

INFER: INtermediate representations for FuturE pRediction

by

Shashank Srikanth, Junaid Ahmed Ansari, Karnik Ram, Sarthak Sharma, Madhava Krishna

in

*IEEE/RSJ International Conference on Intelligent Robots and Systems
(IROS 2019)*

The Venetian Macao, Macau, China

Report No: IIIT/TR/2019/-1



Centre for Robotics
International Institute of Information Technology
Hyderabad - 500 032, INDIA
November 2019

INFER: INtermediate representations for Future pRediction

Shashank Srikanth¹, Junaid Ahmed Ansari¹, R. Karnik Ram¹, Sarthak Sharma¹,
J. Krishna Murthy², and K. Madhava Krishna¹



Fig. 1: **Trajectory forecasting:** *Left:* (top) An interesting scenario from the KITTI dataset [1] where the car (cyan bounding box) will turn left over the next 4 seconds. Our approach—INFER—predicts this future trajectory, conditioned on the previous states of the car (the trail behind the car). This predicted future trajectory (orange) and the ground truth trajectory (blue) are projected on to the image. On the top-right and bottom-right, we show that our approach transfers—without additional learning/finetuning—to other datasets, such as Cityscapes [2] and Oxford Robot Car [3]. Furthermore, while our model is trained on the KITTI [1] dataset, where we use both stereo imagery and lidar scans, when testing on Cityscapes [2], we use only stereo imagery, and when testing on Oxford Robot Car [3], we use only lidar scans. To the best of our knowledge, INFER is the first such approach to demonstrate **cross-sensor** and **varying driving scenario** transferability. In each image: The green 3D bounding box depicts the first sighting of the vehicle of interest (this is when we bootstrap our model); the red 3D bounding box indicates the instant from which our model begins predicting future states. We also register the Lidar/depth information (cyan-road, dark gray-lane, and magenta-road) to demonstrate the accuracy of our prediction.

Abstract—In urban driving scenarios, forecasting future trajectories of surrounding vehicles is of paramount importance. While several approaches for the problem have been proposed, the best-performing ones tend to require extremely detailed input representations (e.g. image sequences). As a result, such methods do not *generalize* to datasets they have not been trained on. In this paper, we propose *intermediate representations* that are particularly well-suited for future prediction. As opposed to using texture (color) information from images, we condition on *semantics* and train an autoregressive model to accurately predict future trajectories of traffic participants (vehicles) (see Fig. above). We demonstrate that semantics provide a significant boost over techniques that operate over raw pixel intensities/disparities. Uncharacteristic of state-of-the-art approaches, our representations and models generalize across different sensing modalities (stereo imagery, LiDAR, a combination of both), and also across completely different datasets, collected across several cities, and also across countries where people drive on opposite sides of the road (left-handed vs right-handed driving). Additionally, we demonstrate an application of our approach in multi-object tracking (data association). To foster further research in transferable representations and ensure reproducibility, we release all our code and data.³

¹Robotics Research Center, KCIS, IIT Hyderabad, India {shashank.s,mkrishna}@research.iit.ac.in

²Mila - Quebec AI Institute, Montreal, Canada.

³More qualitative and quantitative results, along with code and data can be found at <https://rebrand.ly/INFER-results>

I. INTRODUCTION

Deep learning methods have ushered in a new era for computer vision and robotics. With very accurate methods for object detection [4] and semantic segmentation [5], we are now at a juncture where we can envisage the application of these techniques to perform higher-order understanding. One such application which we consider in this work, is forecasting trajectories of traffic participants in urban driving scenarios. Specifically, we argue that constructing *intermediate representations* of the world using off-the-shelf computer vision models for semantic segmentation and object detection, we can train models that account for the multi-modality of future states, and at the same time transfer well across different train and test distributions (datasets) and sensing modalities.

Our approach, dubbed *INFER* (*INtermediate representations For trajEctory foRecasting*), involves training an autoregressive model that takes in an *intermediate representation* of past states of the world, and predicts a multimodal distribution over plausible future states. The key insight is that state-of-the-art approaches [6], [7] heavily rely on images to infer scene context, and as a result suffer from high dataset bias. We mitigate this bias, while preserving scene

context information, by constructing *intermediate* representations of a dynamic scene. Using semantic occupancy grids as our choice of intermediate representations, we demonstrate superior performance to current art [6]–[9]. Furthermore, we believe we are the first to demonstrate the generalization of trajectory forecasting techniques to substantially different datasets (zero-shot). Our approach generalizes across scene layouts, weather conditions, and sensing modalities. We carry out a thorough ablation study on our choice of intermediate representations to capture the influence of different semantic concepts on the accuracy of trajectory forecasting. We also present a use case of trajectory forecasting: multi object tracking and evaluate our performance on select sequences from KITTI [1] and Cityscapes [2].

In summary, we make the following contributions:

- 1) We propose *INFER*, an autoregressive model that uses semantics as an intermediate representation to forecast future trajectories of dynamic traffic participants (vehicles). We significantly improve over current art, and over long horizons.
- 2) Uncharacteristic of prior art, *INFER* generalizes (without additional finetuning) to datasets it has never been trained on. We show results of zero-shot transfer on the Cityscapes [2] and Oxford Robot Car [3] datasets, using a model trained on sequences from KITTI [1].
- 3) We carry out principled ablation studies to gather empirical evidence to answer the question “*What kind of semantics generalize across datasets?*”.

