implicit in the sense that the reward comes as a form of a binary value with 0 corresponding to any unsuccessful run (e.g., crashing into another car so that the car breaks down,) while any successful run (e.g., driving indefinitely without crashing) does not receive the reward. This is the case in which we are interested in this paper. In learning to drive, the reward is simply defined as follows:

$$r(\Omega) = \begin{cases} 1, & \text{if there was no crash,} \\ 0, & \text{otherwise} \end{cases}$$

This reward is implicit, because it is observed only when there is a failure, and no reward is observed with an optimal policy (which never crashes and drives indefinitely.)

Policies

A policy is a function that maps from a state observation $\phi(s)$ to one a of the actions available $A(s)$ at the state s . An underlying state s describes the surrounding environment perfectly, while a policy often has only a limited access to the state via its observation $\phi(s)$. In the context of end-to-end autonomous driving, s summarizes all necessary information about the road (e.g., # of lanes, existence of other cars or pedestrians, etc.) while $\phi(s)$ is, for instance, an image frame taken by a front-facing camera.

We have two separate policies. First, a *primary policy* π is a policy that learns to drive a car. This policy does not observe a full, underlying state s but only has access to the state observation $\phi(s)$, which is in this paper a pixel-level image frame from a front-facing camera. The primary policy is implemented as a function parametrized by a set of parameters θ .

The second one is a *reference policy* π^* . This policy may or may not be optimal, but is assumed to be a good policy which we want the primary policy to *imitate*. In the context of autonomous driving, a reference policy can be a human driver. We use a rule-based controller, which has access to a true, underlying state in a driving simulator, as a reference policy in this paper.

Cost of a Policy Unlike previous works on imitation learning (see, e.g., (Daumé Iii, Langford, and Marcu 2009; Ross, Gordon, and Bagnell 2010; Chang et al. 2015)), we introduce a concept of cost to a policy. The cost of querying a policy for an appropriate action varies significantly based on how the policy is implemented. For instance, it is expensive to query a reference policy, if it is a human driver. On the other hand, it is much cheaper to query a primary policy which is often implemented as a classifier. Therefore, in this paper, we analyze an imitation learning algorithm in terms of *how many queries it makes to a reference policy*.

Driving

A car is driven by querying a policy for an action with a state observation $\phi(s)$ at each time step. The policy, in this paper, observes an image frame from a front-facing camera and returns both the angle of a steering wheel ($u \in [-1, 1]$) and a binary indicator for braking ($b \in \{0, 1\}$). We call this strategy of relying on a single fixed policy a *naive strategy*.

Reachable States With a set of initial state $S_0^\pi \subset S$, each policy π defines a subset of the reachable states S^π . That is, $S^\pi = \cup_{t=1}^\infty S_t^\pi$, where $S_t^\pi = \{s | s = \delta(s', \pi(\phi(s'))), \forall s' \in S_{t-1}^\pi\}$. In other words, a car driven by a policy π will only visit the states in S^π .

We use S^* to be a reachable set by the reference policy. In the case of learning to drive, this reference set is intuitively smaller than that by any other reasonable, non-reference policy. This happens, as the reference policy avoids any state that is likely to lead to a low reward which corresponds to crashing into other cars and road blocks or driving out of the road.

Supervised Learning

Imitation learning aims at finding a primary policy π that *imitates* a reference policy π^* . The most obvious approach to doing so is supervised learning. In supervised learning, a car is first driven by a reference policy while collecting the state observations $\phi(s)$ of the visited states, resulting in $D = \{\phi(s)_1, \phi(s)_2, \dots, \phi(s)_N\}$. Based on this dataset, we define a loss function as

$$l_{\text{supervised}}(\pi, \pi^*, D) = \frac{1}{N} \sum_{n=1}^N \|\pi(\phi(s)_n) - \pi^*(\phi(s)_n)\|^2. \quad (1)$$

Then, a desired primary policy can be written as $\hat{\pi} = \arg \min_{\pi} l_{\text{supervised}}(\pi, \pi^*, D)$.

