

Hierarchical Generative Adversarial Imitation Learning with Mid-level Input Generation for Autonomous Driving on Urban Environments

Gustavo Claudio Karl Couto* and Eric Aislan Antonelo*

Abstract—Deriving robust control policies for realistic urban navigation scenarios is not a trivial task. In an end-to-end approach, these policies must map high-dimensional images from the vehicle’s cameras to low-level actions such as steering and throttle. While pure Reinforcement Learning (RL) approaches are based exclusively on engineered rewards, Generative Adversarial Imitation Learning (GAIL) agents learn from expert demonstrations while interacting with the environment, which favors GAIL on tasks for which a reward signal is difficult to derive, such as autonomous driving. However, training deep networks directly from raw images on RL tasks is known to be unstable and troublesome. To deal with that, this work proposes a hierarchical GAIL-based architecture (hGAIL) which decouples representation learning from the driving task to solve the autonomous navigation of a vehicle. The proposed architecture consists of two modules: a GAN (Generative Adversarial Net) which generates an abstract mid-level input representation, which is the Bird’s-Eye View (BEV) from the surroundings of the vehicle; and the GAIL which learns to control the vehicle based on the BEV predictions from the GAN as input. hGAIL is able to learn both the policy and the mid-level representation simultaneously as the agent interacts with the environment. Our experiments made in the CARLA simulation environment have shown that GAIL exclusively from cameras (without BEV) fails to even learn the task, while hGAIL, after training exclusively on one city, was able to autonomously navigate successfully in 98% of the intersections of a new city not used in training phase. Videos and code available at: <https://sites.google.com/view/hgail>.

Index Terms—Autonomous Driving, Generative Adversarial Imitation Learning, CARLA Simulator, Bird’s-Eye View

I. INTRODUCTION

Commonly, Autonomous Driving (AD) has been implemented using individual modules for perception, planning and control organized in a pipeline [1]–[4]. However, learning approaches have been on the rise, in an attempt to tackle the complexities of AD in different scenarios, in simulation or even in the real world. Most of the approaches are based on Behavior Cloning (BC), which uses supervised learning on a set of expert demonstrations collected offline [5]–[9], e.g., with a human driver generating a set of input (observations) and corresponding desired output pairs. The latter approach suffers from covariate shift [10]–[12], since it can not teach robustly the learning agent a trajectory which does not accumulate errors.

Reinforcement Learning (RL) approaches to AD can learn policies that do not present this covariate shift issue, since

the agent is able to learn in interaction with the environment, considering the whole sample trajectories and not only independent observation-action samples as in BC. However, RL requires the definition of a reward signal, which can be cumbersome to do it considering the complexity of a driver’s behavior and its environment. Although Inverse RL can be used for imitation learning purposes [13], it is expensive to run since it executes RL in a loop. Learning in IRL is computationally more expensive than just learning a policy directly from expert demonstrations.

On the other hand, Generative Adversarial Imitation Learning (GAIL) [14] provides a way to train agents in interaction with their environments, directly from expert demonstrations. This approach has been validated in the CARLA realistic simulator for autonomous driving in urban scenarios previously [15], showing it can scale to large environments. However, in [15] only fixed routes were considered. Although the agent’s architecture was general enough for dynamic routes, with inputs such as the high-level command and the next point of the sparse trajectory in the vehicle’s reference frame, the network had limitations for learning a general policy for dynamic routes, i.e., those that can change on the fly (turn left, right, or go straight at an intersection) from the perspective of the agent.

In [16], the authors show that training complex mathematical functions such as deep neural networks for online policy learning based on complex input patterns from raw camera images is unstable. This is because the reward signal is not informative enough for training such large models. One of the strategies employed by [16] to circumvent this limitation is to decouple feature representation from policy learning. In our context, the decoupling is achieved by learning a mid-level representation of the vehicle’s raw camera images. For instance, Bird’s-Eye View (BEV) representations of the road ahead of the vehicle have been used as mid-level input to trajectory generation and motor control networks in [7] and [17], respectively. In BEV representation, the scene from the camera is projected onto a top-down view, and regions of interest in the image, such as vehicle lanes and sidewalks, are segmented into different colors [2]. Using BEV as input to the control policy, the idea is to make policy learning focus on the navigation problem, since it will use a more abstract and simple view of the road ahead.

One of the advantages of such approach is enabling transfer to real-world by a relatively easy process: it requires mapping the real-world images to the same abstract representation used to train the agent in simulation. Besides, training agents with mid-level visual inputs such as BEV and others (e.g.,

*Gustavo C. K. Couto and Eric A. Antonelo are with the Automation and Systems Engineering Department, Federal University of Santa Catarina, Florianópolis, Brazil gustavo.karl.couto@posgrad.ufsc.br, eric.antonelo@ufsc.br.

This is an author’s version copy. The official publication is available at: <http://dx.doi.org/10.1109/TIV.2024.3436587>.

optical flow, depth, semantic segmentation, and albedo) make policies learn faster, generalize better and achieve higher task performance [18]. Other works employ semantic segmentation for semantic driving [19]–[21], offering supporting evidence that mid-level inputs to agents are useful for realistic downstream active tasks [18]. So far, BEV has been employed for autonomous driving in urban scenarios in [17] and [7]. However, they generate BEV input by using an algorithm which has to access a known map of the city. In contrast, in our work, we consider a mid-level predicted BEV input, feeding the agent’s policy, that is generated by a separate network that learns simultaneously with the policy network in a hierarchical GAIL (hGAIL) architecture. The proposed hGAIL consists of two modules: a learning module for mapping camera images to BEV representations; and a policy learning module based on GAIL that controls the vehicle by using the predicted BEV as mid-level input. The first module uses Conditional Generative Adversarial Nets (GANs) [22] and U-nets to map the images obtained from three frontal vehicle’s cameras to the mid-level BEV input. The latter module, based on GAIL, outputs steer, acceleration and brake signals to drive the vehicle based on the GAN’s output. Thus, although our agent produces a mid-level representation, it still pertains to the class of end-to-end models, since the raw camera images are fed as inputs to the hGAIL architecture. Additionally, the GAIL module’s cost function is augmented with a behavior cloning loss [23] in order to stabilize the policy learning. Both GAN and GAIL networks learn simultaneously, though with their own cost functions, while the agent interacts with the CARLA simulation environment for urban driving. This approach ensures that both the policy and representation networks are trained using on-policy data, learning from the agent’s mistakes.

This work contributes by:

- 1) overcoming the difficulties of policy learning directly from raw pixels for the GAIL agent in autonomous driving. This is achieved by decoupling representation learning from the driving task, through learning a Bird’s-Eye View (BEV) representation of the input, which can be processed more efficiently by GAIL.
- 2) proposing an end-to-end hierarchical architecture based on both GAN and GAIL (hGAIL) for autonomous urban driving. The online nature of GAIL allows the agent to experience out-of-distribution samples not existing in the expert dataset, different from Behavior Cloning approaches [24]. After training, the hGAIL agent only has access to a sparse trajectory, being able to predict a more refined desired trajectory that is used by the policy to directly output the steering and throttle motor actuators, unlike approaches that output the trajectory and use a nonadaptive controller to move the vehicle [7], [25].
- 3) extending previous work that is applicable to fixed routes [15] to a more skilled agent able to follow dynamic routes, making the vehicle able to change the route on the fly. In addition, this work also tackles navigation in cities with pedestrians and other vehicles, and respecting traffic lights, unlike [15].

- 4) proposing a GAN network for BEV mid-level representation generation of three channels (desired route, drivable area, and lane boundaries) based on inputs from the raw vehicle’s frontal camera images, sparse trajectory and high-level command. This approach enables simultaneous representation learning with policy learning in an online manner, without degrading the GAIL training process.

II. RELATED WORKS

A. Behavior Cloning (BC)

The authors of works [8] and [9] utilized Behavior Cloning (BC) for Autonomous Driving (AD) in the CARLA simulator. A large dataset of human driving was collected and augmented using image processing techniques to train end-to-end policies conditioned on the desired route. Further advancements in BC were made in [26], which used a large deep ResNet network for feature extraction and fused data from camera, LiDAR, and radar to generate feature maps.

B. Reinforcement Learning (RL)

In [27], an RL agent was trained and tested in a 2D simulator, CARLA, and a real-world car racing task, highlighting the potential use of auxiliary models to address the weakness of RL in needing large amounts of environment interactions. In [17], RL is employed with a customized reward function to train an agent based on the mid-level BEV representation as input, during the first training phase. Afterwards, using the trained agent as an online expert, they trained a second agent through apprenticeship learning, but now in an end-to-end approach with input directly from the vehicle’s cameras. Their method evaluated using CARLA surpassed the default benchmarks trained by BC.

C. GAIL

The authors in [28] used Generative Adversarial Imitation Learning (GAIL) with simple affordance-style features as inputs to reduce cascading errors in behavior-cloned policies and make them more robust to perturbations. Raw LiDAR readings and simple road features, such as speed and lane center offset, were mapped to turn-rate and acceleration to model human highway driving behavior in a realistic highway simulator. The experiment successfully reproduced human driver behavior while reducing the risk of collisions.

