

LOKI: Long Term and Key Intentions for Trajectory Prediction

Harshayu Girase^{1,2,*†} Haiming Gang^{1,*} Srikanth Malla¹ Jiachen Li^{1,2,†}

Akira Kanehara³ Karttikeya Mangalam² Chiho Choi¹

¹Honda Research Institute USA ²University of California, Berkeley ³Honda R&D Co., Ltd.

{harshayugirase, jiachen_li, mangalam}@berkeley.edu akira.kanehara@jp.honda

{hgang, smalla, cchoi}@honda-ri.com

Abstract

Recent advances in trajectory prediction have shown that explicit reasoning about agents' intent is important to accurately forecast their motion. However, the current research activities are not directly applicable to intelligent and safety critical systems. This is mainly because very few public datasets are available, and they only consider pedestrian-specific intents for a short temporal horizon from a restricted egocentric view. To this end, we propose LOKI (LOng term and Key Intentions), a novel large-scale dataset that is designed to tackle joint trajectory and intention prediction for heterogeneous traffic agents (pedestrians and vehicles) in an autonomous driving setting. The LOKI dataset is created to discover several factors that may affect intention, including i) agent's own will, ii) social interactions, iii) environmental constraints, and iv) contextual information. We also propose a model that jointly performs trajectory and intention prediction, showing that recurrently reasoning about intention can assist with trajectory prediction. We show our method outperforms state-of-the-art trajectory prediction methods by upto 27% and also provide a baseline for frame-wise intention estimation. The dataset is available at <https://usa.honda-ri.com/loki>

1. Introduction

Over the past few years, there has been extensive research into predicting future trajectories of dynamic agents in scenes, such as pedestrians and vehicles. This is an incredibly important and challenging task for safety critical applications such as autonomous vehicles or social robot navigation. While these methods have been significantly advanced over recent years, very few benchmarks specifically test if these models can accurately reason about key

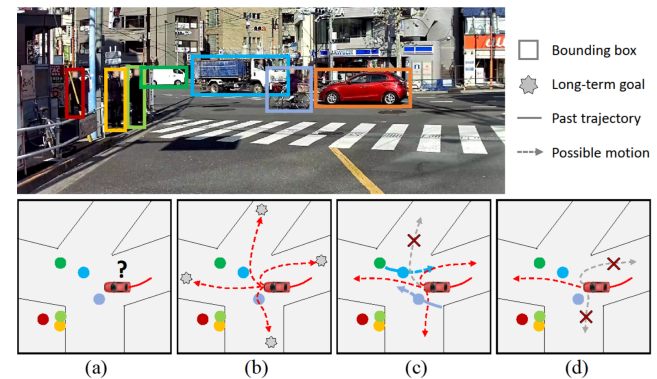


Figure 1: We show that reasoning about long-term goals and short-term intents plays a significant role in trajectory prediction. With a lack of comprehensive benchmarks for this purpose, we introduce a new dataset for intention and trajectory prediction. An example use case is illustrated in (a) where we predict the trajectory of the target vehicle. In (b), long-term goals are estimated from agent's own motion. Interactions in (c) and environmental constraints such as road topology and lane restrictions in (d) influence the agent's short-term intent and thus future trajectories.

maneuvers such as sudden turns and lane changes of vehicles or pedestrians crossing the road. Traditional trajectory error metrics may not capture performance on frame-level maneuvers, which is critical for safe planning.

An intelligent trajectory prediction system should be able to understand and model dynamic human behaviors. The study of human behavior as goal-directed entities has a long and rich interdisciplinary history across the sub-fields of psychology [1], neuroscience [2] and computer vision [3]. The human decision making process is inherently hierarchical, consisting of several levels of reasoning and planning mechanisms that operate in tandem to achieving respective short and long term desires. Recent works have shown that explicitly reasoning about long-term goals [3, 4, 5] and short-term intents [6, 7, 8] can assist with tra-

*Co-first author. The authorship order is randomly determined

†Work done during the internship at Honda Research Institute

jectory prediction.