II. RELATED WORK

Of late, several approaches have been proposed to tackle the problem of future prediction in dynamic scenes. Here, we summarize a few of them, while drawing parallels and contrast to our own.

Classical Methods: The problem has been studied extensively in the classical probabilistic modelling paradigm [10], [11]. However, these approaches typically make strong assumptions about the scene and/or require explicit hand-modelling. In contrast, the recent more powerful learning models seem to show better promise in learning the vehicle models and in accurately forecasting their trajectories [6].

IRL for path prediction: Another set of approaches involve using Inverse Reinforcement Learning to estimate the action taken by an agent at each time step and predict the future paths subsequently by applying the estimated actions sequentially at the current target location. Activity forecasting is done in [12] by combining semantic maps with ideas from optimal control theory. In contrast to our approach, the authors use a stationary surveillance camera to forecast activities of pedestrians, where as we predict future locations of vehicles in highly dynamic scenes.

RNNs for future prediction: RNNs have been used for several sequence modelling tasks and can be used for generating sequential future prediction outputs. DESIRE [6] proposes a stochastic recurrent encoder-decoder network for predicting the future trajectories of agents in dynamic environments. The overall model comprises three components, which includes a conditional VAE followed by an RNN

encoder-decoder to generate a set of diverse plausible future states, followed by an inverse optimal control-based ranking and refinement module to rank the predictions and enforce global consistency. Despite using semantic information, the authors however do not claim transfer across datasets and have different models for the Stanford Drone Dataset [13] and the KITTI dataset [1]. In contrast, we are able to show zero-shot transfer of our approach on different datasets like [2], [3] when trained on [1].

[14] uses an interaction layer and a Kalman filter that is embedded in the architecture to learn high variance sensor inputs. They evaluate on the NGSIM dataset which consists of highway scenes. In [15], the authors leverage scene semantics and the past motion trajectory to predict future trajectories. They evaluate their approach on datasets recorded from a stationary camera, and show transfer across unseen scenes from the datasets, rather than cross dataset transfer as we do. The approach of [16] uses geometric and motion properties of the vehicle in the form of yaw, velocity, acceleration, and bounding box. This approach however does not leverage any scene semantics and does not show any results of transfer across different datasets. All the above approaches rely purely on the autoregressive nature of LSTMs to predict hypotheses for future trajectories of participants.

Generative models: [7], [17] predict pedestrian trajectory and exploit generative adversarial networks (GANs) to regularize the output future trajectories to be more *realistic*. In [17], a novel pooling mechanism was introduced for aggregating information across people, and socially plausible future states are predicted by training adversarially against a recurrent discriminator.

Manoeuvre-based approaches: Another set of approaches like [9], [18] use manoeuvre-based LSTMs for social interaction-aware trajectory prediction. In [9], a convolutional social pooling layer is proposed for robust learning of interdependencies in vehicle motion. The authors define 6 classes of maneuvers (eg. brake, accelerate, etc.) and assign a probability to each maneuver class, to obtain multi-modal outputs. Convolutional social pooling aids the learning process by creating a *social tensor* which is then fed to the decoder to infer a distribution over plausible maneuvers. They showcase results on [19], [20] by forming train and test splits comprising of sequences from both the datasets, rather than showcasing transfer across one another.

Occupancy grid approaches: An approach similar to ours is [21], in which the the future locations of the vehicle is predicted on an occupancy grid map, with results shown on a highway driving dataset. In [8], the authors propose a seq2seq model that takes as input at time t the ego vehicle’s velocity and yaw angle, as well as the surrounding vehicles’ positions and velocities & generates K locally best trajectory candidates over an occupancy grid. The approach is devoid of scene semantics and no transfer across dataset is shown.

Our approach, which predicts future location of the VoI deep into the future (upto 4 sec) in highly dynamic environments leverages semantics, depth information and orthographic mapping to represent the raw sensor data in the form of a novel representation that not only reifies cross-dataset

transfer : from [1] to [2], [3] but is also able to generalize well to cross-sensor modalities (Eg. from LiDAR to stereo, see Sec. V).

III. INTERMEDIATE REPRESENTATIONS

In this section, we describe the intermediate representations used by INFER to describe an urban driving scenario. Our design philosophy is based on the following three desired characteristics that knowledge representation systems must possess:

- 1) **Representational adequacy:** Such representations must have the capability to *adequately* represent task-relevant information.
- 2) **Inferential adequacy:** They must also have the capability to infer traits that cannot otherwise be inferred from the original unprocessed data.
- 3) **Generalizability:** These representations must necessarily generalize to other data distributions (for the same task).

Scene representation for the task of future prediction spans a broad spectrum. On one end is solving the future prediction problem based on raw RGB input and predicting the future locations of the VoI in the image space in the form of heatmaps. On the other end is leveraging the geometric information of the VoI in the form of 3D coordinates w.r.t ego vehicle, rotational parameters, relative velocity etc and then regressing to locations of the VoI vehicle deep into the future. The former representation operates on raw RGB data without reasoning about the scene geometry in any form and hence the predicted location of the VoI in the 2D image space would again need to be interpreted in 3D. Approaches like [22] that reason about depth from single view suffer from dataset dependencies, struggling to show transfer on KITTI [1] after being trained on Cityscapes [2]. The dependence of the network to operate on the RGB input enfeebles the transfer on datasets with different RGB pixel values and sensor modalities, as seen in DESIRE [6] where different networks were used for different datasets.

