

Inverting the Pose Forecasting Pipeline with SPF2: Sequential Pointcloud Forecasting for Sequential Pose Forecasting

Xinshuo Weng¹, Jianren Wang¹, Sergey Levine², Kris Kitani¹, Nicholas Rhinehart²

¹Robotics Institute, Carnegie Mellon University
{xinshuow, jianrenw, kkitani}@cs.cmu.edu

²Berkeley Artificial Intelligence Research Lab, University of California, Berkeley
{svlevine, nrhinehart}@eecs.berkeley.edu

Abstract: Many autonomous systems forecast aspects of the future in order to aid decision-making. For example, self-driving vehicles and robotic manipulation systems often forecast future object poses by first detecting and tracking objects. However, this detect-then-forecast pipeline is expensive to scale, as pose forecasting algorithms typically require labeled sequences of object poses, which are costly to obtain in 3D space. Can we scale performance without requiring additional labels? We hypothesize yes, and propose inverting the detect-then-forecast pipeline. Instead of detecting, tracking and then forecasting the objects, we propose to first forecast 3D sensor data (e.g., point clouds with 100k points) and then detect/track objects on the predicted point cloud sequences to obtain future poses, i.e., a forecast-then-detect pipeline. This inversion makes it less expensive to scale pose forecasting, as the sensor data forecasting task requires no labels. Part of this work’s focus is on the challenging first step – Sequential Pointcloud Forecasting (SPF), for which we also propose an effective approach, SPFNet. To compare our forecast-then-detect pipeline relative to the detect-then-forecast pipeline, we propose an evaluation procedure and two metrics. Through experiments on a robotic manipulation dataset and two driving datasets, we show that SPFNet is effective for the SPF task, our forecast-then-detect pipeline outperforms the detect-then-forecast approaches to which we compared, and that pose forecasting performance improves with the addition of unlabeled data. Our project website is <http://www.xinshuoweng.com/projects/SPF2>.

Keywords: 3D point cloud, sequential forecasting, evaluation metrics

1 Introduction

Many autonomous systems forecast aspects of the future to aid decision-making. For example, self-driving vehicles and robotic manipulation systems often forecast future object poses using pose forecasting algorithms [1, 2, 3, 4, 5, 6, 7, 8, 9]. Given sensor data, these approaches first detect/track objects, and then train a pose forecasting algorithm to predict future poses, e.g., object location and velocity. However, this detect-then-forecast pipeline is expensive to scale, as pose forecasting algorithms typically require labels of future object poses and their correspondences in time, which are costly to obtain, especially in 3D space [10]. Can we create a pose forecasting pipeline that is cheaper to scale by leveraging large unlabeled 3D point cloud sequences?

We hypothesize yes, and propose an inverted forecast-then-detect pipeline (point cloud forecasting followed by detection/tracking). By reserving the order, forecasting is now performed on the point clouds, which does not require any expensive-to-collect labels. Notably, the first step in this pipeline is a challenging problem that no prior work has addressed (the closest is perhaps [3], which predicts flow of the points and requires point-wise annotations). Therefore, we consider this first step an important research direction in its own right, and believe it may prove useful to other downstream tasks. We term this task Sequential Pointcloud Forecasting (SPF): given a sequence of past point clouds as input, output a sequence of future point clouds. Two benefits of the SPF task are that (1) the prediction represents the entire scene, which contains information about both scene background and foreground objects, and (2) it requires only point cloud data, which costs significantly less to collect than labels, and is already prevalent in public datasets [9, 10, 11].

Towards the SPF task, we present SPFNet, which employs an LSTM autoencoder model. To leverage 2D Convolutional Neural Networks (CNNs) while preserving 3D structure of the point cloud, we use a range map representation [12]. For training and evaluation, we use raw LiDAR point clouds from two real-world autonomous driving datasets (KITTI [11] and nuScenes [10]) and depth point clouds from a Baxter robot dataset [3]. We show that SPFNet can reliably forecast future point clouds in different domains and outperform competitive approaches.

We then present our forecast-then-detect pipeline, SPF², which reverses the steps in the conventional detect-then-forecast pipeline that first performs perception (3D object detection & tracking) and then trajectory forecasting. Specifically, our SPF² pipeline first performs forecasting at the sensor data level using SPFNet, and then uses off-the-shelf 3D detection/tracking on the predicted point clouds to obtain future object trajectories. This pipeline inherits the benefits of the SPF task: (1) forecasting at the sensor data level can leverage the geometry of the scene – when forecasting point clouds of the entire scene, locations of every point are implicitly constrained by those of surrounding points; (2) our forecasting module is unsupervised, which facilitates training with data at larger scales than supervised pose forecasting. Note that the detector/tracker in our pipeline may still require labels.