In this work, we propose to couple the tasks of heterogeneous (vehicles, pedestrians, etc.) multi-agent trajectory forecasting and intention prediction. We believe it is critical to explicitly reason about agents' long-term goals as well as their short-term intents. In our work, we define goals to be a final position an agent wants to reach for a given prediction horizon [9, 5], while intent refers to *how* an agent accomplishes their goal [10]. For example, consider a vehicle at an intersection. At the highest level, say they want to reach their ultimate goal of turning left to their final goal point, which in turn might be necessary for some higher-level end (such as going home). However, the exact motion of their trajectory is subject to many factors including i) agent's own will, ii) social interactions, iii) environmental constraints, iv) contextual cues. Thus, when reasoning about the agent's intent to turn left it is important to consider not only agent dynamics but also how intent is subject to change based on map topology or neighboring agents (see Figure 1). We believe this complex hierarchy of short-term intents and long-term goals is ubiquitous and in fact, crucial, for agent motion planning and hence by extension, for motion prediction. We propose an architecture that considers long-term goals similar to [9, 5, 3, 4] but adds a key component of frame-wise intention estimation which is used to condition the trajectory prediction module. By forcing the model to learn discrete short-term intents of agents, we observe improved performance by the prediction module.

Equally rich & successful is the contemporary history of the use of datasets for benchmarking progress in computer vision. Ushered by seminal works such as MNIST [11] and benchmarks such as ImageNet [12], benchmarking progress and learning from data has played a key role in the success of modern deep learning. Currently, there exists no public datasets that allow for explicit frame-wise intention prediction for heterogeneous agents in highly complex environments. Although few datasets are designed to study pedestrian intents or actions [13, 7, 6, 14] from egocentric view, it is an inherent limitation to extensive study of tasks for autonomous driving. Thus, we propose a joint trajectory and intention prediction dataset that contains RGB images with corresponding LiDAR point clouds with detailed, frame-wise labels for pedestrians *and* vehicles. The LOKI dataset allows explicit modeling of agents' future intent and extensive benchmarking for both tasks. It also shows promising directions to jointly reason about intentions and trajectories while considering different external factors such as agents' predilection, social interactions and environmental factors. We show that by modeling short-term intent and long-term goals with explicit supervision via intention labels, better trajectory prediction accuracy can be achieved. In addition, predicting a specific intention at each frame adds a layer of abstraction to our model that improves understanding

prediction decisions, an important step towards maintaining safety critical applications.

In conclusion, the contribution of our work is twofold. **First**, we propose the first publicly available heterogeneous dataset which contains frame-wise intention annotations and captures trajectories of up to 20 seconds containing both 2D and 3D labels with RGB and LiDAR inputs. **Second**, we illustrate the efficacy of separately reasoning about both long-term goals and short-term intents through ablation studies. Specifically, we highlight how the subtask of intention prediction improves prediction performance, and propose a model that outperforms state-of-the-art multimodal benchmarks by upto 27%. We believe our highly flexible dataset will allow the trajectory prediction community to further explore topics within the intention-based prediction space. In addition, the problem of intention estimation is an involved task in and of itself for which our work provides a strong baseline.

2. Related Work

Over the past few years, there has been a rapid improvement in the field of trajectory prediction owing to the success of deep neural networks and larger publicly available datasets [15, 7, 6, 13, 16, 17, 18, 19, 20]. There have been numerous subtopics of interest within the trajectory prediction community including compliant trajectory prediction, multi-modal trajectory prediction, and goal-oriented prediction [21, 22, 23, 24, 25, 3, 9, 26, 4, 27, 28, 29, 30, 7, 31, 5].

2.1. Contextual Trajectory Prediction

Earlier works in the field of trajectory prediction focused on unimodal trajectory prediction – predicting a single future path for each agent. These works underscored the importance of social [32, 27, 33] and scene compliance [34] when making predictions. Over the past few years, trajectory prediction studies have extended these ideas to multimodal frameworks to account for multiple plausible futures each agent can have. In SocialGAN, Gupta et al. [23] introduce a socially-aware multi-modal framework that uses generative adversarial networks to sample a varying number of future trajectories for each agent. Since then, there has been a major emphasis and many interesting approaches to with multimodal forecasting [22, 35, 9, 23, 36, 4, 30, 37].

