

Exploiting Playbacks in Unsupervised Domain Adaptation for 3D Object Detection in Self-Driving Cars

Yurong You^{*†}, Carlos Andres Diaz-Ruiz^{*‡}, Yan Wang[†], Wei-Lun Chao[§],
Bharath Hariharan[†], Mark Campbell[‡], Kilian Q Weinberger[†]

Abstract—Self-driving cars must detect other traffic participants like vehicles and pedestrians in 3D in order to plan safe routes and avoid collisions. State-of-the-art 3D object detectors, based on deep learning, have shown promising accuracy but are prone to over-fit domain idiosyncrasies, making them fail in new environments—a serious problem for the robustness of self-driving cars. In this paper, we propose a novel learning approach that reduces this gap by fine-tuning the detector on high-quality pseudo-labels in the target domain – pseudo-labels that are automatically generated *after driving* based on replays of previously recorded driving sequences. In these replays, object tracks are smoothed forward and backward in time, and detections are interpolated and extrapolated—crucially, leveraging future information to catch hard cases such as missed detections due to occlusions or far ranges. We show, across five autonomous driving datasets, that fine-tuning the object detector on these pseudo-labels substantially reduces the domain gap to new driving environments, yielding strong improvements detection reliability and accuracy.

Index Terms—Object Detection, Segmentation and Categorization; Computer Vision for Automation; Transfer Learning; Deep Learning for Visual Perception

I. INTRODUCTION

Detecting traffic participants such as cars, cyclists, and pedestrians in 3D is a fundamental learning problem for self-driving cars. Typically, inputs consists of LiDAR point clouds and/or images; outputs are sets of tight 3D bounding boxes that envelop detected objects. The problem is challenging, because the detection must be highly accurate and reliable. The current state of the art in 3D object detection is based on deep learning approaches [1], [2], [3], [4], trained on short driving segments with labeled bounding boxes [5], [6], which yield up to 80% average precision on held-out segments [4].

However, as with all machine learning approaches, these techniques succeed when the training and test data distributions match. One way to ensure this is to constrain self-driving cars to a small geo-fenced area, such as with a fleet of similar self-driving taxis collecting and sharing training data about the same area. This approach, however, is not generalizable to consumer self-driving cars which should be able to drive freely anywhere, similar to a human-driven car. This unconstrained scenario introduces an inherent adaptation problem: the car producer cannot foresee where the owner

will ultimately operate the car. For example, the perception system might be trained on urban roads in Germany [5], [6], but the car may be driven in the mountain roads in the USA, where cars are larger and fewer, roads may be snowier, and the environment (trees, roads, etc.) may look different. Past work has shown that such differences can cause >35% drop in detection accuracy [7]. Closing this adaptation gap is one of the biggest challenges for consumer self-driving vehicles.

Formally, this challenge is a problem in *unsupervised domain adaptation (UDA)* [8]: the detector, having been previously trained on labeled data from a *source* domain, must now adapt to a *target* domain where only unlabeled data are available. While the UDA problem is easily cast with training and testing datasets in different locations (Germany vs USA, urban vs rural, etc.), similar problems exist in many self-driving car scenarios. For example, consider the common case of car owners who drive many of the same routes (e.g., commuting), and leave their cars parked (e.g., at night) for extended periods of time. This raises an intriguing possibility: the car can collect sensor data on these trips, and then retrain itself autonomously *offline* to adapt to this new environment, to improve subsequent *online* driving.

In this paper, we present a novel approach for UDA of 3D detectors for self-driving cars. Our approach uniquely uses two key insights. First, data collected over time via a *video* is not simply a bag of independent frames. Second, the *dynamics* of our objects of interest (i.e., cars) can be modeled effectively. We propose to use time correlations between frames and object physics to enable more accurate and efficient solutions to the UDA problem *offline*. More specifically, our approach takes the ‘confident’ detections of nearby objects, estimates their states (e.g., locations, sizes and speeds), and then *extrapolates* the tracks forward and backward in time, discovering challenging cases when the detector made mistakes (e.g. missed detections due to occlusions, detections at far ranges, etc.) The *playback* of the data allows us to go back in time and annotate labels in frames which were previously missed. Although this process cannot be performed in real time (since it uses future information), we can use it *offline* to generate a new training set with pseudo-labels for the target environment. We can then adapt the detector to the target domain using this newly created dataset, thus allowing the detector to generalize to more scenarios. We call our approach *dreaming*, as the car learns by replaying past driving sequences backwards and forwards, potentially while it is parked overnight.

We evaluate our approach on the most challenging of

* Equal contributions

† Computer Science Department, Cornell University {yy785, yw763, bh497, kqw4}@cornell.edu

‡ Mechanical and Aerospace Engineering Department, Cornell University {cad297, mc288}@cornell.edu

§ Department of Computer Science and Engineering, the Ohio State University chao.209@osu.edu

UDA problems, that of 3D object detection across multiple autonomous driving datasets, including KITTI [5], [6], Argoverse [9], Lyft [10], Waymo [11], and nuScenes [12]. We show *across all dataset combinations* that discovering detector mistakes and retraining the detector with our dreaming car procedure strongly reduces the source/target domain gap with high consistency. In fact, the resulting detector after “dreaming” *substantially exceeds* the accuracy of the offline system used to generate the pseudo-labels, which—although able to look into the future—is limited to the extrapolation of confident detections before adaption. Our dreaming procedure can easily be implemented on-device and we believe that it constitutes a significant step towards safely operating autonomous vehicles without geo-restrictions.

II. RELATED WORK

3D object detection can be categorized based on the input: using 3D sensors like LiDAR or 2D images from cameras [13], [14], [15], [16]. We focus on the former due to its higher accuracy. Examples of LiDAR-based detectors include F-PointNet [1], PointRCNN [2], PIXOR [3] and VoxelNet [17]. While these methods have consistently improved the detection accuracy, it has been recently revealed in [7] that they cannot generalize well when trained and tested on different datasets, especially on distant objects with sparse LiDAR points.

Unsupervised domain adaptation (UDA) has been widely studied in machine learning and computer vision, especially on image classification [8], [18], [19]. The common setup is to adapt a model trained from one labeled source domain (e.g., synthetic images) to another unlabeled target domain (e.g., real images). Recent work has extended UDA to driving scenes, but mainly for 2D object detection [20], [21], [22], [23], [24] and semantic segmentation [25], [26], [27], [28], [29]. The mainstream approach is to match the feature distributions or image appearances between domains, e.g., via adversarial learning [18], [25] or image translation [30]. The approaches more similar to ours are [31], [32], [33], [34], [29], [35], [36], [37], [38], which iteratively assign pseudo-labels to (some of) the unlabeled data in the target domain and retrain the models. This procedure, usually named self-training, is proven to be effective in learning with unlabeled data, such as semi-supervised and weakly-supervised learning [39], [40], [41], [42]. For UDA, self-training enables the model to adapt its features to the target domain in a supervised fashion.

