

Conditional Affordance Learning for Driving in Urban Environments

Axel Sauer¹ Nikolay Savinov¹ Andreas Geiger^{1,2}

¹Computer Vision and Geometry Group, ETH Zürich

²Autonomous Vision Group, MPI for Intelligent Systems and University of Tübingen

Abstract: Most existing approaches to autonomous driving fall into one of two categories: modular pipelines, that build an extensive model of the environment, and imitation learning approaches, that map images directly to control outputs. A recently proposed third paradigm, direct perception, aims to combine the advantages of both by using a neural network to learn appropriate low-dimensional intermediate representations. However, existing direct perception approaches are restricted to simple highway situations, lacking the ability to navigate intersections, stop at traffic lights or respect speed limits. In this work, we propose a direct perception approach which maps video input to intermediate representations suitable for autonomous navigation in complex urban environments given high-level directional inputs. Compared to state-of-the-art reinforcement and conditional imitation learning approaches, we achieve an improvement of up to 68 % in goal-directed navigation on the challenging CARLA simulation benchmark. In addition, our approach is the first to handle traffic lights, speed signs and smooth car-following, resulting in a significant reduction of traffic accidents.

Keywords: Autonomous Driving, Sensorimotor Control, Affordance Learning

1 Introduction

An autonomous vehicle is a cognitive system and hence follows the concept of a rational agent: in order to operate safely it must accurately observe its environment, make robust decisions and perform actions based on these decisions [1]. Autonomous navigation can be described as a mapping function from sensory input to control output. To implement this function, three major approaches have been proposed: modular pipelines, behavior reflex, and direct perception.

Most autonomous vehicles are based on *modular pipelines* (MP). An MP splits the autonomous driving problem into smaller, easier subproblems: perception, path planning, and control. The approach often relies on various sensors in order to produce a consistent representation of the surrounding environment. A driving decision is then made based on this representation. While MPs are relatively interpretable due to their modularity, they rely on complex intermediate representations which are manually chosen (e.g., optical flow), often hard to predict with sufficient accuracy and might not necessarily be the optimal choice for solving the sensorimotor control task. Furthermore, MPs require large amounts of annotated data which can be costly to obtain, e.g., pixel-wise semantic segmentation for training neural networks or high-definition maps for localization.

The second approach, *imitation learning* (IL), maps directly from raw input to control output and is learned from data in an end-to-end fashion, skipping the intermediate step of building an explicit environment model. In contrast to MPs, data for training IL-based navigation systems can be collected with relative ease, i.e., by driving around and recording the video footage together with the driven trajectory. However, IL approaches are faced with the problem of learning a very complex mapping (from raw input to control) in a single step. Thus, a model with high capacity as well as a very large amount of training data is required for handling the large variety of real world situations which may be encountered at test time. Besides, IL lacks transparency as it is hard to comprehend the internal decision-making process in a neural network. This raises security concerns: a driving system that does not act transparently might not be trusted or confidently used by humans.

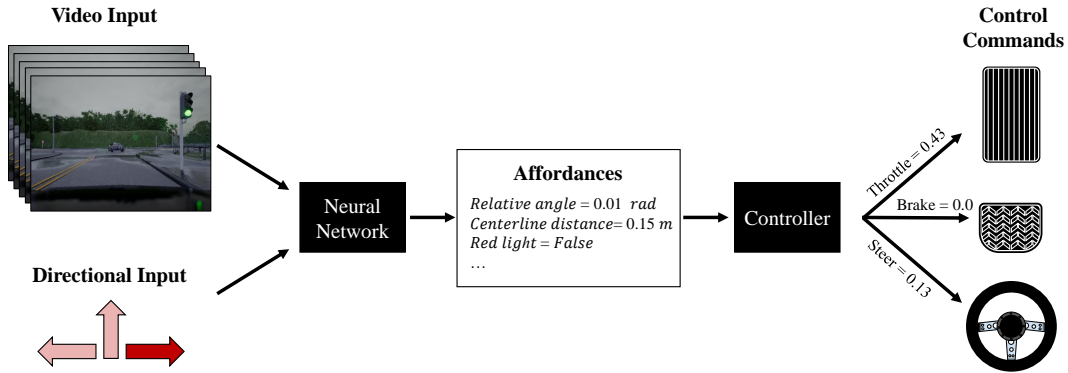


Figure 1: **Conditional Affordance Learning (CAL) for Autonomous Urban Driving.** The input video and the high-level directional input are fed into a neural network which predicts a set of affordances. These affordances are used by a controller to calculate the control commands.