Standard trajectory forecasting evaluation assumes access to the set of ground-truth (GT) past trajectories, and each GT past trajectory is used as input to forecast a future trajectory. As the correspondence between each GT past and GT future trajectory is known, every predicted trajectory has a target GT future trajectory. This *a priori* knowledge of the mapping between each forecasted trajectory and its target GT future trajectory is required to compute the conventional error metrics (ADE and FDE [13]). However, come deployment time, both the conventional detect-then-forecast pipeline and the SPF² pipeline do not have access to GT past trajectories. As a result, there is no known correspondence between each forecasted and GT trajectory, which prevents computation of the ADE/FDE metrics. To evaluate detect-then-forecast and our pipeline in an end-to-end fashion, we propose a procedure that first establishes an association between forecasted and GT trajectories with a matching step. Unlike standard ADE/FDE computation, this association does not provide 100% recall of the GT trajectories – some GT trajectories may be missed, with true positives identified by thresholding the ADE score. Because different methods may achieve different recalls, we propose metrics to *average* ADE/FDE over recalls. Our contributions are summarized as follows:

1. **A novel forecast-then-detect pipeline, SPF²**, that can scale forecasting performance by leveraging unlabeled point cloud data and outperform the detect-then-forecast pipeline.
2. **A new task, Sequential Pointcloud Forecasting (SPF)**, which predicts a 3D representation of the future of the entire scene and does not require any human annotation for training.
3. **An effective approach for SPF**, deemed SPFNet, that outperforms competitive approaches in one robot manipulation dataset and two real-world autonomous driving datasets.
4. **An evaluation procedure with two metrics** that can evaluate performance of the detect-then-forecast pipeline and our SPF² pipeline in an end-to-end fashion.

2 Related Work

Forecasting Tasks. Prior tasks can be categorized into trajectory forecasting, activity forecasting and video forecasting. Trajectory (or pose) forecasting [6, 14, 7, 8, 15, 16] predicts a sequence of future ground positions of target objects. Activity forecasting [1, 2] predicts a person’s future actions (such as walking and running). Both trajectory forecasting and activity forecasting are object-level forecasting, which (1) require object labels; (2) solely predict information about target objects while ignoring everything else in the scene. In contrast, our SPF task predicts a richly informative 3D representation of the entire scene and does not require any human annotation for object labels.

Video forecasting [17, 18] is similar to SPF, as both tasks forecast high-dimensional sensor data. The difference is that video forecasting outputs a geometrically-projected 3D representation, whereas SPF outputs a full 3D representation. Beyond forecasting RGB images, [19, 20] jointly predict the depth image and RGB image in the next frame, which is also related to SPF, as depth forecasting is a 3D forecasting task, although the data representation is different and [19, 20] only perform one-frame prediction. More importantly, prior 3D forecasting methods have primarily focused on the accuracy of the forecasting task itself, while we focus both on the accuracy of the SPF task and on utilizing it to enhance performance of the downstream pose forecasting task.

Point Cloud Generation. Prior work focuses on generating a single object point cloud. [21] proposes to use Chamfer distance [21] and Earth Mover’s distance [22] as loss functions to generate

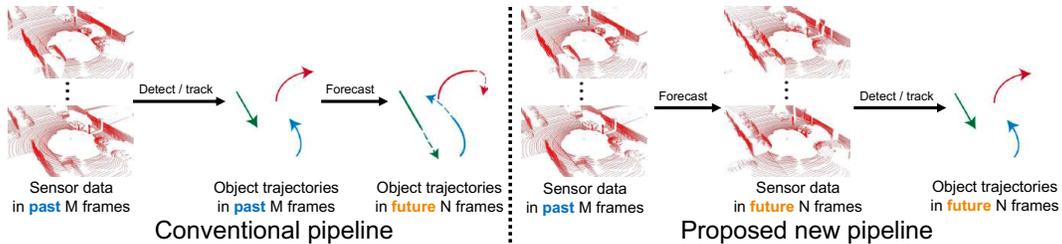


Figure 1: Comparison of perception and trajectory forecasting pipeline. (Left) Conventional detect-then-forecast pipeline. (Right) Our pipeline first forecasts point clouds, and then detects and tracks objects.

object point clouds. [23] generates object point cloud and then project into multiple 2D viewpoints optimized with a 2D re-projection error. Beyond deterministic point cloud generation, prior work employs stochastic frameworks such as flow-based models [24], generative adversarial networks (GANs) [25] and variational autoencoders (VAEs) [26].

Although most prior work focuses on point cloud generation for a single object, a few works attempted to generate point clouds for an entire scene, which we call scene point clouds. [12] uses VAEs and GANs to generate a scene point cloud from a random noise vector. [27] generates a high-resolution scene point cloud conditioned on stereo images and low-resolution LiDAR inputs. [3] predicts a scene point cloud in the next frame by considering the action applied to an object in the scene. Different from [12, 27, 3] which output a single scene point cloud, SPF generates a sequence of scene point clouds. To the best of our knowledge, SPF is the first formulated as a sequence prediction of large-scale scene point clouds (100k points per frame).

Perception and Trajectory Forecasting Pipeline aims to localize objects and predict where they will move in the future, given the sensor data. The conventional detect-then-forecast pipeline has two modules: (1) a perception module (including detection [28, 29] and tracking [30, 31]) to obtain past object trajectories; (2) a trajectory forecasting module to predict the future object trajectories based on the past object trajectories. Prior work often separately evaluates perception and forecasting, e.g., using average precision and CLEAR metrics [32] for perception evaluation, and use ADE and FDE [13] to evaluate trajectory forecasting. However, modularized trajectory forecasting evaluation lacks the ability to assess the performance of the entire pipeline, as it employs the unrealistic assumption that we have access to past GT object trajectories.

Different from the conventional perception and forecasting pipeline, our pipeline first performs the forecasting at the sensor level and then applies detection and tracking to the predicted sensor data to obtain forecasted pose trajectories. Also, our new evaluation procedure can account for the error from detection, tracking and forecasting modules, enabling the evaluation over the entire pipeline. Although [16] also propose to evaluate the entire pipeline in an end-to-end fashion, their evaluation is performed at a single operating point, while our evaluation with two proposed metrics can summarize performance over many operating points.