2.2. Goal-based Prediction

When modeling vehicle and human trajectories, it is natural to formulate the problem as a goal-directed task. Because humans are not completely stochastic agents and have a predilection towards certain actions, very recent trajectory forecasting studies have shown the effectiveness of goal-conditioned predictions [38, 28, 9, 3, 39, 7, 40, 4, 5, 41]. Recently, [9] and [5] showed that considering agents' final goal points can immensely aid in forecasting trajectories.

However, both of these works only consider positional information as their goal states. In our work, we propose and show the effectiveness of considering *both* long-term positional goals as well as short-term intended actions.

2.3. Intention Datasets

To better understand agent intent in traffic scenes, a few works have proposed datasets that contain intention labels to study underlying intent in addition to the traditional trajectory prediction task. The JAAD [7], PIE [13] and STIP [6] datasets are recent datasets designed to study pedestrian intent. The JAAD dataset focuses on traffic scene analysis and behavior understanding of pedestrian at intersection crossing scenarios. The PIE dataset expands on JAAD further and contains more annotations for both intention estimation and trajectory prediction. PIE [7] only predicts intent at the current timestep and focuses on shorter horizon predictions (1.5 seconds). The STIP dataset solves the limitation of only being able to do single-shot intention prediction, as it contains frame-wise intention labels for up to 3 seconds. However, this dataset only contains "crossing/not crossing" labels for pedestrians and does not focus on trajectory prediction. All these datasets only consider intentions of pedestrians at intersections which may not capture the intents of all agents in a highly complex traffic environment with both vehicles and pedestrians.

IntentNet [8] does consider intents for vehicle trajectory prediction; however, they do not consider frame-wise intentions. Furthermore, the dataset and labels are not publicly available. TITAN [14] is another driving action dataset collected from egocentric view. Although it can be potentially used for intention prediction of traffic agents, it only contains ego-view tracklets and lacks environmental and LiDAR information that can be crucial to find agents' intent. Both these works also only focus on short term predictions (less than 3 seconds). Compared with general driving dataset (such as Waymo [42], Nuscenes [43], and INTERACTION [44]), LOKI extends the standard bounding box, track id, etc. to richer intention, contextual and environmental labels.

To the best of our knowledge, currently no publicly available dataset contains detailed, frame-wise annotations to allow for heterogeneous multi-agent trajectory forecasting *and* intention prediction in joint camera and lidar space. Our dataset contains very diverse traffic scenarios through long data collection periods in different locations, weather conditions, roads and lighting. Table 1 shows the details of our LOKI dataset in comparison to other recently available intention datasets (PIE, JAAD, STIP).

3. LOKI Dataset

Exploring predictions in a large traffic environment is a complex problem because the future behavior of each traf-

	PIE [7]	JAAD [13]	STIP [6]	LOKI (ours)
# of scenarios	-	346	556	644
# of agents	1.8K	2.8K	3.3k	28K
# of labeled agents	1.8K	0.6K	3.3	28K
# of classes	1	1	1	8
# of bboxes	740K	391K	350k	886K
# of agent types	1 (Ped)	1 (Ped)	1 (Ped)	8 classes
Avg. agent per frame	2.5	5.2	3.2	21.6
Annotation freq.	-	-	2 FPS	5 FPS
Frame-wise labels	no	✓	✓	✓
RGB Images	✓	✓	✓	✓
LiDAR Point cloud	no	no	no	✓
2D Bounding box	✓	✓	✓	✓
3D Bounding box	no	no	no	✓
Lane Info	no	no	no	✓
Pedestrian attributes	no	✓	no	✓

Table 1: Comparison of LOKI dataset with PIE [7], JAAD [13] and STIP [6].




Figure 2: Distribution of labels sorted according to the different types of intention among the different classes

fic participant is not only indicated by the past behavior, but also highly impacted by the future goals and intentions. With a lack of comprehensive benchmarks for this purpose, we introduce a large scale dataset that is designed for the task of joint intention and trajectory prediction. Our dataset is collected from central Tokyo, Japan using an instrumented vehicle that is equipped with a camera, LiDAR, GPS and vehicle CAN BUS. The recordings are suburban and urban driving scenarios that contain diverse actions and interactions of heterogeneous agents, captured from different times of the day.

