

## SelfD: Self-Learning Large-Scale Driving Policies From the Web

Jimuyang Zhang Ruizhao Zhu Eshed Ohn-Bar  
 Boston University

{zhangjim, rzhu, eohnbar}@bu.edu

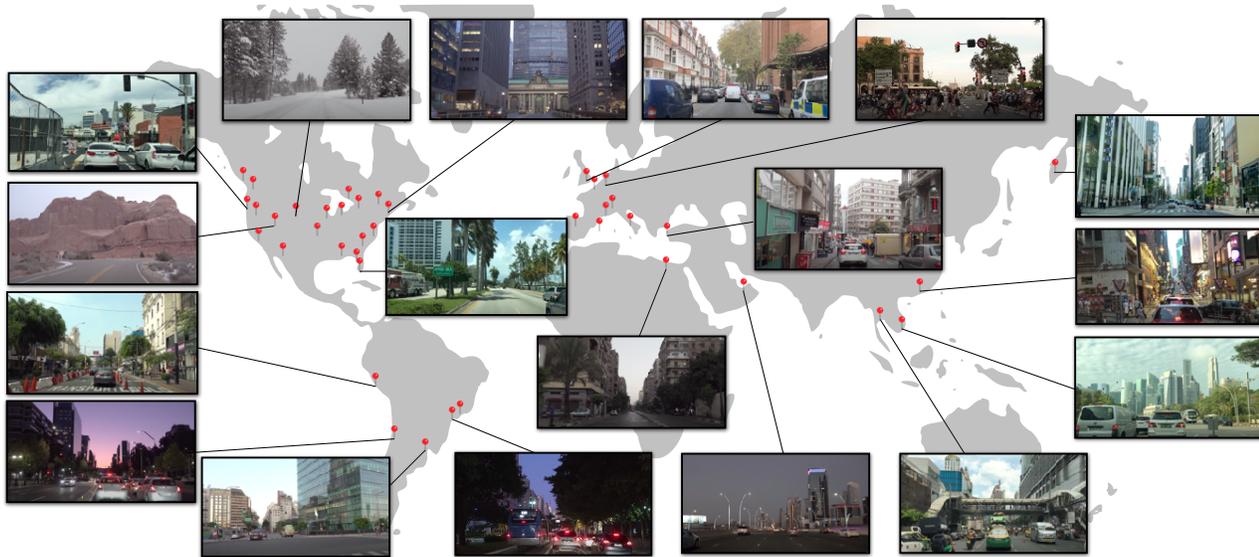


Figure 1. **Learning From Unlabeled Data Facilitates Scalable Navigation Decision-Making.** Our goal is to develop robust, generalized, and easily deployable decision-making policies for navigation. Our key insight is to make use of the hours of freely available and highly diverse navigation data from the web in order to augment the knowledge and robustness of an initially trained navigation policy.

### Abstract

Effectively utilizing the vast amounts of ego-centric navigation data that is freely available on the internet can advance generalized intelligent systems, i.e., to robustly scale across perspectives, platforms, environmental conditions, scenarios, and geographical locations. However, it is difficult to directly leverage such large amounts of unlabeled and highly diverse data for complex 3D reasoning and planning tasks. Consequently, researchers have primarily focused on its use for various auxiliary pixel- and image-level computer vision tasks that do not consider an ultimate navigational objective. In this work, we introduce SelfD, a framework for learning scalable driving by utilizing large amounts of online monocular images. Our key idea is to leverage iterative semi-supervised training when learning imitative agents from unlabeled data. To handle unconstrained viewpoints, scenes, and camera parameters, we train an image-based model that directly learns to plan in the Bird’s Eye View (BEV) space. Next, we use unlabeled

data to augment the decision-making knowledge and robustness of an initially trained model via self-training. In particular, we propose a pseudo-labeling step which enables making full use of highly diverse demonstration data through “hypothetical” planning-based data augmentation. We employ a large dataset of publicly available YouTube videos to train SelfD and comprehensively analyze its generalization benefits across challenging navigation scenarios. Without requiring any additional data collection or annotation efforts, SelfD demonstrates consistent improvements (by up to 24%) in driving performance evaluation on nuScenes, Argoverse, Waymo, and CARLA.

### 1. Introduction

How can we learn generalized models for robust vision-based navigation in complex and dynamic settings? While humans can effortlessly transfer general navigation knowledge across settings and platforms [36, 46, 52, 53], current real-world development of navigation agents generally deploys within a fixed pre-assumed setting (e.g., ge-

ographical location, use-case) and carefully calibrated sensor configurations. Consequently, each autonomy use-case generally requires its own prohibitive data collection and platform-specific annotation efforts [4, 6, 19, 56, 79]. Due to such development bottlenecks, brittle navigation models trained in-house by various developers (e.g., Tesla’s Autopilot [23], Waymo’s Driver [4], Amazon’s Astro [69], Fedex’s Roxo [25], etc.) are easily confounded by the immense complexity of the real-world navigation task, e.g., rare scenarios, novel social settings, geographical locations, and camera mount perturbations. Yet, every minute, a vast amount of highly diverse and freely available ego-centric navigation data containing such scenarios is uploaded to the web. In this paper, we work towards effectively utilizing such freely available demonstration data to improve the efficiency, safety, and scalability of generalized real-world navigation agents.

There are two key challenges to employing the large amounts of unconstrained and unlabeled online data for training robust vision-based navigation policies. First, while online images may be collected in various layouts and camera settings, existing monocular-based prediction and decision-making methods tend to rely on restrictive assumptions of planar scenes [22, 70] and known camera parameters [13, 28, 41, 55, 59, 71, 77]. Towards a dataset and platform-agnostic navigation agent, our proposed architecture does not explicitly rely on such assumptions. Second, due to safety-critical requirements, methods for learning decision-making in complex navigation settings generally also assume access to highly curated benchmarks with clean annotations [4, 7, 10, 13, 17, 20, 26, 33, 64, 74]. Consequently, such methods must be revisited when learning from unlabeled and diverse internet videos, e.g., with various quality demonstrations [74]. To effectively utilize such data, we propose to leverage recent advances in iterative semi-supervised learning [8, 38, 60]. Yet, as such techniques are studied within pixel and image-level tasks [14, 29, 41], their utility for learning complex and safety-critical 3D reasoning and planning tasks is not well-understood.