The third approach, *direct perception* (DP), aims to combine the advantages of MP and IL. Instead of predicting a detailed representation of the environment, the goal in DP is to predict a low-dimensional intermediate representation of the environment which is then used in a conventional control algorithm to maneuver the vehicle. Thus, DP neither requires the network to learn the complex sensorimotor control problem end-to-end, nor does it assume the availability of large annotated datasets for object detection, semantic segmentation, motion estimation or localization. Chen et al. [2] demonstrated good results when applying this approach to highway driving using an open-source simulator. However, highway driving is a rather easy task compared to driving in rural or urban areas. When considering navigation in urban areas, several difficulties are added: the agent must obey traffic rules (speed limits, red lights, etc.), take into account possible obstacles on the road (e.g., pedestrians crossing the street), and handle junctions with more than one possible direction.

At the core of the direct perception approach lies the choice of intermediate representation to be predicted. Ideally, this representation should be of low dimensionality while comprising all the necessary information for making a driving decision. One choice of such representation are *affordances*, attributes of the environment which limit the space of allowed actions. As an example, consider the distance to the vehicle ahead which limits the ability of the ego-vehicle to speed up. Affordances can also be conditioned on high-level driving decisions as provided, e.g., by a conventional navigation system. Consider a vehicle at an intersection: the distance to the centerline changes with the desired turn we would like the vehicle to take (and thus the lane we would like it to follow).

The goal of this work is to generalize the direct perception approach to the urban setting. We develop novel intermediate representations in the form of affordances suitable for urban navigation as well as conditional models where the decision is not based on the image input alone, but also on high-level navigation commands as provided by a consumer grade navigation system (e.g., “turn left at the next intersection”). Moreover, our system is designed to drive more safely and comfortably, alleviating the common problem of jerky driving behavior in existing approaches for autonomous driving [3]. We coin the developed approach “Conditional Affordance Learning” (CAL) as the predicted affordances and the model output are conditioned on high-level navigation commands. Figure 1 shows an overview of our approach. The main contributions of our work are summarized as follows:

- We develop novel intermediate representations which are low-dimensional, yet provide the necessary information to allow for driving in urban environments.
- We deploy a conditional navigation model which allows for taking high-level directional commands prior to intersections. We further implement a control algorithm which allows for a smooth ride while obeying traffic rules.
- We provide an in-depth evaluation of different network architectures and parameterizations. Our network is able to predict all intermediate representations in a single forward pass.
- We demonstrate the benefits of video-based representation learning for navigation. Prior works either operate frame-by-frame or do not evaluate in an online setting.

We will release the code and our trained models upon publication.

2 Related Work

Modular pipelines are the most popular approach to autonomous driving and are employed by most research projects [4, 5, 6]. To build the environmental model, the perception stack needs to detect all aspects of the traffic scene that are likely to be relevant for the driving decision. These detection tasks are usually trained and solved separately [7], most recently by exploiting deep neural networks, e.g., for object detection [8, 9], image segmentation [10, 11], depth [12, 13] or motion estimation [14, 15]. This information can be aggregated into a model of the environment [16, 17] and a planning module generates an obstacle-free trajectory which is executed by the control module [18].

The first implementation of the imitation learning (IL) approach in the real world, ALVINN [19], dates back to the 80s. More recently, Muller et al. [20] employ a similar system to navigate an off-road truck in a variety of terrains, weather, and lighting conditions while avoiding obstacles. Using more computational power and a modern CNN architecture, contemporary approaches such as the one by Bojarski et al. [21] demonstrate impressive performance on real-world tasks like highway following and driving in flat courses. Codevilla et al. [22] propose a conditional imitation learning formulation that is able to traverse intersections based on high-level navigational commands.

The idea of direct perception is to predict several affordance indicators describing the driving scene. The concept of affordances was originally proposed by Gibson [23] in the field of psychology and has been applied to the autonomous driving task by Chen et al. [2], demonstrating strong performance in a racing simulation. Al-Qizwini et al. [24] improve the original approach of Chen et al. [2] by analyzing different CNNs for the mapping from image to indicators. They find GoogLeNet [25] and VGG16 [26] to perform best on this task.

Due to the difficulties of training and testing in real environments, driving simulators are often exploited in related work. Two popular driving simulators for research are the open source car racing game TORCS [2, 24, 27] and the commercial game Grand Theft Auto V (GTA V) [28, 29]. However, TORCS is neither photorealistic nor complex enough, lacking important scene elements such as intersections, pedestrians and oncoming traffic. GTA V can be considered photorealistic, but is closed source and hence limited in its ability to customize and control the environment. In this work, we use the recently released open-source simulator CARLA [30] which provides a trade-off between realism and flexibility, thus addressing some of the problems of previous simulators.