UDA in 3D tackles the domain discrepancy in point clouds. Qin et al. [43] were the first to match point cloud distributions between domains, via adversarial learning. However, they considered point clouds of isolated objects, which are very different from the ones captured in driving scenes. Other approaches project 3D points to the frontal or bird’s-eye view and apply UDA methods in the resulting 2D images [44], [45], [46], which can be sub-optimal in models’ accuracy.

Instead, we follow the self-training paradigm and show that *UDA in 3D object detection can be drastically improved by our high-quality pseudo-labels*. Concurrent work ST3D [47] also tackles UDA in 3D via self-training, but it focuses on addressing the object size discrepancy across domains [7].

It pre-trains a 3D detector with random object scaling and uses pseudo-labels to train student detectors in a curriculum-learning way. In contrast, our work makes use of consecutive LiDAR scans in the target domain to improve pseudo-label quality, especially for faraway or previously missed objects. Our work is thus orthogonal and complementary to theirs.

Leveraging videos for object detection has been explored in [48], [49], [50], [51] to ease the labeling efforts by mining extra 2D bounding boxes from videos. The main idea is to leverage the temporal information to extend weakly-labeled instances or potential object proposals across frames, which are then used as pseudo-labels to retrain the detectors. In the context of UDA, [52], [31], [53] also incorporate object tracks to discover high quality 2D pseudo-labels for self-training.

Our approach is different in two aspects. First, we not only interpolate but also *extrapolate* tracks to infer object locations when they are too far away to be detected accurately. Second, we operate in 3D. We apply a physical-based motion model and exploit the fact that objects in 3D are scale-invariant to correct the detection along tracks. In contrast, the methods above operate in 2D and may disregard faraway objects (that appear too small in images) due to unreliable 2D tracks [53].

Auto-labeling. Our work is also related to concurrent work in 3D auto-labeling [54], [55], which improves the initial detections of a 3D object detector in an offline manner by aggregating them across the whole sequence with a standard tracker [56]. Our work is different from theirs in two aspects. First, our approach not only aggregates detections in a track, but *extrapolates* them with a motion model, which is crucial to recover the missing detections faraway (Table V). Second, our end-goal is to improve the 3D object detector by fine-tuning it with the improved pseudo-labels, while auto-labeling solely focuses on improving pseudo-labels.

III. EXPLOITING PLAYBACKS FOR UDA

Similar to most published work on 3D object detection for autonomous driving [1], [2], [3], [4], [57], we focus on *frame-wise* 3D detectors. A detector is first trained on a source domain and then applied to a target domain (e.g., a new city). [7] revealed a drastic accuracy drop in such a scenario: many of the target objects are either misdetected or mislocalized, especially if they are far away. To aid adaptation, we assume access to an unlabeled dataset of *video sequences* in the target domain, which could simply be recordings of the vehicle’s sensors while it was in operation. Our approach is to generate pseudo-labels for these recordings that can be used to adapt the detector to the new environment during periods when the car is not in use. We assume no access to the source data when performing adaptation—it is unlikely that car producers share data with the customers after the detector is deployed.

A. Tracking for improved detection

One way to improve the test accuracy based on the frame-wise detection outputs is *online* tracking by detection [58], [59], [60]. Here, detected objects are associated across current and *past* frames to derive trajectories, which are used to filter

out false positives, correct false negatives, and adjust the initial detection boxes in the current frame.

Online 3D object tracking. We investigate this idea with a Kalman filter based tracker [61], [62], [56], which has shown promising results in benchmark tracking leader boards [12]. *We opt to not use a learning-based tracker [63] since it would also require adaptation before it can be applied in the target domain.* Specifically, we apply the tracker in [61]. The algorithm estimates the joint probability $p(\mathbf{a}_k, \mathbf{x}_k | \mathbf{z}_k)$ at time k , where \mathbf{x}_k is the set of tracked object states (e.g., cars speeds and locations), \mathbf{z}_k is the set of observed sensor measurements (here each measurement is a frame-wise detection), and \mathbf{a}_k is the assignment of measurements to tracks. The joint probability can be decomposed into continuous estimation $p(\mathbf{x}_k | \mathbf{a}_k, \mathbf{z}_k)$ and discrete data assignment $p(\mathbf{a}_k | \mathbf{z}_k)$. The former is solved recursively via an Extended Kalman Filter (EKF); the latter is solved via Global Nearest-Neighbor (GNN), where the cost to be minimized is the negative BEV IoU between the predicted box and the measurement. The EKF parameterizes the state \mathbf{x} of a single (i th) object as a vehicle (position, velocity, shape) *relative* to the ego-vehicle

$$\mathbf{x}_k^i = [x \ y \ \theta \ s \ l \ w]^T, \quad (1)$$

where x, y are the location of the tracked vehicle's back axle relative to a fixed point on the ego-vehicle, θ is the vehicle orientation relative to the ego-vehicle, s is the absolute ground speed, and l, w are the length and width. The EKF uses a dynamics model of the evolution of the state over time. Here we assume that the tracked vehicle is moving at a *constant* speed and heading in the global coordinate frame, with added noise to represent the uncertainty associated with vehicle maneuvers. This tracker has been shown to work well on tracking moving objects from a self-driving car [61], [64]. We initialize a new track when $c_{\min\text{-hits}}$ measurements of the same tracked object are realized. We end a track when it does not obtain any measurement updates over $c_{\max\text{-age}}$ frames, or the tracked object exits the field of view (FOV). As will be seen in subsection IV-B, applying this tracker can indeed improve the detection accuracy online, via inputting missed detections, correcting mislocalized detections, and rejecting wrong detections in \mathbf{z}_k at current time k by \mathbf{x}_k .

Offline 3D object tracking. Online trackers only use past information to improve current detections. By relaxing this for offline tracking (e.g., to be able to look into the future and come back to the current time), we can obtain more accurate estimates of vehicle states. While such a relaxation is not applicable during test time, higher accuracy tracking on unlabeled driving sequences will be very valuable for adapting the source detector to the target domain in a self-supervised fashion, as we will explain in the following section.