**Contributions:** Motivated by how humans are able to effortlessly learn and adapt various skills using online videos, our SelfD approach introduces three main contributions. *First*, to facilitate learning from unconstrained imagery, we develop a model for mapping monocular observations directly to a Bird’s Eye View (BEV) planning space (i.e., without requiring camera calibration). *Second*, we introduce a novel semi-supervised learning approach incorporating self-training with “hypothetical” data augmentation. We demonstrate the proposed sampling step to be essential for making use of highly diverse demonstration data. *Third*, we perform a set of novel cross-dataset experiments to analyze the ability of an initially trained decision-making model to self-improve its generalization capabilities using minimal

assumptions regarding the underlying data. We evaluate across various datasets with complex navigation and harsh conditions to demonstrate state-of-the-art model generalization performance.

## 2. Related Work

**Observational Imitation Learning:** Our key idea is to leverage the scale and diversity of easily available online ego-centric navigation data to learn a robust conditional imitation learning policy [13, 17]. While learning from labeled demonstrations can significantly simplify the challenging vision-based policy learning task [3, 6, 11, 12, 30, 31, 39, 42, 47, 48, 52, 56, 58, 78, 81, 82], observed images in our settings are not labeled with corresponding actions of a demonstrator. We therefore work to generalize current conditional imitation learning (CIL) approaches [13, 17, 18] to learn, from unlabeled image observations, an agent that can navigate in complex urban scenarios. To address this challenging observational learning task, prior work has recently explored introducing various restrictive assumptions, including access to a hand-designed reward function [9], an interactive environment for on-policy data collection [66], or demonstrator optimality [66, 67]. Moreover, due to unconstrained and dynamic scenes in our settings, learning an inverse model [9, 65–67] (i.e., the visual odometry task of recovering actions from observation sequences) becomes challenging and may result in noisy trajectories. To facilitate scalable training from diverse data sources, our proposed semi-supervised learning approach does not leverage such assumptions. Nonetheless, our resulting model can also be used to bootstrap other methods for policy training, e.g., model-based or model-free reinforcement learning approaches [12, 42, 51, 68].

**Semi-Supervised Learning for Navigation:** CIL generally involves learning from known actions of human experts [17, 18, 58], which is not applicable to our settings. Yet, the recent work of Chen et al. [13] utilizes a multi-stage training step, where a privileged (i.e., ‘teacher’) CIL agent is employed to provide supervision to a non-privileged (i.e., ‘student’) visuomotor CIL agent. As the privileged agent is given access to extensive ground truth information about the world in training and testing, it produces highly plausible and clean trajectories. In contrast, SelfD leverages *the same visuomotor architecture* as teacher and student. We also train in inherently noisy settings, as teacher inference is performed on diverse out-of-distribution image data and not on the original training dataset. Consequently, our approach goes beyond prior policy distillation approaches to handle scenarios where supervision by the teacher model may be potentially noisy and unsafe. We also note the relationship between such distillation and semi-supervised training via pseudo-labeling [8, 38, 43, 60, 62]. However, as far as we are

aware, we are the first to develop a pseudo-labeling based self-training method for learning safe driving policies from complex scenes with diverse navigation data, camera perspectives, geographical locations, and weathers. In contrast, prior work [38, 60] emphasizing semi-supervised learning for image- and object-level recognition tasks [8, 76] has limited utility for complex decision-making tasks as further discussed in Sec. 4.

**Self-Supervised Visual Learning:** Our approach aims to directly leverage experience from large amounts of unlabeled video data to learn complex 3D navigation. An alternative approach to our proposed semi-supervised framework involves learning visual representations from the large unlabeled data [15, 27, 29, 32, 44, 45, 54, 63]. These generic representations can then be transferred to the policy learning task [29]. However, visual representation learning from unlabeled data often rely on various auxiliary image-level data augmentation strategies [2, 50, 72, 80], e.g., Jigsaw or Colorization [29] tasks, that are indirect to the ultimate navigation task. Thus, such methods are limited in utility for elaborate spatial navigation tasks in dynamic settings. Indeed, the utility of image augmentation-based methods beyond navigation in simple static environments (e.g., [29]) has yet to be demonstrated, i.e., for autonomous driving. In contrast, SelfD’s learning and augmentation strategies directly optimize for BEV (map) planning. Moreover, our approach is orthogonal to the aforementioned studies as indirect visual tasks can be leveraged as further pre-training prior to the proposed self-supervised pseudo-labeling step.

### 3. Method

Our goal is to facilitate training driving policies at scale. To efficiently make use of the broad and diverse experience found in large amounts of unlabeled videos, we follow three main steps. We first propose to use a monocular image-based behavior cloning planner that reasons directly in the BEV (Sec. 3.3). Our proposed planner can therefore better generalize across arbitrary perspectives. Next, we introduce a data augmentation step for obtaining multiple plausible pseudo-labels when self-training over unlabeled internet data (Sec. 3.4). Finally, we re-train the model over the larger dataset to learn a more robust and generalized vision-based navigation policy (Sec. 3.5).

#### 3.1. Problem Setting

We consider the task of learning from observations  $\mathbf{x} = (\mathbf{I}, v, c) \in \mathcal{X}$  of a single front camera image  $\mathbf{I} \in \mathbb{R}^{W \times H \times 3}$ , ego-vehicle speed  $v \in \mathbb{R}$ , and a categorical navigational command  $c \in \mathbb{N}$  (e.g., turn left, turn right, and forward [17]). Our agent learns to map such observations to a navigational decision. In general, the decision may either be a low-level vehicle control action [18] (e.g., steering, throt-

tle) or a desired future trajectory relative to the ego-vehicle, i.e., set of  $K$  waypoints  $\mathbf{y} \in \mathcal{Y}$  [13, 47] in the BEV (map) space. In the latter case, future waypoints may be paired with a hand-specified or learned motion controller to produce the low-level action [13, 47]. In this work, we focus on the latter representation due to its interpretability and generalizability [47]. Thus, our goal is to obtain a waypoint prediction function  $f_\theta: \mathcal{X} \rightarrow \mathcal{Y}$ , with learnable parameters  $\theta \in \mathbb{R}^d$ , i.e., the decision-making policy. Without loss of generality, we will slightly augment the output space of this function in Sec. 3.3. Next, we build on recent advances in driving policy learning, in particular CIL [13, 17, 18], to develop a general approach for learning to drive from unlabeled data.