A major issue of this supervised learning approach to imitation learning stems from the imperfection of the primary policy $\hat{\pi}$ even after training. This imperfection likely leads the primary policy to a state s which is not included in the reachable set S^* of the reference policy, i.e., $s \notin S^*$. As this state cannot have been included in the training set $D \subseteq S^*$, the behaviour of the primary policy becomes unpredictable. The imperfection arises from many possible factors, including sub-optimal loss minimization, biased primary policy, stochastic state transition and partial observability.

DAGger: beyond Supervised Learning

A major characteristics of the supervised learning approach described above is that it is only the reference policy π^* that generates training examples. This has a direct consequence that the training set is almost a subset of the reference reachable set S^* . The issue with supervised learning can however be addressed by imitation learning or learning-to-search (Daumé Iii, Langford, and Marcu 2009; Ross, Gordon, and Bagnell 2010).

In the framework of imitation learning, the primary policy, which is currently being estimated, is also used in addition to the reference policy when generating training examples. The overall training set used to tune the primary policy then consists of both the states reachable by the reference policy as well as the intermediate primary policies. This makes it possible for the primary policy to correct its path toward a good state, when it visits a state unreachable by the reference policy, i.e., $s \in S^{\pi} \setminus S^*$.

DAGger is one such imitation learning algorithm proposed in (Ross, Gordon, and Bagnell 2010). This algorithm fine-tunes a primary policy trained initially with the supervised learning approach described earlier. Let D_0 and π_0 be the supervised training set (generated by a reference policy) and the initial primary policy trained in a supervised manner. Then, DAGger iteratively performs the following steps. At each iteration i , first, additional training examples are generated by a mixture of the reference π^* and primary π_{i-1} policies (i.e.,

$$\beta_i \pi^* + (1 - \beta_i) \pi_{i-1} \quad (2)$$

) and combined with all the previous training sets: $D_i = D_{i-1} \cup \{\phi(s)_1^i, \dots, \phi(s)_N^i\}$. The primary policy is then fine-tuned, or trained from scratch, by minimizing $l_{\text{supervised}}(\theta, D_i)$ (see Eq. (1).) This iteration continues until the supervised cost on a validation set stops improving.

DAGger does not rely on the availability of explicit reward. This makes it suitable for the purpose in this paper, where the goal is to build an end-to-end autonomous driving model that drives on a road indefinitely. However, it is certainly possible to incorporate an explicit reward with other imitation learning algorithms, such as SEARN (Daumé Iii, Langford, and Marcu 2009), AggreVaTe (Ross and Bagnell 2014) and LOLS (Chang et al. 2015). Although we focus on DAGger in this paper, our proposal later on applies generally to any learning-to-search type of imitation learning algorithms.

Cost of DAGger At each iteration, DAGger queries the reference policy for each and every collected state. In other words, the cost of DAGger C_i^{DAGger} at the i -th iteration is equivalent to the number of training examples collected, i.e., $C_i^{\text{DAGger}} = |D_i|$. In all, the cost of DAGger for learning a primary policy is $C^{\text{DAGger}} = \sum_{i=1}^M |D_i|$, excluding the initial supervised learning stage.

This high cost of DAGger comes with a more practical issue, when a reference policy is a human operator, or in our case a human driver. First, as noted in (Ross et al. 2013), a human operator cannot drive well without actual feedback, which is the case of DAGger as the primary policy drives most of the time. This leads to suboptimal labelling of the collected training examples. Furthermore, this constant operation easily exhausts a human operator, making it difficult to scale the algorithm toward more iterations.

SafeDAGger: Query-Efficient Imitation Learning with a Safety Classifier

We propose an extension of DAGger that minimizes the number of queries to a reference policy both during training and testing. In this section, we describe this extension, called *SafeDAGger*, in detail.

Safety Classifier

Unlike previous approaches to imitation learning, often as learning-to-search (Daumé Iii, Langford, and Marcu 2009; Ross, Gordon, and Bagnell 2010; Chang et al. 2015), we introduce a classifier c_{safe} , to which we refer as a *safety classifier*. This classifier takes as input both the partial observation of a state $\phi(s)$ and a primary policy π and returns a binary label indicating whether the primary policy π is likely to deviate from a reference policy π^* *without* querying it.