3 Conditional Affordance Learning

Our autonomous driving setup consists of an agent which interacts with the environment as illustrated in Figure 2. A high-level command for the maneuver to be performed (e.g., “go straight”, “turn left”, “turn right”) is provided by a planner, mimicking one of the most common driving scenarios: a human driver (the agent) following the commands of a navigation system (the planner). The actions of the agent are conditioned on this high-level command. The agent receives the observation from the environment and the high-level command, and then chooses the value of throttle, brake and steering wheel. Those values are sent back to the environment model which provides the next observation and high-level command, and the process continues until the goal is reached.

In this work, we focus on visual-based navigation: we only consider observations from a single front-facing camera. Cameras have the advantage that they are ubiquitous and cheap. In addition, they provide information that laser scanners or radar sensors cannot provide, e.g., the color of a traffic light or the type of a speed sign. Thus they are the only sensor that allows to solve the task without any additional sensors. Note that this setup is comparable to how humans solve the driving task as the human visual stereopsis is imprecise at distances more than 10 meters and thus the human visual system can be considered monocular from the viewpoint of autonomous driving.

In the following we will discuss the intermediate representations required for autonomous navigation in urban environments and the submodules of the CAL agent for perception and control.

3.1 Affordances

A good intermediate representation should be low-dimensional, yet expressive enough to meet the necessary requirements for autonomous navigation: (i) the agent should be able to drive from A to B as fast as possible while obeying traffic rules. (ii) infractions against traffic rules should be avoided

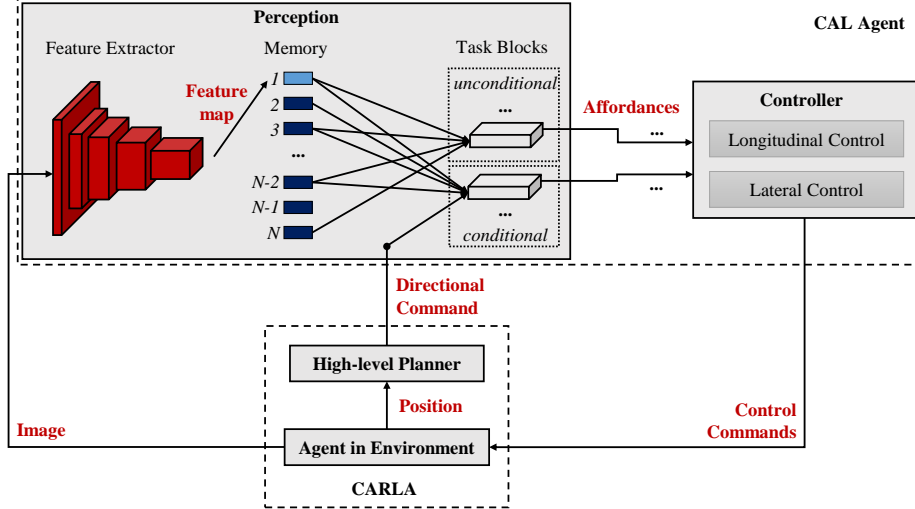


Figure 2: **Overview of our System.** The CAL agent (top) receives the current camera image and a directional command from CARLA [30] (“straight”, “left”, “right”). The feature extractor converts the image into a feature map. The agent stores the last N feature maps in memory, where N is the length of the input sequence required for the perception stack. This sequence of feature maps, together with the directional commands from the planner, are exploited by the task blocks to predict affordances. Different tasks utilize different temporal receptive fields and temporal dilation factors. The control commands calculated by the controller are sent back to CARLA [30] which updates the environment and provides the next observation and directional command.

at all times. In this work, we consider 6 types of infractions: driving on the wrong lane, driving on the sidewalk, running a red light, colliding with other vehicles, hitting pedestrians and hitting static objects. Moreover, the speed limit must be obeyed at all times. (iii) the agent should be able to provide a pleasant driving experience. The car should stay in the center of the road and take turns smoothly. In addition, it should be able to follow a leading car at a safe distance.

Based on these requirements we propose the affordance set shown in Figure 3 (left) and task a neural network to predict these affordances from the monocular input video. The agent should stop for **red traffic lights** and change its speed according to the current **speed sign**. For this purpose we define an *observation area* A_1 in the local (x, y) coordinates of the agent, as shown in Figure 3 (right). If a traffic light or a speed sign is within A_1 , the corresponding affordance indicator switches to *True* or to the speed limit value. The agent keeps an internal memory of the speed limit. Once a speed sign is passed, the current speed limit is updated and the agent adjusts its speed accordingly. If an obstacle is in front of the car — located within an *observation area* A_2 — the agent should perform a **hazard stop** — come to a complete stop as quickly as possible. In CARLA, possible road obstacles are cars and pedestrians. For driving in the center of the road, the position and the orientation of the car on the road must be known. The position is defined by the **distance to centerline** d , the lateral distance of the center of the vehicle’s front axle to the closest point on the centerline. The **relative angle** ψ describes the orientation of the car with respect to the road tangent. To achieve steady vehicle following, the **distance to the vehicle** ahead must be known. If a car is located in an *observation area* A_3 , we measure this distance ℓ as the shortest distance between the bounding boxes of both cars, otherwise set $\ell = 50$. The coordinates of all *observation areas* are listed in the supplementary material.