The other end of the scene representation spectrum deals with representing the VoI in form of its geometric properties viz. 3D location w.r.t ego vehicle, rotational parameters, velocity. Although this form of representation captures the depth and geometric properties, it is not able to reason about the scene layout in any sense. Using such kind of representation makes it infeasible to reason about scene semantics like road, lanes, other vehicles or obstacles present in the scene, forcing the system to reason purely based on the relative geometry of VoI.

To this end, we choose a representation that takes the best of both worlds. The proposed representation does not rely heavily on the camera viewing angle, as camera mounting parameters (height, viewing angle, etc.) vary across datasets, and we want our approach to be robust to such variations. We hence adopt a birds-eye view as a canonical reference frame that we transform sensor data to. Further, this helps get rid of undesirable perspective distortion effects. The proposed approach dispenses the dependency on raw pixel intensities and camera intrinsic matrix by extracting semantics of the

scenes. This also brings into play the crucial role of semantics that the predictions must take into account as shown in the ablation study in Table IV. We encompass scene geometry by using depth sensors to project the scene to an orthographic view. Hence, the intermediate representation reasons about the world, capturing scene layout as well as scene geometry, generalizing well not only across different datasets viz. KITTI [1], Cityscapes [2] and Oxford Robot Car [3] but showcasing cross-modality transfer from Lidar in KITTI [1] (and stereo depth) & Oxford Robot Car [3] to stereo depth in Cityscapes [2].

The scene is represented by a five-channel occupancy grid in the birds-eye view. Each of these channels contains complementary semantic cues from the scene, namely *obstacles*, *road*, *lane markings*, *target vehicle* and *other vehicles* - the intermediate scene concepts. All the five channels are in metric units and are generated from stereo image pairs and discriminative learning methods. The camera is at the center-left of the occupancy grid channels, and faces the right end of the grid with coordinate system being the canonical camera coordinate system (i.e. the X-axis and Z-axis points towards top and right of the grid, respectively). Each generated grid is of the size $512 \times 512 \times 1$, where each pixel has a resolution of 0.25m.

As shown in Fig. 3, we first perform semantic and instance segmentation of the left camera image (considering the left camera to be the coordinate frame of the grid) [4], [24], and generate the disparity maps from the stereo image pairs using PSMNet [25].

The disparity images, segmentation masks, and the camera parameters are used to generate point clouds for static scene categories (viz. *lane markings*, *road* and *obstacles*). These semantically classified point clouds are then projected top-down into their respective discrete and fixed size occupancy grids. The 3D points which do not project within the bounds of the occupancy grids are truncated.

Occupancy grids for vehicles are generated in a similar fashion, except that we use instance segmentation to distinguish between different vehicles. We also track these vehicles across time using an appearance and geometry based multi-object tracker [23]. For each vehicle, we have one channel in the occupancy grid representing its own position - the *target vehicle* channel. The *other vehicles* channel represents the positions of the remaining traffic participants.

The proposed representation is a medley of the two ends of the scene representation spectrum for the task of future prediction as it captures the critical task specific information like scene layout, depth information, and semantic information. Hence, it is *representationally adequate*. The inclusion of depth, orthographic mapping and semantics adorns our representation with traits that cannot be inferred directly from the raw unprocessed data in the form of only RGB images or depth sensor. Having a fine-grained semantic map (different channels for building, kerb etc.) does not add much advanced semantic information than what is already captured by our *obstacle channel* (which represents buildings, kerb, vegetation together), hence making the representation *inferentially adequate*. This representation can be inferred from any raw,

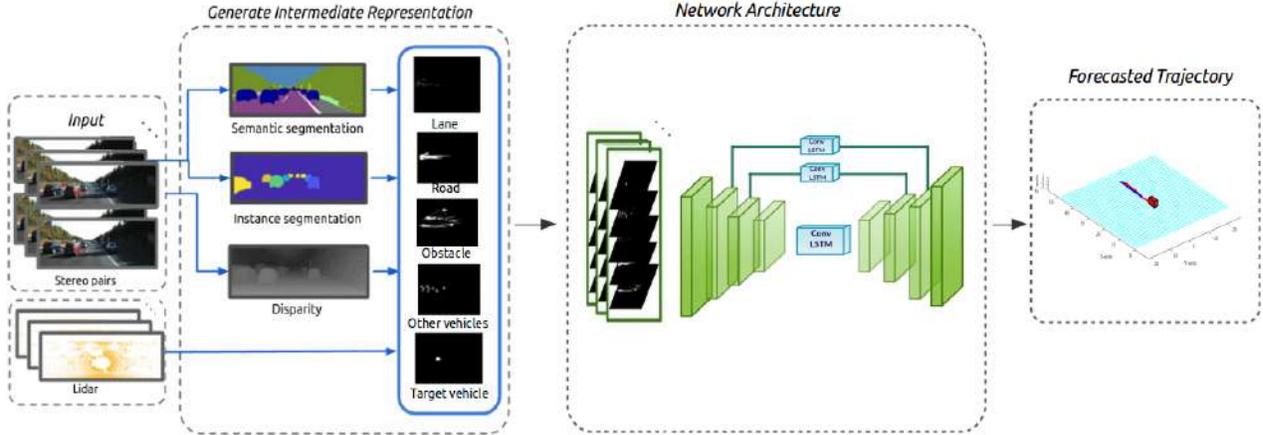


Fig. 3: **Proposed framework:** We first *generate intermediate representations* by fusing monocular images with depth information (from either stereo or Lidar), obtaining semantic and instance segmentation from monocular image, followed by an orthographic projection to bird’s-eye view. We use [4] to detect and [23] to track objects. The generated intermediate representations are fed through the network, and finally we show the resulting prediction of the target vehicle’s trajectory registered in the sensor coordinate frame.

unprocessed data and as we show in Section V, transfers well to a variety of datasets for the task of future prediction, hence is *generalizable*.