From our recordings, we extracted 644 scenarios with average 12.6 seconds length. The synced LiDAR data and RGB image were down sampled to 5HZ for annotation. The total number of agents is over 28K including 8 classes (*i.e.*, Pedestrian, Car, Bus, Truck, Van, Motorcyclist, Bicyclist, Other) of traffic agents, which results in 21.6 average agents in a scene. We annotated all these agents' bounding boxes (total 886K) in the RGB image (2D) as well as LiDAR point cloud (3D) by linking with a same track-ID. The comparison with existing benchmarks is shown in Table 1. The LOKI dataset is annotated with unique attributes that can influence agents' intent such as interaction related labels, environmental constraints and contextual information.

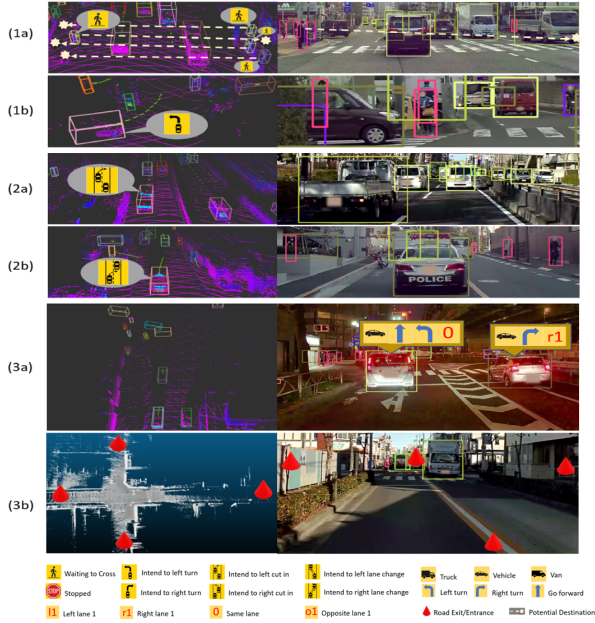


Figure 3: Visualization of three types of labels: (1a-1b) Intention labels for pedestrian; (2a-2b) Intention labels for vehicle; and (3a-3b) Environmental labels. The left part of each image is from laser scan and the right part is from camera. In (1a), the current status of pedestrian is "Waiting to cross", and the potential destination shows the intention of pedestrian. In (3a), the blue arrow indicates the possible action of the current lane where the vehicle is on, and the red words present the lane position related to the ego-vehicle.

3.1. Dataset Annotation

Considering that LiDAR point clouds better capture positional relations among agents than RGB images, we annotate 3D bounding box of agents with their orientation, potential destination of pedestrians, road entrance / exit, and agents' intention as well as action labels in this space. In contrast, in the RGB image space we leverage its contextual clarity to annotate environmental labels such as lane information (what actions can be made from this lane), lane number for vehicles (relative position with respect to the autonomous agent), the gender and age for pedestrian, the state of traffic light, and the type of traffic sign. Note that we also annotate 2D bounding box, potential destination and road entrance / exit information in the RGB space to inspire the potential research in the egocentric view. By using the consistent tracking ID between the same agent in the 3D LiDAR space and 2D image space, our labels can be shared across different spaces.

To dig into more complex prediction researches, our dataset provides denser agents per frame and more meticulous intention attributes compared to other datasets.

We have three types of labels in the LOKI dataset: Intention labels, Environmental labels and Contextual labels

to explore how these can affect the future behavior of agents (details and visuals are in Figure 2 and Figure 3).

Intention labels Intentions are defined to be "how" an actor decides to reach a goal via a series of actions [10]. At each frame, we annotated the current actions of the traffic participants and then used future actions to generate our intention labels. For example, if the current action of vehicle is "Moving" and the future action in 1 second is "Stopped", the vehicle's current intention is to stop. Various intention horizons can be explored; we use 0.8s, as we explore how short-term intent can help guide trajectory prediction.

Environmental labels The environment of driving scene can heavily impact the intention of agent especially for the driving area users, so we include the environmental information such as "Road Exit and Entrance" positions, "Traffic light", "Traffic Sign", "Lane Information" in the LOKI dataset. Those labels determined by the structure of the road and the traffic rules that can be applied to any agent in the scene. The lane information includes the allowed actions of the current lane where the vehicle is on and the relative position between other vehicle and ego-vehicles.