In [29], model-based GAIL (MGAIL) was used with a large-scale expert dataset (100,000 miles) from San Francisco city, and applied to the task of dense urban self-driving. The discriminator and generator in their MGAIL model consists of Transformer networks, which combine different types of inputs: roadgraph points, traffic light signals, and other objects’ trajectories. They introduce an hierarchical model which integrates a high-level graph-based search (to generate the intended trajectory) with a low-level transformer-based MGAIL policy. Besides, they do not learn mid-level input representations from raw camera images as hGAIL in this work, but use a pretrained independent robotic perception

system to generate input features. Additionally, our controller in hGAIL learns with expert samples generated in the CARLA simulator corresponding to only 36 minutes (8 km) of driving. In theory, transferring the trained policy in hGAIL to a real-world setting would require only replacing the BEV network by another network trained on real images.

D. BEV Mid-level representation

Reinforcement learning enables an agent to be trained in a closed-loop fashion, resulting in more robust agents. However, this robustness comes with the cost of instabilities that prevent the use of very deep networks, such as ResNet [30]. In [16], the authors aim to decouple representation learning from the training of reinforcement learning as a solution to train agents that can process high-dimensional raw inputs, such as images from cameras.

In [25], a hierarchical driving model is proposed, consisting of a driving intention module and trajectory generation module. The former learns to generate the future trajectory with a GAN, whose generator is fed with an image from a monocular camera and a local map, cropped from an offline map using the GPS position. The intention map is combined with LiDAR data to generate a potential map, which is their mid-level input representation (similar to BEV), subsequently fed to a controller module trained to imitate a set of demonstration trajectories. While their BEV prediction has only one channel with the desired trajectory (intention), our BEV in hGAIL produces also other two channels for drivable area and lane boundaries. However, in our approach, we do not combine LiDAR data with the BEV prediction. The final output of their model corresponds to a trajectory function that is used by a nonadaptive controller to drive the vehicle, whereas hGAIL outputs directly the values for the steering and throttle motor actuators. The training of their networks are offline, opposed to our online training of hGAIL, which is more challenging with respect to the learning task of directly mapping raw frontal cameras images to motor actuators. They evaluate their method on CARLA and also on a real vehicle.

A similar modular approach is explored in [24] in the CARLA simulation environment, which combines a predicted Bird’s-Eye View representation with a pretrained trajectory of 0.5-meter resolution to serve as inputs to a policy network trained with Behavior Cloning. To generate the BEV, first, the raw image is segmented using a pretrained segmentation network; then, its output is transformed by an U-net into an BEV top-down perspective. It is worth noting that our BEV generation in hGAIL uses a Conditional GAN that maps directly from raw camera images to the final BEV prediction, while [24] uses a segmentation network as an intermediate step. Additionally, our BEV prediction includes the desired trajectory as a BEV channel, while the BEV prediction in [24] does not include this information, needing access to the points of a dense trajectory as inputs to the policy. In contrast, after training, our hGAIL agent only access a sparse trajectory of 50-meter resolution on average, which is 100 times more sparse than the one used in their work. Besides, they train the BEV network offline in a supervised way and before the

policy network is trained by conventional Behavior Cloning, also offline. Notice that, in our work, both BEV and policy networks learn online in an integrated way simultaneously and in an adversarial way. Thus, covariate shift situations are better handled with hGAIL, since the policy can learn with samples not included in the expert dataset.

The BEV input representation in hGAIL is inspired by [7], which trains a system to autonomously drive a vehicle on CARLA using a Bird’s-Eye View representation as input. This BEV representation, like our own, contains a road and route representation from a top-down, agent-centered perspective. However, after training, hGAIL does not need anymore to access the city map to generate the BEV of the vehicle, as this is done by a GAN network, unlike in [7]. Based on this BEV input, in their work, a neural network is trained offline to output the trajectory to be followed by a nonadaptive controller. They also employ data augmentation and auxiliary losses in this training. In contrast to this, our hGAIL agent learns both the policy and the BEV representation online and simultaneously, i.e., in a closed-loop training process, mapping raw frontal camera images directly into motor actuator values (steering, and throttle). Thus, the controller also has to be learned in hGAIL, unlike [7] which learns only the trajectory offline.

The work in [19] uses the output of a scene segmentation network as a mid-level input representation. This enables the transfer of the agent’s policy to the real world, since a similar network can be used to generate the mid-level input from images of real scenes.

Our work is the first instance, to our knowledge, where a GAIL is trained using a mid-level input representation generated by a learning module to tackle the intricate task of autonomous driving navigation. Previous works also show that policies with mid-level representations as input can be trained on a simulator, and subsequently transferred to the real world. This is a cost-effective and safe way to perform closed-loop training of agents (since infractions during training of the agent occur in simulation only), in particular, by the GAIL method which provides an interactive learning based on expert demonstrations.

III. METHODS

A. Conditional Generative Adversarial Networks - CGAN

CGANs are composed of two neural networks, a discriminator D and a generator G . The function of G in a CGAN [22] is to translate an image x into an image y by mapping both x and a random noise vector z into an output image $G : \{x, z\} \rightarrow y$. Both D and C seek to optimize the same objective function:

$$\mathbb{E}_{x,y} [\log(D(x, y))] + \mathbb{E}_{x,z} [\log(1 - D(x, G(x, z)))] \quad (1)$$

where G tries to minimize it, while D seeks to maximize it.

Additionally, a L1 distance loss function is added to the final objective, making the generator network G also learn from the true label y as it would happen in a supervised learning task,

modeling low-frequency characteristics of images. The resulting method, PatchGAN [22], divides the image into multiple patches to be classified by the discriminator individually and, thus, requires fewer parameters.

The CGAN is used in our work to generate the Bird’s-Eye View (BEV) image representation from the agent’s sensors such as frontal cameras and GPS, to be detailed later.

B. Generative Adversarial Imitation Learning - GAIL

Mathematically, GAIL¹ finds a saddle point (π, D) of the expression:

$$\mathbb{E}_{\pi_E} [D(s, a)] - \mathbb{E}_{\pi} [D(s, a)] - \lambda H(\pi) - \lambda_2 L_{gp}, \quad (2)$$

where $D : S \times A \rightarrow (0, 1)$, S is the state space, A is the action space; π_E is the expert policy; π is the agent’s policy; $H(\pi)$ is a policy regularizer controlled by $\lambda \geq 0$ [31]. The discriminator D will try to increase (2), while π seeks to minimize it; and L_{gp} is a loss that penalizes the gradient’s norm, constraining the discriminator network to the 1-Lipschitz function space, according to [32]. The above equation is the Wasserstein loss, used to alleviate vanishing gradient and mode collapse problems, using the Wasserstein distance [32] between the policy distribution and expert distribution, as also done in [33], [34].

C. Bird’s-Eye View (BEV) representation

The BEV of a vehicle represents its position and movement in a top-down coordinate system [7]. The vehicle’s location, heading, and speed are represented by p_t , θ_t , and s_t respectively. The top-down view is defined so that the agent’s starting position is always at a fixed point within an image (the center of it). Furthermore, it is represented by a set of images of size $W \times H$ pixels, at a ground sampling resolution of ϕ meters/pixel. The BEV of the environment moves as the vehicle moves, allowing the agent to see a fixed range of meters in front of it. For instance, the BEV representation for the vehicle whose three frontal cameras are shown in Fig. 1(a) is given in Fig. 1(c), where the desired route, drivable area and lane boundaries form a three-channel image.

IV. AGENT

Our agent’s architecture (Fig. 2) is based on hierarchical Generative Adversarial Imitation Learning (hGAIL) for training policy and mid-level representation simultaneously. There are two main parts of hGAIL: the conditional GAN that generates the BEV representation based on input from the vehicle’s frontal cameras, trajectory and high-level command; and the GAIL that learn the agent’s policy by imitation learning based on input from the BEV representation generated by the first CGAN module, current vehicle’s speed, and the last actuator values.

A. BEV generation with CGANs

The Conditional GAN module, used to transform the images from the frontal cameras into a top-down view representation, has two different networks named Discriminator and Generator, whose architectures can be seen in Appendix VI-B.

1) *Input representation*: The input for the CGANs corresponds to the 192x192 resolution RGB images from the three frontal cameras represented in Fig. 1(a) and the sparse trajectory visual representation from Fig. 1(b) that are stacked to generate a 10x192x192 image, i.e., with 10 channels. The goal of the CGAN’s generator is to translate this stack of images into the 3-channel BEV representation seen in Fig. 1(c). In addition to this RGB input image, the discriminator also receives the 3x192x192 BEV image, which can come from either the generator as *fake* or from the training set as *real*. Other inputs to the generator are the 5 points from the sparse trajectory, one point behind the vehicle and 4 points ahead of it, represented as a vector and the high-level command as a 4-dimensional one-hot encoding vector (“lane follow”, “left”, “right”, “straight”).

The architectures of both Generator and Discriminator are described in Appendix VI-B.

B. Policy learning with GAIL

The generator in the GAIL module iteratively seeks the θ parameters of the policy $\pi_{\theta}(\cdot|s)$ that minimizes (2), while the discriminator seeks to maximize it. To assist the agent’s learning, loss terms for stimulating exploration are added as described after the representations for the input, output, and architecture are presented.

1) *Input representation*: The input s to the agent’s policy is a three-channels 192x192 image generated by the GAN network, corresponding to the mid-level BEV representation of the vehicle in its current position. In addition, the current vehicle’s speed and the last value of the policy actuators (last acceleration and steering) are also fed as input further down in the network layers (to the first fully connected layer).