3 The SPF² Pipeline

As mentioned in Sec. 2, the conventional pipeline is separated into two steps as in Fig. 1 (left): (1) a detection/tracking module to obtain past object trajectories; (2) a forecasting module to predict future trajectories based on the past trajectories. In contrast, we propose a new pipeline, Sequential Pointcloud Forecasting for Sequential Pose Forecasting (SPF²), depicted in Fig. 1 (right). Specifically, we first forecast point cloud sequences using our SPFNet (Sec. 3.2), and then use a 3D detector and tracker on the predicted point cloud sequences to obtain pose sequences (trajectories).

3.1 Sequential Pointcloud Forecasting

Task Formulation. We first define a 3D point cloud enclosing an entire scene as the scene point cloud, which includes points belonging to both objects and scene background. The goal of SPF is to predict a sequence of future scene point clouds given a sequence of past scene point clouds. Specifically, given M frames of past scene point clouds with each frame $S_t = \{(x, y, z)_j\}_{j=1}^{K_t}$, where $t \in [-M + 1, \dots, 0]$ denotes the frame index, $j \in [1, \dots, K_t]$ denotes the index of points and K_t denotes the number of points at frame t , the goal of SPF is to predict N frames of future scene point clouds with each frame $S_t = \{(x, y, z)_j\}_{j=1}^{K_t}$, where $t \in [1, \dots, N]$.

3.2 Approach: SPFNet

Our SPFNNet is shown in Fig. 2, which has four modules: (a) a shared encoder to extract features from past scene point clouds, (b) an LSTM to model temporal dynamics of the features of the scene point clouds, (c) a shared decoder to generate scene point clouds in the future and (d) a combination of the L1 and binary cross-entropy losses is applied to the generated intermediate range maps and a Chamfer distance loss is applied to the predicted point clouds at every future frame.

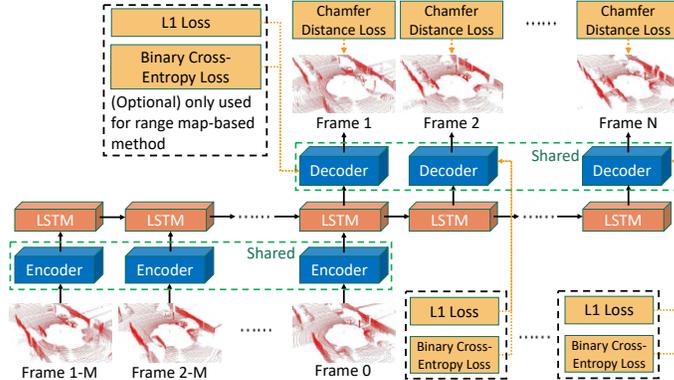


Figure 2: Proposed SPFNNet.

Encoder. To predict the future scene point clouds, it is important to extract useful information from the past through an encoder. In this work, we investigate two types of encoders:

(1) *Point-based encoder.* To process a point cloud, PointNet [33] is a straightforward choice. We use a network (multi-layer perceptron (MLP) plus a max-pooling operator) similar to [33] to obtain the local feature for each point and the global feature for the scene. To leverage both the local per-point feature and global scene feature, we concatenate them so that every point has the feature of its own and the global feature. However, preserving high-dimensional features for each point will reach the GPU memory limit (e.g., RTX 2080 Ti has 11Gb), as SPF requires to process a sequence (e.g., 30 frames) of point clouds each with a large number of points (e.g., 100k). In fact, if maintaining high-dimensional features for each point, the required GPU memory will be massive. As a result, we add another MLP and a max-pooling operator to obtain a single compact feature to represent the entire scene point cloud at every frame as the output of the encoder.

(2) *Range map-based encoder.* In addition to directly extract the feature from point clouds, another popular approach in the literature [34, 12] is to first transform the point cloud into a 2D representation and then use 2D CNNs for feature extraction. Although these methods are fast and efficient, the disadvantage of most of these methods (except for the range map [12]) is to lose significant information. For example, if projecting the point cloud to the bird’s eye view, information in the height direction is lost. As a result, we choose to use the range map as our 2D representation [12], which preserves full 3D information in the LiDAR point cloud. For details about how to obtain the range map from the input scene point cloud, please refer to our supplementary materials. After the scene point cloud is transformed into a 2D range map with a size of $H \times W$ (we use 120×1024 in our experiments to predict 122.88k points per frame), we use a standard 2D CNN with 8 layers (convolution + batch norm + ReLU) to extract features. For the detailed network architecture of our encoders, we provide visualization in our supplementary.

Temporal Dynamics Modeling. As SPF is formulated as a sequence-to-sequence problem, we use a standard LSTM network to learn temporal dynamics of the features. Specifically, we feed the feature obtained from the encoder in the past M frames to the LSTM network and predict features for the future N frames in an autoregressive manner.

Decoder. To predict future scene point clouds, a decoder is needed to convert the predicted features from LSTM to scene point clouds. With two encoders carefully designed, the corresponding decoder follows the encoder design in a symmetric manner: (1) For the point-based decoder, we first increase the dimensionality of the LSTM output feature to $1 \times (K \times 3)$ using MLP, and then reshape it to a scene point cloud with a size of $K \times 3$. (2) For the range map-based decoder, we use a 2D CNNs with eight layers (deconv + batch norm + ReLU) to map the LSTM predicted feature back to a range map with $H \times W$. Additionally, to deal with the holes in the range map (*i.e.*, the pixels that do not have any point in the point cloud), we also predict a binary range mask R_M (0 denotes pixel without a point) with a shape of $H \times W$ using a similar 2D CNNs. Then, when we transform the predicted range map back to a scene point cloud only for the pixels that are not masked out.