**Conditional Imitation Learning:** Navigational demonstrations can be collected by logging the sensor data of a manually operated system, as often done by researchers and developers. Such demonstrations can then be utilized for training decision-making policies through various learning techniques [1, 4, 16, 17, 52, 56, 61]. In its most straightforward implementation, training the mapping function  $f_\theta$  can be simplified to supervised learning via an i.i.d. data assumption [5, 56]. More robust policies can be trained by going beyond such schemes, e.g., through online and interactive learning strategies [52, 57, 61]. In our work, we leverage off-line demonstrations and do not assume access to an interactive environment. Nonetheless, our proposed data augmentation techniques in Sec. 3.4 can significantly improve policy robustness, as will be demonstrated in Sec. 4.

Given a collected dataset of observations and corresponding actions,  $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ , supervised training can be achieved by optimizing an imitation objective

$$\underset{\theta}{\text{minimize}} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\mathcal{L}(\mathbf{y}, f_\theta(\mathbf{x}))] \quad (1)$$

with  $\mathcal{L}$  being a suitable loss function, e.g.,  $L_1$  regression distance to the waypoint targets. Despite widespread use in academia [18, 34, 75] and industry [4, 6], the standard CIL formulation significantly limits scalability. Specifically, approaches for imitation learning generally require having direct access to optimal target action labels. Consequently, they cannot be applied towards learning from the vast and diverse unlabeled data that is uploaded online.

**Conditional Imitation Learning from Observations:** To make full use of unlabeled data in demonstrations containing diverse navigational experience, we develop a framework for Conditional Imitation Learning from Observations (CILfO), i.e., assuming only data of the form  $\mathcal{U} = \{\mathbf{I}_i\}_{i=1}^M$ . Within this more generalized yet difficult learning task, a key challenge lies in recovering suitable label targets  $\hat{\mathbf{y}}$ , navigational command  $\hat{c}$  and speed  $\hat{v}$  to construct a dataset

$$\hat{\mathcal{D}} = \{(\mathbf{I}_i, \hat{v}_i, \hat{c}_i, \hat{\mathbf{y}}_i)\}_{i=1}^M \quad (2)$$

over which Eqn. 1 can then be used to train a policy.

**Initial Data Assumption:** In order to address the challenging CILfO learning task, our key idea is to leverage a small labeled dataset to learn an initial policy mapping using human expert demonstrations. We then sample from this trained function to obtain pseudo-labels [38] over the unlabeled data. This assumption is reasonable considering that there are several publicly available driving datasets with included action labels [7,26,33,73]. Alternatively, a preliminary dataset with a novel platform setup or use-case may be initially collected.

**Sequential Data Assumption:** We note that although sequential video data is used in our ablation, our generic and scalable CILfO formulation does not assume access to temporal data. The reason for this is three-fold. First, while sequential observations may be potentially beneficial for disambiguation of various decision-making factors [21], generalization results for learning temporal sensorimotor driving policies have been inconclusive [21, 48, 70]. Second, while sequential data is assumed in most prior observational imitation learning settings [9,66] recovering the underlying demonstration actions (e.g., using monocular visual odometry [71]) in arbitrary scenes is challenging. Due to the difficulty and noise in sequential action recovery, our proposed single-frame formulation in Sec. 3.2 is shown to significantly outperform such baselines in Sec. 4. Third, we do not make any assumptions regarding the optimality of the demonstrations in the unlabeled dataset  $\mathcal{U}$  [49,67]. Consequently, our generalized method is suitable for scalable learning from videos with complex scenes, environmental conditions, arbitrary viewpoints, and diverse demonstration quality. As our learning task lies well beyond the capabilities of existing CIL and observational methods, we next develop a generalized training method for leveraging unconstrained and unlabeled demonstration data.

### 3.2. SelfD Overview

Through a semi-supervised policy training process, our proposed SelfD navigation policy model can be learned in three summarized steps:

1. Use a small, labeled domain-specific dataset  $\mathcal{D}$  to learn an initial observations-to-BEV policy  $f_\theta$  via imitation.
2. Obtain a large pseudo-labeled dataset  $\hat{\mathcal{D}}$  by leveraging sampling from  $f_\theta$ .
3. Pre-train a generalized policy  $f_\theta$  on  $\hat{\mathcal{D}}$  and fine-tune on the clean labels of  $\mathcal{D}$ .

Note that we re-use the parameter symbol  $\theta$  throughout the steps to simplify notation. Our iterative semi-supervised training enables effectively augmenting the knowledge and robustness of an initially trained policy. As described

next, our proposed initial step facilitates learning a platform and perspective-agnostic policy during subsequent training steps by directly reasoning in a BEV planning space.

### 3.3. BEV Plan Network

In this section, we propose a suitable output representation to account for arbitrary cameras, viewpoints and scene layouts. Current monocular planners generally predict waypoints in the image plane to align with an input image [13,47]. The waypoints are then transformed to a BEV plan using carefully calibrated camera intrinsic and extrinsic (e.g., rotation, height) parameters [13]. Thus, policy models are often trained and evaluated within a fixed pre-assumed setup. In contrast, SelfD predicts a future plan parameterized by waypoints in the BEV plan space directly. Based on our experiments in Sec. 4, we demonstrate this choice to be crucial for real-world planning across settings. The predicted generalized BEV waypoints can be paired with a low-level controller, e.g., a PID controller [13,47].

Due to the difficulty in learning a monocular-to-BEV plan mapping, we follow recent work in confidence-aware learning [37] to train an augmented model  $f_\theta: \mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{R}$  with quality estimates  $\sigma \in \mathcal{R}$ . Our training loss function in Eqn. 1 is defined as

$$\mathcal{L} = \mathcal{L}_{\text{plan}} + \lambda \mathcal{L}_{\text{quality}} \quad (3)$$

where  $\mathcal{L}_{\text{plan}}$  is the  $L_1$  distance between ground-truth and predicted waypoints,  $\mathcal{L}_{\text{quality}}$  is a binary cross-entropy loss [37,76], and the  $\lambda$  hyper-parameter balances the tasks.

### 3.4. “What If” Pseudo-Labeling of Unlabeled Data

Given a set of unlabeled images  $\mathcal{U}$ , we sample from the trained conditional policy  $f_\theta$  in a semi-supervised training process. While the speed and command inputs to  $f_\theta$  can be recovered through visual odometry techniques [71], these result in highly noisy trajectories in our online video settings (discussed in Sec. 4). As the demonstrations in our data may be unsafe or difficult to recover, we propose to leverage a single-frame pseudo-labeling mechanism. Our key insight is to employ the conditional model  $f_\theta$  to generate multiple hypothetical future trajectories in a process referred to as “what if” augmentation. Beyond resolving the missing speed and command inputs, our proposed augmentation provides additional supervision, i.e., a conditional agent that better reasons on what it *might need to do*, for instance, if it had to turn left instead of right at an intersection (Fig. 2). In contrast to related work in policy learning and distillation [13], sampling from the teacher agent is more challenging as the agent is not exposed to extensive 3D perception knowledge about the world and is being evaluated outside of its training settings.