Leveraging prior semantic maps: One aspect of autonomous driving scenarios we wish to leverage is that using prior maps, it is possible to obtain a coarse estimate of the *road* and *lane* semantic channels. Specifically, we adopt the strategy in [26] that uses OpenStreetMap (OSM) [27] in conjunction with GPS, and align the OSM with the current *road channel*. This gives us access to estimates of the *road* and *lane* channels for future frames, which we demonstrate in Section V to boost performance significantly. At test time, however, note that we have no *a priori* information about other vehicles on the road, as those are the very attributes we wish to predict.

IV. INFER: INTERMEDIATE REPRESENTATIONS FOR TRAJECTORY FORECASTING

We formulate the trajectory forecasting problem as a per-cell regression over an occupancy grid. We use the intermediate representations introduced in the previous section to simplify the objective and help the network generalize better. We now detail the network architecture and the training/testing procedure.

A. Problem Formulation

Assume that we are given a sequence of intermediate representations $\tau = \{\mathcal{I}_t\}_{t=1}^M$ for a particular **vehicle of interest (VoI)** \mathcal{V} , where \mathcal{I}_t denotes the intermediate representation for VoI \mathcal{V} at time t . The objective of future state prediction is then to predict a (multi-hypothesis) distribution $\{\mathcal{F}_t\}_{t=M+1}^N$, where $N > M$ ($M, N \in \mathbb{Z}$). Each \mathcal{F}_t is a distribution over a regular grid (x_i, y_i) that represents the likelihood of VoI \mathcal{V} being at (x_i, y_i) , conditioned on τ .

B. Network Architecture

We train an autoregressive model that outputs the VoI’s position on an occupancy grid, conditioned on the previous intermediate representations. We use a simple Encoder-Decoder model connected by a convolutional LSTM to learn temporal dynamics. The proposed trajectory prediction scheme takes as input a sequence of intermediate representations and produces a single channel output occupancy grid.

The input grids to the network are resized to spatial dimension of 256×256 from 512×512 to reduce the network size and training overhead. First, through a series of convolution, pooling and non-linearity operations, the Encoder reduces the input resolution from $5 \times 256 \times 256$ to $64 \times 32 \times 32$. This reduced tensor is passed through the Convolutional LSTM sandwiched between the encoder-decoder. The LSTM consists of 64 convolution filters each with a separate hidden and cell state with kernel dimension 3×3 . The output of the LSTM is upsampled via the decoder to a resolution of $8 \times 256 \times 256$ and then this tensor is convolved with a 1×1 filter to provide the future location of the vehicle, which is a likelihood map of dimension $1 \times 256 \times 256$.

During downsampling of input via the encoder, spatial information is lost in the pooling step. Also, as we are dealing with future prediction of possible location of a vehicle after upsampling, we need to retain the temporal information from each downsampled step in encoder. In order to capture both the spatial as well as temporal information, we add *skip connections* between corresponding encoder and decoder branches. We also experiment without convolutional LSTMs over the skip connections.

C. Training objective

Our training objective comprises two terms: a reconstruction loss term, and a safety loss term. More formally,

the reconstruction loss \mathcal{L}_{rec} penalizes the deviation of the predicted future distribution $\hat{\mathcal{F}}$ from the actual future state \mathcal{F}^* . Mathematically,

$$\mathcal{L}_{rec} = \|\hat{\mathcal{F}} - \mathcal{F}^*\|^2 \quad (1)$$

We also add a *safety* loss term, that penalizes all predicted states of vehicles that lie in an obstacle cell.

$$\mathcal{L}_{safe} = \|\mathcal{O} \odot \hat{\mathcal{F}}\| \quad (2)$$

Here, \mathcal{O} denotes the *obstacle* channel, and \odot denotes the elementwise matrix product.

D. Training Phase

To train the model, we feed intermediate representations corresponding to the first 2 second of a sequence to provide sufficient context to the convolutional LSTM. Thereafter, we operate akin to sequence-to-sequence models, i.e., we obtain an output from the network, construct an intermediate representation using this output and the next incoming frame, and feed this into the network as a subsequent input. We train the model by truncated backpropagation through time, once 20 frames are predicted.

E. Test Phase

During the test phase, intermediate representations from the first 2 sec is used to initialize the LSTM and the remaining frames are predicted one step at a time. The final output predicted by the network is upsampled to a resolution of 512×512 using bilinear interpolation and the point with the highest activation in the heatmap is chosen as the predicted location of the target vehicle. Each grid cell of a grid size of 512×512 corresponds to 25×25 cm² area which is suitable resolution in real world for autonomous driving scenarios.

V. RESULTS

A. Dataset

Most approaches to future trajectory prediction [6], [28] demonstrate results over the KITTI [1] autonomous driving benchmark. ([28] deals only with pedestrians). But, the KITTI benchmark alone does not address many challenges that trajectory forecasting algorithms face in real-world operation. In us humans, one would expect to learn a prediction policy in a particular city (or a toy driving environment/park) and expect it to generalize to newer scenarios, even across cities.

