

# Action-based Representation Learning for Autonomous Driving

Yi Xiao<sup>1</sup>, Felipe Codevilla<sup>2</sup>, Christopher Pal<sup>2</sup>, Antonio M. López<sup>1</sup>

<sup>1</sup>Computer Vision Center (CVC) and Computer Science Dpt.

at the Universitat Autònoma de Barcelona (UAB), Barcelona, Spain

<sup>2</sup>Montreal Institute for Learning Algorithms (MILA), Montreal, Canada

**Abstract:** Human drivers produce a vast amount of data which could, in principle, be used to improve autonomous driving systems. Unfortunately, seemingly straightforward approaches for creating end-to-end driving models that map sensor data directly into driving actions are problematic in terms of interpretability, and typically have significant difficulty dealing with spurious correlations. Alternatively, we propose to use this kind of action-based driving data for learning representations. Our experiments show that an affordance-based driving model pre-trained with this approach can leverage a relatively small amount of weakly annotated imagery and outperform pure end-to-end driving models, while being more interpretable. Further, we demonstrate how this strategy outperforms previous methods based on learning inverse dynamics models as well as other methods based on heavy human supervision (ImageNet).

**Keywords:** Representation Learning, Autonomous Driving, Imitation Learning

## 1 Introduction

The development of autonomous vehicles (AVs) is a significant multidisciplinary challenge. Currently, the main paradigm being pursued for developing AVs follows a traditional divide-&-conquer engineering strategy. In particular, modular pipelines are proposed with key modules for perception, route planning and maneuver control, among others [1]. In turn, these modules may be composed to deal with different tasks, *e.g.*, perception encompasses object detection and tracking, semantic class/instance segmentation, etc. [2]. These tasks rely on models trained from data using modern deep learning techniques [3]. Following such a data-driven approach is not a problem in itself since it is possible to collect petabytes of on-board data (raw sensor data, vehicle state variables, etc.) continuously, not only from fleets of AVs under development, but also from sensorized human-driven vehicles under naturalistic driving. However, in practice, the best performing models arise from supervised deep learning, and this means that the raw data must be augmented with ground truth, which is collected through time consuming and costly human annotations (*e.g.*, bounding boxes, object silhouettes, etc).

The data annotation bottleneck associated with these approaches has caused the idea of end-to-end driving [4, 5] to receive renewed interest [6, 7, 8, 9, 10, 11]. In this paradigm a deep model is trained to directly control an AV from input raw sensor data (mainly images), *i.e.*, without a clear separation between perception and maneuver planning, and without explicit intermediate perceptual tasks to be solved. In this pure data-centered approach, the supervision required to train deep end-to-end driving models does not come from human annotation; instead, the vehicle’s state variables, which can be automatically collected from fleets of human-driven vehicles, are used as supervision (*e.g.* speed, steering angle, acceleration, braking). These models are mainly trained by behaviour cloning (BC) of human driving experiences. However, despite the undeniable good performance shown by end-to-end driving models, their reliability is controversial, due in particular to the difficulty of interpreting the relationship between inferred driving actions and image content [12], as well as training instabilities [11].

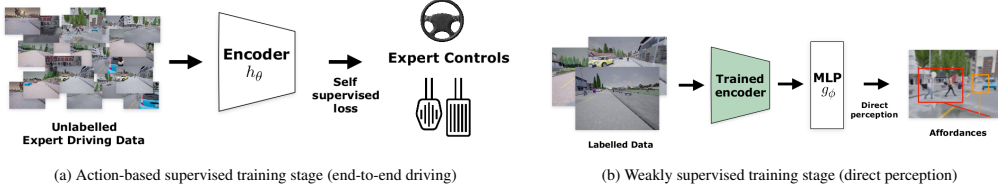


Figure 1: Approach overview: (a) an encoder is trained following an end-to-end driving setting (e.g. using BC or inverse model); (b) this pre-trained encoder together with a multi-layer perceptron (MLP) are used for predicting affordances. The affordances are used as input to a simple PID controller to drive the vehicle.

A different paradigm, conceptually midway between pure modular and end-to-end driving, is the so-called direct perception approach [13, 14], which focuses on learning deep models to predict driving affordances, from which an additional controller can maneuver the AV. In general, such affordances can be understood as a relatively small set of interpretable variables describing events that are relevant for an agent acting in an environment [15]. Driving affordances bring interpretability while only requiring weak supervision, in particular, human annotations just at the image level (*i.e.*, not pixel-wise).

In this paper, we show that action-based methods, that focus on predicting the control actions, such as end-to-end driving trained with BC, can be an effective pre-training strategy for learning a direct perception model (Fig. 1). This strategy enables a significant reduction on the number of annotated images required to train such a model. Overall, this means that we can leverage the data collected by fleets of human-driven vehicles for training interpretable driving models, thus, keeping a major advantage of modular pipelines while reducing data supervision (*i.e.*, human annotation). Further, our approach, improves over other recent pre-training proposals such as contrastive methods [16] and even over ImageNet (supervised) pre-training. We also show that learning from expert data in our approach leads to better representations compared to training inverse dynamics models using the approach in [17]. This shows that expert driving data (*i.e.* coming from human drivers) is an important source for representation learning. As is common practice nowadays, we run our experiments in the CARLA simulator [18]. To the best of our knowledge, this work is the first to show that expert demonstrations can act as an effective action-based representation learning technique. This constitutes the primary contribution of this paper.

## 2 Related Work

Since human-based data annotation is a general problem for all kinds of new data-intensive applications, not only for autonomous driving, learning representations (deep models) with the support of weak supervision and self-supervision are open challenges that attract great interest.

In the autonomous driving context, the use of driving affordances [13, 14] allows for weak supervision since only annotations at the image level are required. Based on these interpretable affordances a controller is tuned to drive. Since [14] focuses on urban driving using CARLA simulator, inspired from this work, we have defined four affordances to consider the explicit detection of hazards involving pedestrians and vehicles, respecting traffic lights and considering the heading of the vehicle within the current lane. Defining the best set of affordances to drive is not the focus of this paper, but we have chosen a reasonable set.

In order to solve visual tasks, we can find self-supervision based on auxiliary and relatively simple (pretext) tasks such as learning colorization [19], rotations [20, 21, 22], shuffling cues [23], or solving a jigsaw puzzle of image parts [24]; it has been shown that self-supervision can match traditional ImageNet (supervised) pre-training provided one works with large enough CNNs [25], although it has been argued that these proxy tasks are not sufficiently hard so as to fully exploit large unsupervised datasets [26]. Another branch of self-supervised learning is based on contrastive methods [27, 28, 29], which learn representations by comparing data pairs. While, in these methods supervision is based on different ways of transforming or comparing the input data itself, in this paper, supervision comes in the form of expert driver actions. In fact, we include in our study a recent contrastive method, ST-DIM, designed in the context of playing Atari games [16], adapted

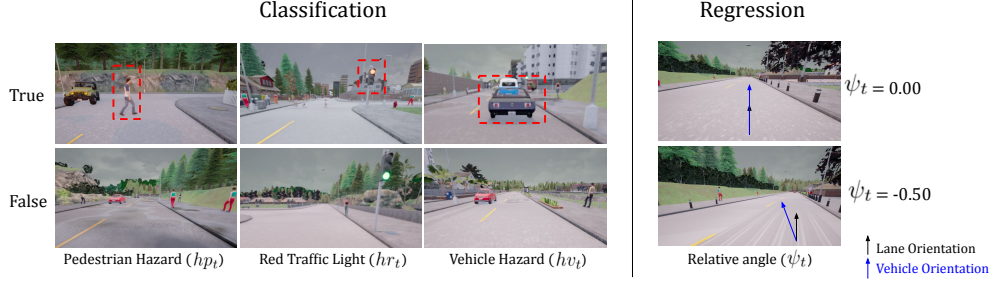


Figure 2: Our affordances illustrated on images from CARLA. Classification ones are binary variables (t/f), and regression runs on  $[-\pi, \pi]$ rad. See Sect. 3.3 for details.

to actions required for driving. We will see how action-based representation learning outperforms ST-DIM as a representation learning strategy to infer affordances.