We repeatedly sample  $\hat{v}$  and  $\hat{c}$  uniformly and rely on the conditional model to provide pseudo-labels  $(\hat{y}, \hat{\sigma}) =$



Figure 2. “What If” Hypothetical Pseudo-Labeling. We generate multiple plausible future trajectories (units are meters in BEV) as pseudo-labels for each unlabeled frame by sampling the conditional planner from Sec. 3.3. We depict two scenarios from our dataset with inference shown for various inputs. Speed is in meters per second and conditional commands are either left (1), forward (2), or right (3).

$f_{\theta}(\mathbf{I}, \hat{v}, \hat{c})$  for additional supervision across all conditional branches and speed observations. In this manner, querying the “teacher” model  $f_{\theta}$  enables us to generate various scenarios beyond the original demonstration. In particular, as discussed in Sec. 4, we find self-training strategies to provide limited generalization gains without this “what if” data augmentation step. This augmentation strategy enables our single-frame pseudo-labeling approach to significantly outperform approaches that are more elaborate to train at scale, as they may involve additional modules relying on approximating  $\hat{y}$ ,  $\hat{c}$ , and  $\hat{v}$  from video. Finally, to avoid incorporating potentially noisy trajectories, the corresponding quality estimates  $\hat{\sigma}$  can be used to process and filter examples in the pseudo-labeled dataset  $\hat{\mathcal{D}}$ .

### 3.5. Model Pre-Training and Fine-Tuning

As a final training step, we re-train the waypoint network  $f_{\theta}$  from scratch over the large and diverse dataset  $\hat{\mathcal{D}}$ . The pre-trained policy can then be further fine-tuned over the original dataset  $\mathcal{D}$ , thus leveraging the additional knowledge gained from  $\hat{\mathcal{D}}$  to improve its performance. We note that we employ separate training over the two datasets  $\hat{\mathcal{D}}$  and  $\mathcal{D}$  and rely on knowledge transfer through learned representations as it reduces the need for any careful hyperparameter tuning beyond the overall learning rate. For instance, Caine et al. [8] empirically demonstrate the importance of delicately optimizing the ratio of labeled to pseudo-labeled data when mixing the datasets for a 3D object detection task, while also showing it to vary among object categories, e.g., pedestrians vs. vehicles. Thus, through a pre-training mechanism, we avoid the need to carefully mix the cleanly

labeled and pseudo-labeled datasets [8, 63, 76].

## 3.6. Implementation Details

We implement our BEV waypoint prediction network  $f_{\theta}$  leveraging a state-of-the-art conditional sensorimotor agent [13]. However, as discussed in Sec. 3.2, we do not assume a fixed known BEV perspective transform. Thus, we remove the fixed perspective transformation layer which restricts scalability and replace it with a per-branch BEV prediction module (see supplementary for additional implementation and architecture details). During training we use a learning rate of  $1e-3$  and resize images to  $400 \times 225$ .

## 4. Experiments

### 4.1. Experimental Setup

To obtain large amounts of driving data, we downloaded 100 hours of front view driving videos from popular driving channels on YouTube (across cities, weathers, and times of day, as shown in Fig. 1). In our real-world evaluation, we use the nuScenes [7], Waymo [64] and Argoverse [10] datasets. While nuScenes is a highly instrumented and annotated dataset, it is primarily used for perception tasks without an official benchmark for future waypoint planning. The recent work by Hu et al. [35] uses a random split, yet this is not suitable for studying generalization (the method also employs LiDAR input whereas we do not). We create a geography-based split into nuScenes Boston (nS-Boston) and nuScenes Singapore (nS-Singapore) which is challenging due to the significant domain shift. To further evaluate generalization, we also utilize the Waymo and Argoverse datasets which were collected across 8 different cities. Specifically, we construct a future waypoint prediction benchmark from the Waymo perception testing set and the Argoverse 3D tracking training set. Note that in our evaluation we do not assume any access to the target domain’s pseudo-labels, as often done in related research. Overall, our split results in 13K, 11K, 11K, and 26K samples from nS-Boston, nS-Singapore, Argoverse and Waymo, respectively.

The open-loop real-world evaluation compares predictions to an expert driver in complex interactions (i.e., diverse maneuvers, yielding, merging, irregular intersections). To analyze the impact of the proposed approach during closed-loop driving, we further perform interactive on-policy evaluation in CARLA [18, 24]. We replicate our training settings in simulation and predict a final low-level action by employing a PID controller [13]. In particular, we train on labeled data from Town 1 and evaluate on Town 2 in regular traffic (following NoCrash [18]). To understand the impact of pseudo-labeling on the close-loop driving metrics we keep a portion of the training town data for pseudo-labeling. The evaluation is conducted over 25 pre-

defined routes in Town 2 under four different weathers, one of which is not seen in annotated data. Additional details regarding the experimental setup can be found in the supplementary.

**Evaluation Metrics:** We follow standard open-loop evaluation and use Average and Final  $L_2$  Displacement Error (ADE and FDE, respectively) over future waypoints in the BEV [40]. We also compute a collision rate [34] which measures collisions along the predicted waypoints with other vehicles (we are only able to compute this metric on nuScenes and Argoverse due to the provided annotations for the surrounding vehicles). In our CARLA experiments we also report agent Success Rate (SR), Route Completion (RC), and Collision (Coll.) frequency per 10km.

## 4.2. Results

**Model Architecture:** We first analyze the proposed BEV planning model architecture on the nuScenes cross-town split. The architecture is then fixed for subsequent experiments. As shown in Table 1, the state-of-the-art CIL model [13] achieves a 1.86 ADE on the cross-town training and testing split. We emphasize that the CIL baseline [13] predicts the waypoints in 2D images and transform the predicted waypoints into BEV through a perspective transformation. However, such transformation may not be accurate in the real-world. We can therefore see the benefits of the BEV planner, even before pseudo-labeling techniques have been applied. Our proposed BEV planner achieves a 1.14 ADE (a nearly 40% reduction in error compared to CIL). Given prior work on nuScenes primarily relies on non-conditional approaches based on LiDAR [34], we also ablate over different inputs to the model. Intuitively, we find the speed measurements to be critical to BEV waypoint prediction. Our proposed BEV planner with both speed and conditional command input gives the best ADE performance on nS-Singapore. We now continue to analyze the benefits of a semi-supervised learning step to improve policy performance.