Losses. We use the Chamfer distance [21] loss \mathcal{L}_{cd} to minimize the distance between ground truth $S_{1:N}^*$ and predicted point cloud sequences $\hat{S}_{1:N}$. For the range map-based method, we apply two additional losses on the intermediate range map outputs to form our full objective in Eq. 1: (1)

a L_1 distance loss between the ground truth, R^* , and predicted, \hat{R} , 2D range maps; (2) a binary cross-entropy loss L_{bce} between the predicted \hat{R}_M and ground truth R_M^* range mask.

$$\mathcal{L} \doteq \sum_{i=1}^N \mathcal{L}_{cd}(\hat{S}_i, S_i^*) + \lambda_1 \mathcal{L}_1(\hat{R}_i, R_i^*) + \lambda_2 \mathcal{L}_{bce}(\hat{R}_{Mi}, R_{Mi}^*), \quad (1)$$

$$\text{where } \mathcal{L}_{cd}(\hat{S}, S^*) \doteq \sum_{x \in \hat{S}} \min_{y \in S^*} \|x - y\|_2^2 + \sum_{y \in S^*} \min_{x \in \hat{S}} \|x - y\|_2^2,$$

3.3 3D Object Detection and 3D Multi-Object Tracking

Once we predict future point clouds using SPFNet, the last step of our SPF² pipeline is 3D detection and tracking to obtain objects’ future poses. In theory, our SPF² pipeline is generic so any detector and tracker can be used, e.g., using unsupervised detectors/trackers to make our full pipeline unsupervised. However, existing unsupervised 3D detectors often have lower accuracy than supervised 3D detectors. To ensure accuracy of our SPF² pipeline, we use a supervised 3D detector [35] while use an unsupervised 3D tracker [30], which can achieve strong tracking performance in KITTI.

4 End-to-End Perception and Trajectory Forecasting Evaluation

As mentioned in Sec. 1 and 2, standard modularized evaluation for trajectory forecasting relies on the unrealistic assumption that we have access to GT past trajectories. As a result, standard modularized evaluation is not applicable for evaluating either the detect-then-forecast or our pipeline. Consider two common forecasting metrics: Average Displacement Error (ADE) and Final Displacement Error (FDE) [13], which compute the distance between each pair of prediction and its corresponding GT. In the modularized evaluation setting, as GT past trajectories are used, computing ADE and FDE is straightforward as each predicted future trajectory has its corresponding future GT trajectory. However, in the end-to-end evaluation setting, computing ADE and FDE requires an additional matching step between the forecasted trajectories and GT trajectories as their correspondences are not known as *a priori*: see Fig. 3 (left). To obtain the correspondences, we match forecasted trajectories to GT trajectories based on ADE scores using the Hungarian algorithm [36]. Please refer to our supplementary material for details regarding this matching step.

As errors can be accumulated in detection/tracking/forecasting, the final forecasted trajectories are not perfect and some GT objects might not be matched to any predicted trajectory (i.e., a false negative). We identify true positives by thresholding the ADE score. In general, the recall value will not be 100% and it might vary substantially across methods. This is in contrast to the modularized evaluation where the recall is always 100%. As an example, we show the ADE-over-recall curve computed on the KITTI dataset in Fig. 3 (right) for our method and two baselines (see Sec. 5.3 for details). We can see that ADE tends to increase as recall increases, i.e., a looser ADE threshold will include more high-ADE forecasted trajectories. Therefore, to achieve a relatively fair comparison, methods must compare ADE either at an identical recall or by averaging ADE across the same range of recalls. We choose the latter solution because evaluating at a single recall can produce less robust numbers¹, e.g., some methods might perform better at a recall point but worse at other recalls. To that end, we propose two metrics called AADE/AFDE (Average ADE/FDE) to summarize ADE/FDE performance across a range of recalls. Specifically, AADE/AFDE are computed by integrating ADE/FDE. Similar to other integral metrics such as average precision in object detection, we approximate the integration by a summation of the ADE/FDE over a set of recalls.

5 Experiments

Datasets. We evaluate on two real-world autonomous driving datasets: KITTI-Raw [11], nuScenes [10], and a robot manipulation dataset of the Baxter robot [4].

5.1 Evaluating SPFNet in the Autonomous Driving Domain

Evaluation and Settings. We use standard Earth Mover’s Distance (EMD) [22] and Chamfer Distance (CD) [21] to measure the distance between GT and predicted future point clouds. We evaluate two settings for all methods: (1) observe 1.0s past and predict the next 1.0s; (2) observe 3.0s past and predict the next 3.0s. For KITTI, 1.0s and 3.0s means 10 and 30 frames. For nuScenes key-frames only, 1.0s and 3.0s means 2 and 6 key-frames.

Baselines. As this paper is the first towards the SPF task, there are no existing methods for comparison. Therefore, we devise baselines based on existing techniques: (1) *Identity*. We duplicate the

¹[16] identify true positives using a 3D IoU threshold of 0.5, i.e., a single operating point. As a result, the reported numbers do not account for performance at other operating points, e.g., at a higher recall including objects below the IoU threshold of 0.5, e.g., those difficult to detect.