In fact, in a perceive-&-act context, dynamics learning [30, 31] and inverse dynamics [17, 32, 33] can be used as action-based supervision strategy. Broadly speaking, being able to predict the next states of an agent, or the action between state transitions, yields useful representations. This action-centric approach to supervision is in line with our work. Thus, our study includes experiments with different inverse and forward dynamics supervision strategies. We show the importance of those strategies in the autonomous driving context. However, different than previous work, we empirically demonstrate that expert actions yields a better representation learning than random actions used in [17].

Finally, it is also worth mentioning teacher-student strategies [34, 35] which allow one to train an end-to-end driving student model from a teacher model. In this case, even the student is end-to-end, the data annotation bottleneck arises during the supervised training of the teacher, which requires bounding boxes and/or semantic segmentation. Since the student is still an end-to-end driving model, the issue of interpretability once this model is deployed in the AV still would remain open.

### 3 Action-based Representation Learning

#### 3.1 Overall Approach

As can be seen in Fig. 1, we study our action-based representation learning strategy by learning affordances in two stages. The first stage relies on non-manually annotated data to learn a representation (encoder). We will consider different methods to learn this representation (Sect. 3.2), all of them based on predicting driving actions from on-board data. The second stage uses this pre-trained representation together with a multi-layer perceptron (MLP) to learn the considered affordances (Sect. 3.3).

Therefore, for the first stage of our approach, we assume that we have access to a sequence of  $N_u$  data samples  $\mathcal{D}^u = \{d_t\}_{t=1}^{N_u}$ , which have been acquired on-board a human-driven sensorized vehicle but without human annotations. Thus, we have  $d_t = \{o_t, a_t\}$ , where  $o_t$  and  $a_t$  are respectively, at a certain time  $t$ , the observation acquired by the vehicle’s sensors and the driving action taken by the expert driver. We must understand that  $a_t$  is the expert reaction to the environment when  $o_t$  was acquired. In general, we will have different  $\mathcal{D}^u$  sequences acquired at different driving runs, however, without losing generality and for the sake of keeping a simpler notation, we can assume all of them appended in one single sequence. In this paper, we assume that each observation  $o_t$  contains an *image* capture of the driving environment, the *vehicle speed* ( $v_t$ ) at the moment the image is acquired, and a high level *navigation command* such as *continue* in the same lane, or in the next intersection *go straight/left/right*, *i.e.* as introduced in the so-called conditional imitation learning [10] in the form of one-hot vector  $c_t$ . The corresponding action  $a_t$  is defined in terms of the *steering angle*, *acceleration*, and *break* values that must be applied to maneuver the vehicle.

For the second stage of our approach, we assume a relatively small dataset of on-board images with image-level affordance annotations (weak supervision), *i.e.*,  $\mathcal{D}^l = \{x_j\}_{j=1}^{N_l}$ , with  $x_j = \{o_j, y_j\}$ ,

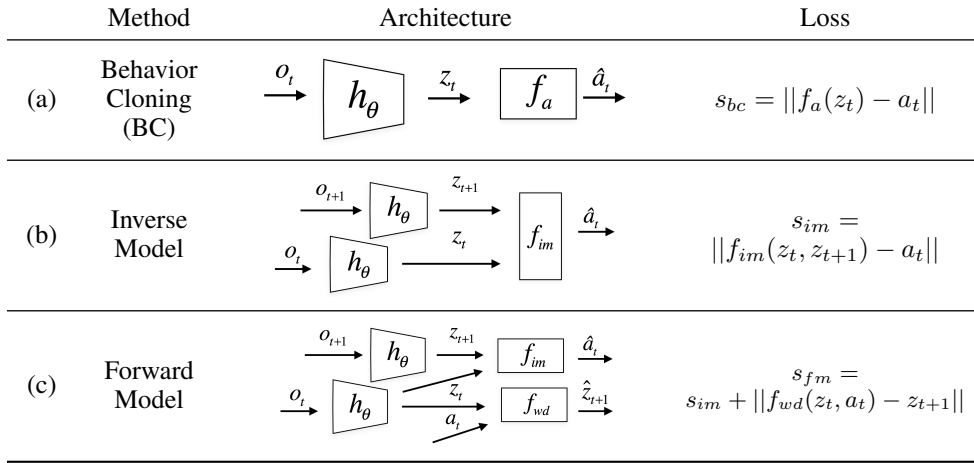


Figure 3: Proposed action-based supervised losses based on expert actions. To train the encoder  $h_\theta$ , these are minimized over the dataset  $\mathcal{D}^u$ .

being  $o_j$  the observation and  $y_j$  the corresponding affordance annotation. In particular,  $y_j$  contains variables indicating situations such as a *pedestrian hazard*, a *vehicle hazard*, a *red traffic light*, and a *relative heading angle* (Fig. 2). In this setting, we can assume that the images used in  $\mathcal{D}^l$  come from sub-sequences of  $\mathcal{D}^u$ , but selected so that  $N_u \gg N_l$ .

Accordingly, the two stages can be summarized as follows: (1) use  $\mathcal{D}^u$  to train a deep encoder  $h_\theta$ ; (2) use  $h_\theta$  and  $\mathcal{D}^l$  to train a projection network,  $g_\phi$ , for predicting affordances. For actual driving, we develop a controller  $C : g_\phi(h_\theta(o_t)) \rightarrow \hat{a}_t$ ; i.e., given the affordances  $g_\phi(h_\theta(o_t))$  predicted from an observation  $o_t$ ,  $C$  estimates the action  $\hat{a}_t$  to maneuver the vehicle. In order to show driving results, we will use a simple PID controller.

### 3.2 Action-based Supervised Stage

At the action-based supervised stage the objective is to train an encoder  $h_\theta$  to produce a set of features  $z_t$  (encoder’s bottleneck) given an input observation  $o_t$ . With this purpose, we have studied some alternatives illustrated in Fig. 3. At training time, all of them rely on  $\mathcal{D}^u$ , but they use different inputs and losses to be minimized. We summarize these alternatives in the following.