**Pseudo-Labeling Approaches:** We consider various approaches for leveraging the unlabeled YouTube data (we select 10 driving hours for the ablation). To emphasize generalization, we do not pseudo-label the unseen test datasets or incorporate their unlabeled samples into the self-training in any way. We report waypoint prediction performance before and after the final fine-tuning step on nS-Boston across all test sets. As shown in Table 2, pre-training over the pseudo-labeled YouTube data and fine-tuning on the clean dataset consistently improves ADE, FDE, and collision metrics across evaluation settings. We see how leveraging a state-of-the-art Visual Odometry (VO) model [71] as a teacher for SelfD instead of the proposed pseudo-labeling mechanism results in reduced performance. As recovering

Table 1. **Model Architecture.** The CIL baseline [13] predicts waypoints in the image and relies on perspective transform with known camera calibration, while our proposed monocular-to-BEV planner does not have access to camera intrinsic or extrinsic parameters. All models are trained on nuScenes Boston and evaluated on nuScenes Singapore.

Models	Command	Speed	ADE (meters)
CIL Baseline [13]	✓	✓	1.86
BEV Planner	✓		4.77
		✓	4.64
	✓	✓	<b>1.14</b>

the camera motion in dynamic scenes with unknown camera parameters is difficult, leveraging this model results in noisy and unsafe trajectories. While we find limited utility in leveraging temporal data, this may be alleviated by improved visual odometry methods in the future. As such methods become more robust to arbitrary settings, they may be integrated to complement our method, e.g., to leverage both rollouts as well as the proposed data augmentation. We also find our pre-training pipeline to improve over mixing-based approaches [63].

**What If Augmentation:** Results in Table 2 demonstrate how self-training without the proposed data augmentation results in limited gains. Here, we experiment with pseudo-labeling using various sampling strategies for speed measurements and conditional commands. In addition to uniform and prior-based sampling (determined empirically based on the training data), we also report findings using the visual odometry [71] model estimates for the speed and conditional command. We find uniform sampling to outperform more informed prior and visual odometry-based sampling due to the extensive speed profile and multi-branch supervision it provides. We can see how a model learned through self-training but without such powerful augmentation (i.e., SelfD without What if in Table 2) results in worse performance compared to the baseline BEV planning model, e.g., 1.18 vs. 1.14 ADE and 2.23 vs. 2.16 FDE. We also find the augmentation to improve safety of the predicted trajectories, as shown by reduced collision rates. This is affirmed in our closed-loop evaluation shown in Table 5.

**Generalization to Harsh Settings:** We stress test SelfD in Table 3 under the most challenging subsets across the datasets, as selected by video episodes collected in certain environmental conditions. Interestingly, we find SelfD to efficiently reduce error on nighttime conditions, which is only seen in YouTube videos. However, as nS-Boston already contains some scenes with cloudy and light rain settings, the improvement on rainy conditions on Waymo is

Table 2. **Pseudo-Labeling Approaches.** We compare various pseudo-labeling approaches by varying the underlying sampling mechanism for the speed and conditional inputs. These may either be estimated with visual odometry [71] (i.e., SelfD without “what if” augmentation) or sampled from a prior distribution (i.e., uniformly or with a prior). We find uniform sampling to work best as a form of data augmentation. We also compare to a baseline using a state-of-the-art visual odometry method for pseudo-labels. Note that DVO-based approaches utilize sequential video data while the rest are single-frame. “Coll.” refers to collision rate and “nS” to nuScenes (see Sec. 4.1).

Method		nS-Singapore			Argoverse			Waymo	
		ADE	FDE	Coll. (%)	ADE	FDE	Coll. (%)	ADE	FDE
BEV Planner (nS-Boston)		1.14	2.16	9.50	1.07	2.16	16.50	2.17	3.87
Pre-Trained (YouTube)	SelfD (VO)	4.63	8.07	12.91	4.35	7.54	31.87	5.16	9.02
	SelfD (w/o What If)	1.32	2.45	10.02	1.27	2.46	<b>13.92</b>	2.43	4.28
	SelfD (What If-Prior)	1.19	2.40	10.01	1.12	2.34	16.51	2.10	3.94
	SelfD (What If-Uniform)	1.19	2.28	10.10	1.13	2.28	15.70	2.24	4.01
Fine-Tuned (YouTube & nS-Boston)	FixMatch [63]	1.18	2.22	12.02	1.10	2.24	15.20	2.18	3.88
	SelfD (VO)	1.09	2.09	9.33	1.03	2.11	17.52	2.01	3.62
	SelfD (w/o What If)	1.18	2.23	9.34	1.09	2.18	16.26	2.03	3.62
	SelfD (What If-Prior)	1.09	2.10	9.47	1.00	2.06	15.71	2.10	3.73
	SelfD (What If-Uniform)	<b>1.00</b>	<b>1.93</b>	<b>9.30</b>	<b>0.99</b>	<b>2.05</b>	14.20	<b>1.65</b>	<b>3.02</b>

Table 3. **Generalization to Harsh Conditions.** We analyze the impact of SelfD on generalization from nS-Boston to specific environmental conditions. We compute the performance over domain subsets within each dataset. While insightful for the larger nuScenes and Waymo splits, we note that this splits leads to a small dataset on Argoverse with limited insights.

Method	nS-Singapore Night		Argoverse Rainy		Argoverse Night		Waymo Rainy		Waymo Night	
	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE
BEV Planner (nS-Boston)	1.67	3.05	0.81	<b>1.71</b>	<b>0.78</b>	<b>1.61</b>	0.70	<b>1.28</b>	1.81	3.25
SelfD (Pre-Trained)	1.73	3.17	0.84	1.85	0.94	1.94	0.77	1.42	1.96	3.55
SelfD (Fine-Tuned)	<b>1.28</b>	<b>2.42</b>	<b>0.79</b>	<b>1.71</b>	0.91	1.80	<b>0.67</b>	1.32	<b>1.53</b>	<b>2.80</b>

Table 4. **Impact of Initial Training Dataset.** We analyze the impact of different datasets for training the initial BEV planner model.