B. Self-training for UDA

Self-training is a simple yet fairly effective way to improve a model with unlabeled data [39], [40], [41], [65]. The basic idea is to apply an existing model to an unlabeled dataset and use the high confidence predictions (here detections), which are likely to be correct, as “pseudo-labels” for fine-tuning.

One key to success for self-training is the quality of the pseudo-labels. In particular, we desire two qualities out of the detections we use as pseudo-labels. First, they should be *correct*, i.e., they should not include false positives. Second, they should have *high coverage*, i.e., they should cover all cases of objects. Choosing high confidence detections as pseudo-labels satisfies the first criterion but not the second. For 3D object detection, we find that most of the high confidence examples are easy cases: unoccluded objects near the self-driving car. This is where offline tracking becomes a crucial component to include the more challenging cases (far away, or partially occluded objects) in the pseudo-label pool.

C. High quality pseudo-labels via 3D offline tracking

How do we obtain pseudo-labels for far-away, hard-to-detect objects that the detector cannot reliably detect? We propose to exploit tracking by leveraging two facts in the autonomous driving scenario. First, the available unlabeled data is in the form of *sequences* (akin to videos) of point clouds over time. Second, the objects of interest and the self-driving car move in fairly constrained ways. We will run the object detector on *logged* data, so that we can easily analyze both forwards and backwards in time. The object detector will detect objects accurately only when they are close to the self-driving car. Once detected over a few frames, we can estimate the object's motion either towards the self-driving car or away from it, and then both *interpolate* the object's positions in frames where it was missed, or *extrapolate* the object into frames where it is too far away for accurate detection. We show an example of this procedure in Figure 1. Through dynamic modeling, tracking, and smoothing over time we can correct noisy detections. Further with extrapolation and interpolation, we can recover far away and missed detections.

Concretely, we proposed to augment the following functionalities that utilize the future information into the online tracker introduced in subsection III-A, turning it into an offline tracker specifically designed to improve detection, i.e., generating higher quality pseudo-labels for self-training.

State smoothing. Frame-wise 3D object detectors can generate inconsistent, noisy detection across time (e.g., frame 4, 7 and 9 in Figure 1). The model-based tracking approach in subsection III-A reduces this noise, but we can go further by smoothing tracks back and forth over time, since our data is offline. In this work, we use a fixed-point Rauch-Tung-Striebel (RTS) smoother [66] to smooth the tracked state estimates. Smoothing requires a backward iteration ($k = N, N-1, \dots, 1$) that is performed after the forward filtering, where the a-posteriori state and state error covariance estimates ($\bar{\mathbf{x}}_{k|k}$ and $P_{k|k}$) and the a-priori state and state error covariance estimates ($\bar{\mathbf{x}}_{k+1|k}$ and $P_{k+1|k}$) have been calculated. The smoothed gain, C_k , is obtained from

$$C_k = P_{k|k} F_k^T P_{k+1|k}^{-1}, \quad (2)$$

where F_k is the Jacobian of the dynamics model evaluated at $\bar{\mathbf{x}}_{k|k}$. The smoothed state is then evaluated as

$$\bar{\mathbf{x}}_{k|N} = \bar{\mathbf{x}}_{k|k} + C_k [\bar{\mathbf{x}}_{k+1|N} - \bar{\mathbf{x}}_{k|k}] \quad (3)$$

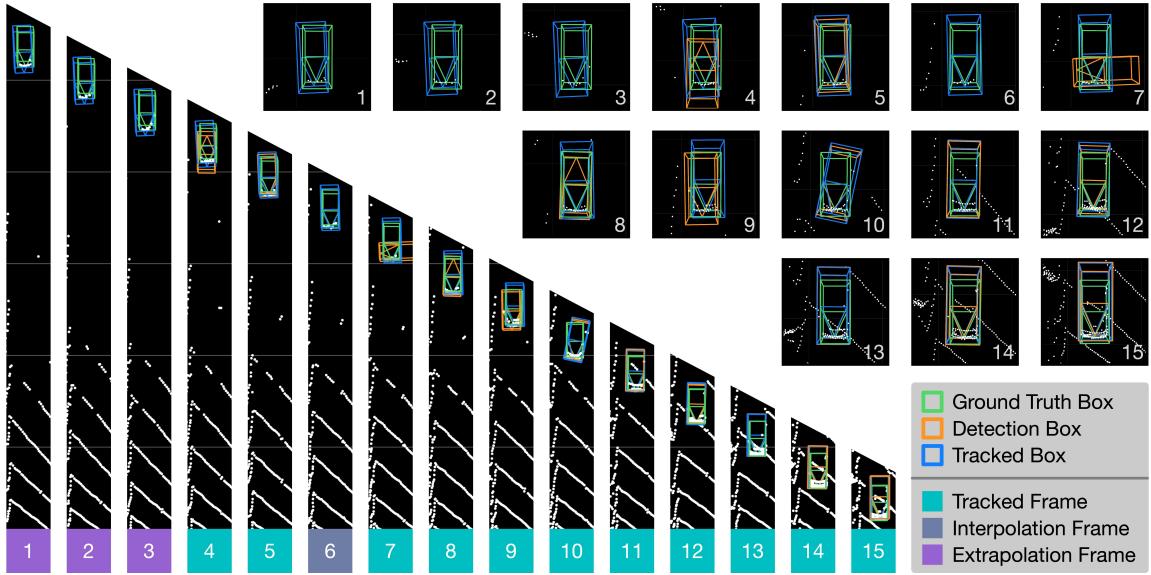


Fig. 1: An example tracked vehicle moves towards ego-vehicle (at the bottom), where pseudo-labels are recovered through extrapolation, interpolation, and smoothing. The improvements of pseudo-labels via offline tracking are instances where estimated bounding boxes (blue) are observed, while the frame-wise detections (orange) are missing or poorly aligned with ground truths (green). Better viewed in color.

while the covariance of the smoothed state is evaluated as

$$P_{k|N} = P_{k|k} + C_k [P_{k+1|N} - P_{k+1|k}] C_k^T. \quad (4)$$

Adjusting object sizes. As shown in [7], the distribution of car sizes in different domains (e.g. different cities) can be different. As such, when tested on a novel domain, detectors often predict incorrect object sizes. This is especially true when the LiDAR signal is too sparse for correct size information. We can also use our tracking to correct such systematic error. Assuming that the most confident detections are more likely to be accurate, we estimate the size of the object by averaging the size of the three highest confidence detections. We use this size for all objects in this track.