2) *Output representation*: The vehicle in CARLA has three actuators as: *steering* $\in [-1, 1]$, *throttle* $\in [0, 1]$, and *brake* $\in [0, 1]$. Our agent’s action space is $\mathbf{a} \in [-1, 1]^2$, where the two components of \mathbf{a} correspond to steering and acceleration. Braking occurs when acceleration is negative. In this way, by modeling brake and throttle with one dimension, the agent is not allowed to brake and accelerate simultaneously [35]. Instead of using the Gaussian distribution for the policy’s actions, common choice in model-free RL, we employ the Beta distribution $\mathcal{B}(\alpha, \beta)$ due to its bounded support, which allows us to model bounded continuous action distributions, usually found in real-world applications such as autonomous driving [35], where the action space is not unbounded (i.e., the gas pedal can be actuated up to a certain limit). Besides, the policy loss \mathcal{L}_P can be explicitly computed since clipping or squashing is not used to enforce input constraints (in the case of Gaussian distribution). Furthermore, the Beta distribution allows the policy to act in extreme situations of vehicle driving, where sharp turns and sudden braking are necessary, as its parameters α and β , which are defined as outputs of the policy

¹More details on GAIL are given in Appendix VI-A.

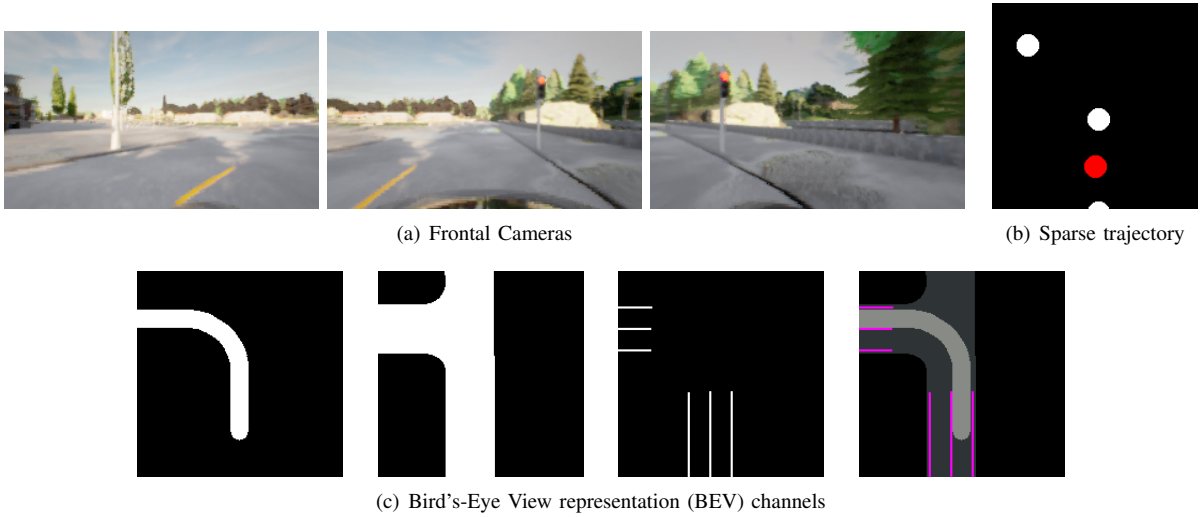


Fig. 1. a) Images from three frontal cameras located at the left, central, and right part of the vehicle, respectively. They were taken after the first few interactions of the agent in the CARLA simulation environment. Each camera produces a 256x144 RGB image. b) The corresponding sparse trajectory visual input captured at the same frame. The points from the sparse trajectory and the highlighted vehicle position are plotted as circles with a radius of 10 pixels, using the same scale (pixels per meter) and perspective as the BEV representation. When the image is fed to the CGAN, it is represented with only one channel and a size of 192x192 pixels. c) The three channels of BEV image that our agent employs, computed at the same instant shown in (a). From left to right, the channels correspond to: desired route, drivable area, and lane boundaries. The last image shows all three channels combined in different colors.

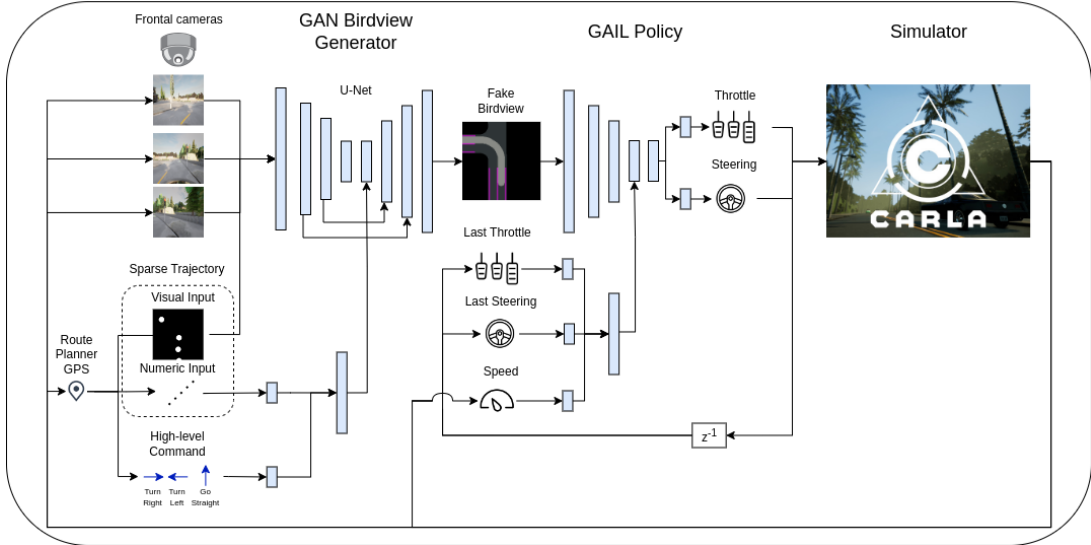


Fig. 2. Hierarchical Generative Adversarial Imitation Learning (hGAIL) for policy learning with mid-level input representation. It basically consists of chained CGAN and GAIL networks, where the first one (CGAN) generates BEV representation from the vehicle’s three frontal cameras, sparse trajectory and high-level command, while the latter (GAIL) outputs the acceleration and steering based on the predicted BEV input (generated by CGAN), the current speed and the last applied actions. Both CGAN and GAIL learn simultaneously while the agent interacts to the CARLA environment. The discriminator parts of both networks are not shown for the sake of simplicity.

neural network π_θ and control the shape of the distribution, can be tuned to produce such characteristic vehicle behaviors.

The architectures of both Generator and Discriminator are described in Appendix VI-C.

3) *BC augmentation* : The generator loss in GAIL is also augmented with a Behavior Cloning (BC) loss \mathcal{L}_{BC} ($-\mathbb{E}_{\pi_E}[\log(\pi(a|s))]$) using the available expert samples. The underlying idea is to provide the training process of the policy with useful and informative gradients, especially when the discriminator is not yet fully trained. This concept and the practical implementation follows [23].

4) *Encouraging Exploration*: During training, the agent is encouraged to explore the environment through two objectives, as in [17]:

$$\mathcal{L}_{\text{ent}} + \mathcal{L}_{\text{exp}} \quad (3)$$

where: the first loss function corresponds to the entropy loss commonly used to promote exploration:

$$\mathcal{L}_{\text{ent}} = -\lambda_{\text{ent}} \cdot \mathbb{H}(\pi_\theta(\cdot|s)), \quad (4)$$

Minimizing \mathcal{L}_{ent} means maximizing entropy and thus uncertainty for the policy distribution π_θ , which stimulates the

agent try more diverse actions since the policy distribution for a certain state s does not become too certain too quickly in the process. It also drives the action (policy) distribution towards a uniform prior (which represents maximum entropy and uncertainty) since it is equivalent to minimizing the KL-divergence to the uniform distribution defined in the support of the Beta policy $[-1, 1]$:

$$H(\pi_\theta) = -\text{KL}(\pi_\theta || \mathcal{U}(-1, 1)), \quad (5)$$

We can also bias the agent’s learning with priors that signify meaningful behaviors for an autonomous vehicle and helps to improve and speed up the overall agent’s training from scratch. This is accomplished with the following exploration loss \mathcal{L}_{exp} [17]:

$$\mathcal{L}_{\text{exp}} = \lambda_{\text{exp}} \cdot \mathbb{1}_{\{T-N_z+1, \dots, T\}}(k) \cdot \text{KL}(\pi_\theta(\cdot | s) || p_z), \quad (6)$$

where $\mathbb{1}$ is the indicator function and $z \in \mathcal{Z}$ is the terminal event that finishes the episode. Some examples of events in \mathcal{Z} would be collision, route deviation or the car being still or blocked for too long. \mathcal{L}_{exp} imposes a prior p_z to the policy during the last N_z steps of an episode ending with one of the events in \mathcal{Z} . The indicator function serves as a selection mechanism of the last steps in the episode. This p_z promotes exploration as follows: if z is a collision, $p_z = \mathcal{B}(1, 2.5)$ for the acceleration actuator, which encourages slowing down behavior; if the car is still, the acceleration prior is $p_z = \mathcal{B}(2.5, 1)$, favoring increasing the vehicle’s speed; if the vehicle deviates from the trajectory, a uniform prior $\mathcal{B}(1, 1)$ is employed for the steering actuator [17].