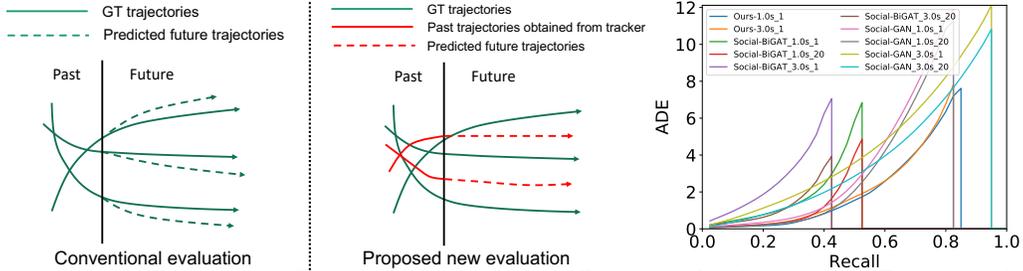


Figure 3: (Left) Comparison of the evaluation procedure. Conventional evaluation uses GT past to predict the future, which has the prior knowledge of the mapping between each predicted and GT future trajectory. In the real world, people use tracking outputs as past trajectories so the predicted trajectories do not have known mapping with GT, and thus a matching step is needed. **(Right) ADE-over-recall curve.** Methods can have different maximum recalls so it is not fair to compare ADE of each method performing at its maximum recall. Therefore, integral metrics that summarize performance across the same range of recalls are preferred.

Table 1: Quantitative evaluation for the proposed SPF task on the KITTI and nuScenes datasets.

Datasets	Metrics	Identity	GT-Ego	Est-Ego	Align-ICP	Align-[39]	SceneFlow	Ours+Point	Ours+RM
KITTI-1.0s	CD↓	12.82	5.47	9.18	6.13	6.02	3.15	2.37	0.92
	EMD↓	526.87	391.03	495.21	418.25	439.17	291.63	211.47	128.81
KITTI-3.0s	CD↓	13.31	7.91	11.31	9.14	9.57	5.08	3.91	1.57
	EMD↓	602.89	452.81	502.83	470.25	493.26	351.46	267.42	175.54
nuScenes-1.0s	CD↓	8.42	2.16	4.91	4.04	3.50	1.93	1.25	0.40
	EMD↓	461.63	168.37	299.13	281.53	270.81	117.41	135.94	78.37
nuScenes-3.0s	CD↓	10.16	2.85	6.52	7.13	5.27	3.81	2.97	0.69
	EMD↓	494.81	190.14	370.91	419.37	332.97	294.53	128.26	91.83

scene point cloud from the last past frame to the future for N times; (2) *GT-Ego*. We use ground truth ego-motion (rotation and translation) between adjacent past frames and compute an average motion. Then, we use the average motion to warp the last past frame of point cloud to the future frames; (3) *Est-Ego*. Instead of using the ground truth ego-motion, we estimate the ego-motion using [37]; (4) *Align*. We use the point cloud alignment methods (ICP [38] and Deep-ICP [39]) to compute the global rigid motion between adjacent past frames and obtain the average motion for warping; (5) *SceneFlow*. We use scene flow methods [40] to estimate the point-wise motion between the last two past frames and use the motion to warp every single point.

Results. We summarize the results in Table 1, where +RM and +Point denote our range map-based and point-based method. Our SPFNet consistently outperforms baselines in both metrics. The results are reasonable because our SPFNet (1) is designed for sequence prediction and (2) can learn the future location for each individual point, while the baselines are weak because they are originally designed for one-frame prediction (i.e., warping between two frames) and also most baselines only estimate global motion. However, it is a common difficulty to find strong baselines when proposing a new task. Also, we observed that Ours+RM performs better than Ours+Point. We hypothesize that predicting in the range map space may be relatively easier than in the point cloud space for our SPF task, although it is possible to see further improvements for the point-based method if we use a more advanced point-based encoder. Additionally, we find that the results on nuScenes seem to consistently better than KITTI, which we believe are primarily due to two reasons: 1) nuScenes LiDAR sensor has a smaller sensing range (e.g., 70m v.s. 120m of the KITTI LiDAR sensor), and forecasting for the nearby range is generally considered easier; 2) nuScenes keyframes have a lower frame rate (2Hz v.s. 10Hz in KITTI), leading to fewer frames of prediction at the same time horizon.

Performance scaling with unlabeled data. As our SPFNet can be trained with massive unlabeled point clouds without any human annotation, it is useful to know if its performance scales with the amount of training data. To that end, we train our SPFNet on the KITTI dataset multiple times, each time a different number of training sequences are withheld. This experiment is conducted in the 1.0s prediction setting using the range map-based method. The results are shown in Fig. 4 with two sub-figures, one with number of training samples and the other with number of training sequences as the x axis. Clearly, with more training sequences or more training samples, CD and EMD are decreasing, which confirms our hypothesis that it is cheaper to scale performance of the proposed SPFNet by leveraging unlabeled point cloud data.

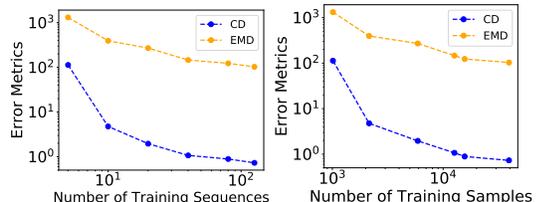


Figure 4: Error v.s. amount of training data.

Table 2: Quantitative evaluation on the Baxter dataset. Note that the numbers within the parentheses for the PPFE metric are what we have reproduced, which are very close to the reported results by the authors of the baselines. This demonstrates that we successfully reproduced the baselines for comparison.