**Behavior cloning (BC).** Common deep architectures trained by BC consist of an encoder  $h_\theta$  extracting features  $z_t$  from observations  $o_t$  (i.e.  $z_t = h_\theta(o_t)$ ), and a fully-connected projection network,  $f_a(z_t)$ , which predicts an expert action  $\hat{a}_t$  from such features. In this paper, we follow the architecture presented in [11]; however, instead of using the high level command ( $c_t$ ) for branching to different projection functions, since our goal is to pre-train a useful representation not performing actual driving with it, we use  $c_t$  as part of  $o_t$ . Therefore, the encoder of [11] is modified to input the high level command the same way as the speed variable ( $v_t$ ). The input image is processed by a ResNet34 backbone. The following alternatives, also rely on this encoder architecture.

**Inverse model.** By considering not only  $o_t$  but also the subsequent observation  $o_{t+1}$  as input, we turn BC into an Inverse model [17, 32]. In this case, thinking of the encoder’s bottleneck as encoding an agent (driver) internal state, the problem to solve consists of predicting the action that transforms the state  $z_t$  into the state  $z_{t+1}$ . Differently than BC, with an inverse model the actions can come from either an expert driver or just random roaming (or poor driving).

**Forward model.** In this case, we want to learn an encoder that is able to output a state  $z_t = h_\theta(o_t)$  such that, given an action  $a_t$ , we can predict the future state as  $z_{t+1} = f_{wd}(z_t, a_t)$ . However, this can lead to the trivial solution  $z_t = \mathbf{0}$  from which  $f_{wd}$  can still produce  $z_{t+1}$ . Such degenerated encoders  $h_\theta$  are of course not of interest. Therefore, we use the regularization strategy of [17], consisting of adding also the Inverse model so that the encoded observation  $z_t$  is able to predict the action too. Again, the actions can come from either an expert driver or just random roaming.

### 3.3 Weakly Supervised Stage: Learning Affordances

The affordances used in this paper (Fig. 2) consider explicit detection of hazards involving pedestrians and vehicles, respecting traffic lights and considering the heading of the vehicle within the current lane. More specifically, we have considered the following four affordances:

**Pedestrian hazard** ( $hp_t$ ). This variable is set to one if there is a pedestrian in our lane at a distance lower than 10 m; otherwise, is set to zero.

**Vehicle hazard** ( $hv_t$ ). This variable is set to one if there is a vehicle in our lane at a distance lower than 10 m; otherwise, is set to zero. Vehicles refers to cars, vans, motorbikes, and cyclists.

**Red traffic light** ( $hr_t$ ). This variable is set to one if there is a traffic light in red affecting our lane at a distance lower than 10 m; otherwise, is set to zero.

**Relative heading angle** ( $\psi_t$ ). This variable accounts for the relative angle of the longitudinal vehicle axis with respect to the lane in which it is navigating. The variable runs on  $[-\pi, \pi]$  rad with  $\psi_t = 0$  when vehicle and lane are aligned (no matter the lateral vehicle position within the lane).

Note that  $\{hp_t, hv_t, hr_t\}$  are binary variables, so predicting them is a binary classification problem, while  $\psi_t$  is a real number, thus, predicting it is a regression problem. These binary variables are critical to perform stop-&-go maneuvers by any controller relying on these affordances, while the regressed angle is critical to properly navigating without going out of the lane. Overall, the idea behind these affordances is that relevant visual competences for driving emerge when training the corresponding models; for instance, some kind of pedestrian and vehicle detection, red traffic light detection, and localization of the vehicle within the lane for proper navigation.

The affordance prediction model,  $g_\phi$ , that predicts  $\{hp_t, hv_t, hr_t, \psi_t\}$  is obtained by training a MLP, which receives the output from the pre-trained  $h_\theta(o_t)$  as input.

## 4 Experimental Results

**Environment.** As in most recent works addressing autonomous driving, we perform our experiments and data collection in the CARLA simulator [18], in particular, using version 0.9.6. We rely on the widely used Town01 as training and Town02, the new town, for testing, from now on denoted as T1 and T2. More details about the environment and the benchmarks are provided on the supplementary material.

**Dataset.** In order to collect the action-based supervised dataset ( $\mathcal{D}^u$ ) and the weakly supervised dataset ( $\mathcal{D}^l$ ), we modified the default CARLA’s autopilot for not only recording  $o_t$  and  $a_t$ , but also our image-level affordances ( $y_t$ ). We collected  $\sim 50$  hours of image sequences in T1 for training purposes, balancing the training weather conditions, at 20 fps. In this data collection process, we have three cameras, a forward-facing (central) camera from which we will drive at testing time, and two lateral cameras only used for training purposes as in [6, 10]. Thus, in terms of samples to train  $h_\theta$ , we have  $N_u \sim 108,000,000$ . This dataset plays the role of  $\mathcal{D}^u$ , while to play the role of  $\mathcal{D}^l$  we selected  $\mathcal{D}^u$ ’s sub-sequences corresponding to 1% and 10% of the total amount. In this case, we only consider images acquired by the central camera; thus, totalling  $N_l \sim 36,000$  for 1% and  $N_l \sim 360,000$  for 10%. These sub-sequences were selected semi-randomly to ensure that jointly form a dataset where the relative heading angle approximates a Gaussian distribution centered at  $\psi_t = 0$ . Finally, for testing purposes, two new datasets were collected, namely, by driving  $\sim 1$  hour in T1 and also  $\sim 1$  hour in T2. This driving was balanced among all weather conditions, and only the central camera is considered; thus, for each town we have  $\sim 72,000$  images.

**Baselines.** In Sect. 3.2 we have presented the action-based pre-training strategies for  $h_\theta$  that we want to study. In addition, we have incorporated ST-DIM [16], a contrastive representation learning baseline used by agents playing Atari games. We have modified the code provided by the authors just to include ResNet34 as backbone, *i.e.* as for the rest of pre-training strategies. In short, ST-DIM is trained to answer if two frames are consecutive or not, without any action-related information involved. Moreover, for the Inverse, Forward, and ST-DIM strategies, we have included seldom variants which require to collect additional  $\sim 20$  hours of image sequences in T1. However, in this case, instead of relying on our expert driver autopilot, the driving was *random*; thus, eventually run-



Pre-training	Binary Affordances			Relative Angle ( $\psi_t$ )		
	Pedestrian ( $hp$ )	Vehicle ( $hv$ )	Red T.L. ( $hr$ )	Left Turn	Straight	Right Turn
No pre-training	26 $\pm$ 0	50 $\pm$ 1	42 $\pm$ 0	11.38 $\pm$ 0.18	1.85 $\pm$ 0.03	24.68 $\pm$ 0.03
ImageNet	37 $\pm$ 2	75 $\pm$ 0	47 $\pm$ 0	11.69 $\pm$ 0.57	2.83 $\pm$ 0.07	25.55 $\pm$ 0.10
Contrastive (ST-DIM)	47 $\pm$ 1	53 $\pm$ 0	53 $\pm$ 0	10.43 $\pm$ 0.21	2.62 $\pm$ 0.03	18.75 $\pm$ 0.23
Forward	<b>50 <math>\pm</math> 0</b>	63 $\pm$ 0	60 $\pm$ 0	5.35 $\pm$ 0.03	0.52 $\pm$ 0.00	6.61 $\pm$ 0.03
Inverse	49 $\pm$ 0	78 $\pm$ 0	70 $\pm$ 0	<b>3.57 <math>\pm</math> 0.03</b>	<b>0.46 <math>\pm</math> 0.00</b>	<b>3.78 <math>\pm</math> 0.06</b>
Behavior Cloning (BC)	47 $\pm$ 0	<b>81 <math>\pm</math> 0</b>	<b>75 <math>\pm</math> 0</b>	4.89 $\pm$ 0.03	1.24 $\pm$ 0.03	6.25 $\pm$ 0.10

Table 1: Linear probing results. Left: F1 score for the binary affordances (higher is better). Results are scaled by 100 for visualization purpose. Right: MAE of the relative angle (lower is better), shown for different navigation maneuvers. MAE is shown in degrees for an easier understanding.