Thus, uniting the Proximal Policy Optimization (PPO) loss \mathcal{L}_P [36] that minimizes (2), the Behavior Cloning loss \mathcal{L}_{BC} [23] and the entropy and exploration terms in (3), the total loss function for policy learning of all GAIL-based agents in this work is as follows:

$$\alpha \mathcal{L}_{BC} + (1 - \alpha) \mathcal{L}_P + \mathcal{L}_{\text{ent}} + \mathcal{L}_{\text{exp}}, \quad (7)$$

where α controls the weight of the BC term \mathcal{L}_{BC} with respect to the PPO loss \mathcal{L}_P during training.

V. EXPERIMENTAL RESULTS

The goal of the vehicle is to learn to navigate autonomously in the city shown in Fig. 3 using the hGAIL agent’s architecture with mid-level BEV input generation, and subsequently generalize its driving skill to a new town.

A. Collected data

The environment and trajectories are obtained from the CARLA Leaderboard evaluation platform². In particular, the *town01* environment from this platform along with ten predefined trajectories are employed to generate the expert training set.

The expert dataset is constructed using a deterministic agent that navigates using a dense point trajectory and a classic



Fig. 3. *Town01* environment of the agent, with one of the routes used to collect data by the expert. The highlighted path has 740 meters, 20 points in the sparse trajectory (shown as yellow dots) and 762 points in the dense point trajectory (not shown).

PID controller [37]. The dense point trajectory provides many points at a fine resolution, whereas a sparse point trajectory consists of considerably fewer points, providing only a general sense of direction to the agent. As a result, the dense point trajectory is utilized to generate training data by the expert, whereas the sparse point trajectory is employed by the agent for more general guidance.

In Fig. 3, one of the 10 routes executed by the expert to form the labeled training set of demonstrations is shown, where the line starting in yellow and ending in red represents the desired trajectory (not observable to the agent as it is). The sparse trajectory can be seen as yellow dots, generated every 50 meters traveled or when the vehicle is about to start a different movement (from *straight* to *turn* and vice-versa).

The ten trajectories of the training set were recorded at a rate of 10 hertz, resulting in 10 observation-action pairs per second. For the shortest route of 1480 samples (average route of 2129 samples), it represents 2.5 minutes (3.5 minutes) of simulated driving. All the ten trajectories yielded a total of 21,287 training samples (30 GB of uncompressed data). The total set corresponds approximately to 36 minutes or 8km of driving.

B. Training

The training was conducted using six parallel actors in a synchronous manner, with each actor running its own instance of the CARLA simulator. In the simulation, each episode begins with the vehicle at zero speed at a random starting point. The episode concludes upon the occurrence of any infraction, collision, or lane invasion, and a new episode begins with the vehicle located at a random point of the map to provide diversified experiences for each policy update.

At every 12,288 environment interactions (steps), the agent’s architecture is updated in a central computer: the parametrized policy using loss function (7) is trained for 20 epochs using PPO (K=20), while the GAIL’s Discriminator is trained for 2 epochs (J=2) on these 12,288 samples; and the CGAN’s Generator for BEV with the loss function (1) is trained for 4 epochs, while its Discriminator for 4 epochs.

²CARLA Autonomous Driving Leaderboard available at: <https://leaderboard.carla.org/>

This process corresponds to one training cycle of the full hGAIL. A new cycle will collect the next 12,288 samples from all the actors, and execute the training as described above again. As six parallel actors are used, 2,048 steps or environment interactions per actor are recorded, totalling the 12,288 environment interactions. Thus, the episode does not have to end for a policy update to happen. It is important to note that at any given moment, any of the six actors may be interacting with the environment in different parts of the environment. Additional hyperparameters’s values can be found in Appendix VI-D.

It is worth noting that the GAN’s generator of hGAIL is pretrained on the fixed set of the ten expert trajectories (with 21,287 pairs of input and BEV targets) for 4 epochs in a supervised way. After this pretraining, the batch of the *real* BEV samples for training the GAN’s discriminator evolves with the agent’s training and corresponds exactly to the batch of 12,288 samples collected by the six parallel actors. These samples include the topdown view of the vehicle, computed by the simulator, which becomes the real BEV images which the GAN should learn to generate. This happens at every 12,288 steps executed by all actors together and, thus, the training of the GAN for generating a mid-level input representation is accomplished as the agent interacts with the environment, similarly to a strategy for decoupling representation learning from reinforcement learning in [38]. The idea is that, once the GAN’s training is turned off, the predicted BEV from the GAN’s generator could be used in new city where the real BEV is not available. The detailed algorithm for hGAIL training is shown in algorithm (1). It is relevant to notice that the training of the BEV representation will benefit from the online training samples collected as the agent interacts with the environment. In this setting, many examples are generated where the agent deviates from the correct trajectory, automatically enriching the training of the BEV network.

C. Evaluation

The main evaluation environment is *town2*, shown in Appendix VI-F. It was used to test how well the hGAIL agent can generalize its driving skills to unseen, new environments. All experiments below consider agents trained exclusively in *town01* environment.

The training progress can be seen in Fig. 4 for *hGAIL*, *GAIL* w/ *real BEV*, and *GAIL* from *cameras* agents. The second agent is trained directly on the real Bird’s-Eye View image computed from the simulator, while the last one is trained with input coming directly from the three frontal cameras and sparse trajectory, disregarding any Bird’s-Eye View representation. The plot shows the average and standard deviation of the number of infractions for three runs for each agent’s stochastic policy.

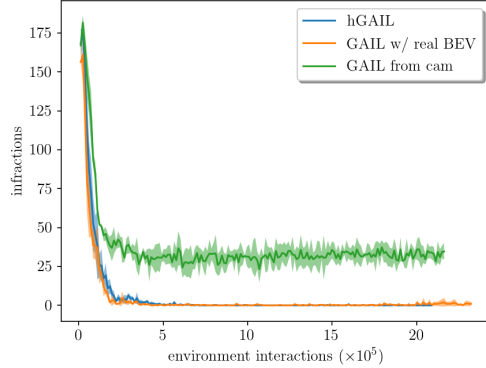
The resulting deterministic policies of each agent trained in *town1* for all three runs are also evaluated in *town2* as training evolves, as shown in Fig. 5. It shows the average percentage of completed routes from a total of six Leaderboard routes in *town2* as learning proceeds. Each run uses a different

Algorithm 1: Hierarchical Generative Adversarial Imitation Learning

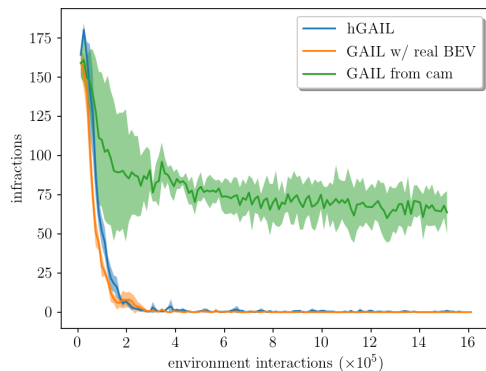
Input: Expert transitions buffer \mathcal{B}_E , T , J , K , L ;
Input parameters: Actor θ , Critic ϕ , Discriminator ω ;
 Pretrain CGAN BEV network $G(\cdot)$ w/ samples from \mathcal{B}_E ;
for $episode = 1, 2, \dots$ **do**
 for $t = 1, 2, \dots, T$ **do** // Collect env. samples
 // x_t are CGAN $G(\cdot)$ ’s inputs; y_t is
 the true BEV
 Choose action $a_t \sim \pi_\theta(G(x_t))$; $v_t \leftarrow V_\phi(G(x_t))$;
 $x_{t+1}, y_{t+1} \leftarrow act(a_t)$;
 Add $(x_t, y_t, G(x_t), a_t, v_t)$ to the buffer \mathcal{B}_π ;
 end
 for $j = 1, 2, \dots, J$ **do** // Update GAIL discrim.
 Sample $\{(y^{(j)}, a^{(j)})^\pi\}_{i=1}^m$ and $\{(y^{(j)}, a^{(j)})^E\}_{i=1}^m$
 from policy transitions buffer \mathcal{B}_π and expert
 transitions \mathcal{B}_E , respectively;
 Update the policy discriminator parameters w to
 increase (2);
 end
 Compute advantage $A_{t \in \{1, 2, \dots, T\}}$ and add to policy
 transitions buffer \mathcal{B}_π ;
 for $k = 1, 2, \dots, K$ **do** // Update agent w/ PPO
 Sample $\left\{ \left(G(x^{(k)}), a^{(k)}, A^{(k)} \right) \right\}_{i=1}^m$ from \mathcal{B}_π ;
 Update policy parameters θ to minimize (7);
 end
 Train CGAN BEV network w/ samples from \mathcal{B}_π for L
 epochs (both generator $G(\cdot)$ and its discriminator);
end

agent trained exclusively in *town1*. Both hGAIL and GAIL with real BEV are able to generalize the learning in *town1* to *town2*. Besides, the ablation on hGAIL agent of the visual input of the sparse trajectory negatively affect the evaluation performance. In this context, it is worth noticing that if we would have pretrained its GAN network for more epochs on the expert dataset, we would improve the performance of the *hGAIL ablated* (results not shown). This means that the sparse trajectory given as extra visual stimuli makes the GAN train faster.