We define the deviation of a primary policy π from a reference policy π^* as

$$\epsilon(\pi, \pi^*, \phi(s)) = \|\pi(\phi(s)) - \pi^*(\phi(s))\|^2.$$

Note that the choice of error metric can be flexibly chosen depending on a target task. For instance, in this paper, we simply use the L_2 distance between a reference steering angle and a predicted steering angle, ignoring the brake indicator.

Then, with this defined deviation, the optimal safety classifier c_{safe}^* is defined as

$$c_{\text{safe}}^*(\pi, \phi(s)) = \begin{cases} 0, & \text{if } \epsilon(\pi, \pi^*, \phi(s)) > \tau \\ 1, & \text{otherwise} \end{cases}, \quad (3)$$

where τ is a predefined threshold. The safety classifier decides whether the choice made by the policy π at the current state can be trusted with respect to the reference policy. We emphasize again that this determination is done without querying the reference policy.

Learning A safety classifier is *not* given, meaning that it needs to be estimated during learning. A safety classifier c_{safe} can be learned by collecting another set of training examples:¹

¹It is certainly possible to simply set aside a subset of the original training set for this purpose.

$D' = \{\phi(s)'_1, \phi(s)'_2, \dots, \phi(s)'_N\}$. We define and minimize a binary cross-entropy loss:

$$l_{\text{safe}}(c_{\text{safe}}, \pi, \pi^*, D') = -\frac{1}{N} \sum_{n=1}^N c_{\text{safe}}^*(\phi(s)'_n) \log c_{\text{safe}}(\phi(s)'_n, \pi) + (1 - c_{\text{safe}}^*(\phi(s)'_n)) \log(1 - c_{\text{safe}}(\phi(s)'_n, \pi)), \quad (4)$$

where we model the safety classifier as returning a Bernoulli distribution over $\{0, 1\}$.

Driving: Safe Strategy Unlike the naive strategy, which is a default go-to strategy in most cases of reinforcement learning or imitation learning, we can design a *safe strategy* by utilizing the proposed safety classifier c_{safe} . In this strategy, at each point in time, the safety classifier determines whether it is safe to let the primary policy drive. If so (i.e., $c_{\text{safe}}(\pi, \phi(s)) = 1$), we use the action returned by the primary policy (i.e., $\pi(\phi(s))$.) If not (i.e., $c_{\text{safe}}(\pi, \phi(s)) = 0$), we let the reference policy drive instead (i.e., $\pi^*(\phi(s))$.)

Assuming the availability of a good safety classifier, this strategy avoids any dangerous situation arisen by an imperfect primary policy, that may lead to a low reward (e.g., breakdown by a crash.) In the context of learning to drive, this safe strategy can be thought of as letting a human driver take over the control based on an automated decision.² Note that this driving strategy is applicable regardless of a learning algorithm used to train a primary policy.

Discussion The proposed use of safety classifier has a potential to address this issue up to a certain point. First, since a separate training set is used to train the safety classifier, it is more robust to unseen states than the primary policy. Second and more importantly, the safety classifier finds and exploits a simpler decision boundary between safe and unsafe states instead of trying to learn a complex mapping from a state observation to a control variables. For instance, in learning to drive, the safety classifier may simply learn to distinguish between a crowded road and an empty road and determine that it is safer to let the primary policy drive in an empty road.

Relationship to a Value Function A value function $V^\pi(s)$ in reinforcement learning computes the reward a given policy π can achieve in the future starting from a given state s (Sutton and Barto 1998). This description already reveals a clear connection between the safety classifier and the value function. The safety classifier $c_{\text{safe}}(\pi, s)$ determines whether a given policy π is likely to *fail* if it operates at a given state s , in terms of the deviation from a reference policy. By assuming that a reward is only given at the very end of a policy run and that the reward is 1 if the current policy acts exactly like the reference policy and otherwise 0, the safety classifier precisely returns the value of the current state.

A natural question that follows is whether the safety classifier can drive a car on its own. This perspective on the safety classifier as a value function suggests a way to use the safety

²Such intervention has been done manually by a human driver (Pomerleau 1992).