Figure 4: **Attention heatmaps.** Top: RGB input images from a town unseen during training. Mid: attention heatmaps of ImageNet pre-trained encoder. Bottom: attention heatmaps of a BC encoder pre-trained with 50 hours of expert driving data. From left to right, we show cases involving different affordances: pedestrian hazard, vehicle hazard, red traffic light detection and navigation.

ning into accidents, driving over the sidewalk, in the wrong lane, etc. In short, navigating by random actions. As additional baselines, we have used ImageNet and no pre-training (random initialization).

**Training details.** We train the encoder  $h_\theta$  for 100K iterations (mini-batches) using  $\mathcal{D}^u$ . Then, using  $\mathcal{D}^l$  we train the affordance prediction model,  $g_\phi$ , for 20K iterations, no matter if this stage rely on linear classification/regression or fine-tuning. These iteration values were found by preliminary experiments where we monitored training convergence. When  $\mathcal{D}^l$  assumes 10% of the dataset, we iterated 100K. In all cases, we used ADAM optimizer with an initial learning rate of 0.0002 and batch size of 120. Moreover, for any kind of training, after 75K iterations the learning rate becomes half. For obtaining reliable results, we repeat both the encoder and affordance training over three different random seeds, and pick up the model with best performance on training town for driving. As observed in [11] we saw a meaningful random seed variation for training the encoders.

**Controller.** In order to perform driving evaluations, we have tuned a PID controller that takes the estimated affordances as input and outputs the action commands ( $a_t$ ) to control the AV. Given perfect affordances (*i.e.* those annotated as ground truth), this controller has been tuned to drive well in T1 (lateral and longitudinal control).

## 4.1 Results

**Linear probing.** We start by evaluating the representation learning capabilities of action-based methods using the commonly applied linear probing technique [36, 28, 29, 37]. More specifically, using a frozen  $h_\theta$  as feature extractor, we train a linear classifier to predict affordances with an affordance dataset  $\mathcal{D}^l = 1\%$ . Each trained model is tested in the  $\sim 1$  hour testing set for T2, not seen during training. In order to assess the performance of binary-affordance models we use F1-score, while for assessing the performance of the relative heading angle, we use MAE. For better analysis, we divide the relative heading angle into three cases, left turn, straight and right turn, according to the navigation situation. We consider as left regime those cases where the relative heading angle ground truth is lower than  $-0.1$  rad, as right regime when it is larger than  $0.1$  rad, and

Pre-training	Binary Affordances			Relative Angle ( $\psi_t$ )		
	Pedestrian ( $hp$ )	Vehicle ( $hv$ )	Red T.L. ( $hr$ )	Left Turn	Straight	Right Turn
No pre-training	26 $\pm$ 0	50 $\pm$ 1	42 $\pm$ 0	11.38 $\pm$ 0.18	1.85 $\pm$ 0.03	24.68 $\pm$ 0.03
Contrastive (ST-DIM)	41 $\pm$ 0	62 $\pm$ 1	63 $\pm$ 1	9.01 $\pm$ 0.46	2.77 $\pm$ 0.18	18.37 $\pm$ 0.45
Contrastive Random (ST-DIM)	39 $\pm$ 1	<b>73 <math>\pm</math> 1</b>	47 $\pm$ 0	9.70 $\pm$ 0.41	2.98 $\pm$ 0.11	15.89 $\pm$ 0.41
Forward	<b>50 <math>\pm</math> 0</b>	51 $\pm$ 0	58 $\pm$ 0	4.87 $\pm$ 0.00	0.52 $\pm$ 0.00	6.07 $\pm$ 0.06
Forward Random	20 $\pm$ 1	38 $\pm$ 0	16 $\pm$ 0	11.54 $\pm$ 0.03	1.20 $\pm$ 0.00	19.14 $\pm$ 0.00
Inverse	45 $\pm$ 0	66 $\pm$ 0	<b>73 <math>\pm</math> 0</b>	<b>3.02 <math>\pm</math> 0.03</b>	<b>0.42 <math>\pm</math> 0.03</b>	<b>5.06 <math>\pm</math> 0.17</b>
Inverse Random	26 $\pm$ 0	49 $\pm$ 0	59 $\pm$ 0	8.50 $\pm$ 0.53	1.45 $\pm$ 0.03	13.14 $\pm$ 0.34

Table 2: Linear probing results comparing encoders trained with random policy training data versus expert demonstration data. Note that, to provide fair comparison, the encoders here were trained with 20 hours data, which are different than the ones from Table 1.

as straight regime otherwise. For each pre-training strategy, we repeat linear classifier training with three random seeds, and compute its mean and standard deviation in Tables 1 and 2.

Table 1 shows the F1/MAE scores for the affordances prediction. Note that those results consist of zero-shot generalization to an unseen town (T2). The main observation is that action-based pre-training (Forward/Inverse/BC) outperforms all the other reported pre-training strategies. However, we see that the action-based pre-training is mostly beneficial to help on reliably estimating the vehicle’s relative angle with respect to the road. The contrastive method, ST-DIM, shows promising results for the binary affordances but results on very poorly relative angle estimations. We also observed a poor generalization capability for ImageNet pre-training. These results suggest that a useful scene representation is learned by training encoders with expert demonstration data. Additional evidence is presented in Fig. 4, which shows examples of attention heatmaps from an ImageNet encoder, and a BC encoder that was trained with 50 hours of data. These attention maps are calculated by the simple average of the feature maps from the third ResNet34’s block. We can see that the necessity to imitate expert demonstrations creates activations on useful objects such as pedestrians, vehicles, traffic lights and lane markings; which is in agreement with the fact that a linear classifier can predict well the set of affordances with the action-based pre-training.

We also further study if the source data needs to come from expert driving or from random actions as in [17]. The general intuition is that the Inverse model can learn a good representation by learning the dynamics of the scene. We show on Table 2 that inverse model can, indeed, outperform the no pre-training condition even when using random actions. However, we show that there is a lot more benefit for representation learning obtained from the expert action information than from random action. For ST-DIM, as expected, the difference between random and expert policy is smaller since it is not based on action.