Interpolation and extrapolation. We use estimation (forward in time) and smoothing (backward in time) to recover missed detections, and in turn, to increase the recall rate of pseudo-labels (e.g., frame 1–3 and 6 in Figure 1). If a detection is missed in the middle of a track, we restore it by taking the estimated state from smoothing. We also extrapolate the tracks both backward and forward in time, so that tracks that were prematurely terminated due to missing detections, can be recovered. More concretely, we are able to recover detections of vehicles that were lost as they moved away from the ego-vehicle because the sensor signals became sparser (or in turn, the vehicles started far away and were then only detected when they got close enough). Extrapolations are performed by first using dynamics model predictions of the EKF to predict potential bounding boxes; measurements are obtained by performing a search and detection in the vicinity of the prediction. We apply the detector in a 3 m^2 area around the extrapolated prediction, yielding several 3D bounding box candidates. After filtering out candidates with confidences lower than some threshold, we select the candidate with the highest BEV IoU with the prediction as the measurement. If a track loses such a measurement for three consecutive

frames, extrapolations are stopped. With this targeted search, we are able to recover objects that were missed due to low confidence. After extrapolating and interpolating detections for all tracks, we perform Non Maximum Suppression (NMS) over bounding boxes in BEV, where more recent extrapolations/interpolations are prioritized.

Discussion. The tracker we apply is standard and simple. We opt it to show the power of our dreaming approach for UDA—exploiting offline, forward and backward information to derive high-quality pseudo-labels for adapting detectors. More sophisticated trackers will likely improve the results further. While we focus on frame-wise 3D detectors, our algorithm can be applied to adapt video-based 3D object detectors [67] as well. One particular advantage of fine-tuning on the pseudo-labeled target data is to allow the detector adapting not only its predictions (e.g., the box regression) but also its features (e.g., early layers in the neural networks) to the target domain. The resulting detector thus can usually lead to more accurate detections than the pseudo-labels it has been trained on.

IV. EXPERIMENTS

Datasets. We experiment with five autonomous driving data sets: KITTI [5], [6], Argoverse [9], nuScenes [12], Lyft [10] and Waymo [11]. All datasets provide LiDAR data in sequences and ground-truth bounding box labels for either all or part of the data. We briefly summarize these five datasets in the supplementary. We follow a setup similar to [7], but with different splits on Lyft, nuScenes, and Waymo in order to keep the training and test sets non-overlapping with sequences. For nuScenes and Waymo, we only use data from a single location (Boston for nuScenes and San Francisco for Waymo).

UDA settings. We train models in the source domain using labeled frames. We split each dataset into two non-overlapping parts, a training set and a test set. *We use the train set (and*

its pseudo-labels) of the target domain to adapt the source detector, and evaluate the adapted detector on the **test** set.

Metric. We follow KITTI to evaluate object detection in 3D and BEV metrics. We focus on the *Car* category as it is the main focus of existing work. We report average precision (AP) with the intersection over union (IoU) thresholds at 0.5 or 0.7, i.e., a car is correctly detected if the IoU between it and the detected box is larger than 0.5 or 0.7. We denote AP in 3D and BEV by AP_{3D} and AP_{BEV} , respectively. Because on the other datasets there is no official separation on the difficulty levels like in KITTI, we split AP by depth ranges.

3D object detection models. We use two LiDAR-based models POINTRCNN [2] and PIXOR [3] to detect objects in 3D. They represent two different but popular ways of processing point cloud data. POINTRCNN uses PointNet++ [68] to extract point-wise features, while PIXOR applies 2D convolutions in BEV of voxelized point clouds. Neither relies on images. We mainly report and discuss results of POINTRCNN except for the last study in [subsection IV-B](#).

Hyper-parameters. To train detectors on source domain, we use the hyper-parameters provided by [2] for POINTRCNN. For PIXOR, we follow [7] to train it using RMSProp with momentum 0.9 and learning rate 5×10^{-5} (decreased by a factor of 10 after 50 and 80 epochs) for 90 epochs. For self-training on the target domain, we initialize from the pre-trained model on the source domain. For POINTRCNN, we fine-tune it with learning rate 2×10^{-4} and 40 epochs in RPN and 10 epochs in RCNN. For PIXOR, we use RMSProp with momentum 0.9 and learning rate 5×10^{-6} (decreased by a factor of 10 after 10 and 20 epochs) for 30 epochs.

We developed and tuned our dreaming method with Argoverse as the source domain and KITTI as the target domain (in the target domain, we only use the training set). We then fixed all hyper-parameters for all subsequent experiments.

A. Baselines

We compare against two baselines under the UDA setting.

Self-Training (ST). We apply a self-training scheme similar to that typically used in the 2D problems [40]. When adapting the model from the source to the target, we apply the source model to the target training set. We then keep the detected cars of confidence scores > 0.8 (label-sharpening) and use them as pseudo-labels to fine-tune the model. We select the threshold following our hyper-parameter selection procedure and apply it to all the experiments.

Statistical Normalization (SN). [7] showed that car sizes vary between domains: popular cars at different areas can be different. When the mean bounding box size in the target domain is accessible, either from limited amount of labeled data or statistical data, we can apply *statistical normalization (SN)* [7] to mitigate such a systematic difference in car sizes. SN adjusts the bounding box sizes and corresponding point clouds in the source domain to match those in the target domain, and fine-tunes the model on such “normalized” source data, with no need to access target sensor data. We follow the exact setting in [7] to apply SN.

TABLE I: **Pseudo-label quality.** We compare the quality of the pseudo-labels on KITTI **train** set generated by a detector trained on Argoverse (PL) and those after smoothing, resizing, interpolation, and extrapolation (PL (AFTER)). We show the AP_{BEV} / AP_{3D} of the **car** category at IoU = 0.7 and 0.5 across different depth range.

Method	IoU 0.5			IoU 0.7		
	0-30	30-50	50-80	0-30	30-50	50-80
PL	80.5 / 80.0	63.7 / 60.7	23.9 / 18.6	63.9 / 36.6	33.0 / 12.1	10.0 / 0.9
PL (AFTER)	80.5 / 80.2	68.0 / 61.1	29.5 / 22.6	64.9 / 38.0	38.4 / 15.1	12.1 / 0.9

B. Empirical Results

Pseudo-label quality. In [Table I](#) we evaluate the quality of pure pseudo-labels and pseudo-labels after smoothing, resizing, interpolation, and extrapolation under the Argoverse to KITTI setting. It can be seen that the dreaming process significantly improves the pseudo-label quality across all ranges, which leads to further improvement after self-training.