Algorithm 1 SafeDagger Blue fonts are used to highlight the differences from the vanilla DAgger.

```

1: Collect  $D_0$  using a reference policy  $\pi^*$ 
2: Collect  $D_{\text{safe}}$  using a reference policy  $\pi^*$ 
3:  $\pi_0 = \arg \min_{\pi} l_{\text{supervised}}(\pi, \pi^*, D_0)$ 
4:  $c_{\text{safe},0} = \arg \min_{c_{\text{safe}}} l_{\text{safe}}(c_{\text{safe}}, \pi_0, \pi^*, D_{\text{safe}} \cup D_0)$ 
5: for  $i = 1$  do  $M$ 
6:   Collect  $D'$  using the safe strategy  $\pi_{i-1}$  and  $c_{\text{safe},i-1}$ 
7:   Subset Selection:
8:      $D' \leftarrow \{\phi(s) \in D' \mid c_{\text{safe},i-1}(\pi_{i-1}, \phi(s)) = 0\}$ 
9:      $D_i = D_{i-1} \cup D'$ 
10:     $\pi_i = \arg \min_{\pi} l_{\text{supervised}}(\pi, \pi^*, D_i)$ 
11:     $c_{\text{safe},i} = \arg \min_{c_{\text{safe}}} l_{\text{safe}}(c_{\text{safe}}, \pi_i, \pi^*, D_{\text{safe}} \cup D_i)$ 
12: end for
13: return  $\pi_M$  and  $c_{\text{safe},M}$ 

```

classifier directly to drive a car. At a given state s , the best action \hat{a} can be selected to be $\arg \max_{a \in A(s)} c_{\text{safe}}(\pi, \delta(s, a))$. This is however not possible in the current formulation, as the transition function δ is unknown. We may extend the definition of the proposed safety classifier so that it considers a state-action pair (s, a) instead of a state alone and predicts the safety in the next time step, which makes it closer to a Q function.

SafeDagger: Safety Classifier in the Loop

We describe here the proposed SafeDagger which aims at reducing the number of queries to a reference policy during iterations. At the core of SafeDagger lies the safety classifier introduced earlier in this section. The SafeDagger is presented in Alg. 1. There are two major modifications to the original DAgger.

First, we use the safe strategy, instead of the naive strategy, to collect training examples (line 6 in Alg. 1). This allows an agent to simply give up when it is not safe to drive itself and hand over the control to the reference policy, thereby collecting training examples with a much further horizon without crashing. This would have been impossible with the original DAgger unless the manually forced take-over measure was implemented (Ross et al. 2013).

Second, the subset selection (line 7 in Alg. 1) drastically reduces the number of queries to a reference policy. Only a small subset of states where the safety classifier returned 0 need to be labelled with reference actions. This is contrary to the original DAgger, where all the collected states had to be queried against a reference policy.

Furthermore, this subset selection allows the subsequent supervised learning to focus more on difficult cases, which almost always correspond to the states that are problematic (i.e., $S \setminus S^*$.) This reduces the total amount of training examples without losing important training examples, thereby making this algorithm data-efficient.

Once the primary policy is updated with D_i which is a union of the initial training set D_0 and all the hard examples collected so far, we update the safety classifier. This step ensures that the safety classifier correctly identifies which states are difficult/dangerous for the latest primary policy.

This has an effect of automated curriculum learning (Bengio et al. 2009) with a mix strategy (Zaremba and Sutskever 2014), where the safety classifier selects training examples of appropriate difficulty at each iteration.

Despite these differences, the proposed SafeDagger inherits much of the theoretical guarantees from the DAgger. This is achieved by gradually increasing the threshold τ of the safety classifier (Eq. (3)). If $\tau > \epsilon(\pi, \phi(s))$ for all $s \in S$, the SafeDagger reduces to the original DAgger with β_i (from Eq. (2)) set to 0. We however observe later empirically that this is not necessary, and that training with the proposed SafeDagger with a fixed τ automatically and gradually reduces the portion of the reference policy during data collection over iterations.