**Driving results.** We evaluate the driving performance of our method on the CARLA NoCrash benchmark [11], which mainly focuses on the capabilities of models to drive under the presence of pedestrians and vehicles. The objective is to complete a set of goal-oriented episodes without crashing. The CARLA simulator provides high level planning commands to navigate towards a targeted town location (goal) from current one; these commands are part of the observations used to train  $h_\theta$ . For this evaluation, we updated the CARLA NoCrash benchmark to CARLA 0.9.6 version augmented with the new pedestrian crossing algorithms (recently incorporated to last CARLA version). For driving, we fine tune the whole network using three layers as the projection  $g_\phi$ . We report the success rate (higher is better) on the driving tasks and the percentage of traffic lights crossed in red (lower is better). For each model, we repeat driving for three times, and compute its mean and standard deviation in Tables 3 and 4.

Table 3 compares the performance of the action-based methods with the baselines and other methods from the literature. For all implemented methods, we considered fine-tuning with  $\mathcal{D}^l = 10\%$ . Firstly, we see that Inverse model and BC are the best representation learning strategies to pre-train the encoder, specially in new town. Both models also clearly outperform the contrastive-based baseline (ST-DIM) in new town. However, in the training town the contrastive method obtained relevant results, specially under dense traffic. As reference, we report results from Learning by Cheating (LBC) [34] and the CILRS method [11], presented on the bottom of Table 3. Shown results are copied from the corresponding papers. Our proposed approaches also achieved very close results to the LBC method [34] and outperforms CILRS specially when reacting to traffic lights. Note that the LBC method requires high supervision for training a teacher network, which teaches a student

Technique	Empty	Training Town			T.L.	New Town			T.L.
		Regular	Dense			Regular	Dense		
No pre-training	78 $\pm$ 4	79 $\pm$ 7	48 $\pm$ 4	12 $\pm$ 1	57 $\pm$ 1	37 $\pm$ 4	10 $\pm$ 1	18 $\pm$ 2	
Image Net	86 $\pm$ 1	<u>89 <math>\pm</math> 3</u>	66 $\pm$ 1	12 $\pm$ 0	21 $\pm$ 3	19 $\pm$ 3	13 $\pm$ 5	23 $\pm$ 2	
Contrastive (ST-DIM)	73 $\pm$ 2	84 $\pm$ 4	<b>71 <math>\pm</math> 3</b>	10 $\pm$ 1	66 $\pm$ 5	49 $\pm$ 2	17 $\pm$ 0	21 $\pm$ 1	
Forward	68 $\pm$ 3	84 $\pm$ 3	68 $\pm$ 8	<u>9 <math>\pm</math> 1</u>	49 $\pm$ 5	37 $\pm$ 3	18 $\pm$ 5	13 $\pm$ 2	
Inverse	73 $\pm$ 4	82 $\pm$ 2	61 $\pm$ 3	<b>8 <math>\pm</math> 1</b>	<b>83 <math>\pm</math> 2</b>	<b>67 <math>\pm</math> 8</b>	<b>26 <math>\pm</math> 8</b>	<b>8 <math>\pm</math> 0</b>	
Behavior Cloning (BC)	<b>91 <math>\pm</math> 1</b>	<u>91 <math>\pm</math> 4</u>	68 $\pm$ 5	<b>8 <math>\pm</math> 1</b>	<u>83 <math>\pm</math> 3</u>	61 $\pm$ 4	<u>25 <math>\pm</math> 4</u>	<b>8 <math>\pm</math> 1</b>	
CILRS 0.8.4 [11]	97 $\pm$ 2	83 $\pm$ 0	42 $\pm$ 2	47	66 $\pm$ 2	49 $\pm$ 5	23 $\pm$ 1	64	
LBC [34]	97 $\pm$ 1	93 $\pm$ 1	71 $\pm$ 5	N/A	100 $\pm$ 0	94 $\pm$ 3	51 $\pm$ 3	N/A	

Table 3: Comparison of action-based pre-training with baselines (top) and other methods from the literature (bottom). We are able to surpass ImageNet pre-training and the CILRS baseline. Results are from the CARLA 0.9.6 NoCrash benchmark.

$(D^l)$	$(D^u)$	Training Town			T.L.	New Town			T.L.
		Empty	Regular	Dense		Empty	Regular	Dense	
0.5 hours	No Data	25 $\pm$ 4	11 $\pm$ 3	3 $\pm$ 2	51 $\pm$ 2	1 $\pm$ 1	1 $\pm$ 1	0 $\pm$ 0	57 $\pm$ 4
	5 hours	68 $\pm$ 5	58 $\pm$ 5	27 $\pm$ 5	<b>10 <math>\pm</math> 1</b>	14 $\pm$ 1	7 $\pm$ 1	3 $\pm$ 1	23 $\pm$ 2
	20 hours	<b>98 <math>\pm</math> 1</b>	<b>93 <math>\pm</math> 2</b>	<b>61 <math>\pm</math> 6</b>	<u>10 <math>\pm</math> 1</u>	17 $\pm$ 2	12 $\pm$ 2	2 $\pm$ 2	15 $\pm$ 2
	50 hours	95 $\pm$ 2	87 $\pm$ 2	47 $\pm$ 2	12 $\pm$ 1	<b>25 <math>\pm</math> 3</b>	<b>20 <math>\pm</math> 4</b>	<b>6 <math>\pm</math> 1</b>	<b>10 <math>\pm</math> 1</b>
5 hours	No Data	78 $\pm$ 4	79 $\pm$ 7	48 $\pm$ 4	12 $\pm$ 1	57 $\pm$ 1	37 $\pm$ 4	10 $\pm$ 1	18 $\pm$ 2
	5 hours	58 $\pm$ 6	80 $\pm$ 4	69 $\pm$ 5	<u>9 <math>\pm</math> 1</u>	68 $\pm$ 6	55 $\pm$ 5	23 $\pm$ 3	12 $\pm$ 2
	20 hours	74 $\pm$ 2	81 $\pm$ 4	<u>70 <math>\pm</math> 5</u>	10 $\pm$ 1	<u>82 <math>\pm</math> 2</u>	<b>66 <math>\pm</math> 5</b>	<b>34 <math>\pm</math> 0</b>	12 $\pm$ 1
	50 hours	<b>91 <math>\pm</math> 1</b>	<b>91 <math>\pm</math> 4</b>	<u>68 <math>\pm</math> 5</u>	<b>8 <math>\pm</math> 1</b>	<b>83 <math>\pm</math> 3</b>	<u>61 <math>\pm</math> 4</u>	25 $\pm$ 4	<b>8 <math>\pm</math> 1</b>

Table 4: Driving performance ablation study of the BC pre-training encoder investigating the quantity of expert driving data ( $D^u$ ) given a smaller amount of annotated affordances ( $D^l$ ).

network to drive end-to-end. Our method uses much less densely annotated data and does not use dataset aggregation (Dagger [38]). Moreover, note that reported results from CILRS are on CARLA version 0.8.4, the benchmark from the newer version is considerably more difficult. We compare both on the supplementary material.

**Ablation study.** In Table 4, we analyse the data amount impact of both supervised data and expert demonstrations. We used the simplest pre-training encoding, BC. We observe a clear correlation between performance improvement and the quantity of action-based supervised data  $D^u$ . We can see that with only 30 minutes (0.5 hours) of annotated data, the pre-trained network has a performance close to our best reported results in training town. For the generalization conditions, however, only 0.5 hours of annotated data is not enough to obtain satisfactory results. Pre-training is useful also when using higher amounts of annotated data (5 hours), as shown on the bottom of Table 4. However, the impact of pre-training when more annotated data is available is clearer in New Town. Note that when comparing 50 hours with 20 hours, on pre-training, the results are similar, since 20 hours driving seems to be sufficient to capture the inherent variability of T1.