Adaptation from Argoverse to KITTI. We compare the UDA methods under Argoverse to KITTI in [Table III](#) and observe several trends: 1) models experience a smaller domain gap when objects are closer (0-30 m vs 30-80 m); 2) though directly applying online tracking can improve the detection performance, models improve more after just self-training; 3) the offline tracking is used to provide extra pseudo-labels for fine-tuning, and interestingly, models fine-tuned from pseudo-labels can outperform pseudo-labels themselves; 4) DREAMING improves over ST and SN by a large margin, especially on IoU at 0.5; 5) DREAMING has a large gain in AP for faraway objects, e.g., on range 50-80 m, compared to ST, it boosts the AP_{BEV} on IoU at 0.5 from 28.2 to 30.1.

Adaptation among five datasets. We further applied our methods to adaptation tasks among the five datasets. Due to limited space, we show the results of AP_{BEV} and AP_{3D} on range 50-80 m and 0-80 m at IoU = 0.5 in [Table II](#). Our method consistently improves the adaptation performance on faraway ranges, while having mostly equal or better performance over the full range. We include detailed evaluation results across all ranges, at IoU = 0.7, and with SN in the supplementary material. We observe a consistent and clear trend as in [Table II](#).

Adaptation between different locations inside the same dataset. Different datasets not only come from different locations but also use different sensor configurations. To isolate the effects of the former (which is our motivating application), in [Table IV](#) we evaluate our method’s performance for domain adaptation within the KITTI dataset. In this case, the source and target domain are all scenes from KITTI, while the source is composed of city and campus scenes (38 sequences, 9,556 frames, 3,420 labeled frames) and the target consists of residential and road scenes (23 sequences, 10,448 frames, 3619 labeled frames). Our method consistently outperforms no fine-tuning and ST, especially on 30-80 m range.

Ablation Study. We show ablation results in [Table V](#). Here we fine-tune models using ST and adding smoothing (S), resizing (R), interpolation (I) and extrapolation (E) to the pseudo-label generation. It can be observed that ST alone already boosts performance considerably. Through selecting high confidence detections, smoothing and adjusting the object size we ensure that the pseudo-labels provided are mostly

TABLE II: **Dreaming results on UDA among five auto-driving datasets.** We report AP_{BEV} and AP_{3D} of the Car category on far-away range (50-80m) and full range (0-80m) at IoU= 0.5. On each entry (*row, column*), we report AP of UDA from *row* to *column* in the order of No Re-training / Self-Training / Dreaming. At the diagonal entries is the AP of in-domain model. Our method is marked in **blue**.

(a) 50 – 80 m							(b) 0 – 80 m						
AP _{BEV}	KITTI	Argoverse	Lyft	nuScenes	Waymo		AP _{BEV}	KITTI	Argoverse	Lyft	nuScenes	Waymo	
KITTI	40.4	20.1/26.9/ 28.4	49.6/56.3/ 56.4	1.4/ 9.1 / 4.5	42.8/48.5/ 50.2		KITTI	87.1	56.8/58.8/ 59.4	68.2/68.3/ 68.6	27.7/27.8/ 28.9	62.5/68.8/ 69.0	
Argoverse	18.0/28.2/ 30.1	37.8	46.3/48.8/ 54.5	0.6/ 9.1 / 3.0	50.4/50.9/ 56.0		Argoverse	82.3/83.2/ 83.3	68.5	66.7/67.0/ 67.6	26.9 /25.2/ 25.4	69.8/69.6/ 70.0	
Lyft	26.1/30.8/ 33.9	29.3/30.7/ 35.2	67.2	4.5/ 9.1 / 6.0	51.6/51.9/ 56.9		Lyft	82.6/84.9/ 85.3	60.2/65.7/ 66.0	79.2	28.8/28.2/ 29.1	70.6/70.9/ 71.1	
nuScenes	9.6/16.4/ 21.7	3.0/12.4/ 17.7	22.9/39.1/ 46.4	3.5	24.9/42.5/ 50.1		nuScenes	61.7/76.5/ 79.5	22.4/38.0/ 46.6	41.1/59.0/ 65.5	37.7	51.5/62.2/ 68.7	
Waymo	14.3/25.6/ 27.8	24.7/23.4/ 28.3	45.3/54.0/ 55.5	0.2/ 3.0/ 9.1	58.2		Waymo	81.2/82.2/ 83.1	56.9/58.3/ 59.2	67.3/68.9/ 69.4	23.1/27.2/ 29.1	71.9	
AP _{3D}	KITTI	Argoverse	Lyft	nuScenes	Waymo		AP _{3D}	KITTI	Argoverse	Lyft	nuScenes	Waymo	
KITTI	36.3	15.4/19.5/ 20.8	40.0/46.3/ 47.7	0.9/ 0.4 / 0.4	33.4/39.3/ 41.0		KITTI	86.7	49.0/55.0/ 55.9	60.7/65.2/ 66.4	20.9/22.3/ 23.1	59.4/60.5/ 61.4	
Argoverse	13.9/22.5/ 25.6	30.0	42.9/46.0/ 47.6	0.1/ 2.0 / 3.0	47.8/48.2/ 49.1		Argoverse	77.4/ 81.1 / 81.1	65.9	64.9/65.5/ 66.4	22.8 /22.1/ 21.7	62.1/62.4/ 67.7	
Lyft	20.4/24.0/ 26.4	25.3/26.7/ 27.3	65.5	0.4/ 9.1 / 1.1	49.6/50.6/ 50.8		Lyft	77.5/82.2/ 82.3	56.7/58.6/ 58.9	78.3	24.1/23.6/ 26.3	63.0/ 69.1 / 68.7	
nuScenes	5.5/ 8.7/ 13.4	3.0/ 9.1/ 13.0	15.1/30.2/ 37.8	2.8	22.5/39.6/ 45.5		nuScenes	45.5/67.5/ 70.3	19.5/36.1/ 42.7	32.4/56.7/ 58.6	36.8	42.6/60.2/ 61.1	
Waymo	8.5/18.1/ 19.7	19.6/21.3/ 22.3	43.0/48.1/ 49.6	0.0/ 0.3/ 9.1	50.8		Waymo	74.1/76.6/ 77.3	54.1/55.9/ 56.2	66.0/68.1/ 68.7	21.6/23.5/ 24.3	71.3	

TABLE III: **UDA from Argoverse to KITTI.** We report AP_{BEV}/ AP_{3D} of the **car** category at IoU = 0.7 and IoU = 0.5 across different depth range on the test set. NR stands for No-Retrain baseline, ST stands for Self-Training [40], SN stands for Statistical Normalization [7]. Our method *Dream* is marked in **blue**. We show the performance of in-domain model, i.e., the model trained and evaluated on KITTI, at the first row in **gray**. We also show results by directly applying online and offline (not feasible in real-time) tracking. Best viewed in color.