Metrics	SE3-Pose-Nets + Joint Angles	SE3-Pose-Nets [4]	SE3-Nets [41]	Ours+Point	Ours+RM
PPFE↓	0.017 (0.019)	0.027 (0.025)	0.022 (0.024)	0.035	0.030
CD↓	0.315	0.313	0.357	0.153	0.094
EMD↓	52.7	50.9	64.1	33.5	18.2

5.2 Evaluating SPFNet in the Robot Manipulation Domain

To show that our proposed SPFNet can be applicable to not only the autonomous driving domain but also other domains such as robot manipulation, we evaluate on the Baxter dataset collected by [4] using the Baxter simulator. Specifically, the dataset contains a Baxter robot (shown in Figure 5) which moves its right arm with seven joints towards the target position in front of an RGBD camera.

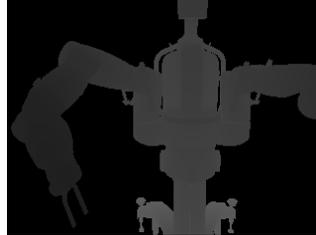


Figure 5: A example depth image of the Baxter robot.

Evaluation Metrics. As the Baxter dataset provides camera matrix and pixel correspondences between depth images across frames, computing flow in 3D space is possible. Therefore, we follow [4] and report a per-point flow error (PPFE) to measure the differences between predicted and ground truth flow in the 3D space, in addition to the EMD and CD metrics.

Baselines. We compare against S.O.T.A. methods (SE3-Nets [41], SE3-Pose-Nets [4], and its one variant). Note that these baselines are different from our SPFNet in two-fold: 1) they only support one-frame prediction (i.e., not sequence-to-sequence forecasting) and it is nontrivial to adapt them for sequence forecasting. Thus, to obtain a direct comparison, we only test the first predicted frame of our SPFNet, though we train our SPFNet with 10 input and 10 output frames to better learn scene dynamics; 2) instead of predicting the next frame of the point cloud, these baselines predict the flow of each point and then warp the past point cloud to the next frame with the predicted flow. As a result, these baselines require point correspondences to obtain ground truth flow for training and cannot be applied to LiDAR point clouds which generally do not have point correspondences.

Results. We summarize the results in Table 2. Note that only results for the PPFE metric is reported in [4]. In order to compute the results of baselines for CD and EMD, we first need to re-train the baselines (no pre-trained model is released by the authors of the baselines). To make sure that we successfully re-trained the baselines, we show the reproduced results on the PPFE metric within the parentheses, which are very close to the reported results in [4]. In fact, our re-trained SE3-Pose-Net obtained a slightly lower error in PPFE than the reported number in [4], i.e., 0.025 v.s. 0.027. With successfully re-trained baselines, we can now additionally compute for CD and EMD metrics. We can see that, our SPFNet achieved much better performance on the CD and EMD metrics than all baselines. This is probably because the baselines are trained to optimize the flow instead of the point cloud distance (e.g., CD or EMD). On the other hand, though our SPFNet is not optimized for the flow, our PPFE is competitive to the baselines such as SE3-Pose-Net (0.030 v.s., 0.027). Overall, the results show that our SPFNet is able to learn scene dynamics on a robot dataset.

5.3 Evaluating Perception and Trajectory Forecasting Pipeline

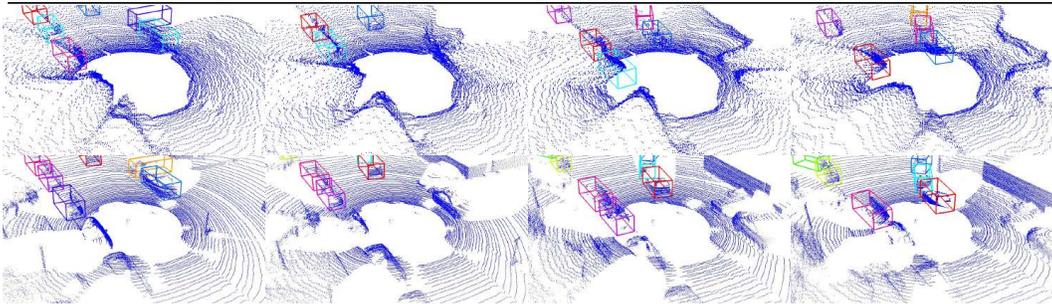
Evaluation and Settings. We use the new evaluation procedure with the AADE/AFDE metrics. To fairly compare AADE/AFDE for all methods, we only integrate ADE/FDE up to a recall that all methods can reach. This is because we cannot integrate ADE/FDE for a method beyond its maximum recall, which is undefined. We use the 1.0 and 3.0 seconds prediction. As trajectory forecasting modules used in the conventional pipeline often perform stochastic prediction, we evaluate them by sampling 1 and 20 times. For our pipeline with deterministic forecasting module SPFNet, it can be viewed as 1 sample. We use the range map-based encoder here for our method.

Baselines. We employ different trajectory forecasting modules [6, 7, 42, 14] in the conventional perception and trajectory forecasting pipeline. For a fair comparison, we use the same detector [35] and tracker [30] for both our and conventional pipeline. Although [6, 7] are originally proposed for predicting the trajectory of people, we adapt them to work with the car.

Results. We summarize performance on the KITTI and nuScenes datasets in Table 3. We can see that our SPF² outperforms the conventional detect-then-forecast pipeline using different forecasting

Table 3: Evaluation for the perception and trajectory forecasting pipeline on the KITTI and nuScenes datasets.