Adaptation to Other Imitation Learning Algorithms

The proposed use of a safety classifier is easily adaptable to other more recent cost-sensitive algorithms. In AggreVaTe (Ross and Bagnell 2014), for instance, the roll-out by a reference policy may be executed not from a uniform-randomly selected time point, but from the time step when the safety classifier returns 0. A similar adaptation can be done with LOLS (Chang et al. 2015). We do not consider these algorithms in this paper and leave them as future work.

Experimental Setting

Simulation Environment

We use TORCS (tor accessed May 12 2016), a racing car simulator, for empirical evaluation in this paper. We chose TORCS based on the following reasons. First, it has been used widely and successfully as a platform for research on autonomous racing (Loiacono et al. 2008), although most of the previous work, except for (Koutník et al. 2013; Chen et al. 2015), are not comparable as they use a radar instead of a camera for observing the state. Second, TORCS is a light-weight simulator that can be run on an off-the-shelf workstation. Third, as TORCS is an open-source software, it is easy to interface it with another software which is Torch in our case.³

Tracks To simulate a highway driving with multiple lanes, we modify the original TORCS road surface textures by adding various lane configurations such as the number of lanes, the type of lanes. We use ten different tracks in total for our experiments. We split those ten tracks into two disjoint sets: seven training tracks and three test tracks. All training examples as well as validation examples are collected from the training tracks only, and a trained primary policy is tested on the test tracks. See Fig. 1 for the visualizations of the tracks and the first section for the types of information collected as examples, in the supplementary material.

Reference Policy π^* We implement our own reference policy which has access to an underlying state configuration. The state includes the position, heading direction, speed, and distances to others cars. The reference policy follows a simple traffic rule. The cars driven by the reference policy either

follow the current lane (accelerating up to the speed limit), change the lane if there is a slower car in the front and a lane to the left or right is available, or brakes.

Data Collection

We use a car in TORCS driven by the pre-defined reference policy to collect data. For each training track, we add 40 cars driven by the same reference policy to simulate traffic. We run up to three iterations in addition to the initial supervised learning stage. In the case of SafeDagger, we collect 30k, 30k and 10k of training examples (after the subset selection in line 6 of Alg. 1.) In the case of the original DAgger, we collect up to 390k data each iteration and uniform-randomly select 30k, 30k and 10k samples as our training examples.

This data collection strategy was designed in order to keep the amount of data used for training a primary policy network at the same level. Note that the number of queries to the reference policy is much higher (up to 39 folds) with the original DAgger, as it queries the reference policy for every single image frame simultaneously with a primary policy.

Policy Networks

Primary Policy π_θ We use a deep convolutional network that has five convolutional layers followed by a set of fully-connected layers. This convolutional network takes as input the pixel-level image taken from a front-facing camera. It predicts the angle of steering wheel ($[-1, 1]$) and whether to brake ($\{0, 1\}$). Furthermore, the network predicts as an auxiliary task the car’s affordances, including the existence of a lane to the left or right of the car and the existence of another car to the left, right or in front of the car. We have found this multi-task approach to easily outperform a single-task network, confirming the promise of multi-task learning from (Caruana 1997).

Safety Classifier c_{safe} We use a feedforward network to implement a safety classifier. The activation of the primary policy network’s last hidden convolutional layer is fed through two fully-connected layers followed by a softmax layer with two categories corresponding to 0 and 1. We choose $\tau = 0.0025$ as our safety classifier threshold so that approximately 20% of initial training examples are considered unsafe, as shown in Fig. 2. See Fig. 3 in the supplementary material for some examples of which frames were determined safe or unsafe. For more details, see the second section in the supplementary material.

Evaluation

Training and Driving Strategies We mainly compare three training strategies; (1)Supervised Learning, (2) DAgger (with $\beta_i = \mathbf{I}_{i=0}$) and (3) SafeDagger. For each training strategy, we evaluate trained policies with both of the driving strategies; (1) naive strategy and (2) safe strategy.

Evaluation Metrics We evaluate each combination by letting it drive on the three test tracks up to three laps. All these runs are repeated in two conditions; without traffic and with traffic, while recording three metrics. The first metric is *the number of completed laps* without going outside a track, averaged over the three tracks. When a car drives out of the

³We will release a patch to TORCS that allows seamless integration between TORCS and Torch.