In addition to the above three agents, other two agents were tested: the Behavior Cloning (BC) and GAIL from cameras agents. The BC agent consists of basically substituting the GAIL policy for a BC policy, which receives the BEV prediction from a GAN. Both BC and its associated GAN were trained on the same set of ten demonstration trajectories in an offline manner. Both of these agents are not shown in the plot, since they fail to learn to complete any route (staying at 0% if shown in Fig. 5). After training, the agent was also evaluated at a given T intersection and compared to the target given by the expert. Fig. 6 shows the resulting trajectories, with blue and orange denoting the agent’s and expert’s trajectories, respectively. It is worth noting that the policy’s network in the hGAIL agent receives as input only the generated (fake) BEV mid-level image, the current speed, and last applied actions for throttle and steering. For instance, this BEV image corresponds to the topdown image with three channels from Fig. 1(c). It is important to observe that the only information denoting the desired movement for the agent comes from the yellow desired route in the drivable red area.



(a) With visual trajectory input

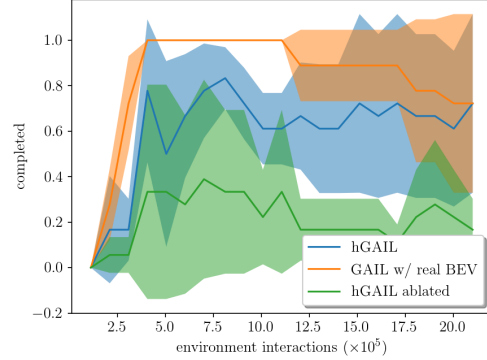


(b) Without visual trajectory input

Fig. 4. Number of committed infractions vs. environment interactions during training in *town1* environment. The top (bottom) plot shows the results for the agents receiving (disregarding) the sparse trajectory as visual input (*trajectory 1x192x192* in Fig.10) in the corresponding CNN. For each method (hGAIL, GAIL with real Bird’s-Eye View, GAIL from cameras), the average performance of three runs is depicted considering a stochastic policy. The shaded area represents the standard deviation. The GAIL from cameras agent fails to learn the task and keep the sum of committed infractions close to zero, while the goal of zero infractions is achieved by both hGAIL and GAIL with real BEV.

This yellow route occupies the whole lane in the BEV image, which leaves open how the agent will learn to turn at certain intersections. In other words, the agent’s policy can not see directly the points in the sparse trajectory, as these points are fed to the GAN part of the architecture and not to the policy. This means that how we terminate the episode, such as through infractions and lane invasion, will influence to a great extent the type of behavior the agent learns. Such an example can be seen in the turns of Fig. 6, where the agent’s trajectory does not match exactly with the expert’s one. We conjecture that the behavior of the agent could have been made more similar to the expert if: the network of the agent’s policy would have received all the information that the expert receives as input (e.g., dense trajectory); the infractions would have been more strict; additional reward terms would have been added in the reward function to punish certain behaviors.

The agent trained only on *town1* was also evaluated at every T intersection in *town2* environment, i.e., 8 different T inter-



(a) Evaluation

Fig. 5. Evaluation of agents in *town2*, trained exclusively in *town1*. The plot shows the percentage of completed routes from a total of six Leaderboard routes in *town2* vs. environment interactions, averaged over three different runs, where each run entails a different agent trained only in *town1*. For each method (hGAIL, GAIL with real Bird’s-Eye View, hGAIL ablated), the average performance of three runs is depicted considering a deterministic policy. The shaded area represents the standard deviation. Not shown in the plot, Behavior Cloning (BC) and GAIL from cameras agents fail to learn the task and complete any route (staying at 0% if shown in the plot). Both hGAIL and GAIL with real BEV agents are able to generalize the learning in *town1* to *town2*. The latter agent does not have to learn BEV, as it has always access to the true BEV. The *hGAIL ablated* agent receives no visual input of the sparse trajectory, but only its numeric vector.

TABLE I
EVALUATION PERFORMANCE FOR 8 T INTERSECTIONS AND 6 TYPE OF TURNS IN TOWN2

Turn type	BC	hGAIL	GAIL from cam.	hGAIL ablated
Top-right	0%	88%	0%	100%
Top-left	75%	100%	0%	75%
Right-left	50%	100%	25%	100%
Right-top	88%	100%	13%	63%
Left-right	88%	100%	13%	100%
Left-top	0%	100%	0%	88%
All types	50%(24)	98%(47)	8%(4)	88%(42)

sections, and compared to Behavior Cloning and GAIL from cameras agents. The latter corresponds to a GAIL agent with input directly from the vehicle’s three frontal cameras instead of the mid-level BEV input. The results are summarized in Table I, whose lines presents the results for each possible turn out of 6 in total at a given T intersection (as shown in Fig. 6). Thus, each turn, covering around 100 meters, was evaluated in 8 different T intersections, totalling 48 experiments for each agent. The success percentage for each turn type is given in this table, where we can see that hGAIL can turn without failing in all intersections and for all turn types except for one *top-right* intersection, while BC fails 50% of the times, and GAIL from cameras fails to learn most of the required driving behavior, succeeding only in 4 turns out of 48. This ablation of the GAN from hGAIL (which is the GAIL from cameras) shows the need for learning the mid-level input representation to succeed in this complex task. Additionally, notice that only hGAIL can finish complete trajectories (of around 1 km, as shown in Fig. 5), as both BC and GAIL from cameras fail in at least one turn type.

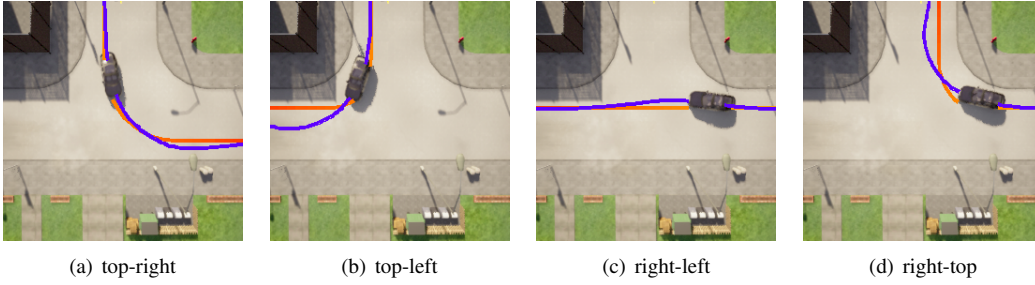


Fig. 6. Agent’s trajectories in *town2* in blue color generated by the deterministic policy after training in *town1* (at epoch 100) superimposed on the expert trajectory in orange color. At the same T intersection, 6 possible movements are possible: from top to right, top to left, right to left, right to top, left to right (omitted) and left to top (omitted).

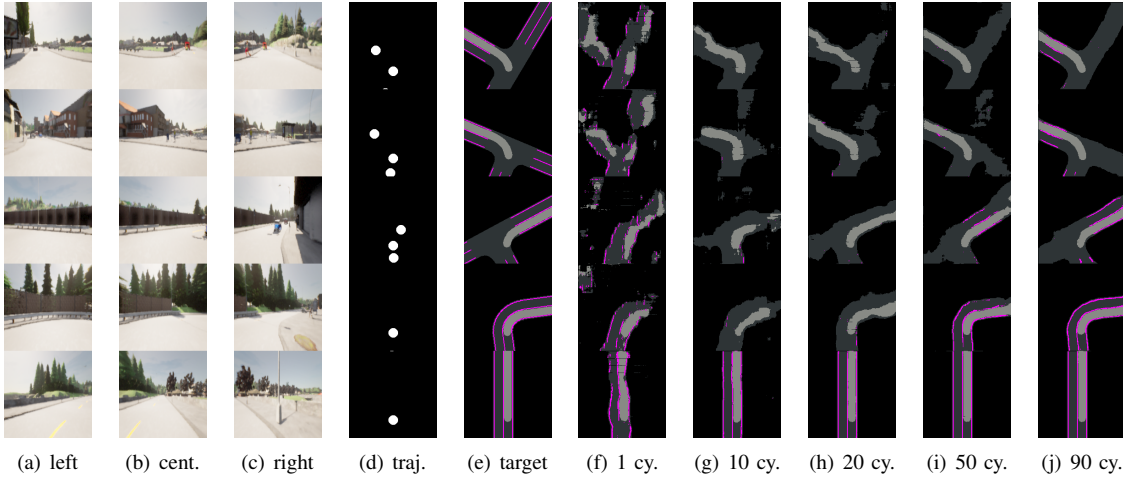


Fig. 7. BEV generation in *town2* as the agent goes through training in *town1*. The first three columns shows the images from the cameras attached to the front of the vehicle. The fourth column shows the sparse trajectory from the route planner. The fifth column shows five BEV images computed by the simulator and are considered the target output. The following columns show the BEV images generated by the GAN from the agent’s architecture as it undergoes training, at: 12,288 environment steps (1 cycle), 122,880 environment steps (10 cycles), 245,760 environment steps (20 cycles), 614,400 environment steps (50 cycles), and 1,105,920 environment steps (90 cycles). One cycle is similar to the concept of epoch, and consists of the full training of hGAIL using the last 12,288 steps collected; see Section V-B for more details.