3.2 Perception

We formulate perception as a multi-task learning (MTL) problem: using a single neural network we predict all affordances in a single forward pass. This yields a computationally efficient solution which conforms to the real-time requirements of autonomous systems. Furthermore, MTL has been shown to learn internal representations which are able to improve the network’s generalization abilities [31]. Xu et al. [32] demonstrate this effect in the context of autonomous driving. In our network, all convolutional layers are shared between tasks. The output of the last convolutional layer is used as input for the task-specific branches which we will call “task blocks” in the following.

| Type | Conditional | Affordances | Acronym | Range of values |
|------------|-------------|----------------------------|---------|----------------------------|
| discrete | No | Hazard stop | - | $\in \{True, False\}$ |
| | | Red Traffic Light | - | $\in \{True, False\}$ |
| | | Speed Sign [km/h] | - | $\in \{None, 30, 60, 90\}$ |
| continuous | No | Distance to vehicle [m] | ℓ | $\in [0, 50]$ |
| continuous | Yes | Relative angle [rad] | ψ | $\in [-\pi, \pi]$ |
| | | Distance to centerline [m] | d | $\in [-2, 2]$ |

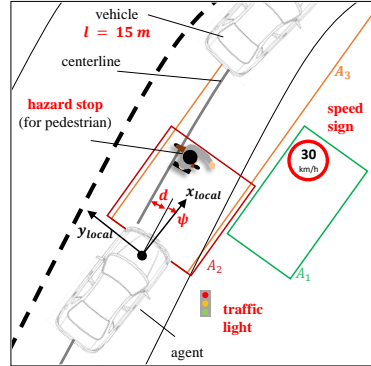


Figure 3: **Affordances.** Left: We categorize affordances according to their type (discrete/continuous) and whether they are conditional (dependent on directional input) or unconditional. Right: Illustration of the affordances (red) and observation areas used by our model.

Neural network architecture. To extract features from raw input images, we use the convolutional layers of a VGG16 network [26] pretrained on ImageNet [33]. The simple structure of the network make it an appealing choice for a feature extractor with real-time requirements. The feature extractor is applied to every image in the input sequence to extract a sequence of feature maps. The task blocks are shallow networks and consist of a task-specific layer, followed by a batch normalization layer [34] and a dropout layer. For the task-specific layers we experiment with dense layers, LSTMs [35], GRUs [36], and temporal convolutional layers. Each task block has a different temporal receptive field and temporal dilation factors (e.g., using only every other image). The output nodes of the conditional task blocks are divided into three equally sized groups. The directional command $c \in \{straight, left, right\}$ is used to switch between these groups of nodes. Only the active group is used to predict the target label. During training the gradients are only backpropagated through this group. We observed that training specialized submodules for each directional command leads to better performance compared to using the directional command as an additional input to the task networks. This observation agrees with Codevilla et al. [22]. However, note that in contrast to the Conditional Imitation Learning (CIL) approach of Codevilla et al. [22], we use specialized task networks to predict conditional affordances rather than conditional steering commands.

Training. As our approach decouples affordance prediction and control, we are able to train our perception stack independently from the chosen control algorithm. To train the affordance prediction network, we recorded a dataset with the CARLA simulator [30]. CARLA provides two towns: Town 1 for training and Town 2 for testing. For collecting training data in Town 1, we navigate through the city using our controller (described in section 3.3) based on the ground truth affordances. During data collection, the directional input at intersections is chosen at random. All other traffic participants and all traffic lights are controlled by the in-game AI. The full dataset contains approximately 240 thousand images and the corresponding ground truth labels for each of the six affordances. Five percent of the dataset is used for validation. The network is trained on mini-batches of size 32. For optimization we use Adam [37] with a base learning rate of $5e-5$. For classification we use the class-weighted categorical cross-entropy H where the weights are chosen inversely proportional to the class label occurrence in the training set. For regression, we use the mean average error MAE . All six affordance loss functions are added to yield the loss function for the whole network:

$$\mathcal{L} = \sum_{j=1}^3 H_j + \sum_{k=1}^3 MAE_k. \quad (1)$$

Data augmentation is crucial to achieve generalization. In this work, we used the following augmentations: change in color, contrast, and brightness as well as Gaussian blur and salt-and-pepper noise. Common techniques like vertical and horizontal flipping are not used, because in the driving context vertical flipping would correspond to switching from right- to left-hand traffic and horizontal flipping — to driving upside down. In addition, we employ camera pose randomization. It is well known that for imitation learning utilizing a single fixed camera pose leads to unstable behavior [21, 38] due to control errors and unexpected observations reinforcing each other. Thus, we use a three-camera setup that records the same episode at 3 different baselines (-50 cm, 0 cm, $+50$ cm). In addition, each camera is randomly rotated around its vertical axis by up to $\pm 15^\circ$.