Contextual labels There are some other factors may also affect the future behavior of agent. We define the "Weather", "Road condition", "Gender", "Age" as external contextual labels. These factors are the characters of the agent or environment which can cause the different intentions even under similar environment condition.

4. Proposed Method

4.1. Problem Formulation

In this work, we tackle the problem of multi-agent trajectory forecasting while concurrently predicting agent intentions. The type of intentions vary between agent classes: vehicles and pedestrians. We formulate the problem as follows. Suppose in a given scene, \mathcal{S} , we have N agents, $A_{1:N}$. Given the past $t_{obs} = 3s$ of trajectory history in BEV coordinates, the problem requires forecasting the future $t_{pred} = 5s$ coordinates of the agent in top-down image space. Since our dataset allows for frame-wise intention predictions depending on agent type (pedestrians vs. vehicles), we define another task to predict discrete intentions for each agent at each timestep, in addition to the traditional trajectory prediction problem.

4.2. Model Design

4.2.1 Long-term Goal Proposal Network

Intuitively, agents have a predetermined, long-term goal that they want to reach. Many recent goal-directed works have focused on modeling this through estimating final "end-point" or "goal state" distributions as done in [9, 3, 5, 28, 4]. Inspired by agents' rational decision-making process and

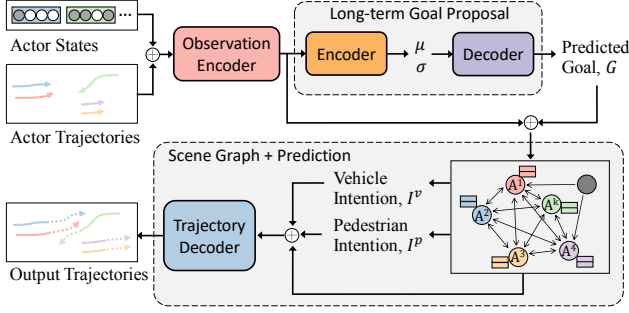


Figure 4: Our model first encodes past observation history of each agent to propose a long-term goal distribution over potential final destinations for each agent independently. A goal, G is then sampled and passed into the Joint Interaction and Prediction module. A scene graph is constructed to allow agents to share trajectory information, intentions, and long-term goals. Black nodes denote road entrance/exit information which provides agents with map topology information. At each timesteps, current scene information is propagated through the graph. We then predict an intent (the action will the agent take in the near future) for each agent. Finally, the trajectory decoder is conditioned on predicted intentions, goals, past motion, and scene before forecasting the next position. This process is recurrently repeated for the horizon length. Note that \oplus denotes concatenation.

the success of prior works, we design a goal network similar to the method proposed in [9]. For each agent, A_k , we use a Conditional Variational Autoencoder (CVAE) to estimate the final long term goal G_k that is simply the estimated position in BEV $u_{k_f} = (x_{k_f}, y_{k_f})$ where f indicates the final frame. The inputs into the CVAE are the encodings from the Observation RNN Encoder. The goal network only consider agents’ own history, as agents have a predetermined long term goal irrespective of other agents.

4.2.2 Scene Graph + Trajectory Decoder

Our main insight and promising directions from our proposed dataset comes from agents’ short-term intentions. As described earlier, we have different intentions for pedestrians and vehicles. Without loss of generality, we will refer to agents, A , and intentions, I , without specifying the type of agent. We believe agents’ have intermediate stochastic intents that can change depending on agent behavior, agent-agent interaction, or environmental factors. To account for this, we construct a traffic scene graph \mathcal{G} to account for social and environmental factors that may affect intent and trajectory prediction.

More concretely, suppose we have a scene graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where vertices, \mathcal{V} , denote agents and road entrances/exits and edges, \mathcal{E} , capture agent-agent and agent-map influence. In a given scene, for neighboring agents

v_i and v_j , there is a directional edge e_{ij} if agent i affects agents j (within a certain distance threshold). Static road entrance/exit nodes can affect agents but without incoming edges. We connect a directional edge e_{ij} if road entrance/exit node i is within a certain distance from agent j .