D. Mid-level representation learning

Here, we present some results of the Bird’s-Eye-View representation learned by the GAN’s generator of the hGAIL agent’s architecture. While this GAN learns in *town01*, Fig. 7 shows the evolution of the representations of five different vehicle’s positions taken in *town02* at different training epochs of the agent in *town1*, where each row corresponds to a different particular position of the vehicle in the *town02* environment. The first four columns are the input to the GAN’s generator, consisting of the images from the three frontal cameras and the sparse trajectory given as an image. The fifth column corresponds to the targets (labels), i.e., the BEV generated by the simulator, which is used to train the GAN’s generator. The other columns show the mid-level BEV representation evolving from a poor prediction at cycle 1 (after 12,288 environment steps) to a good enough prediction at cycle 32. It is worth noticing that the GAN has never seen *town02*, and was trained only on *town01*.

E. Dynamic environment

Here, we extend the hGAIL agent’s architecture to allow for environments with traffic lights, pedestrians and other vehicles.

Initially, to make the agent able to learn to respect traffic lights, it was necessary to employ importance weighting in the BC term (\mathcal{L}_{BC}) of the loss function in (7) so that each sample in the expert set is weighted according to its importance. These weights were found by computing the inverse of the output of a Kernel Density Estimator (KDE) applied on the target outputs (action values), so that less frequent actions (e.g., accelerating as the light turns green) are given more weight in the BC loss. The KDE used a Gaussian Kernel with bandwidth equal to 0.2. Additionally, three extra binary inputs were used for the GAIL policy in hGAIL, denoting the detection of traffic lights, pedestrians and vehicles in front of the cameras. Also, the weight of the BC term in (7) was set to $\alpha = 0.012$, and the number of discriminator epochs (J) to 4.

The expert dataset was built similarly to the previous experiment, with the same trajectories, as described in Section V-A, except that now the expert respects traffic lights and breaks when pedestrians and other vehicles appear close enough ahead of its vehicle, which increases the dataset size. We trained two agents from scratch for each method, i.e., hGAIL, GAIL with real BEV, and GAIL from camera. The last two methods also employed importance weighting and three extra

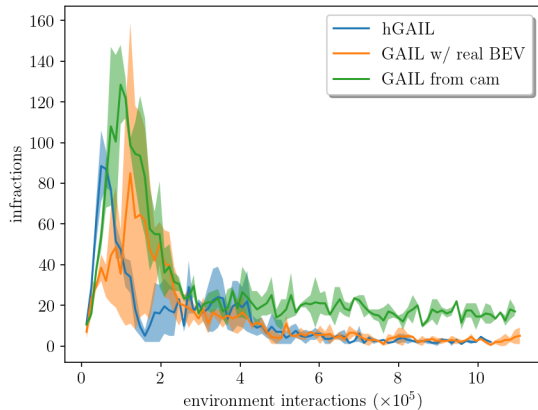


Fig. 8. Number of infractions for hGAIL, GAIL from real BEV, and GAIL from camera agents during training in *town1* with pedestrians, vehicles and traffic lights.

inputs as done in hGAIL, but processed inputs directly from the real BEV of the scene for the second method, and from the three frontal cameras for the last method. The training evolution can be seen in Fig. 8, where the first two agents converge to a low infraction rate in *town01*. The evaluation results can be seen in Table V-E for both training *town1* and test *town2* environments, showing the mean and the standard deviations over all routes in each town and the two agents trained for each method. Even though the GAIL agent has a slightly better *Route Completion*, its *Driving Score* (22%) is worse than the score (45%) of the hGAIL agent in *town02* because the former commits more infractions and collisions. It could be that the noisy predicted BEV that hGAIL uses as input to the policy makes it more robust when learning to avoid infractions in comparison to the GAIL with real BEV agent. We can also notice that the expert fails sometimes. Particularly, hGAIL surpass the expert for the *Red Light Infraction* in *town01*, and it gets similar performance to the expert in other metrics. The GAIL agent trained with input directly from camera presented the worst performance, with 28% and 5% average driving scores for *town01* and *town02*, respectively.

In Fig. 9, we can observe that the hGAIL agent can generate the three BEV channels for navigation in spite of the pedestrians and vehicles that may pass in front of the cameras. Thus, it learns to predict what is the road ahead in the BEV representation as if the vehicle or pedestrian would not be in front of the camera blocking the passage and the view.

VI. CONCLUSION

In this work, the hGAIL architecture was proposed to solve the autonomous navigation of a vehicle in the CARLA simulation environment in an end-to-end approach, connecting sensory perceptions to low-level actions directly with neural networks (sensory-motor coupling), while learning mid-level input representations of the agent’s environment. hGAIL is an hierarchical Adversarial Imitation Learning architecture composed of two main modules: the CGAN which generates the Bird’s-Eye View (BEV) representation from the three

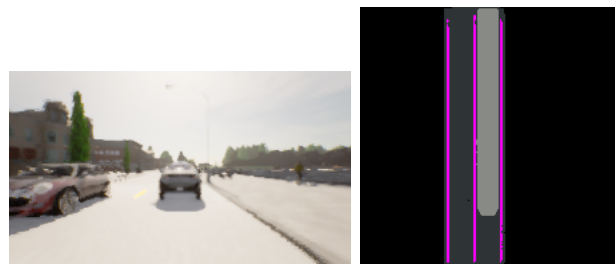


Fig. 9. BEV generation in a scene with two cars captured by the vehicle’s camera. Left: central camera image input to hGAIL. Right: predicted BEV output by the CGAN in hGAIL.

frontal cameras of the vehicle, desired trajectory and high-level command, which is a mid-level (more abstract) input representation of the scene in front of the vehicle; and the GAIL which learns to control the vehicle based mainly on the BEV predictions from the CGAN as input.

Both GAIL and CGAN in hGAIL learns simultaneously to control the agent and generate the input representations, respectively. The learning takes place in an urban city initially without pedestrians or other cars, but with dynamic routes that can change the path on the fly. In the following experiment, hGAIL is extended so that the agent can learn to respect traffic lights and deal with the detection of pedestrians and vehicles. Our experiments have shown that the mid-level input generated by CGAN is essential for the learning task as the GAIL exclusively from cameras (without BEV) fails to even learn the task, keeping a high-infraction rate through training. The BC agent with its associated GAN can complete some turns in the trajectories, but not consistently as hGAIL can. In fact, hGAIL, after training exclusively on ten expert trajectories from one city, was able to generalize its navigation skills to a new, unseen city, and complete 98% of all 48 possible turns in the eight T intersections from that new city.

This work has demonstrated the usefulness of mid-level BEV input for realistic navigation scenarios, but also that this input representation can be learned concomitantly with the agent’s policy training. Thus, the BEV generation is learned with the same data distribution used to train the agent’s policy, i.e., the expert dataset and the data generated by the agent. In future work, we are interested in extending the method making the policy generalize to other weather conditions and other new cities; and sim2real experiments. Also, we would like to investigate more closely the rarer situations that lead to infractions, proposing a more refined architecture to handle these difficult cases, as well as improve the quality of the expert.

APPENDIX

A. Generative Adversarial Imitation Learning (GAIL)

In GAIL [14], basically, there are two components that are trained iteratively in a min-max game: a discriminative classifier D is trained to distinguish between samples generated by the learning policy π and samples generated by the expert policy π_E (i.e., the labelled training set); and the learning policy π is optimized to imitate the expert policy π_E .

TABLE II
DRIVING PERFORMANCE AND INFRACTION ANALYSIS FOR DIFFERENT AGENTS ON TOWN01 AND TOWN02.

Agent	Town	Metrics						
		Driving Score (%, \uparrow) ^a	Route Completion (%, \uparrow)	Infraction Penalty (%, \uparrow)	Collisions Vehicle (#/Km, \downarrow)	Collisions Pedestrian (#/Km, \downarrow)	Red Light Infraction (#/Km, \downarrow)	Agent Blocked (#/Km, \downarrow)
Expert	town01	86 \pm 9	98 \pm 2	87 \pm 8	0.0 \pm 0.0	0.0 \pm 0.0	0.62 \pm 0.43	0.21 \pm 0.21
	town02	46 \pm 16	82 \pm 1	60 \pm 18	0.31 \pm 0.14	0.0 \pm 0.0	2.03 \pm 1.32	1.17 \pm 0.29
hGAIL	town01	82 \pm 1	95 \pm 6	86 \pm 4	0.14 \pm 0.14	0.09 \pm 0.09	0.48 \pm 0.11	0.26 \pm 0.26
	town02	45 \pm 2	66 \pm 8	65 \pm 3	0.98 \pm 0.49	0.17 \pm 0.17	2.24 \pm 0.95	2.00 \pm 0.51
GAIL	town01	67 \pm 1	97 \pm 0	70 \pm 1	0.76 \pm 0.27	0.05 \pm 0.05	0.53 \pm 0.09	0.13 \pm 0.0
	town02	22 \pm 2	72 \pm 1	31 \pm 5	3.93 \pm 1.13	0.08 \pm 0.08	2.24 \pm 0.93	1.62 \pm 0.13
GAIL from cam	town01	28 \pm 1	50 \pm 6	61 \pm 10	0.36 \pm 0.0	0.35 \pm 0.35	2.63 \pm 0.46	8.32 \pm 4.02
	town02	5 \pm 2	10 \pm 3	48 \pm 4	10.76 \pm 6.80	0.62 \pm 0.62	0.62 \pm 0.62	18.00 \pm 3.75

^a The *Driving Score*, a metric from CARLA Leaderboard, is calculated by the product of *Route Completion* (the percentage of the route distance completed) and an *Infraction Penalty* which accounts for all violations committed. For instance, if an agent ran a red light and collided with a pedestrian during a route and the penalty for running one red light was 0.7 and for colliding with a pedestrian is 0.5, the penalty would be $0.7 \times 0.5 = 0.35$.