During 10 % of the training episodes, we additionally set the distance to the preceding vehicle ℓ to zero to provoke rear-end collisions with other cars which otherwise would not occur in our training set. This way the network learns to correctly predict these situation as hazard stops.

Hyperparameter search. The task block hyperparameters entail the *type* of layer (Dense, LSTM, GRU, Temporal Convolution), the number of *nodes* within these layers, the dropout ratio p , the sequence length *seq*, and the dilation value *dil*. The architecture of our network is optimized in a two-step procedure. First, we randomly choose *type* and sample values for *nodes*, p , *seq*, and *dil* from predefined value ranges (see supplementary material for the respective ranges). We initialize the network with the same set of parameters for each task block and train it until the error on the validation set converges. Second, the overall best performing task blocks are combined into a new model. The combined model is then trained jointly until the validation error converges. The hyperparameters of the best performing task blocks are shown in the supplementary material.

3.3 Controller

The control algorithm is split into longitudinal control and lateral control since it is possible to decouple these two tasks [39]. The longitudinal controller calculates the throttle and brake outputs, the lateral controller — the steering output.

Longitudinal control. The longitudinal controller is subdivided into several states, see the supplementary material for an illustration. The states are (in ascending importance): *cruising*, *following*, *over_limit*, *red_light*, and *hazard_stop*. All states are mutually exclusive. For *cruising* the target speed v^* is the current speed limit. Given the current speed $v(t)$ we minimize the error term $e = v^*(t) - v(t)$ using a PID controller [40] for the throttle value. The brake is set to zero. The value of v^* also depends on the current road type. If the agent receives a directional input to turn left or right, we reduce v^* by 10 km/h to enable smooth and safe navigation through turns. The state switches to *following* if another car is less than 35 m ahead. We will denote the distance to the front car as ℓ and the current speed limit as v_{max} . We use the optimal car-following model of Chen et al. [2] and minimize the error term

$$e(t) = v^*(t) - v(t) = v_{max} \left(1 - \exp \left(-\frac{c}{v_{max}} \ell(t) \right) - d \right) - v(t) \quad (2)$$

with a second PID controller. If the car is driving faster than the current speed limit, the state switches to *over_limit*, for a detected red light it switches to *red_light*, for a detected obstacle it switches to *hazard_stop*. For *over_limit*, *red_light*, and *hazard_stop*, the throttle and brake value depend on the vehicle speed, the prediction probability and the urgency to stop. The details as well as a description of the PID tuning procedure are provided in the supplementary material.

Lateral control. For lateral control we use the Stanley Controller (SC) [4]. The SC uses two error metrics: the distance to centerline $d(t)$ and the relative angle $\psi(t)$. The control law to calculate the steering angle $\delta_{SC}(t)$ at the current vehicle speed $v(t)$ is given by

$$\delta_{SC}(t) = \psi(t) + \arctan \left(\frac{k d(t)}{v(t)} \right) \quad (3)$$

where k is a gain parameter. To calculate the final steering output δ , we add a damping term δ_d which reduces swaying on straight roads. $\delta_d(t)$ is calculated by using the current output of the control law $\delta_{SC}(t)$, the steering angle at the previous time step $\delta(t-1)$, and a positive damping constant D :

$$\delta(t) = \delta_{SC}(t) + \delta_d(t) \quad \text{where} \quad \delta_d(t) = -D * (\delta_{SC}(t) - \delta(t-1)) \quad (4)$$

4 Results

We use the driving benchmark and evaluation protocol of [30] as the basis of our evaluation. The benchmark is composed of four different driving tasks with increasing difficulty: driving straight, driving through a single turn, navigating through the town taking several turns, and navigating through the town with additional dynamic obstacles (cars, pedestrians). The agent is initialized at a pre-defined point in the town and has to reach a goal under a task-specific time limit. The time limit equals the time needed to reach the goal when driving along the optimal path at 10 km/h.