Hence, we expand our test dataset to comprise sequences from the KITTI [1], Cityscapes [2], and Oxford RobotCar [3] datasets, for they span a number of cities, provide for enough weather variations, and also exhibit a switch from right-side (Germany) to left-side (the United Kingdom) driving.

The proposed approach is trained on KITTI dataset. We take 21 sequences from the train split of the KITTI Tracking benchmark and 5 sequences from the KITTI Raw dataset. We use ONLY these 26 sequences for training and validating our models. We extract a total of 223 trajectories from these sequences, and then divide them into train and test : 178 trajectories for train and 45 trajectories for test, comprising

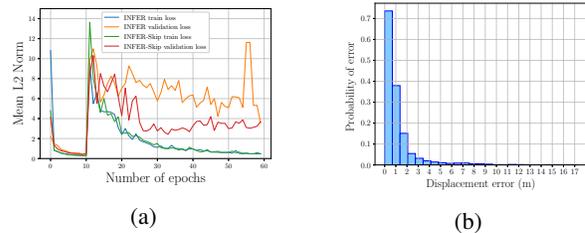


Fig. 4: (a) The average pixel error vs number of epochs for a given split. For the first 10 epochs, we predict one frame into the future. Subsequently, deep future prediction is done and hence there’s a sudden rise in loss. (b) Histogram of per frame $L2$ -norm for all frames in KITTI, for the *INFER-Skip* model. 86% of all our predictions lie within a 2m threshold.

of over 11K frames. We perform a 5 fold cross validation on these 26 sequences. The length of these sequences vary, from a minimum of 3 sec (30 frames) to a max of 6 sec (60 frames, the frame rate for KITTI is 10 fps).

To highlight the core idea of learning transferrable representations for future prediction, we test the best performing model from the above split (i.e., the KITTI dataset) on sequences from Cityscapes [2] and Oxford RobotCar [3].

Specifically in Cityscapes [2], we choose 26 trajectories spread over 15 different cities. The Cityscapes dataset [2] provides sequences of length 30 frames, recorded at 17 fps, resulting in trajectories of length upto 1.76 seconds, comprising of approx 800 frames.

The Oxford RobotCar [3] consists of several sequences recorded over different routes through Oxford, UK. The dataset was recorded at 16 fps and we choose a few trajectories with a duration of 4 sec or more, comprising of approx 500 frames. For both, [3] & [2] we pass every alternate frame while testing.

We evaluate all our models using the average displacement error (ADE) metric which is defined as the average $L2$ distance between the ground truth and predicted trajectories, over all vehicles and all time steps. We compare our methods with the following baselines:

- *Markov-Baseline*: A simple discrete Bayes filter implementation over a grid that uses a constant velocity motion model [29].
- *RNN Encoder Decoder* [16]: A RNN encoder-decoder model that regresses to future locations based on the past trajectories and vehicle information in the form of yaw, velocity, acceleration, bounding box coordinates etc.
- *DESIRE-S* [6]: The best-performing variant from [6], that uses a scene context fusion (SCF) module.
- *ConvLSTM-Baseline*: Our ConvLSTM architecture, taking a 4 channel input, first 3 being the RGB scene in bird’s eye view and the 4th channel being the target vehicle channel. That is, a variant of our model that operates on pixel intensities, as opposed to semantics.
- *INFER*: Our proposed model but with a single convolutional LSTM.
- *INFER-Skip*: The same model as *INFER* but with 2 additional convolutional LSTMs which serve as skip connections.

B. Training Details

All the models were implemented in PyTorch and trained on a single NVIDIA GeForce GTX 1080 GPU for a maximum of 60 epochs each. The training process takes about 6-8 hours. The models were trained using the ADAM optimizer with a learning rate of 0.0001 and gradient clipping with L2 norm of 10.0. The loss trends across train & validation sets for *INFER-Skip* & *ConvLSTM Baseline* is shown in Fig. 4a. The distribution of L2 norm of all frames across the KITTI dataset is shown in Fig. 4b.

Method	1 sec	2 sec	3 sec	4 sec
Markov-Baseline	0.70	1.41	2.12	2.99
RNN Encode-Decoder ¹ [16]	0.68	1.94	3.2	4.46
ConvLSTM-Baseline (Top 5)	0.76	1.23	1.60	1.96
DESIRE-SI (Best) [6]	0.51	1.44	2.76	4.45
DESIRE-SI (10%) [6]	0.28	0.67	1.22	2.06
INFER (Top 5)	0.61	0.87	1.16	1.53
INFER-Skip (Top 1)	0.75	0.95	1.13	1.42
INFER-Skip (Top 3)	0.63	0.82	1.00	1.30
INFER-Skip (Top 5)	0.56	0.75	0.93	1.22

TABLE I: Quantitative results of baseline models vs. *INFER-Skip* across KITTI dataset for the task of predicting upto 4 sec into the future. Error metrics reported are ADE in metres. We have a single model that predicts 1, 2, 3 and 4 seconds into the future.

C. Performance evaluation and transfer results

The results of the *Markov baseline* are shown in Fig. 6a. It performs significantly better than the *RNN Encoder Decoder* [16] due to the use of grid based representation for the target vehicle. Note that both these approaches output a single prediction. The *ConvLSTM* baseline outperforms these two while predicting deep into the future, while degrading slightly when predicting upto 1 sec into the future.