Thus, in this game, both D and π have opposite interests: D feeds on state-action pair (s, a) and its output seeks to detect whether (s, a) comes from learning policy π or expert policy π_E ; and π maps state s to a probability distribution over actions a , learning this mapping by relying on D 's judgements on state-action samples (i.e., D informs how close π is from π_E). Mathematically, GAIL finds a saddle point (π, D) of the expression:

$$\mathbb{E}_{\pi_E} [D(s, a)] - \mathbb{E}_{\pi} [D(s, a)] - \lambda H(\pi) - \lambda_2 L_{gp}, \quad (8)$$

where $D : S \times A \rightarrow (0, 1)$, S is the state space, A is the action space; π_E is the expert policy; $H(\pi)$ is a policy regularizer controlled by $\lambda \geq 0$ [31]. The discriminator will try to increase (8), while π seeks to minimize it; and L_{gp} is a loss that penalizes the gradient's norm, constraining the discriminator network to the 1-Lipschitz function space, according to [32]. The above equation is the Wasserstein loss, used to alleviate vanishing gradient and mode collapse problems, using the Wasserstein distance [32] between the policy distribution and expert distribution, as also done in [33], [34]. It measures the minimum effort to move one distribution to the place of the other, yielding a better feedback signal than the Jensen-Shannon divergence. Both D and π can be represented by deep neural networks. In practice, a training iteration for D uses Adam gradient-based optimization [39] to increase (8), and in the next iteration, π is trained with any on-policy gradient method such as Proximal Policy Optimization (PPO) [36] to decrease (8).

B. CGAN Network Architecture

Both CGAN networks' architectures are presented in Fig. 10, where the common layers in orange refer to layers existing in both the Generator and the Discriminator. Notice that they are separate networks which do not share parameters: the figure was made to not repeat equivalent layers when describing both networks.

Generator: It can be seen in this figure and also in Fig. 2 that the CGAN's generator is a U-Net [40], usually employed for image translation or segmentation. Further, while the image is processed by convolution layers, the other perceptual inputs

(trajectory and command, second column in the figure) are processed by two fully connected layers followed by two transposed convolution layers which upsample their input to reach the desired resolution so that it can be merged with the last orange $256 \times 10 \times 10$ layer in the left column. The next transposed convolution grey layer ($256 \times 22 \times 22$) merges information coming from the frontal cameras's RGB images (left column) and the trajectory points plus the command (right command) for the generator network. Its final output is $3 \times 192 \times 192$, corresponding to the three-channels BEV translated image.

Discriminator: The discriminator is also conditioned on the RGB images from the frontal cameras and the trajectory visual representation, which are merged to the (fake/real) BEV image, totalling $13 \times 192 \times 192$ input to the first convolutional layer of the discriminator. The other perceptual inputs (right column) are processed similarly to the generator until it merges in a new $384 \times 11 \times 11$ layer (in blue) with information coming from the images ($256 \times 10 \times 10$, left column). The final output corresponds to the one given by PatchGAN.

C. GAIL policy architecture

The agent's policy part, which corresponds to the right side of Fig. 2, has the architecture shown in Fig. 11. The discriminator layers are also shown, even though both network's weights are not shared, as in the previously presented CGAN architecture. The only shared part corresponds to the layers between the Generator and the value function $V_{\phi}(\cdot)$ until the main branch splits into two heads: one for the actions *steering* and *throttle* for the generator (with 2 *softplus* units that outputs the α and β parameters of a Beta distribution, for each action); and another for the value of state s , given by a linear unit. The discriminator $D(s, a)$ receives an action a in addition to the observation s and maps to a linear output unit, whose output value is employed as reward when training with PPO.

D. Hyperparameters

Tables III and Table IV show the hyperparameter values for the GAIL and GAN parts of hGAIL, respectively.

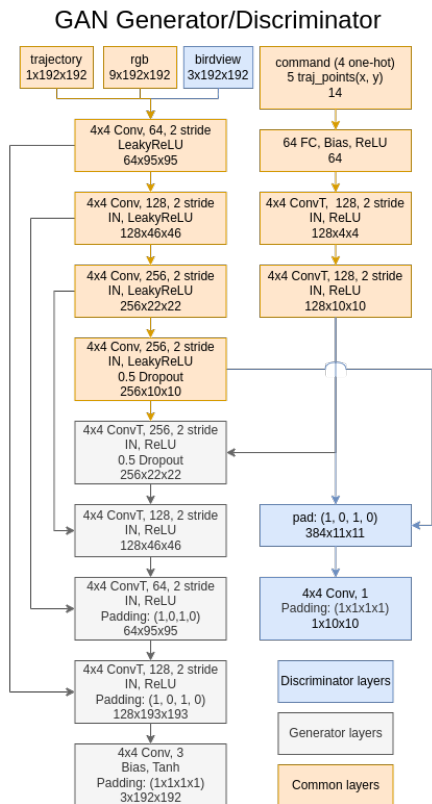


Fig. 10. Conditional GAN architecture for generating the BEV input representation. The Generator and the Discriminator are separate networks which do not share parameters: the figure was made to not repeat equivalent layers when describing both networks. The generator corresponds to the U-net at the left side of Fig. 2 and aims at translating RGB 10x192x192 images from the vehicle’s frontal cameras and sparse trajectory to BEV mid-level input representation (3x192x192 images).

TABLE III
HYPERPARAMETERS FOR GAIL

Description	Value
Parallel environments (N)	6
Initial adam step size (lr)	2.0×10^{-5}
Adam step size exponential decay (λ_{lr})	0.96
Number of PPO epochs (K)	20
Mini-batch size (m)	256
Discount (γ)	0.99
GAE parameter (λ)	0.9
Clipping parameter (ϵ)	0.2
Value Function clipping parameter (ϵ_{vf})	0.2
Value Function coefficient (c_1)	0.5
Entropy coefficient (λ_{ent})	0.01
Exploration coefficient (λ_{exp})	0.05
Timesteps per epoch (T)	12288
Weight of BC in GAIL loss (α)	0.004
Discriminator adam step size (lr)	2.5×10^{-4}
Number discriminator epochs (J)	2

E. Training progress

The evolution of training for the hGAIL agent in *town1* can also be seen in Fig. 12, where the whole trajectory throughout the city is plot at three different moments in training. Early in the training process, the infractions or errors, given by red triangles, are frequent. These infractions decrease as learning proceeds.

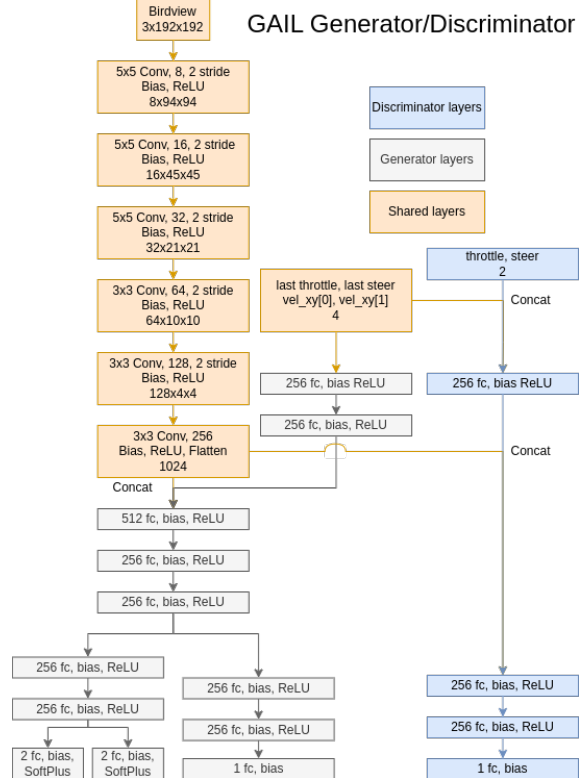


Fig. 11. GAIL architecture for policy learning, corresponding to the Generator network at the right side of Fig. 2 and the Discriminator responsible for producing the reward signal. The Generator, $\pi_{\theta}(a|s)$, receives the predicted BEV image from the GAN’s generator, the last agent’s actions (throttle, steer), and the current speed as input (which forms the observation s of the policy), and outputs the α and β parameters of the Beta distribution for both steering and action with the SoftPlus activation function. The Value function $V_{\phi}(s)$ shares the Generator network’s layers until it branches into a separate head with more two hidden layers and a linear output unit. The Discriminator $D_{\omega}(s, a)$ receives the actions *throttle* and *steer* in addition to the observation s and has a linear output unit. Notice that features from the last convolutional layer are flattened before they are merged (*concat*) with other information into FC (fully connected) layers.