## 5 Conclusions

In this paper, we have shown that representations learned by using action-based methods (BC or Inverse model) are promising as pre-trained representations for autonomous driving controllers based on affordances. Moreover, we have found that most of the benefit comes when the driving experiences (actions) are captured from proper driving (humans). In other words, expert driving outperforms random roaming for representation learning. While considerable future research is needed to improve the raw performance of the methods explored here, the fact that the required data can be easily obtained by simply recording the actions of good drivers, highlights the potential of action-based methods for learning representations for autonomous vehicles beyond pure end-to-end autonomous driving models.

It would be relevant to explore if the pre-training strategy presented here can be also helpful for training annotation-intensive visual models, such as those for semantic class/instance segmentation or object detection. Further, it would be interesting to examine other strategies to use the expert driver data in order to further improve performance.



## Acknowledgments

Yi Xiao and Antonio M. López acknowledge the financial support received for this work from the Spanish TIN2017-88709-R (MINECO/AEI/FEDER, UE) project. Antonio acknowledges the financial support to his general research activities given by ICREA under the ICREA Academia Program. Yi Xiao acknowledges the Chinese Scholarship Council (CSC) grant No.201808390010. Yi Xiao and Antonio M. López acknowledge the support of the Generalitat de Catalunya CERCA Program as well as its ACCIO agency to CVC's general activities. Felipe Codevilla and Christopher Pal would like to thank the Natural Sciences and Engineering Research Council and the Programme d'innovation en Cybersécurité du Quebec for their support.

## References

- [1] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda. A survey of autonomous driving: Common practices and emerging technologies. arXiv:1906.05113v2, 2019.
- [2] J. Janai, F. Güney, A. Behl, and A. Geiger. Computer vision for autonomous vehicles: Problems, datasets and state of the art. arXiv:1704.05519v2, 2019.
- [3] S. Grigorescu, B. Trasnea, T. Cocias, and G. Macesanu. A survey of deep learning techniques for autonomous driving. arXiv:1906.05113v2, 2019.
- [4] D. Pomerleau. ALVINN: An autonomous land vehicle in a neural network. In *Neural Information Processing Systems (NIPS)*, 1989.
- [5] Y. LeCun, U. Muller, J. Ben, E. Cosatto, and B. Flepp. Off-road obstacle avoidance through end-to-end learning. In *Neural Information Processing Systems (NIPS)*, 2005.
- [6] M. Bojarski, D. D. 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. arXiv:1712.00409, 2016.
- [7] H. Xu, Y. Gao, F. Yu, and T. Darrell. End-to-end learning of driving models from large-scale video datasets. In *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [8] S. Hecker, D. Dai, and L. Van Gool. End-to-end learning of driving models with surround-view cameras and route planners. In *European Conference on Computer Vision (ECCV)*, 2018.
- [9] A. Amini, W. Schwarting, G. Rosman, B. Araki, S. Karaman, and D. Rus. Variational auto-encoder for end-to-end control of autonomous driving with novelty detection and training de-biasing. In *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2018.
- [10] F. Codevilla, M. Müller, A. M. López, V. Koltun, and A. Dosovitskiy. End-to-end driving via conditional imitation learning. In *International Conference on Robotics and Automation (ICRA)*, 2018.
- [11] F. Codevilla, E. Santana, A. M. López, and A. Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *International Conference on Computer Vision (ICCV)*, 2019.
- [12] P. Hamm, D. Jayaraman, and S. Levine. Causal confusion in imitation learning. In *Neural Information Processing Systems (NIPS)*, 2018.
- [13] C. Chen, A. Seff, A. L. Kornhauser, and J. Xiao. DeepDriving: Learning affordance for direct perception in autonomous driving. In *International Conference on Computer Vision (ICCV)*, 2015.
- [14] A. Sauer, N. Savinov, and A. Geiger. Conditional affordance learning for driving in urban environments. In *Conference on Robot Learning (CoRL)*, 2018.
- [15] J. J. Gibson. *The ecological approach to visual perception: classic edition*. Psychology Press, 2014.

- [16] A. Anand, E. Racah, S. Ozair, Y. Bengio, M.-A. Côté, and R. D. Hjelm. Unsupervised state representation learning in Atari. In *Neural Information Processing Systems (NIPS)*, 2019.
- [17] P. Agrawal, A. V. Nair, P. Abbeel, J. Malik, and S. Levine. Learning to poke by poking: Experiential learning of intuitive physics. In *Neural Information Processing Systems (NIPS)*, 2016.
- [18] A. Dosovitskiy, G. Ros, F. Codevilla, A. López, and V. Koltun. CARLA: An open urban driving simulator. In *Conference on Robot Learning (CoRL)*, 2017.
- [19] G. Larsson, M. Maire, and G. Shakhnarovich. Colorization as a proxy task for visual understanding. In *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [20] S. Gidaris, P. Singh, and N. Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations (ICLR)*, 2018.
- [21] Z. Feng, C. Xu, and D. Tao. Self-supervised representation learning by rotation feature decoupling. In *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [22] J. Xu, L. Xiao, and A. López. Self-supervised domain adaptation for computer vision tasks. *IEEE Access*, 7:156694–156706, October 2019.
- [23] I. Misra, C. L. Zitnick, and M. Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision (ECCV)*, 2016.
- [24] M. Noroozi and P. Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision (ECCV)*, 2016.
- [25] A. Kolesnikov, X. Zhai, and L. Beyer. Revisiting self-supervised visual representation learning. In *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [26] P. Goyal, D. Mahajan, A. Gupta, and I. Misra. Scaling and benchmarking self-supervised visual representation learning. In *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [27] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2006.
- [28] A. v. d. Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. arXiv:1807.03748v2, 2018.
- [29] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. arXiv:12002.05709, 2020.
- [30] K. Chua, R. Calandra, R. McAllister, and S. Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Neural Information Processing Systems (NIPS)*, 2018.
- [31] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics for planning from pixels. In *International Conference on Computer Vision (ICCV)*, 2019.
- [32] E. Shelhamer, P. Mahmoudieh, M. Argus, and T. Darrell. Loss is its own reward: Self-supervision for reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2017.
- [33] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven exploration by self-supervised prediction. In *Int. Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2017.
- [34] D. Chen, B. Zhou, V. Koltun, and P. Krähenbühl. Learning by cheating. In *Conference on Robot Learning (CoRL)*, 2019.
- [35] A. Zhao, T. He, Y. Liang, H. Huang, G. V. den Broeck, and S. Soatto. Lates: Latent space distillation for teacher-student driving policy learning. arXiv:1912.02973, 2019.

- [36] G. Alain and Y. Bengio. Understanding intermediate layers using linear classifier probes. In *International Conference on Learning Representations (ICLR) Workshops*, 2017.
- [37] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick. Momentum contrast for unsupervised visual representation learning. arXiv:1911.05722, 2019.
- [38] 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*, pages 627–635, 2011.