Datasets	Metrics	Samples	Conv-Social [14]	Social-GAN [6]	Social-BiGAT [7]	TraPHic [42]	Ours
KITTI-1.0s	AADE↓	1	0.792	0.524	1.099	0.470	0.317
		20	0.623	0.340	0.443	0.382	–
	AFDE↓	1	1.285	0.886	1.708	0.889	0.405
		20	1.152	0.511	0.546	0.613	–
KITTI-3.0s	AADE↓	1	1.692	1.362	2.720	1.432	0.408
		20	1.593	0.984	1.231	0.725	–
	AFDE↓	1	2.670	2.267	3.938	2.536	0.504
		20	2.385	1.512	1.405	1.118	–
nuScenes-1.0s	AADE↓	1	1.186	1.117	2.030	1.214	0.821
		20	0.907	0.762	0.826	0.881	–
	AFDE↓	1	1.490	1.310	2.337	1.563	0.825
		20	1.231	0.763	0.849	1.197	–
nuScenes-3.0s	AADE↓	1	1.794	2.224	4.954	2.417	1.044
		20	1.658	1.426	1.760	1.938	–
	AFDE↓	1	2.850	3.224	6.765	3.479	1.043
		20	2.538	1.652	1.845	2.766	–

**Figure 6:** Qualitative comparison between our prediction and GT on KITTI. **(Top)** Predicted future scene point clouds by our SPFNet and detection/tracking outputs on the predicted point clouds by our SPF² pipeline. Bounding boxes with the same color across frames represent a predicted object trajectory by our pipeline; **(Bottom)** GT future point clouds and GT bounding boxes. Note that KITTI only provides object labels in the frontal view, *i.e.*, objects that are behind the ego-vehicle do not have bounding box labeled. Also, the color of the same object might be different in prediction and GT as their absolute IDs (identities) are different.

modules with 1 sample. When sampling 20 times for stochastic forecasting modules in the conventional pipeline, we do see improvements, but surprisingly, our deterministic SPF² with 1 sample still achieves the best performance in most settings. We hypothesize it is because that (1) our forecasting at the sensor level is indeed effective in our SPF² pipeline, and (2) the data used for evaluation might not contain much uncertainty, *i.e.*, future is quite certain as vehicles are moving straight most of the time. As a result, evaluation on other datasets with diverse futures might be valuable in the future.

Visualization. We visualize the results of our SPFNet on the KITTI dataset for the 3.0s prediction setting in Fig. 6 (top). We plot 4 frames of results from a prediction of 30 frames with an interval of 7 frames, showing that our predicted scene point clouds reasonably reconstruct the objects and scene, and align well with GT shown in Fig. 6 (bottom). However, the current limitation is that the shape of the objects in predicted point clouds are not sharp enough. Moreover, we visualize detection/tracking outputs (colored bounding boxes) on top of our predicted scene point clouds, showing that the tracking outputs of our SPF² pipeline are very close to the GT bounding boxes.

6 Conclusion

We proposed a new perception and trajectory forecasting pipeline, SPF², that inverts the detect-then-forecast pipeline. Towards our SPF² pipeline, we proposed a new forecasting task, SPF, and method, SPFNet, which predicts a richly informative 3D representation of the future and does not need any human annotation for training. Also, an evaluation procedure is proposed to support end-to-end evaluation of the perception and forecasting pipeline. We verified the performance of our SPFNet and SPF² pipeline on two real-world autonomous driving datasets and one robotic manipulation dataset, which demonstrates that our forecasting performance can scale by leveraging unlabeled data and the proposed SPF is useful to the downstream task of trajectory forecasting. As our inverted pose forecasting pipeline is generic, it might be possible to extend it to other instances of pose forecasting (not only trajectory forecasting) domains, such as future activity forecasting.

Acknowledgments

This research was supported in part by the Office of Naval Research and ARL DCIST CRA W911NF-17-2-0181, and JST AIP Acceleration, Grant Number JPMJCR20U1, Japan.

References

- [1] T. Lan, T. C. Chen, and S. Savarese. A Hierarchical Representation for Future Action Prediction. *ECCV*, 2014.
- [2] H. S. Koppula and A. Saxena. Anticipating Human Activities Using Object Affordances for Reactive Robotic Response. *TPAMI*, 2016.
- [3] A. Byravan and D. Fox. SE3-Nets: Learning Rigid Body Motion Using Deep Neural Networks. *ICRA*, 2017.
- [4] A. Byravan, F. Leeb, F. Meier, and D. Fox. SE3-Pose-Nets: Structured Deep Dynamics Models for Visuomotor Planning and Control. *ICRA*, 2018.
- [5] N. Rhinehart, M. Kris, and P. Vernaza. R2P2: A Reparameterized Pushforward Policy for Diverse, Precise Generative Path Forecasting. *ECCV*, 2018.
- [6] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi. Social GAN: Socially Acceptable Trajectories with Generative Adversarial Networks. *CVPR*, 2018.
- [7] V. Kosaraju, A. Sadeghian, R. Martín-Martín, I. Reid, S. H. Rezatofighi, and S. Savarese. Social-BiGAT: Multimodal Trajectory Forecasting using Bicycle-GAN and Graph Attention Networks. *NeurIPS*, 2019.
- [8] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine. PRECOG: PREDiction Conditioned On Goals in Visual Multi-Agent Settings. *ICCV*, 2019.
- [9] M.-f. Chang, J. Lambert, P. Sangkloy, J. Singh, B. Sławomir, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, and J. Hays. Argoverse: 3D Tracking and Forecasting with Rich Maps. *CVPR*, 2019.
- [10] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, and Q. Xu. nuScenes: A Multimodal Dataset for Autonomous Driving. *CVPR*, 2020.
- [11] A. Geiger, P. Lenz, and R. Urtasun. Are We Ready for Autonomous Driving? the KITTI Vision Benchmark Suite. *CVPR*, 2012.
- [12] L. Caccia, Herke Van Hoof, A. Courville, and J. Pineau. Deep Generative Modeling of LiDAR Data. *IROS*, 2019.
- [13] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social LSTM: Human Trajectory Prediction in Crowded Spaces. *CVPR*, 2016.
- [14] N. Deo and M. M. Trivedi. Convolutional Social Pooling for Vehicle Trajectory Prediction. *CVPRW*, 2018.
- [15] B. Ivanovic and M. Pavone. The Trajectron: Probabilistic Multi-Agent Trajectory Modeling With Dynamic Spatiotemporal Graphs. *ICCV*, 2019.
- [16] M. Liang, B. Yang, W. Zeng, Y. Chen, R. Hu, S. Casas, and R. Urtasun. Pnpnet: End-to-end perception and prediction with tracking in the loop. *CVPR*, 2020.
- [17] X. Y. Jingwei Xu, Bingbing Ni, Zefan Li, Shuo Cheng. Structure Preserving Video Prediction. *CVPR*, 2018.
- [18] M. Oliu, J. Selva, and S. Escalera. Folded Recurrent Neural Networks for Future Video Prediction. *ECCV*, 2018.
- [19] X. Qi, Z. Liu, D. corporation Qifeng Chen, J. Jia CUHK, and Y. Lab. 3D Motion Decomposition for RGBD Future Dynamic Scene Synthesis. *CVPR*, 2019.