Our model *INFER-Skip* which consists of convolutional LSTMs along skip connections outperforms current state of the art DESIRE [6] in KITTI [1] while predicting deep into the future (3 sec, 4 sec). It outperforms all other models on all the evaluation metrics. *INFER* which consists of only a single convolutional LSTM also beats the current state of the art method [6] for predictions deep into the future (3 sec & 4 sec). These models do not perform as well for shorter timesteps in the future as we are limited by the resolution of the intermediate representation, hence making the comparison favourable for [6]. Being multi-modal in nature, we can predict multiple future trajectories for a given track history by sampling the top K samples at each time step. We report the performance of *INFER-Skip* with several values of $K = 1, 3$ & 5 respectively. Using the top 5 predictions significantly improves performance by upto 20 cms, demonstrating the multi-modal nature of the model.

Transfers: We test the transfer ability of our model by training on KITTI [1] & testing on Cityscapes & Oxford Robotcar dataset [2], [3]. We precondition our model for a total of 0.8 sec and test upto 1 sec into the future for Cityscapes dataset [2] as it provides only sequences of

¹ [16] predicts only upto 2 sec into the future, hence we extrapolate the values to get predictions for 3 & 4 seconds into the future.

Method	1 sec (*)	1 sec (**)
ConvLSTM-Baseline (Top 1)	1.5	1.23
ConvLSTM-Baseline (Top 3)	1.36	1.09
ConvLSTM-Baseline (Top 5)	1.28	1.021
INFER-skip (Top 1)	1.11	1.12
INFER-skip (Top 3)	0.99	0.98
INFER-skip (Top 5)	0.91	0.91

TABLE II: Transfer results of *INFER-Skip* & the *ConvLSTM Baseline* models on Cityscapes [2]. We report the ADE in metres.

Method	1 sec	2 sec	3 sec	4 sec
INFER-Skip	0.85	1.14	1.29	1.50

TABLE III: Transfer results of *INFER-Skip* model tested on a few sequences of Oxford Robotcar. We report the ADE in metres.

Method	KITTI				Cityscapes
	1 sec	2 sec	3 sec	4 sec	1 sec
INFER-Skip	0.56	0.75	0.93	1.22	1.12
INFER-Skip w/o road	0.70	1.20	1.80	2.49	5.62
INFER-Skip w/o obstacles	0.54	0.80	1.00	1.24	1.21
INFER-Skip w/o lane	0.57	0.76	0.94	1.21	1.33

TABLE IV: Ablation results across KITTI & Cityscapes by removing semantics corresponding to road, lane & obstacles from the intermediate representation. We report the ADE in metres.

length 1.76s. We present the performance of the *INFER-Skip* across Cityscapes in Table II. In (*), the results of *INFER-Skip* without vehicle channel are shown. This model transfers well to the Cityscapes dataset which differs a lot from the KITTI dataset in terms of scene layout, weather condition & vehicle motion. The performance across KITTI & Cityscapes is off by only 34 cm for 1 sec into the future. We find *ConvLSTM-Baseline* transfers well to the Cityscapes [2] dataset which suggests that using a birds eye view helps in transfer, as such a representation does not rely heavily on the camera viewing angle and camera mounting parameters. However, *INFER-Skip* performs better than the *ConvLSTM Baseline* highlighting the potent of intermediate representations over RGB data. All our models were trained with a preconditioning of 2 sec and then predict upto 4 sec into the future. Thus, predicting the future trajectories in Cityscapes with 1 sec preconditioning leads to higher than expected error. In (**), we train our main model *INFER-Skip* & the *ConvLSTM Baseline* on another split of the KITTI dataset that consists of trajectories of 2 sec length only. These models were trained with only 1 sec preconditioning and the results of their transfer on Cityscapes dataset is shown in the 2nd column of Table II.

We further test transfer of *INFER-Skip* on a few trajectories of the Oxford Robot Car [3] dataset and the performance does not degrade while predicting deep into the future as can be seen in the table III. Our method is able to generalize well to drastic change in scene layout (Germany to UK) and right-side to left-side driving.

D. Ablation study

We conduct detailed ablation studies on *INFER-Skip* to determine the type of semantics that help the model perform well across KITTI [1] and Cityscapes [2]. For [2], we directly do ablation on the transfer rather than training on it first. We experiment with 3 major variants of *INFER-Skip* with the road, obstacles & lane channels removed. The results in

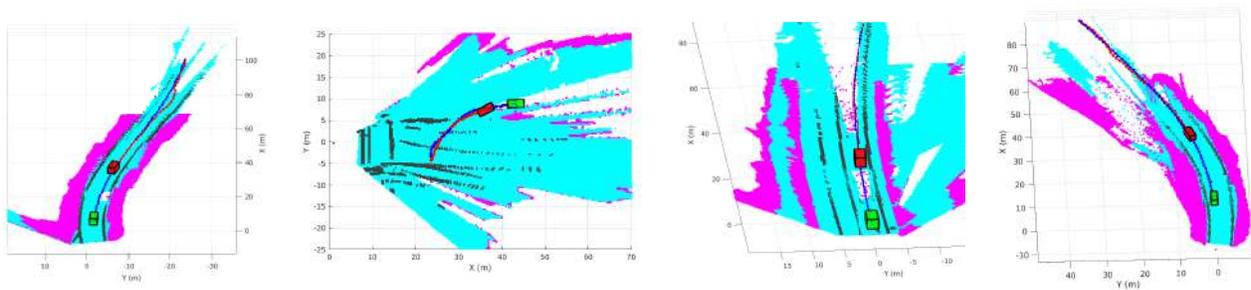


Fig. 5: The above qualitative results from the validation fold of KITTI [1] showcase the efficacy of *INFER-Skip* in using the intermediate representation to predict complex trajectories. For example, in the second plot from left, the network is able to accurately predict the turn taken by the target vehicle (vehicle of interest) at an intersection (predicted and ground truth trajectories are shown in red and blue color, respectively). The green and red 3D bounding boxes indicate start of preconditioning and start of prediction of the vehicle of interest, respectively. It is worth noting that the predicted trajectories are well within the lane (dark gray) and road region (cyan), while avoiding collisions with the obstacles (magenta).