Method	IoU 0.5			IoU 0.7		
	0-30	30-50	50-80	0-30	30-50	50-80
in-domain	90.0 / 89.9	81.0 / 79.9	40.4 / 36.3	89.0 / 78.0	70.3 / 51.5	26.6 / 9.8
NR	89.4 / 88.8	71.0 / 65.2	18.0 / 13.9	72.6 / 47.8	35.8 / 14.6	4.9 / 3.0
NR + online	89.3 / 88.5	71.9 / 66.1	18.3 / 14.0	72.3 / 47.5	38.2 / 14.9	5.5 / 1.2
NR + offline	89.3 / 88.8	72.7 / 67.7	18.9 / 15.0	74.5 / 49.6	43.5 / 18.1	7.3 / 1.8
ST	89.5 / 89.2	73.2 / 68.5	28.2 / 22.5	76.8 / 52.9	44.2 / 20.6	13.1 / 2.1
Dream	89.3 / 89.2	74.6 / 72.4	30.1 / 25.6	77.6 / 54.7	49.9 / 24.3	14.5 / 3.4
SN	89.3 / 88.2	69.6 / 65.4	14.6 / 13.3	83.8 / 59.1	53.5 / 27.2	9.3 / 3.6
SN + online	89.4 / 88.2	68.8 / 65.4	19.4 / 16.2	83.5 / 60.1	50.7 / 27.0	13.4 / 9.5
SN + offline	89.4 / 88.2	68.9 / 66.1	21.8 / 19.3	83.3 / 59.2	50.9 / 26.9	13.8 / 9.8
SN + ST	89.8 / 89.3	74.4 / 72.8	22.3 / 21.3	87.0 / 70.4	62.2 / 37.7	16.4 / 7.1
SN + Dream	89.8 / 89.4	75.4 / 73.8	29.4 / 25.4	87.0 / 73.5	62.8 / 41.9	17.2 / 10.3

TABLE IV: **UDA from KITTI (city, campus) to KITTI (road, residential).** Naming is as in **Table III**.

Method	IoU 0.5			IoU 0.7		
	0-30	30-50	50-80	0-30	30-50	50-80
NR	89.4 / 89.4	79.7 / 77.9	36.2 / 30.9	88.3 / 77.9	66.0 / 47.5	19.2 / 6.2
ST	89.4 / 89.4	78.0 / 76.7	35.9 / 31.2	88.2 / 77.8	66.3 / 48.3	19.0 / 7.7
Dream	89.6 / 89.6	80.1 / 78.9	41.4 / 35.3	88.5 / 78.2	68.0 / 49.6	23.2 / 12.7

TABLE V: **Ablation study of UDA from Argoverse to KITTI.** We report AP_{BEV}/ AP_{3D} of the **car** category at IoU = 0.5 and IoU = 0.7 across different depth range, using POINTRCNN model. Naming is as in **Table III**. S stands for smoothing, R for resizing, I for interpolation and E for extrapolation.

Method	IoU 0.5			IoU 0.7		
	0-30	30-50	50-80	0-30	30-50	50-80
NR	89.4 / 88.8	71.0 / 65.2	18.0 / 13.9	72.6 / 47.8	35.8 / 14.6	4.9 / 3.0
ST	89.5 / 89.2	73.2 / 68.5	28.2 / 22.5	76.8 / 52.9	44.2 / 20.6	13.1 / 2.1
ST + S	89.5 / 89.3	73.7 / 68.7	28.3 / 22.3	76.8 / 53.4	45.2 / 21.5	12.9 / 4.7
ST + S + R	89.5 / 89.3	74.0 / 71.6	28.3 / 23.9	78.0 / 55.2	50.8 / 23.9	10.3 / 2.7
ST + S + R + I	89.3 / 89.1	73.9 / 71.6	28.1 / 23.3	77.6 / 54.5	50.4 / 24.0	11.2 / 3.4
ST + S + R + I + E	89.4 / 89.2	74.9 / 72.5	31.0 / 25.7	77.8 / 55.1	50.4 / 24.1	14.3 / 3.3

correct. But just these do not address the second criteria for desired pseudo-labels: *high coverage*. We observe noticeable boosts when interpolation and extrapolations are added, specially for far away objects. This is due to extrapolations and interpolations recovering pseudo-labels for low confidence or missed detections for distant vehicles.

Adaptation results using PIXOR. To show the generality of our approach, we further apply it to another detector PIXOR [3] from Argoverse to KITTI in **Table VI**. Dreaming improves the accuracy at farther ranges (30-80 m) while maintaining the accuracy at close range (0-30 m). Interestingly,

TABLE VI: **UDA from Argoverse to KITTI using PIXOR.** We report AP_{BEV} of the **car** category at IoU = 0.5 and 0.7. Naming is as in **Table III**.

Method	IoU 0.5			IoU 0.7		
	0-30	30-50	50-80	0-30	30-50	50-80
in-domain	88.7	62.6	21.4	79.6	49.9	10.0
NR	85.7	57.2	12.9	54.2	23.6	4.7
ST	86.0	56.3	12.2	55.3	24.6	2.5
Dream	87.1	61.1	20.2	58.0	28.1	4.8
SN	86.7	58.7	15.1	76.2	38.7	5.1
SN + ST	87.4	58.9	12.4	78.0	42.4	3.8
SN + Dream	87.4	64.2	22.2	77.9	42.5	4.5

at IoU 0.5 in the 30-80 m ranges, we are able to surpass the in-domain performance, which uses models trained only in the target domain with the ground-truth labels. This results showcases the power of unsupervised domain adaptation (UDA): with a suitably designed algorithm, UDA that leverages both the source and target domain could outperform models trained only in a single domain.

Others. We show more results and qualitative visualizations in the supplementary material.