# Supplementary Material for “Action-based Representation Learning for Autonomous Driving”

## 1 NoCrash Benchmark: CARLA 0.9.6 (modified)

For performing the driving evaluation, we updated the *NoCrash* benchmark to work with version 0.9.6. This version has some differences compared with CARLA 0.8.4, the main one is related to pedestrians that now have a very different crossing pattern. In addition, they are now able to cross roads in groups and are more equally spread over the town. Thus, it is necessary to change the number of pedestrians spawned on the town in order to reproduce the benchmark from version 0.8.4. For the regular task, we increased the number of pedestrians from 50 to 125 for Town01, and 50 to 100 for Town02. In dense task, we increased from 250 to 400 for Town01, and from 150 to 300 for Town02.

The routes have been also changed since the new API allowed a much more controlled sampling of the routes to be used on the benchmark. Thus, we made a few changes in the routes to guarantee they followed the restrictions described on [1]. We restrict at least 1000 meters distance between start and end point for Town01, and 500 for Town02.

The visuals also have been slightly changed especially with respect to the dynamic obstacles. We can see on Figure 3 some new features of the benchmark that are present on version 0.9.6. The vehicles now include bikes, motorbikes and infant pedestrians. Our update was analogous to the one done by [2], however, instead of implementing it ourselves, we added to version 0.9.6 the pedestrian navigation algorithm provided by the official CARLA 0.9.7 version. Note that for this paper, the version used was still CARLA 0.9.6, while only the pedestrian navigation from 0.9.7 was added. We decided to stay on version 0.9.6 since newer versions of CARLA have incorporated major differences on traffic management that drastically changed vehicle behavior.

## 2 Data distribution

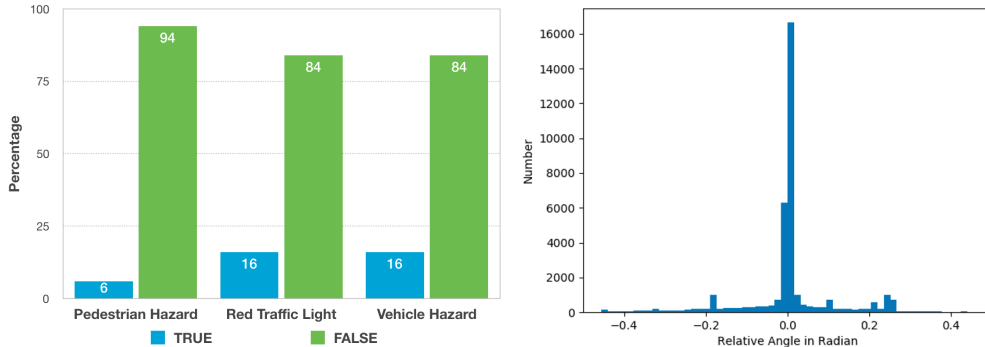


Figure 1: The distributions of affordances on the  $\sim 30$  minutes training dataset.

In Figure 1, we show the distributions of the 30 minutes annotated dataset, which was used for training the affordances prediction network  $g_\phi$ . This dataset is a subset of the full 50 hours dataset, carefully sampled to maintain the same distribution as the full dataset.

### 3 Network Architectures

In Figure 2, we detail the architecture of our encoder,  $h_\phi$ . We follow the architecture presented in [3] using the high level command,  $c_t$ , the image input,  $I_t$ , and the speed variable  $v_t$ , as part of  $o_t$ . In Table 1, we detail the parameters of network architectures for different action-based representation learning approaches. We also detail the affordance projection network.

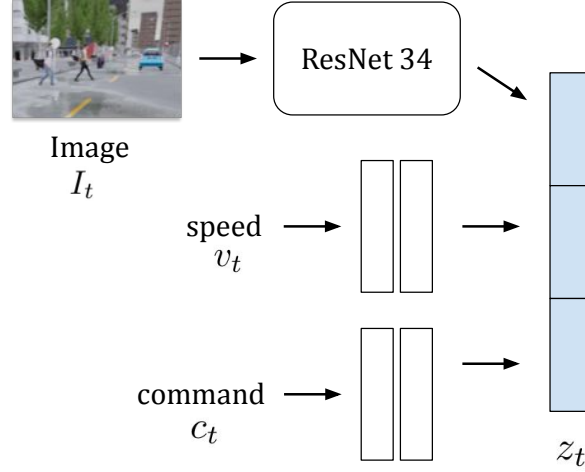


Figure 2: The architecture of the encoder  $h_\theta$ . Blue rectangles indicate that the features are appended. White rectangles are fully connected layers.

### 4 Controller

To evaluate our models in the NoCrash benchmark, we tuned a controller using our affordances. Given a set of affordances at time  $t$ ,  $\{\psi(t), hv(t), hp(t), hr(t)\}$ , the controller outputs an action  $a(t)$  defined by  $\{S(t), T(t), B(t)\}$ , *i.e.*, steering, throttle, and break, respectively. In particular, for lateral control (*i.e.*,  $S(t)$ ) and for longitudinal control (*i.e.*,  $T(t)$  and  $B(t)$ ) we use the following PID-based equations:

$$\begin{aligned} S(t) &= PID(\psi(t)) = K_p\psi(t) + K_i \int_0^t \psi(\tau) d\tau + K_d \frac{\partial \psi(t)}{\partial t}, \\ B(t) &= \max(hr(t), hp(t), hv(t)), \\ T(t) &= PID(v \text{ if } (B(t) > 0), 0 \text{ otherwise}), \end{aligned}$$

where the hazard functions either equal to 1 or 0, and  $v$  is the target maximum speed, 20Km/h in these experiments. We tuned the constants  $K_p$ ,  $K_i$  and  $K_d$  in town 1 (T1) to obtain a perfect driving (no errors, all episodes completed) for the *dense* condition of the NoCrash benchmark provided we use ground truth (perfect) affordances. We did it in that way to provide a driving evaluation directly depending on the quality of the affordance predictions, not in the controller itself, since it is not the focus of this paper.



Figure 3: New types of vehicles and pedestrians present on the updated version of the benchmark. Left: motorbikes were not included on the previous benchmark. Middle: kids are now also part of pedestrians. Right: now pedestrians are able to better agglomerate when crossing roads.