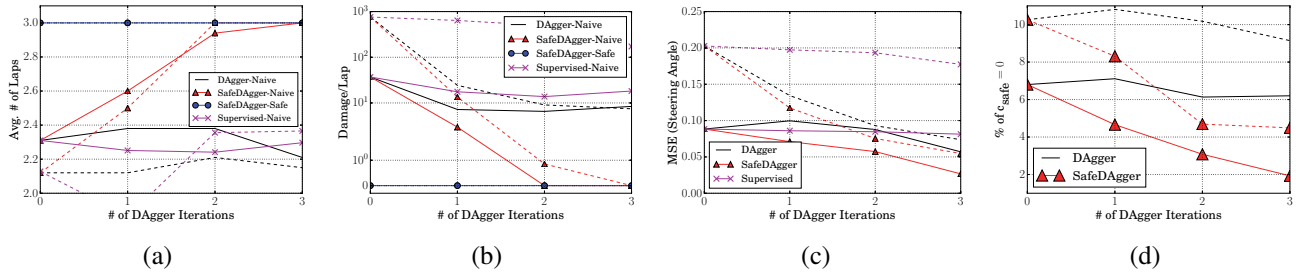


Figure 1: (a) Average number of laps (\uparrow), (b) damage per lap (\downarrow), (c) the mean squared error of steering angle for each configuration (training strategy–driving strategy) over the iterations and (d) the portion of time driven by a reference policy during test. We use solid and dashed curves for the cases without and with traffic, respectively.




Figure 2: The histogram of the log square errors of steering angle after *supervised* learning only. The dashed line is located at $\tau = 0.0025$. 77.70% of the training examples are considered safe.

track, we immediately halt. Second, we look at the damage accumulated while driving. Damage happens each time the car bumps into another car. Instead of a raw, accumulated damage level, we report the *damage per lap*. Lastly, we report the *mean squared error of steering angle*, computed while the primary policy drives.

Results and Analysis

In Fig. 1, we present the result in terms of both the average laps and damage per lap. The first thing we notice is that a primary policy trained using supervised learning (the 0-th iteration) alone works perfectly when a safety classifier is used together. The safety classifier switched to the reference policy for 7.11% and 10.81% of time without and with traffic during test.

Second, in terms of both metrics, the primary policy trained with the proposed SafeDagger makes much faster progress than the original DAgger. After the third iteration, the primary policy trained with the SafeDagger is perfect. We conjecture that this is due to the effect of automated curriculum learning of the SafeDagger. Furthermore, the examination of the mean squared difference between the primary policy and the reference policy reveals that the SafeDagger more rapidly brings the primary policy closer to the reference policy.

As a baseline we put the performance of a primary policy trained using purely supervised learning in Fig. 1 (a)–(b). It clearly demonstrates that supervised learning alone cannot

train a primary policy well even when an increasing amount of training examples are presented.

In Fig. 1 (d), we observe that the portion of time the safety classifier switches to the reference policy while driving decreases as the SafeDagger iteration progresses. We conjecture that this happens as the SafeDagger encourages the primary policy’s learning to focus on those cases deemed difficult by the safety classifier. When the primary policy was trained with the original DAgger (which does not take into account the difficulty of each collected state), the rate of decrease was much smaller. Essentially, using the safety classifier and the SafeDagger together results in a virtuous cycle of less and less queries to the reference policy during both training and test.

Lastly, we conduct one additional run with the SafeDagger while training a safety classifier to predict the deviation of a primary policy from the reference policy one second in advance. We observe a similar trend, which makes the SafeDagger a realistic algorithm to be deployed in practice.

Conclusion

In this paper, we have proposed an extension of DAgger, called SafeDagger. We first introduced a safety classifier which prevents a primary policy from falling into a dangerous state by automatically switching between a reference policy and the primary policy *without* querying the reference policy. This safety classifier is used during data collection stages in the proposed SafeDagger, which can collect a set of progressively difficult examples while minimizing the number of queries to a reference policy. The extensive experiments on simulated autonomous driving showed that the SafeDagger not only queries a reference policy less but also trains a primary policy more efficiently.