Table IV show that the scene semantics do indeed play a vital role in the performance of our proposed model. The removal of the road channel reduces the performance of the model drastically in KITTI. The performance of the model degraded severely when tested on Cityscapes as the error reaches as high as 5.62m. While the other channels like obstacles & lane do not seem to affect the performance of the model in the KITTI dataset greatly, the corresponding transfer error in Cityscapes increases significantly, showcasing how semantics play a critical role in transferring across datasets. Thus, our representations are *representationally adequate* and *generalizes* well to other datasets. We also show ablation on the frame rate in Fig. 6b varying it while testing on KITTI [1]. We show the loss trend for 1, 2, 3, and 4 seconds into the future. The variation in frame rate captures the change in relative velocity of VoI w.r.t ego vehicle. We observe that even when the frame rate is dropped to 60%, the performance does not degrade significantly, highlighting the potent of the representations to generalize to different relative velocities.

E. Qualitative results

We showcase some qualitative results of our approach on challenging KITTI [1], Cityscapes [2], and Oxford Robotcar [3] sequences in figures 5, 7a & 7b respectively. These results illustrate the effectiveness of our *intermediate representations* and models to predict complex trajectories and transfer

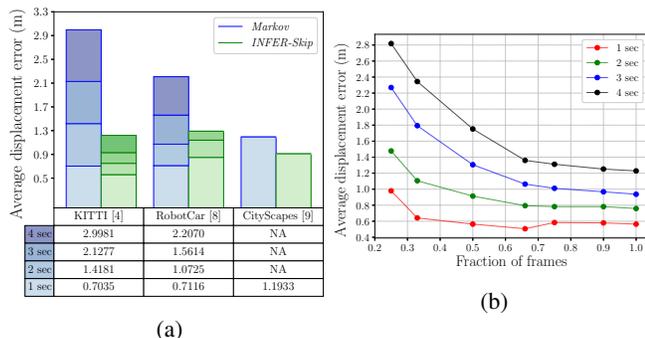


Fig. 6: (a) Shows the performance of the *Markov* baseline across three different datasets, for predictions upto 4 seconds, and is compared with *INFER-Skip*. (b) Shows the performance of *INFER-Skip* for different frame rates.

zero-shot across datasets [2], [3]. The green 3D bounding box depicts the first sighting of the vehicle of interest which is also when we start preconditioning, and the red 3D bounding box indicates the start of prediction. The plots clearly show that *INFER-Skip*, using the proposed *intermediate representations*, is accurately predicting the trajectories. It can also be seen that even non-trivial trajectories constituting of multiple turns are predicted very close to their respective ground truth while being within the road (cyan) and lane (dark gray) regions; obstacles are shown in magenta color. For the purpose of visualization, all the trajectories along with the intermediate representations (road, lane, and obstacle) are registered in the ego vehicle coordinate frame corresponding to the start of prediction.

	KITTI [1]	Cityscapes [2]
Association accuracy	85.71%	75%

TABLE V: Object association for Multi-Object tracking

F. Summary of results

The cornerstone of this effort is that intermediate representations are well apt for the task of future prediction. In Table I we highlight that using semantics provides a significant boost over techniques that operate over raw pixel intensities/disparities. In Table II, III we showcase the efficacy of intermediate representations to transfer across different datasets collected across different cities. We perform an extensive ablation study on which semantics help in transfer and report the results of this ablation on KITTI [1] and Cityscapes [2] in Table IV, highlighting how semantics are critical in transfer across dataset. We do an ablation study on frame rates in Fig. 6b and showcase the generalizability of our approach to the change in relative velocity of VoI and ego vehicle. We highlight a usecase of our approach for object association in multi-object tracking based on future predicted locations in Table V. Association is done based on the minimum L_2 distance the future location of VoI has with all the vehicles present in the scene. We showcase comparable performance on [1], [2] to popular approaches like [30].

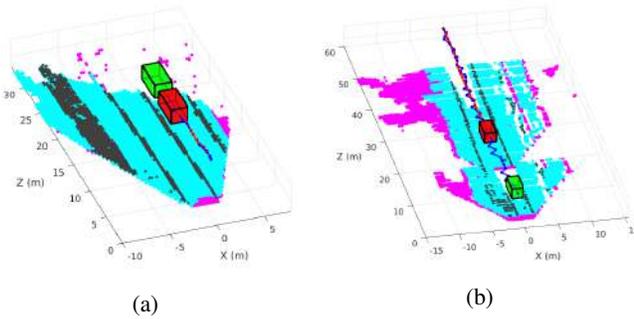


Fig. 7: Above figure qualitatively shows the result of zero-shot transfer of *INFER-Skip*, trained on KITTI dataset [1], to (a) Cityscapes [2] and (b) Oxford robot car [3] datasets. The color conventions are same as that of the plots in Fig. 5.