Module	Input Dimension	Output Channels	Num. of Dropout
ResNet 34	$200 \times 88 \times 3$	512	0.0
Speed	1	128	0.0
	128	128	0.0
Command	4	128	0.0
	128	128	0.0
Join	512+128+128	512	0.0
Action Branch	512	256	0.0
	256	256	0.5
	256	3	0.0
Speed Branch	512	256	0.0
	256	256	0.5
	256	1	0.0

(a) Behavior Cloning (BC)

Module	Input Dimension	Output Channels	Num. of Dropout
ResNet 34	$200 \times 88 \times 3$	512	0.0
Speed	1	128	0.0
	128	128	0.0
Command	4	128	0.0
	128	128	0.0
Join	512+128+128	512	0.0
Join ( $Z_t, Z_{t+1}$ )	512 + 512	512	0.0
Action ( $a_t$ )	512	256	0.0
	256	256	0.0
	256	3	0.0

(b) Inverse model

Module	Input Dimension	Output Channels	Num. of Dropout
ResNet 34	$200 \times 88 \times 3$	512	0.0
Speed	1	128	0.0
	128	128	0.0
Command	4	128	0.0
	128	128	0.0
Action ( $a_t$ )	3	512	0.0
	512	256	0.0
	256	512	0.0
Join	512+128+128	512	0.0
Join ( $Z_t, Z_{t+1}$ )	512 + 512	512	0.0
Action ( $a_t$ )	512	256	0.0
	256	256	0.5
	256	3	0.0
Join ( $Z_t$ , Action)	512 + 512	512	0.0

(c) Forward Model

Module	Input Dimension	Output Channels	Num. of Dropout
Affordances	512	512	0.0
	512	256	0.0
	256	M	0.0

(M: Classification - 2; Regression - 1)

(d) Affordances Network

Table 1: Network architecture details for the encoders  $h_\theta$  (a,b and c) and the affordance projection network  $g_\phi$  (d) when fine-tuned for driving.

Pre-training	Binary Affordances			Relative Angle ( $\psi_t$ )		
	Pedestrian ( $hp$ )	Vehicle ( $hv$ )	Red T.L. ( $hr$ )	Left Turn	Straight	Right Turn
No pre-training	$38 \pm 1$	$59 \pm 0$	$45 \pm 0$	$10.03 \pm 0.06$	$1.80 \pm 0.03$	$17.57 \pm 0.03$
ImageNet	$35 \pm 1$	$67 \pm 0$	$54 \pm 1$	$14.46 \pm 0.35$	$2.31 \pm 0.09$	$19.06 \pm 0.18$
Contrastive (ST-DIM)	$38 \pm 0$	$57 \pm 1$	$76 \pm 1$	$6.72 \pm 0.03$	$2.65 \pm 0.03$	$13.12 \pm 0.11$
Forward	$56 \pm 0$	$62 \pm 0$	$55 \pm 0$	$6.97 \pm 0.12$	<b><math>0.17 \pm 0.00</math></b>	$5.88 \pm 0.03$
Inverse	<b><math>57 \pm 1</math></b>	<b><math>82 \pm 0</math></b>	<b><math>89 \pm 0</math></b>	<b><math>3.84 \pm 0.06</math></b>	<b><math>0.21 \pm 0.03</math></b>	<b><math>2.96 \pm 0.09</math></b>
Behavior Cloning (BC)	$46 \pm 0$	<b><math>83 \pm 0</math></b>	$86 \pm 0$	<b><math>3.76 \pm 0.03</math></b>	$0.63 \pm 0.00$	$4.35 \pm 0.06$

Table 2: Linear probing results on Town01 testing set. Left: F1 score for the binary affordances (higher is better). Results are scaled by 100 for visualization purpose. Right: MAE of the relative angle (lower is better), shown for different navigation maneuvers. MAE is shown in degrees for an easier understanding.

## 5 Additional Results

**Linear Probing.** Tables 2 and 3 show the linear probing evaluation results of models tested on the  $\sim 1$  hour Town01 testing set. This testing set has similar appearance to the training data. We observe a similar tendency to the results obtained on Town02, those shown in the main paper.

**Driving Results.** An important observation is that using expert demonstration as pre-training seems to be more beneficial than training a model end-to-end to directly perform control. In Table 4, we show a behavior cloning encoder trained with 50 hours of expert demonstrations. We compare two different uses of this encoder: directly producing driving controls and serving as representation learning for an affordance prediction model. With our pre-training strategy, and the complementary

Pre-training	Binary Affordances			Relative Angle ( $\psi_t$ )		
	Pedestrian ( $hp$ )	Vehicle ( $hv$ )	Red T.L. ( $hr$ )	Left Turn	Straight	Right Turn
No pre-training	$38 \pm 1$	$59 \pm 0$	$45 \pm 0$	$10.03 \pm 0.06$	$1.80 \pm 0.03$	$17.57 \pm 0.03$
Contrastive (ST-DIM)	$36 \pm 1$	$60 \pm 2$	$67 \pm 1$	$7.43 \pm 0.17$	$2.79 \pm 0.18$	$12.11 \pm 0.18$
Contrastive Random (ST-DIM)	$34 \pm 2$	<b><math>78 \pm 1</math></b>	$52 \pm 2$	$11.10 \pm 0.52$	$2.02 \pm 0.09$	$14.65 \pm 0.29$
Forward	<b><math>52 \pm 0</math></b>	$53 \pm 0$	$60 \pm 0$	$3.95 \pm 0.00$	$0.11 \pm 0.00$	$4.37 \pm 0.03$
Forward Random	$5 \pm 0$	$27 \pm 0$	$2 \pm 0$	$19.10 \pm 0.03$	$0.63 \pm 0.00$	$17.97 \pm 0.03$
Inverse	$48 \pm 1$	$71 \pm 0$	<b><math>90 \pm 0</math></b>	<b><math>2.03 \pm 0.07</math></b>	<b><math>0.21 \pm 0.03</math></b>	<b><math>2.54 \pm 0.03</math></b>
Inverse Random	$34 \pm 0$	$53 \pm 0$	$55 \pm 0$	$12.49 \pm 0.45$	$0.84 \pm 0.03$	$13.22 \pm 0.54$

Table 3: Linear probing results on Town01 testing set, comparing encoders trained with random policy training data versus expert demonstration data. Note that, to provide fair comparison, the encoders here were trained with 20 hours data.

	Technique	Empty	Regular	Dense	T.L.
Training	BC driving	$72 \pm 5$	$47 \pm 2$	$20 \pm 3$	$74 \pm 2$
	BC pre-training	<b><math>91 \pm 1</math></b>	<b><math>91 \pm 4</math></b>	<b><math>68 \pm 5</math></b>	<b><math>8 \pm 1</math></b>
New Town	BC driving	$71 \pm 2$	$42 \pm 7$	$12 \pm 3$	$63 \pm 1$
	BC pre-training	<b><math>83 \pm 3</math></b>	<b><math>61 \pm 4</math></b>	<b><math>25 \pm 4</math></b>	<b><math>8 \pm 1</math></b>

Table 4: We compare the behavior cloning (BC) technique used as driving technique (BC driving) to it used as a pre-training technique (BC pre-training) for an affordances based model.

affordance training, our model (BC pre-training) is able to greatly outperform the end-to-end driving results (BC driving). This difference is expressive especially when comparing the capability to stop on red traffic lights. Finally, note that the “BC driving” results from Table 4 are in practice a re-training of the CILRS [3] baseline to work on version 0.9.6.

## References

- [1] A. Dosovitskiy, G. Ros, F. Codevilla, A. López, and V. Koltun. CARLA: An open urban driving simulator. In *Conference on Robot Learning (CoRL)*, 2017.
- [2] D. Chen, B. Zhou, V. Koltun, and P. Krähenbühl. Learning by cheating. In *Conference on Robot Learning (CoRL)*, 2019.
- [3] F. Codevilla, E. Santana, A. M. López, and A. Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *International Conference on Computer Vision (ICCV)*, 2019.