Table 1: **Quantitative Evaluation on Goal-Directed Navigation Tasks.** We show the percentage of successfully completed episodes per task and compare CAL to a Modular Pipeline (MP) [30], Conditional Imitation Learning (CIL) [22] and Reinforcement Learning (RL) [30].

| Task | Training conditions | | | | New weather | | | | New town | | | | New town and new weather | | | |
|--------------|---------------------|-----------|----|------------|-------------|-----|----|------------|----------|-----------|----|-----------|--------------------------|-----|----|-----------|
| | MP | CIL | RL | CAL | MP | CIL | RL | CAL | MP | CIL | RL | CAL | MP | CIL | RL | CAL |
| Straight | 98 | 95 | 89 | 100 | 100 | 98 | 86 | 100 | 92 | 97 | 74 | 93 | 50 | 80 | 68 | 94 |
| One turn | 82 | 89 | 34 | 97 | 95 | 90 | 16 | 96 | 61 | 59 | 12 | 82 | 50 | 48 | 20 | 72 |
| Navigation | 80 | 86 | 14 | 92 | 94 | 84 | 2 | 90 | 24 | 40 | 3 | 70 | 47 | 44 | 6 | 68 |
| Nav. dynamic | 77 | 83 | 7 | 83 | 89 | 82 | 2 | 82 | 24 | 38 | 2 | 64 | 44 | 42 | 4 | 64 |

The simulation engine provides two different towns (Town 1, Town 2) and different weather conditions. Only Town 1 has been seen during training. The evaluation distinguishes between weather conditions that were seen during training and those that were not. For each combination of task, town, and weather condition, 25 different episodes are evaluated. We compare our Conditional Affordance Learning (CAL) approach to Conditional Imitation Learning (CIL) [22] as well as the Modular Pipeline (MP) and Reinforcement Learning (RL) approach presented in [30]. Note that the baselines ignore all red lights and target a general cruising speed of 20 km/h. In contrast, our agent stops for red lights and recognizes speed signs. However, to provide a fair comparison to prior art, we limit our algorithm to a cruising speed of 20 km/h for the quantitative evaluation against the baselines. We use the same model for all tasks and weather conditions without parameter fine-tuning.

4.1 Goal-Directed Navigation

Table 1 shows our quantitative comparison with respect to the state-of-the-art on CARLA on the task of goal-directed navigation. The table contains percentages of successfully completed episodes per task and per condition. Our agent outperforms the baselines in most of the tasks and under most conditions. As expected, performance decreases for all methods with increasing task difficulty. CAL particularly excels in generalizing to the new town, a scenario where most baselines did not perform well. For the new town under training weather conditions, the performance of the CAL agent is up to 68 % better than the best baseline.

CAL outperforms MP on almost all of the tasks except for the tasks in the training town under new weather conditions where MP performs slightly better. In almost all tasks and conditions, CAL shows better performance compared to CIL. Prediction of the relative angle ψ especially profits from the temporal structure of the CAL network as shown in the supplementary material. Also, the additional parameters of the controller (e.g., damping D) improve driving performance. CAL is outperformed by CIL only on one task: driving straight under the training weather conditions in the test town. Our analysis revealed that this happens due to a systematic failure of the CAL perception stack in one specific weather condition where several trees cast shadows onto the street, affecting the prediction of the distance d to centerline such that the agent occasionally drifts off the street.

4.2 Infraction Analysis

Table 2 shows the average distances driven between two infractions during the hardest task of the benchmark: navigation with dynamic obstacles. Under test conditions (New weather, new town, new weather & new town), our approach is able to outperform the baselines in almost every measure. This indicates the strong generalization ability of our approach in comparison to the baselines. CAL performs particularly well in avoiding collisions with other vehicles and pedestrians. In the new environment, our agent performs more than 10 times better than the best baseline (RL). During testing, pedestrians are crossing the street more frequently than in the training scenario. The fact that our agent is still able to perform well demonstrates that the hazard stop network was able to predict this affordance robustly despite its sparsity.

Occasionally, the CAL agent drives on the sidewalk. This infraction occurs primarily during right turns. Since rights turns in CARLA are generally sharper than left turns, the relative motion of pixels between consecutive frames is larger during a right turn. This may pose a bigger challenge to our perception network which is operating on sequential input. Driving on the sidewalk sometimes also results in crashes with static objects, e.g., traffic lights. Note that the infraction count for static

Table 2: **Infraction Analysis.** We show the average distance (in km) driven between two infractions. ‘>’ indicates that no infraction occurred over the whole course of testing. Note that the total driven distance depends on the amount and length of the episodes that an agent finishes.