DISCUSSION

Learning a suitable representation from the dataset makes the whole point of transferability moot as the learned representation would again be dependent on the dataset it was learned from, hence defying the whole idea of transferability. These learned representations would prove out to be insignificant when faced with a new dataset. On the other hand, the proposed representation is agnostic to the dataset and can be incorporated in any future trajectory prediction problem. It is simple and easy to compute rather than relying on complex handcrafted features. Future work can look at trying to learn representations that are data agnostic. Predicting future semantics to aid the prediction of future locations can also be looked at.

VI. CONCLUSIONS

We propose *intermediate representations* that are apt for the task of future trajectory prediction of vehicles. As opposed to using raw sensor data, we condition on *semantics* and train an autoregressive network to accurately predict future trajectories of vehicles. We outperform the current state of the art approaches, demonstrating that semantics provide a significant boost over techniques that operate solely over raw pixel intensities or depth information. We show that our representations and models transfer zero-shot to completely different datasets, collected across different cities, weather conditions, and driving scenarios. We carry out a thorough ablation study on the importance of our semantics and show generalization of our approach on different frame rates. Additionally, we demonstrate an application of our approach in data association in multi-object tracking.

REFERENCES

- [1] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *CVPR*, 2012.
- [2] M. Cordts, M. Omran, *et al.*, "The cityscapes dataset for semantic urban scene understanding," in *CVPR*, 2016.
- [3] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000km: The oxford robotcar dataset," *IJRR*, 2017.
- [4] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, "Mask r-cnn," *ICCV*, 2017.
- [5] Y. Zhuang, F. Yang, *et al.*, "Dense relation network: Learning consistent and context-aware representation for semantic image segmentation," in *ICIP*, 2018.
- [6] N. Lee, W. Choi, *et al.*, "Desire: Distant future prediction in dynamic scenes with interacting agents," in *CVPR*, 2017.

- [7] A. Sadeghian, V. Kosaraju, *et al.*, "Sophie: An attentive gan for predicting paths compliant to social and physical constraints," *arXiv preprint*, 2018.
- [8] S. Park, B. Kim, *et al.*, "Sequence-to-sequence prediction of vehicle trajectory via lstm encoder-decoder architecture," *arXiv preprint*, 2018.
- [9] N. Deo and M. M. Trivedi, "Convolutional social pooling for vehicle trajectory prediction," *arXiv preprint*, 2018.
- [10] A. Barth and U. Franke, "Estimating the driving state of oncoming vehicles from a moving platform using stereo vision," *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 4, pp. 560–571, 2009.
- [11] R. Danescu, F. Oniga, and S. Nedeveschi, "Modeling and tracking the driving environment with a particle-based occupancy grid," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 4, pp. 1331–1342, 2011.
- [12] K. M. Kitani, B. D. Ziebart, J. A. Bagnell, and M. Hebert, "Activity forecasting," in *ECCV*. Springer, 2012.
- [13] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *ECCV*, 2016.
- [14] C. Ju, Z. Wang, and X. Zhang, "Socially aware kalman neural networks for trajectory prediction," *arXiv preprint*, 2018.
- [15] A. Sadeghian, F. Legros, *et al.*, "Car-net: Clairvoyant attentive recurrent network," in *ECCV*, 2018.
- [16] J. Virdi, "Using deep learning to predict obstacle trajectories for collision avoidance in autonomous vehicles," Ph.D. dissertation, UC San Diego, 2017.
- [17] A. Gupta, J. Johnson, *et al.*, "Social gan: Socially acceptable trajectories with generative adversarial networks," *CVPR*, 2018.
- [18] N. Deo and M. M. Trivedi, "Multi-modal trajectory prediction of surrounding vehicles with maneuver based lstms," *IV*, 2018.
- [19] J. Colyar and J. Halkias, "Us highway 101 dataset," *Federal Highway Administration (FHWA), Tech. Rep. FHWA-HRT-07-030*, 2007.
- [20] X.-Y. Lu and A. Skabardonis, "Freeway traffic shockwave analysis: exploring the ngsim trajectory data," in *86th Annual Meeting of the Transportation Research Board, Washington, DC*, 2007.
- [21] B. Kim, C. M. Kang, *et al.*, "Probabilistic vehicle trajectory prediction over occupancy grid map via recurrent neural network," *ITSC*, 2017.
- [22] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," in *CVPR*, 2017.
- [23] E. Bochinski, V. Eiselein, and T. Sikora, "High-speed tracking-by-detection without using image information," in *International Workshop on Traffic and Street Surveillance for Safety and Security at IEEE AVSS*, 2017.
- [24] S. R. Bulò, L. Porzi, and P. Kotschieder, "In-place activated batch-norm for memory-optimized training of dnns," *arXiv preprint*, 2017.
- [25] J.-R. Chang and Y.-S. Chen, "Pyramid stereo matching network," in *CVPR*, 2018.
- [26] S. Schuler, M. Zhai, N. Jacobs, and M. Chandraker, "Learning to look around objects for top-view representations of outdoor scenes," in *ECCV*, 2018.
- [27] OpenStreetMap contributors, "Planet dump retrieved from <https://planet.osm.org>," <https://www.openstreetmap.org>, 2017.
- [28] V. Karasev, A. Ayvaci, B. Heisele, and S. Soatto, "Intent-aware long-term prediction of pedestrian motion," in *ICRA*, 2016.
- [29] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.
- [30] B. Lee, E. Erdenee, *et al.*, "Multi-class multi-object tracking using changing point detection," in *ECCV*. Springer, 2016.