V. CONCLUSION AND DISCUSSION

In this paper, we have introduced a novel method towards closing the gap between source and target in unsupervised domain adaptation for LiDAR-based 3D object detection. Our approach is based on self-training, while leveraging vehicle dynamics and offline analysis to generate pseudo-labels. Importantly, we can generate high quality pseudo-labels even for difficult cases (i.e. far-away objects), which the detector tends to miss before adaptation. Fine-tuning on these pseudo-labels improves detection performance drastically in the target domain. It is hard to conceive an autonomous vehicle manufacturer that could collect, label, and update data for every consumer environment, meeting the requirements to allow self-driving cars to operate everywhere freely and safely. By significantly reducing the adaptation gap between domains, our approach takes a significant step towards making this vision a reality nevertheless.

ACKNOWLEDGMENT

This research is supported by grants from the National Science Foundation NSF (III-1618134, III-1526012, IIS-1149882, IIS-1724282, TRIPODS-1740822, IIS-2107077, OAC-2118240, OAC-2112606), the Office of Naval Research DOD (N00014-17-1-2175), the Bill and Melinda Gates Foundation, and the Cornell Center for Materials Research with funding from the NSF MRSEC program (DMR-1719875).

REFERENCES

[1] C. R. Qi, W. Liu, C. Wu, H. Su, and L. J. Guibas, “Frustum pointnets for 3d object detection from rgb-d data,” in *CVPR*, 2018.

[2] S. Shi, X. Wang, and H. Li, “Pointrcnn: 3d object proposal generation and detection from point cloud,” in *CVPR*, 2019.

[3] B. Yang, W. Luo, and R. Urtasun, “Pixor: Real-time 3d object detection from point clouds,” in *CVPR*, 2018.

[4] S. Shi, C. Guo, L. Jiang, Z. Wang, J. Shi, X. Wang, and H. Li, “Pv-rcnn: Point-voxel feature set abstraction for 3d object detection,” in *CVPR*, 2020.

[5] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” in *CVPR*, 2012.

[6] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The kitti dataset,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.

[7] Y. Wang, C. Xiangyu, Y. Yurong, E. L. Li, B. Hariharan, M. Campbell, K. Q. Weinberger, and C. Wei-Lun, “Train in germany, test in the usa: Making 3d object detectors generalize,” in *CVPR*, 2020.

[8] B. Gong, Y. Shi, F. Sha, and K. Grauman, “Geodesic flow kernel for unsupervised domain adaptation,” in *CVPR*, 2012.

[9] M.-F. Chang, J. W. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, and J. Hays, “Argoverse: 3d tracking and forecasting with rich maps,” in *CVPR*, 2019.

[10] R. Kesten, M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinsky, W. Jiang, and V. Shet, “Lyft level 5 av dataset 2019,” [urlhttps://level5.lyft.com/dataset/](https://level5.lyft.com/dataset/), 2019.

[11] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettlinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov, “Scalability in perception for autonomous driving: Waymo open dataset,” 2019.

[12] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liang, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nuscenes: A multimodal dataset for autonomous driving,” *preprint, arXiv:1903.11027*, 2019.

[13] Y. Wang, W.-L. Chao, D. Garg, B. Hariharan, M. Campbell, and K. Q. Weinberger, “Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving,” in *CVPR*, 2019.

[14] Y. You, Y. Wang, W.-L. Chao, D. Garg, G. Pleiss, B. Hariharan, M. Campbell, and K. Q. Weinberger, “Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving,” in *ICLR*, 2020.

[15] R. Qian, D. Garg, Y. Wang, Y. You, S. Belongie, B. Hariharan, M. Campbell, K. Q. Weinberger, and W.-L. Chao, “End-to-end pseudo-lidar for image-based 3d object detection,” in *CVPR*, 2020.

[16] P. Li, X. Chen, and S. Shen, “Stereo r-cnn based 3d object detection for autonomous driving,” in *CVPR*, 2019.

[17] Y. Zhou and O. Tuzel, “Voxelnet: End-to-end learning for point cloud based 3d object detection,” in *CVPR*, 2018.

[18] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, “Domain-adversarial training of neural networks,” *JMLR*, vol. 17, no. 1, pp. 2096–2030, 2016.

[19] K. Saito, K. Watanabe, Y. Ushiku, and T. Harada, “Maximum classifier discrepancy for unsupervised domain adaptation,” in *CVPR*, 2018.

[20] Y. Chen, W. Li, C. Sakaridis, D. Dai, and L. Van Gool, “Domain adaptive faster r-cnn for object detection in the wild,” in *CVPR*, 2018.

[21] Z. He and L. Zhang, “Multi-adversarial faster-rcnn for unrestricted object detection,” in *ICCV*, 2019.

[22] T. Kim, M. Jeong, S. Kim, S. Choi, and C. Kim, “Diversify and match: A domain adaptive representation learning paradigm for object detection,” in *CVPR*, 2019.

[23] K. Saito, Y. Ushiku, T. Harada, and K. Saenko, “Strong-weak distribution alignment for adaptive object detection,” in *CVPR*, 2019.

[24] X. Zhu, J. Pang, C. Yang, J. Shi, and D. Lin, “Adapting object detectors via selective cross-domain alignment,” in *CVPR*, 2019.

[25] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. A. Efros, and T. Darrell, “Cycada: Cycle-consistent adversarial domain adaptation,” in *ICML*, 2018.

[26] F. S. Saleh, M. S. Aliakbarian, M. Salzmann, L. Petersson, and J. M. Alvarez, “Effective use of synthetic data for urban scene semantic segmentation,” in *ECCV*. Springer, 2018.

[27] Y.-H. Tsai, W.-C. Hung, S. Schulter, K. Sohn, M.-H. Yang, and M. Chandraker, “Learning to adapt structured output space for semantic segmentation,” in *CVPR*, 2018.

[28] Y. Zhang, P. David, H. Foroosh, and B. Gong, “A curriculum domain adaptation approach to the semantic segmentation of urban scenes,” *TPAMI*, vol. 42, no. 8, pp. 1823–1841, 2019.

[29] Y. Zou, Z. Yu, B. Vijaya Kumar, and J. Wang, “Unsupervised domain adaptation for semantic segmentation via class-balanced self-training,” in *ECCV*, 2018.

[30] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *ICCV*, 2017.

[31] A. RoyChowdhury, P. Chakrabarty, A. Singh, S. Jin, H. Jiang, L. Cao, and E. Learned-Miller, “Automatic adaptation of object detectors to new domains using self-training,” in *CVPR*, 2019.

[32] Q. Tao, H. Yang, and J. Cai, “Zero-annotation object detection with web knowledge transfer,” in *ECCV*, 2018.