| Task | Training conditions | | | | New weather | | | | New town | | | | New town and new weather | | | |
|-----------------------|---------------------|-------------|------|-------------|-------------|-------------|------|-------------|-------------|------|------|--------------|--------------------------|------|------|-------------|
| | MP | CIL | RL | CAL | MP | CIL | RL | CAL | MP | CIL | RL | CAL | MP | IL | RL | CAL |
| Opposite lane | 10.2 | 33.4 | 0.18 | 6.7 | 16.1 | 57.3 | 0.09 | > 60 | 0.45 | 1.12 | 0.23 | 2.21 | 0.40 | 0.78 | 0.21 | 2.24 |
| Sidewalk | 18.3 | 12.9 | 0.75 | 6.1 | 24.2 | > 57 | 0.72 | 6.0 | 0.46 | 0.76 | 0.43 | 0.88 | 0.43 | 0.81 | 0.48 | 1.34 |
| Collision: Static | 10.0 | 5.38 | 0.42 | 2.5 | 16.1 | 4.05 | 0.24 | 6.0 | 0.44 | 0.40 | 0.23 | 0.36 | 0.45 | 0.28 | 0.25 | 0.31 |
| Collision: Car | 16.4 | 3.26 | 0.58 | 12.1 | 20.2 | 1.86 | 0.85 | > 60 | 0.51 | 0.59 | 0.41 | 2.04 | 0.47 | 0.44 | 0.37 | 1.68 |
| Collision: Pedestrian | 18.9 | 6.35 | 17.8 | 30.3 | 20.4 | 11.2 | 20.6 | > 60 | 1.40 | 1.88 | 2.55 | 26.49 | 1.46 | 1.41 | 2.00 | 6.72 |



Figure 4: **Attention Analysis.** Left: The network focuses at the periphery for recognizing traffic lights. Middle: The network is anticipating a hazard stop by observing the shadow of a pedestrian. Right: The network focuses its attention on the pedestrian to correctly predict a hazard stop.

collisions is generally higher than that for driving on the sidewalk as the agent sometimes slides along a railing located next to the road, continuously increasing the infraction count.

4.3 Attention Analysis

In order to gain more insight into the decision making process of our perception stack we explore gradient-weighted class activation maps (Grad-CAMs) [41] which expose the implicit attention of the CNN with respect to the input image. Figure 4 shows the results for the affordances *hazard_stop* and *red_light*. We encode the prediction probability for a particular class with different colors (green means low probability, red means high probability). The attention given to a particular image region is encoded in the intensity channel. The stronger the attention, the more intense the color. Regions with very low attention are shown transparent for clarity.

Figure 4 (left) shows the agent driving through town 1 with the attention map for *red_light* overlaid. Note that the network focuses its attention on the road outlines and structures on the roadside where traffic lights are expected to be located. In Figure 4 (middle) the shadow of a pedestrian is visible at the bottom left of the image. Interestingly, the network anticipates an impending *hazard_stop* and focuses its attention to this region despite the fact that the pedestrian itself is not visible in the image. Finally, in Figure 4 (right) we illustrate a hazard stop situation. The network focuses all its attention on the pedestrian’s lower body. A possible explanation is that the network is able estimate the pedestrian’s walking direction by observing its legs and feet. Note that no pixel-wise supervision has been provided to the network during training. Instead, this behavior has been learned from binary image-level labels ($hazard_stop = True/False$) alone.

5 Conclusion

We have proposed Conditional Affordance Learning which combines the advantages of both modular pipelines and end-to-end approaches for navigating complex urban environments using high-level directional commands as input. We tested our approach extensively in simulation and demonstrated significant performance gains with respect to the state-of-the-art on CARLA. In the future, we plan to extend our work in several directions. Training on more diverse environments will likely improve the agent’s performance in all domains. An extended study of network architectures can improve the quality of the predictions and therefore improve the driving capabilities of the agent. We also plan to extend the sensor setup, e.g., using a stereo camera to improve distance prediction. More sophisticated control algorithms like model predictive control may further improve the performance, in particular when training the parameters of the controller jointly with the perception stack. Finally, inspired by recent work [42], we plan to transfer our results from simulation to the real world.

Acknowledgments

We would like to thank Vanessa Diedrichs for many helpful discussions about autonomous driving and data visualization and for her general support.