Initial Dataset	Method	nS-Singapore	
		ADE	FDE
Argoverse	BEV Planner	1.42	2.58
	SelfD (w/o FT)	1.38	2.52
	SelfD (w/ FT)	<b>1.18</b>	<b>2.23</b>
Waymo	BEV Planner	1.08	2.03
	SelfD (w/o FT)	1.13	2.14
	SelfD (w/ FT)	<b>0.91</b>	<b>1.81</b>

limited. While we show results on similar Argoverse settings for consistency and completeness, our harsh settings split results in very few samples (in particular for Argoverse Night). Hence, results become less meaningful. While insightful, this experiment demonstrates the difficulty in obtaining diverse annotated data, further motivating our SelfD approach for development and real-world scalability.

**Impact of Initial Training Data:** We further investigate the role of the initial training dataset on the pseudo-labeling step and resulting model generalization performance in

Table 5. **Closed-Loop Evaluation in CARLA.** We analyze the impact of the proposed self-training step on closed-loop metrics of Success Rate (SR), Route Completion (RC), and Collision frequency (Coll.) per 10km.

Method	ADE ↓	FDE ↓	SR (%) ↑	RC (%) ↑	Coll. ↓
CIL Baseline [13]	0.57	1.30	16	66.88	10.71
BEV Planner	0.66	1.32	12	65.66	11.33
SelfD (Town 1)	0.56	1.30	25	71.84	11.51
SelfD (Town 1 & 2)	<b>0.55</b>	<b>1.24</b>	<b>26</b>	<b>74.50</b>	<b>5.18</b>

Table 4. Here, we repeat the experiments by replacing nS-Boston with Argoverse and Waymo and test on nS-Singapore (i.e., for direct comparison with Table 2). We keep all hyper-parameters fixed. While models trained on different dataset may bias differently the pseudo-labels on the YouTube dataset, overall we find similar trends in performance regardless of the original training dataset.

**Closed-Loop Evaluation in CARLA:** We show our analysis in CARLA in Table 5. We report both open-loop and closed-loop metrics for completeness. In this experiment, we train our BEV planner using an annotated portion of Town 1, and test on Town 2. We note that while our CIL

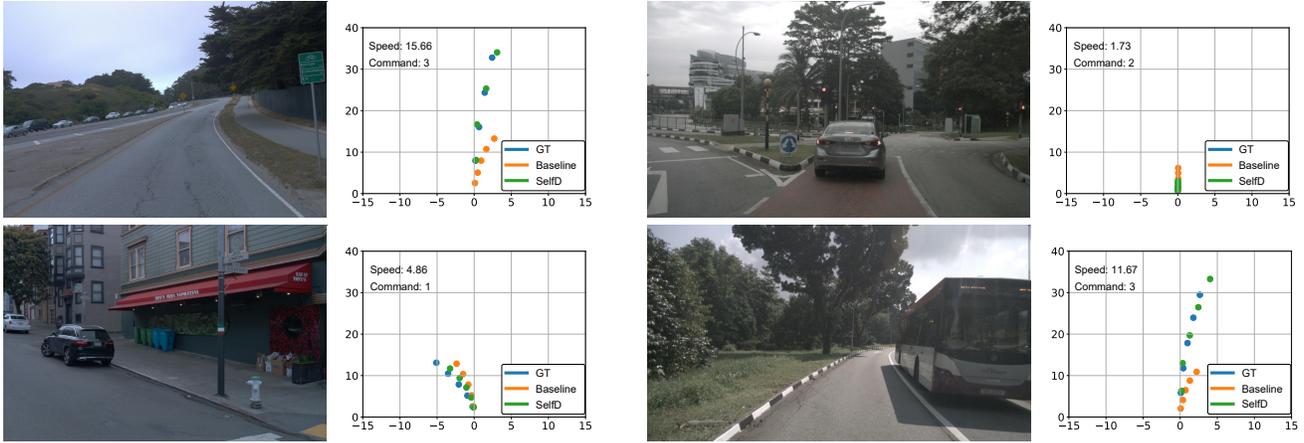


Figure 3. **Qualitative Results.** We compare predicted waypoints in BEV for the ground truth trajectory, the baseline model (trained on nS-Boston), and SelfD (pre-trained on YouTube, fine-tuned on nS-Boston). Results are shown on Waymo (left) and nS-Singapore (right). The SelfD model (green) is shown to generalize well across unusual scenarios due to the proposed pseudo-labeling mechanism. We also observe improved robustness to rare command and speed inputs. **Units are meters** for the BEV plan plots.

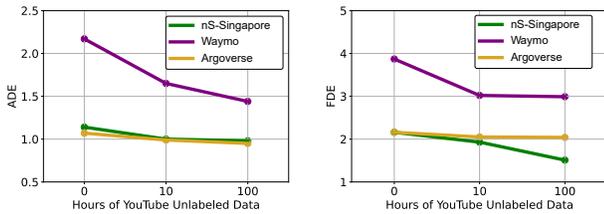


Figure 4. **Additional Unlabeled Data.** Results when increasing the amount of pseudo-labeled YouTube driving data up to 100 hours. As our proposed self-training pipeline facilitates scalable learning from large amounts of data, the model can continue to improve with additional YouTube data.

Baseline benefits from the known fixed perspective transform assumption the BEV planner has no knowledge of camera parameters and must learn it from the data. When training SelfD using pseudo-labeled data from Town 01, we achieve a near doubling in the success rate, from 12% to 25%. This increase in model performance using the most difficult closed-loop evaluation metric highlights the benefits of the proposed approach. We further pseudo-label and use in pre-training an unlabeled dataset from Town 2 (i.e., SelfD - Town 1 & 2), showing additional gains in driving performance across metrics. This experiment suggests multiple ways in which the proposed approach may be leveraged in real-world generalization and adaptation settings.

**Additional YouTube Pre-Training Data:** While we leverage 10 hours of YouTube data in our ablation analysis, we provide results with additional YouTube data of up to 100 hours. The results, shown in Fig. 4, suggest further gains can be made with larger and more diverse pre-training

datasets from the web. As our proposed self-training pipeline facilitates scalable learning from large amounts of data, the model may continue to improve with additional YouTube data. However, this requires further study in the future, e.g., with respect to larger network capacity.