[33] J. Liang, R. He, Z. Sun, and T. Tan, “Distant supervised centroid shift: A simple and efficient approach to visual domain adaptation,” in *CVPR*, 2019.

[34] W. Zhang, W. Ouyang, W. Li, and D. Xu, “Collaborative and adversarial network for unsupervised domain adaptation,” in *CVPR*, 2018.

[35] S. Kim, J. Choi, T. Kim, and C. Kim, “Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection,” in *ICCV*, 2019.

[36] G. French, M. Mackiewicz, and M. Fisher, “Self-ensembling for visual domain adaptation,” in *ICLR*, 2018.

[37] N. Inoue, R. Furuta, T. Yamasaki, and K. Aizawa, “Cross-domain weakly-supervised object detection through progressive domain adaptation,” in *CVPR*, 2018.

[38] M. Khodabandeh, A. Vahdat, M. Ranjbar, and W. G. Macready, “A robust learning approach to domain adaptive object detection,” in *ICCV*, 2019.

[39] D. McClosky, E. Charniak, and M. Johnson, “Effective self-training for parsing,” in *ACL*, 2006.

[40] A. Kumar, T. Ma, and P. Liang, “Understanding self-training for gradual domain adaptation,” *preprint, arXiv:2002.11361*, 2020.

[41] D.-H. Lee, “Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks,” in *Workshop on challenges in representation learning, ICML*, 2013.

[42] R. G. Cinbis, J. Verbeek, and C. Schmid, “Weakly supervised object localization with multi-fold multiple instance learning,” *TPAMI*, vol. 39, no. 1, pp. 189–203, 2016.

[43] C. Qin, H. You, L. Wang, C.-C. J. Kuo, and Y. Fu, “Pointdan: A multi-scale 3d domain adaption network for point cloud representation,” in *NeurIPS*, 2019.

[44] K. Saleh, A. Abobakr, M. Attia, J. Iskander, D. Nahavandi, M. Hossny, and S. Nahvandi, “Domain adaptation for vehicle detection from bird’s eye view lidar point cloud data,” in *ICCVW*, 2019.

[45] Z. Wang, S. Ding, Y. Li, M. Zhao, S. Roychowdhury, A. Wallin, G. Sapiro, and Q. Qiu, “Range adaptation for 3d object detection in lidar,” in *ICCVW*, 2019.

[46] B. Wu, X. Zhou, S. Zhao, X. Yue, and K. Keutzer, “Squeezesegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud,” in *ICRA*, 2019.

[47] J. Yang, S. Shi, Z. Wang, H. Li, and X. Qi, “S3d: Self-training for unsupervised domain adaptation on 3d object detection,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10 368–10 378.

[48] X. Liang, S. Liu, Y. Wei, L. Liu, L. Lin, and S. Yan, “Towards computational baby learning: A weakly-supervised approach for object detection,” in *ICCV*, 2015.

[49] A. Osep, P. Voigtlaender, J. Luiten, S. Breuers, and B. Leibe, “Large-scale object mining for object discovery from unlabeled video,” in *ICRA*, 2019.

[50] I. Misra, A. Shrivastava, and M. Hebert, “Watch and learn: Semi-supervised learning for object detectors from video,” in *CVPR*, 2015.

[51] K. Kumar Singh, F. Xiao, and Y. Jae Lee, “Track and transfer: Watching videos to simulate strong human supervision for weakly-supervised object detection,” in *CVPR*, 2016.

[52] K. Tang, V. Ramanathan, L. Fei-Fei, and D. Koller, “Shifting weights: Adapting object detectors from image to video,” in *NeurIPS*, 2012.

[53] A. Osep, P. Voigtlaender, J. Luiten, S. Breuers, and B. Leibe, “Large-scale object discovery and detector adaptation from unlabeled video,” *preprint, arXiv:1712.08832*, 2017.

[54] C. R. Qi, Y. Zhou, M. Najibi, P. Sun, K. Vo, B. Deng, and D. Anguelov, “Offboard 3d object detection from point cloud sequences,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6134–6144.

- [55] B. Yang, M. Bai, M. Liang, W. Zeng, and R. Urtasun, “Auto4d: Learning to label 4d objects from sequential point clouds,” *arXiv preprint arXiv:2101.06586*, 2021.
- [56] X. Weng, J. Wang, D. Held, and K. Kitani, “3d multi-object tracking: A baseline and new evaluation metrics,” *preprint, arXiv:1907.03961*, 2020.
- [57] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, “Pointpillars: Fast encoders for object detection from point clouds,” in *ICCV*, 2019, pp. 12 697–12 705.
- [58] M. D. Breitenstein, F. Reichlin, B. Leibe, E. Koller-Meier, and L. Van Gool, “Online multiperson tracking-by-detection from a single, uncalibrated camera,” *TPAMI*, vol. 33, no. 9, pp. 1820–1833, 2010.
- [59] ———, “Robust tracking-by-detection using a detector confidence particle filter,” in *ICCV*. IEEE, 2009, pp. 1515–1522.
- [60] Y. Hua, K. Alahari, and C. Schmid, “Online object tracking with proposal selection,” in *ICCV*, 2015, pp. 3092–3100.
- [61] C. Diaz-Ruiz, Y. Wang, W. Chao, K. Weinberger, and M. Campbell, “Vision-only 3d tracking for self-driving cars,” in *CASE*, 2019, pp. 1029–1034.
- [62] H.-k. Chiu, A. Prioletti, J. Li, and J. Bohg, “Probabilistic 3d multi-object tracking for autonomous driving,” *preprint, arXiv:2001.05673*, 2020.
- [63] T. Yin, X. Zhou, and P. Krähenbühl, “Center-based 3d object detection and tracking,” *preprint, arXiv:2006.11275*, 2020.
- [64] I. Miller, M. Campbell, and D. Huttenlocher, “Efficient unbiased tracking of multiple dynamic obstacles under large viewpoint changes,” *IEEE Transactions on Robotics*, vol. 27, no. 1, pp. 29–46, 2011.
- [65] M. Chen, K. Q. Weinberger, and J. Blitzer, “Co-training for domain adaptation,” in *NeurIPS*, 2011.
- [66] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2001.
- [67] J. Yin, J. Shen, C. Guan, D. Zhou, and R. Yang, “Lidar-based online 3d video object detection with graph-based message passing and spatiotemporal transformer attention,” in *CVPR*, 2020.
- [68] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “Pointnet++: Deep hierarchical feature learning on point sets in a metric space,” in *NeurIPS*, 2017.