Imitation learning, in the form of the SafeDagger, allows a primary policy to learn without any catastrophic experience. The quality of a learned policy is however limited by that of a reference policy. More research in finetuning a policy learned by the SafeDagger to surpass existing, reference policies, for instance by reinforcement learning (Silver et al. 2016), needs to be pursued in the future.

Acknowledgments We thank the support by Facebook, Google (Google Faculty Award 2016) and Nvidia (GPU Center of Excellence 2015-2016).

References

- Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, 41–48. ACM.
- Bojarski, M.; Testa, D. D.; Dworakowski, D.; Firner, B.; Flepp, B.; Goyal, P.; Jackel, L. D.; Mathew Monfort, U. M.; Zhang, J.; Zhang, X.; Zhao, J.; and Zieba, K. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*.
- Caruana, R. 1997. Multitask learning. *Machine learning* 28(1):41–75.
- Chang, K.-w.; Krishnamurthy, A.; Agarwal, A.; Daume, H.; and Langford, J. 2015. Learning to search better than your teacher. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, 2058–2066.
- Chen, C.; Seff, A.; Kornhauser, A.; and Xiao, J. 2015. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, 2722–2730.
- Daumé Iii, H.; Langford, J.; and Marcu, D. 2009. Search-based structured prediction. *Machine learning* 75(3):297–325.
- Judah, K.; Fern, A. P.; Dietterich, T. G.; et al. 2014. Active Imitation learning: formal and practical reductions to iid learning. *The Journal of Machine Learning Research* 15(1):3925–3963.
- Kim, B., and Pineau, J. 2013. Maximum mean discrepancy imitation learning.
- Koutník, J.; Cuccu, G.; Schmidhuber, J.; and Gomez, F. 2013. Evolving large-scale neural networks for vision-based reinforcement learning. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, 1061–1068. ACM.
- Laskey, M.; Staszak, S.; Hsieh, W. Y.-S.; Mahler, J.; Pokorny, F. T.; Dragan, A. D.; and Goldberg, K. 2016. Shiv: Reducing supervisor burden in dagger using support vectors for efficient learning from demonstrations in high dimensional state spaces. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 462–469. IEEE.
- Loiacono, D.; Togelius, J.; Lanzi, P. L.; Kinnaird-Heather, L.; Lucas, S. M.; Simmerson, M.; Perez, D.; Reynolds, R. G.; and Saez, Y. 2008. The wcci 2008 simulated car racing competition. In *CIG*, 119–126. Citeseer.
- Muller, U.; Ben, J.; Cosatto, E.; Flepp, B.; and Cun, Y. L. 2005. Off-road obstacle avoidance through end-to-end learning. In *Advances in neural information processing systems*, 739–746.
- Pomerleau, D. A. 1989. Alvin: An autonomous land vehicle in a neural network. Technical report, DTIC Document.
- Pomerleau, D. A. 1992. Progress in neural network-based vision for autonomous robot driving. In *Intelligent Vehicles' 92 Symposium., Proceedings of the*, 391–396. IEEE.
- Ross, S., and Bagnell, J. A. 2014. Reinforcement and imitation learning via interactive no-regret learning. *arXiv preprint arXiv:1406.5979*.
- Ross, S.; Melik-Barkhudarov, N.; Shankar, K. S.; Wendel, A.; Dey, D.; Bagnell, J. A.; and Hebert, M. 2013. Learning monocular reactive uav control in cluttered natural environments. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, 1765–1772. IEEE.
- Ross, S.; Gordon, G. J.; and Bagnell, J. A. 2010. A reduction of imitation learning and structured prediction to no-regret online learning. *arXiv preprint arXiv:1011.0686*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529(7587):484–489.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement learning: An introduction*, volume 28. MIT press.
- accessed May 12, 2016. *The Open Racing Car Simulator*.
- Zaremba, W., and Sutskever, I. 2014. Learning to execute. *arXiv preprint arXiv:1410.4615*.