## 5. Conclusion

We envision broad and easily deployable autonomous navigation systems. However, access to resources and data limits the scope of the brittle autonomous systems today. Our SelfD approaches enables to significantly improve an initially trained policy without incurring additional data collection or annotation efforts, i.e., for a new platform, perspective, use-case, or ambient settings. Crucially, due to the proposed underlying model architecture, we do not incorporate camera parameters or configuration assumptions into the monocular inference. As SelfD is self-improving, a future direction could be to continue and learn from increasingly larger online datasets beyond what is described in our study. While we emphasized efficient large-scale training in our model development, how to best extend SelfD to more explicitly leverage temporal demonstration data is still an open question which could be studied further in the future. Finally, beyond complex 3D navigation, it would be interesting to explore the applicability of our proposed training framework for learning various embodied tasks from unlabeled web data.

**Acknowledgments:** We thank the Red Hat Collaboratory and the BU Center for Information and Systems Engineering for research awards.

## References

- [1] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *ICML*, 2004. 3
- [2] Yuki M Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision. *arXiv preprint arXiv:2006.13662*, 2020. 3
- [3] Michael Bain and Claude Sammut. A framework for behavioural cloning. In *Machine Intelligence*, 1996. 2
- [4] Mayank Bansal, Alex Krizhevsky, and Abhijit Ogale. ChauffeurNet: Learning to drive by imitating the best and synthesizing the worst. In *RSS*, 2019. 2, 3
- [5] Aseem Behl, Kashyap Chitta, Aditya Prakash, Eshed Ohn-Bar, and Andreas Geiger. Label-efficient visual abstractions for autonomous driving. In *IROS*, 2020. 3
- [6] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016. 2, 3
- [7] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *CVPR*, 2020. 2, 4, 5
- [8] Benjamin Caine, Rebecca Roelofs, Vijay Vasudevan, Jiquan Ngiam, Yuning Chai, Zhifeng Chen, and Jonathon Shlens. Pseudo-labeling for scalable 3d object detection. In *arXiv preprint arXiv:2103.02093*, 2021. 2, 3, 5
- [9] Matthew Chang, Arjun Gupta, and Saurabh Gupta. Semantic visual navigation by watching YouTube videos. In *NeurIPS*, 2020. 2, 4
- [10] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *CVPR*, 2019. 2, 5
- [11] Chenyi Chen, Ari Seff, Alain L. Kornhauser, and Jianxiong Xiao. DeepDriving: Learning affordance for direct perception in autonomous driving. In *ICCV*, 2015. 2
- [12] Dian Chen, Vladlen Koltun, and Philipp Krähenbühl. Learning to drive from a world on rails. In *ICCV*, 2021. 2
- [13] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *CoRL*, 2020. 2, 3, 4, 5, 6, 7
- [14] Xinlei Chen and Abhinav Gupta. Webly supervised learning of convolutional networks. In *ICCV*. 2
- [15] Yuhua Chen, Cordelia Schmid, and Cristian Sminchisescu. Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. In *CVPR*, 2019. 3
- [16] Sanjiban Choudhury, Mohak Bhardwaj, Sankalp Arora, Ashish Kapoor, Gireeja Ranade, Sebastian Scherer, and Debadepta Dey. Data-driven planning via imitation learning. *International Journal of Robotics Research*, 37(13-14):1632–1672, 2018. 3
- [17] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *ICRA*, 2018. 2, 3
- [18] Felipe Codevilla, Eder Santana, Antonio M López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *ICCV*, 2019. 2, 3, 5
- [19] Benjamin Coors, Alexandru Paul Condurache, and Andreas Geiger. Nova: Learning to see in novel viewpoints and domains. In *3DV*, 2019. 2
- [20] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 2
- [21] Pim de Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. In *NeurIPS*, 2019. 4
- [22] Ernst Dieter Dickmanns and Volker Graefe. Dynamic monocular machine vision. *Machine vision and Applications*, 1(4):223–240, 1988. 2
- [23] Murat Dikmen and Catherine M Burns. Autonomous driving in the real world: Experiences with tesla autopilot and summon. In *AutomotiveUI*, 2016. 2
- [24] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *CoRL*, 2017. 5
- [25] Fedex. Delivering the future: Fedex unveils autonomous delivery robot. <https://newsroom.fedex.com/>, 2029. 2
- [26] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *CVPR*, 2012. 2, 4
- [27] Clément Godard, Oisin Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *CVPR*, 2019. 3
- [28] Ariel Gordon, Hanhan Li, Rico Jonschkowski, and Anelia Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *CVPR*, 2019. 2
- [29] Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. *arXiv preprint arXiv:1905.01235*, 2019. 2, 3
- [30] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *CVPR*, 2017. 2
- [31] Jeffrey Hawke, Richard Shen, Corina Gurau, Siddharth Sharma, Daniele Reda, Nikolay Nikolov, Przemysław Mazur, Sean Micklethwaite, Nicolas Griffiths, Amar Shah, et al. Urban driving with conditional imitation learning. In *ICRA*, 2020. 2
- [32] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 3
- [33] John Houston, Guido Zuidhof, Luca Bergamini, Yawei Ye, Long Chen, Ashesh Jain, Sammy Omari, Vladimir Igloukov, and Peter Ondruska. One thousand and one hours: Self-driving motion prediction dataset. *arXiv preprint arXiv:2006.14480*, 2020. 4