References

- [1] S. J. Russell and P. Norvig. *Artificial intelligence: a modern approach*. 2016.
- [2] C. Chen, A. Seff, A. Kornhauser, and J. Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *ICCV*, 2015.
- [3] R. Metz. And the award for most nauseating self-driving car goes to ... *MIT Technology Review*, Jan 2018.
- [4] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, and G. Hoffmann. Stanley: The robot that won the DARPA grand challenge. *Journal of Field Robotics*, 23(9):661–692, 2006.
- [5] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. Clark, J. Dolan, D. Duggins, T. Galatali, and C. Geyer. Autonomous driving in urban environments: Boss and the urban challenge. *Journal of Field Robotics*, 25(8):425–466, 2008.
- [6] J. Ziegler, P. Bender, M. Schreiber, H. Lategahn, T. Strauss, C. Stiller, T. Dang, U. Franke, N. Appenrodt, and C. G. Keller. Making Bertha drive — An autonomous journey on a historic route. *IEEE Intelligent Transportation Systems Magazine*, 6(2):8–20, 2014.
- [7] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [8] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016.
- [9] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [10] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *ICCV*, 2017.
- [11] Y. Li, H. Qi, J. Dai, X. Ji, and Y. Wei. Fully convolutional instance-aware semantic segmentation. In *CVPR*, 2017.
- [12] A. Kendall, H. Martirosyan, S. Dasgupta, and P. Henry. End-to-end learning of geometry and context for deep stereo regression. In *ICCV*, 2017.
- [13] J. Žbontar and Y. LeCun. Stereo matching by training a convolutional neural network to compare image patches. *Journal of Machine Learning Research (JMLR)*, 17(65):1–32, 2016.
- [14] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume. *arXiv.org*, 2017.
- [15] F. Güney and A. Geiger. Deep discrete flow. In *ACCV*, 2016.
- [16] H. Zhang, A. Geiger, and R. Urtasun. Understanding high-level semantics by modeling traffic patterns. In *ICCV*, 2013.
- [17] A. Geiger, M. Lauer, C. Wojek, C. Stiller, and R. Urtasun. 3d traffic scene understanding from movable platforms. *IEEE transactions on pattern analysis and machine intelligence*, 36(5):1012–1025, 2014.
- [18] W. Schwarting, J. Alonso-Mora, and D. Rus. Planning and decision-making for autonomous vehicles. *Annual Review of Control, Robotics, and Autonomous Systems*, 1(1):187–210, 2018.
- [19] D. A. Pomerleau. Alvin: An autonomous land vehicle in a neural network. In *NIPS*, 1989.

- [20] U. Muller, J. Ben, E. Cosatto, B. Flepp, and Y. L. Cun. Off-road obstacle avoidance through end-to-end learning. In *NIPS*, 2006.
- [21] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, and J. Zhang. End to end learning for self-driving cars. *arXiv.org*, 2016.
- [22] F. Codevilla, M. Müller, A. Dosovitskiy, A. Lopez, and V. Koltun. End-to-end driving via conditional imitation learning. In *ICRA*, 2018.
- [23] J. J. Gibson. *The senses considered as perceptual systems*. 1966.
- [24] M. Al-Qizwini, I. Barjasteh, H. Al-Qassab, and H. Radha. Deep learning algorithm for autonomous driving using GoogLeNet. In *Intelligent Vehicles Symposium*, 2017.
- [25] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [26] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv.org*, 2014.
- [27] J. Huang, I. Tanev, and K. Shimohara. Evolving a general electronic stability program for car simulated in TORCS. In *CIG*, 2015.
- [28] S. R. Richter, V. Vineet, S. Roth, and V. Koltun. Playing for data: Ground truth from computer games. In *ECCV*, 2016.
- [29] S. Ebrahimi, A. Rohrbach, and T. Darrell. Gradient-free policy architecture search and adaptation. In *CoRL*, 2017.
- [30] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. CARLA: An open urban driving simulator. In *CoRL*, 2017.
- [31] R. Caruana. Multitask learning: A knowledge-based source of inductive bias. In *ICML*, 1993.
- [32] H. Xu, Y. Gao, F. Yu, and T. Darrell. End-to-end learning of driving models from large-scale video datasets. In *CVPR*, 2017.
- [33] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.
- [34] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv.org*, 2015.
- [35] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [36] K. Cho, B. van Merriënboer, Ç. Gülçehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv.org*, 2014.
- [37] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv.org*, 2014.
- [38] M. Müller, V. Casser, J. Lahoud, N. Smith, and B. Ghanem. Sim4cv: A photo-realistic simulator for computer vision applications. *International Journal of Computer Vision*, 2018.
- [39] J. Kosecka, R. Blasi, C. J. Taylor, and J. Malik. A comparative study of vision-based lateral control strategies for autonomous highway driving. In *ICRA*, 1998.
- [40] K. H. Ang, G. Chong, and Y. Li. Pid control system analysis, design, and technology. *IEEE transactions on control systems technology*, 13(4):559–576, 2005.
- [41] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. *arXiv.org*, 2016.
- [42] M. Müller, A. Dosovitskiy, B. Ghanem, and V. Koltun. Driving policy transfer via modularity and abstraction. *arXiv.org*, 2018.