- [20] R. Mahjourian, M. Wicke, and A. Angelova. Geometry-Based Next Frame Prediction from Monocular Video. *IV*, 2017.
- [21] H. Fan, H. Su, and L. Guibas. A Point Set Generation Network for 3D Object Reconstruction from a Single Image. *CVPR*, 2017.
- [22] Y. Rubner, C. Tomasi, and L. J. Guibas. The Earth Mover’s Distance as a Metric for Image Retrieval. *IJCV*, 2000.
- [23] C.-H. Lin, C. Kong, and S. Lucey. Learning Efficient Point Cloud Generation for Dense 3D Object Reconstruction. *AAAI*, 2018.
- [24] G. Yang, X. Huang, Z. Hao, M.-Y. Liu, S. Belongie, and B. Hariharan. PointFlow: 3D Point Cloud Generation with Continuous Normalizing Flows. *ICCV*, 2019.
- [25] D. W. Shu, S. W. Park, and J. Kwon. 3D Point Cloud Generative Adversarial Network Based on Tree Structured Graph Convolutions. *ICCV*, 2019.
- [26] R. Klokov, J. Verbeek, and E. Boyer. Probabilistic Reconstruction Networks for 3D Shape Inference from a Single Image. *BMVC*, 2019.
- [27] P. Tomasello, S. Sidhu, A. Shen, M. W. Moskewicz, N. Redmon, G. Joshi, R. Phadte, P. Jain, and F. N. Iandola. DSCnet: Replicating LiDAR Point Clouds with Deep Sensor Cloning. *CVPRW*, 2019.
- [28] W. Luo, B. Yang, and R. Urtasun. Fast and Furious: Real Time End-to-End 3D Detection, Tracking and Motion Forecasting with a Single Convolutional Net. *CVPR*, 2018.
- [29] X. Weng and K. Kitani. Monocular 3D Object Detection with Pseudo-LiDAR Point Cloud. *ICCVW*, 2019.
- [30] X. Weng, J. Wang, D. Held, and K. Kitani. 3D Multi-Object Tracking: A Baseline and New Evaluation Metrics. *IROS*, 2020.
- [31] X. Weng, Y. Wang, Y. Man, and K. Kitani. GNN3DMOT: Graph Neural Network for 3D Multi-Object Tracking with 2D-3D Multi-Feature Learning. *CVPR*, 2020.
- [32] K. Bernardin and R. Stiefelhagen. Evaluating Multiple Object Tracking Performance: The CLEAR MOT Metrics. *Journal on Image and Video Processing*, 2008.
- [33] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. *CVPR*, 2017.
- [34] M. Liang, B. Yang, S. Wang, and R. Urtasun. Deep Continuous Fusion for Multi-Sensor 3D Object Detection. *ECCV*, 2018.
- [35] S. Shi, X. Wang, and H. Li. PointRCNN: 3D Object Proposal Generation and Detection from Point Cloud. *CVPR*, 2019.
- [36] H. W. Kuhn. The Hungarian Method for the Assignment Problem. *Naval Research Logistics Quarterly*, 1955.
- [37] T. Zhou, M. Brown, G. Noah, S. Google, and D. G. Lowe Google. Unsupervised Learning of Depth and Ego-Motion from Video. *CVPR*, 2017.
- [38] P. Besl and N. McKay. A Method for Registration of 3-D Shapes. *TPAMI*, 1992.
- [39] Y. Wang and J. M. Solomon. Deep Closest Point: Learning Representations for Point Cloud Registration. *ICCV*, 2019.
- [40] X. Liu, C. R. Qi, and L. J. Guibas. FlowNet3D: Learning Scene Flow in 3D Point Clouds. *CVPR*, 2019.
- [41] A. Byravan and D. Fox. SE3-Nets: Learning Rigid Body Motion Using Deep Neural Networks. *ICRA*, 2017.
- [42] R. Chandra, U. Bhattacharya, A. Bera, and D. Manocha. TraPHic: Trajectory Prediction in Dense and Heterogeneous Traffic Using Weighted Interactions. *CVPR*, 2019.