- [34] Peiyun Hu, Aaron Huang, John Dolan, David Held, and Deva Ramanan. Safe local motion planning with self-supervised freespace forecasting. In *CVPR*, 2021. 3, 6
- [35] Peiyun Hu, Jason Ziglar, David Held, and Deva Ramanan. What you see is what you get: Exploiting visibility for 3D object detection. In *CVPR*, 2020. 5
- [36] Elia Kaufmann, Mathias Gehrig, Philipp Foehn, René Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza. Beauty and the beast: Optimal methods meet learning for drone racing. In *ICRA*, 2019. 1
- [37] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*, 2018. 4
- [38] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, 2013. 2, 3, 4
- [39] Guohao Li, Matthias Mueller, Vincent Casser, Neil Smith, Dominik L Michels, and Bernard Ghanem. Oil: Observational imitation learning. In *RSS*, 2019. 2
- [40] Lingyun Luke Li, Bin Yang, Ming Liang, Wenyuan Zeng, Mengye Ren, Sean Segal, and Raquel Urtasun. End-to-end contextual perception and prediction with interaction transformer. In *IROS*, 2020. 6
- [41] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T Freeman. Learning the depths of moving people by watching frozen people. In *CVPR*, 2019. 2
- [42] Xiaodan Liang, Tairui Wang, Luona Yang, and Eric Xing. CIRL: Controllable imitative reinforcement learning for vision-based self-driving. In *ECCV*, 2018. 2
- [43] Zhiqiu Lin, Deva Ramanan, and Aayush Bansal. Streaming self-training via domain-agnostic unlabeled images. *arXiv preprint arXiv:2104.03309*, 2021. 2
- [44] Pengpeng Liu, Irwin King, Michael R Lyu, and Jia Xu. Flow2stereo: Effective self-supervised learning of optical flow and stereo matching. In *CVPR*, 2020. 3
- [45] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Self-low: Self-supervised learning of optical flow. In *CVPR*, 2019. 3
- [46] Dmytro Mishkin, Alexey Dosovitskiy, and Vladlen Koltun. Benchmarking classic and learned navigation in complex 3d environments. *arXiv preprint arXiv:1901.10915*, 2019. 1
- [47] Matthias Müller, Alexey Dosovitskiy, Bernard Ghanem, and Vladlen Koltun. Driving policy transfer via modularity and abstraction. In *CoRL*, 2018. 2, 3, 4
- [48] Urs Muller, Jan Ben, Eric Cosatto, Beat Flepp, and Yann L Cun. Off-road obstacle avoidance through end-to-end learning. In *NeurIPS*, 2006. 2, 4
- [49] Adithyavairavan Murali, Siddarth Sen, Ben Kehoe, Animesh Garg, Seth McFarland, Sachin Patil, W Douglas Boyd, Susan Lim, Pieter Abbeel, and Ken Goldberg. Learning by observation for surgical subtasks: Multilateral cutting of 3D viscoelastic and 2D orthotropic tissue phantoms. In *ICRA*, 2015. 4
- [50] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 3
- [51] Eshed Ohn-Bar, Aditya Prakash, Aseem Behl, Kashyap Chitta, and Andreas Geiger. Learning situational driving. In *CVPR*, 2020. 2
- [52] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, Jan Peters, et al. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2):1–179, 2018. 1, 2, 3
- [53] Yunpeng Pan, Ching-An Cheng, Kamil Saigol, Keuntaek Lee, Xinyan Yan, Evangelos A Theodorou, and Byron Boots. Imitation learning for agile autonomous driving. *International Journal of Robotics Research*, 39(2-3):286–302, 2020. 1
- [54] Mandela Patrick, Yuki M Asano, Polina Kuznetsova, Ruth Fong, João F Henriques, Geoffrey Zweig, and Andrea Vedaldi. On compositions of transformations in contrastive self-supervised learning. In *CVPR*, 2021. 3
- [55] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3D. In *ECCV*, 2020. 2
- [56] Dean A Pomerleau. ALVINN: An autonomous land vehicle in a neural network. In *NeurIPS*, 1989. 2, 3
- [57] Aditya Prakash, Aseem Behl, Eshed Ohn-Bar, Kashyap Chitta, and Andreas Geiger. Exploring data aggregation in policy learning for vision-based urban autonomous driving. In *CVPR*, 2020. 3
- [58] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. In *CVPR*, 2021. 2
- [59] Rene Ranftl, Vibhav Vineet, Qifeng Chen, and Vladlen Koltun. Dense monocular depth estimation in complex dynamic scenes. In *CVPR*, 2016. 2
- [60] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *ICLR*, 2021. 2, 3
- [61] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, 2011. 3
- [62] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *ICML*, 2017. 2
- [63] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020. 3, 5, 6, 7
- [64] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2020. 2, 5
- [65] Wen Sun, Anirudh Vemula, Byron Boots, and J. Andrew Bagnell. Provably efficient imitation learning from observation alone. In *ICML*, 2019. 2
- [66] Faraz Torabi, Garrett Warnell, and Peter Stone. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*, 2018. 2, 4

- [67] Faraz Torabi, Garrett Warnell, and Peter Stone. Recent advances in imitation learning from observation. *arXiv preprint arXiv:1905.13566*, 2019. 2, 4
- [68] Marin Toromanoff, Emilie Wirbel, and Fabien Moutarde. End-to-end model-free reinforcement learning for urban driving using implicit affordances. In *CVPR*, 2020. 2
- [69] Charlie Tritschler. Meet astro, a home robot unlike any other. <https://www.aboutamazon.com/news/devices/meet-astro-a-home-robot-unlike-any-other>, 2021. 2
- [70] Dequan Wang, Coline Devin, Qi-Zhi Cai, Philipp Krähenbühl, and Trevor Darrell. Monocular plan view networks for autonomous driving. In *IROS*, 2019. 2, 4
- [71] Sen Wang, Ronald Clark, Hongkai Wen, and Niki Trigoni. Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks. In *ICRA*, 2017. 2, 4, 6, 7
- [72] Jamie Watson, Oisín Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In *CVPR*, 2021. 3
- [73] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, et al. Argoverse 2.0: Next generation datasets for self-driving perception and forecasting. In *NeurIPS*, 2021. 4
- [74] Yueh-Hua Wu, Nontawat Charoenphakdee, Han Bao, Voot Tangkaratt, and Masashi Sugiyama. Imitation learning from imperfect demonstration. In *International Conference on Machine Learning*, pages 6818–6827. PMLR, 2019. 2
- [75] Huazhe Xu, Yang Gao, Fisher Yu, and Trevor Darrell. End-to-end learning of driving models from large-scale video datasets. In *CVPR*, 2017. 3
- [76] Jihan Yang, Shaoshuai Shi, Zhe Wang, Hongsheng Li, and Xiaojuan Qi. ST3D: Self-training for unsupervised domain adaptation on 3d object detection. In *CVPR*, 2021. 3, 4, 5
- [77] Nan Yang, Lukas von Stumberg, Rui Wang, and Daniel Cremers. D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In *CVPR*, 2020. 2
- [78] Jimuyang Zhang and Eshed Ohn-Bar. Learning by watching. In *CVPR*, 2021. 2
- [79] Jimuyang Zhang, Minglan Zheng, Matthew Boyd, and Eshed Ohn-Bar. X-World: Accessibility, vision, and autonomy meet. In *ICCV*, 2021. 2
- [80] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016. 3
- [81] Albert Zhao, Tong He, Yitao Liang, Haibin Huang, Guy Van den Broeck, and Stefano Soatto. Lates: Latent space distillation for teacher-student driving policy learning. *arXiv preprint arXiv:1912.02973*, 2019. 2
- [82] Brady Zhou, Philipp Krähenbühl, and Vladlen Koltun. Does computer vision matter for action? *Science Robotics*, 4(30), 2019. 2