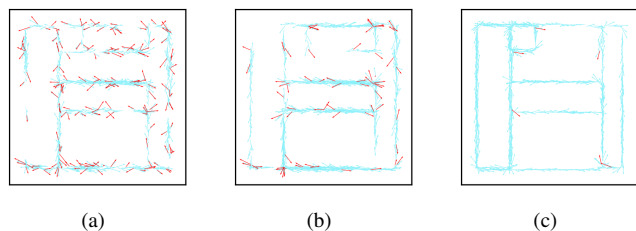


Fig. 12. The vehicle’s trajectory in *town1*, in blue, during different moments of the training process. In the early training iterations, errors, marked in red color, are common. As training proceeds, less and less mistakes happen: in (a), until 86,000 environment interactions; in (b) from 61,000 to 147,000 interactions; in (c) from 245,000 to 331,000 interactions.

TABLE IV
HYPERPARAMETERS FOR CGAN

Description	Value
Adam step size (lr)	2.0×10^{-4}
Number of GAN epochs (L)	4
Mini-batch size (m)	32
Patch size (γ)	(10, 10)
Resize (λ)	(192, 192)
Lambda pixel (ϵ)	100

F. Evaluation environment

Town02 in Fig. 13 is the environment used for testing the generalization capabilities of the trained agents.



Fig. 13. Town02 evaluation city, with one of the routes used to test agents trained in town01. The highlighted path has 1010 meters, 29 points in the sparse trajectory (shown as yellow dots) and 1030 points in the dense point trajectory (not shown).

ACKNOWLEDGMENT

The authors acknowledge CAPES (Coordination for the Improvement of Higher Education Personnel) - Brazil for financial support.

REFERENCES

- [1] M. Montemerlo, J. Becker, S. Bhat, H. Dahlkamp, D. Dolgov, S. Ettinger, D. Haehnel, T. Hilden, G. Hoffmann, B. Huhnke, D. Johnston, S. Klumpp, D. Langer, A. Levandowski, J. Levinson, J. Marcil, D. Orenstein, J. Paefgen, I. Penny, and S. Thrun, "Junior: The stanford entry in the urban challenge," *Journal of Field Robotics*, vol. 25, pp. 569 – 597, 09 2008.
- [2] H. Zhu, K.-V. Yuen, L. Mihaylova, and H. Leung, "Overview of environment perception for intelligent vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 10, pp. 2584–2601, 2017.
- [3] B. Paden, M. Cap, S. Z. Yong, D. S. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, pp. 33–55, 2016.
- [4] L. M. S. Parizotto and E. A. Antonelo, "Cone detection with convolutional neural networks for an autonomous formula student race car," in *26th ABCM International Congress of Mechanical Engineering (COBEM 2021)*, 2021.
- [5] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba, "End to end learning for self-driving cars," 2016. [Online]. Available: <https://arxiv.org/abs/1604.07316>

- [6] H. Xu, Y. Gao, F. Yu, and T. Darrell, "End-to-end learning of driving models from large-scale video datasets," 2016. [Online]. Available: <https://arxiv.org/abs/1612.01079>
- [7] M. Bansal, A. Krizhevsky, and A. S. Ogale, "Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst," in *Robotics: Science and Systems XV*, 2019, pp. 1–8.
- [8] F. Codevilla, M. Muller, A. Dosovitskiy, A. M. Lopez, and V. Koltun, "End-to-end driving via conditional imitation learning," *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1–9, 2017. [Online]. Available: <https://api.semanticscholar.org/CorpusID:3672301>
- [9] F. Codevilla, E. Santana, A. Lopez, and A. Gaidon, "Exploring the limitations of behavior cloning for autonomous driving," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 9328–9337.
- [10] S. Ross and D. Bagnell, "Efficient reductions for imitation learning," *Journal of Machine Learning Research - Proceedings Track*, vol. 9, pp. 661–668, 01 2010.
- [11] A. Prakash, A. Behl, E. Ohn-Bar, K. Chitta, and A. Geiger, "Exploring data aggregation in policy learning for vision-based urban autonomous driving," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11 760–11 770.
- [12] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, G. Gordon, D. Dunson, and M. Dudık, Eds., vol. 15. Fort Lauderdale, FL, USA: PMLR, 11–13 Apr 2011, pp. 627–635. [Online]. Available: <https://proceedings.mlr.press/v15/ross11a.html>
- [13] A. Y. Ng, S. J. Russell *et al.*, "Algorithms for inverse reinforcement learning," in *ICML*, 2000, pp. 663–670.
- [14] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Advances in Neural Information Processing Systems*, 2016, pp. 4565–4573.
- [15] G. C. Karl Couto and E. A. Antonelo, "Generative adversarial imitation learning for end-to-end autonomous driving on urban environments," in *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2021, pp. 1–7.
- [16] K. Ota, D. K. Jha, and A. Kanezaki, "Training larger networks for deep reinforcement learning," *arXiv preprint arXiv:2102.07920*, 2021.
- [17] Z. Zhang, A. Liniger, D. Dai, F. Yu, and L. Van Gool, "End-to-end urban driving by imitating a reinforcement learning coach," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 15 202–15 212.
- [18] A. Sax, J. O. Zhang, B. Emi, A. Zamir, S. Savarese, L. Guibas, and J. Malik, "Learning to navigate using mid-level visual priors," in *Conference on Robot Learning*. PMLR, 2020, pp. 791–812.
- [19] M. Mueller, A. Dosovitskiy, B. Ghanem, and V. Koltun, "Driving policy transfer via modularity and abstraction," in *Conference on Robot Learning*. PMLR, 2018, pp. 1–15.
- [20] A. Mousavian, A. Toshev, M. Fišer, J. Košecka, A. Wahid, and J. Davidson, "Visual representations for semantic target driven navigation," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8846–8852.
- [21] W. Yang, X. Wang, A. Farhadi, A. Gupta, and R. Mottaghi, "Visual semantic navigation using scene priors," in *International Conference on Learning Representations (ICLR)*, 2019, p. 2.
- [22] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 5967–5976.
- [23] R. Jena, C. Liu, and K. Sycara, "Augmenting gail with bc for sample efficient imitation learning," in *Conference on Robot Learning*. PMLR, 2021, pp. 80–90.
- [24] S. Teng, L. Chen, Y. Ai, Y. Zhou, Z. Xuanyuan, and X. Hu, "Hierarchical interpretable imitation learning for end-to-end autonomous driving," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 1, pp. 673–683, 2023.
- [25] Y. Wang, D. Zhang, J. Wang, Z. Chen, Y. Li, Y. Wang, and R. Xiong, "Imitation learning of hierarchical driving model: From continuous intention to continuous trajectory," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2477–2484, 2021.
- [26] P. Cai, S. Wang, Y. Sun, and M. Liu, "Probabilistic end-to-end vehicle navigation in complex dynamic environments with multimodal sensor fusion," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4218–4224, 2020.

- [27] P. Cai, H. Wang, H. Huang, Y. Liu, and M. Liu, "Vision-based autonomous car racing using deep imitative reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7262–7269, 2021.
- [28] A. Kuefler, J. Morton, T. Wheeler, and M. Kochenderfer, "Imitating driver behavior with generative adversarial networks," in *Intelligent Vehicles Symposium (IV), 2017 IEEE*. IEEE, 2017, pp. 204–211.
- [29] E. Bronstein, M. Palatucci, D. Notz, B. White, A. Kuefler, Y. Lu, S. Paul, P. Nikdel, P. Mougin, H. Chen *et al.*, "Hierarchical model-based imitation learning for planning in autonomous driving," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 8652–8659.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [31] M. Bloem and N. Bambos, "Infinite time horizon maximum causal entropy inverse reinforcement learning," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014, pp. 4911–4916.
- [32] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, "Improved training of wasserstein gans," 2017.
- [33] M. Zhang, Y. Wang, X. Ma, L. Xia, J. Yang, Z. Li, and X. Li, "Wasserstein distance guided adversarial imitation learning with reward shape exploration," *2020 IEEE 9th Data Driven Control and Learning Systems Conference (DDCLS)*, Nov 2020. [Online]. Available: <http://dx.doi.org/10.1109/DDCLS49620.2020.9275169>
- [34] Y. Li, J. Song, and S. Ermon, "Infogail: Interpretable imitation learning from visual demonstrations," 2017.
- [35] I. G. Petrazzini and E. A. Antonelo, "Proximal policy optimization with continuous bounded action space via the beta distribution," in *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 2021, pp. 1–8.
- [36] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *CoRR*, vol. abs/1707.06347, 2017. [Online]. Available: <http://arxiv.org/abs/1707.06347>
- [37] D. Chen, B. Zhou, V. Koltun, and P. Krähénbühl, "Learning by cheating," 2019.
- [38] A. Stooke, K. Lee, P. Abbeel, and M. Laskin, "Decoupling representation learning from reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2021, pp. 9870–9879.
- [39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [40] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.



Gustavo Claudio Karl Couto received the master's degree in automation and systems engineering at Federal University of Santa Catarina in 2023 and graduated from the Militar Institute of Engineering in 2018 with a degree in Electronic Engineering. His main areas of interest are robotics and intelligent systems. His current research focuses on applications of imitation learning algorithms in autonomous driving systems.



Eric Aislan Antonelo received the Ph.D. degree in Computer Engineering from Ghent University, Belgium, in 2011. He is currently a Faculty Member with the Department of Automation and Systems Engineering, Federal University of Santa Catarina, Florianópolis, Brazil. His research interests include imitation learning for autonomous vehicles, reservoir computing and machine learning for robotic or industrial applications.