We then predict agents’ future locations via a daisy chained process described as follows. At each frame, m , our model first shares information between agents via the attention mechanism used in [45]:

$$x_i^{t+1} = \gamma(x_i^t) + \sum_{x_j \in \mathcal{N}(x_i)} \alpha_{ij} * \phi(x_j^t, e_{ij}),$$

where x_i^{t+1} represents the updated node features following attention-based feature aggregation with all of its neighbors $x_j \in \mathcal{N}(x_i)$. We use agents’ velocities and relative positions as edge features. These features are encoded by a 2-layer MLP prior to message passing at each timestep. We use the scaled dot-product attention [45] formulation:

$$a_{ij} = \text{softmax}\left(\frac{\psi(x_i)^T \xi(x_j, e_{ij})}{\sqrt{d}}\right)$$

Here, a_{ij} represents the attention coefficient between two nodes i and j and d represents the degree of the node. We use a single-layer for ϕ , γ , ψ , and ξ .

After message passing which allows agents to share their past trajectory, goal, and intention information along with road information through the road entrance/exit nodes, our model then predicts agent intent, which we define to be the agent’s future action $m + q$ frames ahead. In our experiments, we set $q = 4$, thus predicting short-term intent $0.8s$ in the future. We then condition trajectory prediction for frame $m + 1$ based on agent intent at frame m . This process of information sharing and intention conditioning is recurrently repeated for the next $f - ob$ timesteps where f denotes the last prediction frame number and ob denotes the last observation frame. Formally, at each frame, m , we first estimate the probability distribution over a discrete set of intentions (different set of intentions for pedestrian vs. vehicle) for an agent, A_i :

$$P(I_{i_m} | I_{i_{ob:m-1}}, U_{i_{0:m-1}}, G_i, a_{i_{0:ob}}, \cup_{A_j \in \mathcal{N}(A_i)} I_{j_{ob:m-1}}, U_{j_{0:m-1}}, G_j, a_{j_{0:ob}}, R_{ee})$$

where I refers to intention, U is position, G is long-term positional goal, a is action, and R_{ee} refers to road entrances/exits labels. The intention networks are two-layer MLPs which predicts intention using each actor’s updated hidden states from the most recent message passing. Following this, we then predict the next position of each agent, U , conditioned as follows:

$$P(U_{i_{m+1}} | I_{i_{0:m}}, U_{i_{0:m}}, G_i, a_{i_{0:ob}}, \cup_{A_j \in \mathcal{N}(A_i)} I_{j_{0:m}}, U_{j_{0:m}}, G_j, a_{j_{0:ob}}, R_{ee})$$

The trajectory decoder module consists of a GRU that updates each actor’s current hidden state followed by a 2-layer MLP used to predict positions at each step. The overview of our model is illustrated in Figure 4. Specific model architecture details will be provided in the supplementary material.

4.2.3 Loss Functions

Our goal proposal network (GPN) follows the approach introduced in [9] and is trained via the following loss function:

$$\mathcal{L}_{GPN} = \alpha_1 D_{KL}(\mathcal{N}(\mu, \sigma) \| \mathcal{N}(0, I)) + \alpha_2 \|\hat{G} - G\|_2^2$$

Here α_1 and α_2 are tunable parameters to weight the KL Divergence loss and goal reconstruction loss for training the CVAE. We observed that training via conditioning with ground-truth goal positions helps with model convergence because the intentions are dependent on the long-term goal.

Our decoder module which is responsible for both intention and trajectory prediction is composed of separate loss terms for each. Our intention loss is defined as follows:

$$\mathcal{L}_{int} = - \sum_{j=t_{ob}+1}^{t_f} \sum_{i=0}^n w_{ij} * y_{ij} * \log(\hat{y}_{ij})$$

Due to heavy class imbalance, we not only augment rare trajectories such as lane changes and turning but also weight the cross entropy loss by w_i , which is the inverse frequency of the class.

Since we predict offsets in position (velocity) rather than position directly for better model convergence, our loss is on the predicted velocity V for all timesteps:

$$\mathcal{L}_{traj} = \|V - \hat{V}\|_2$$

We train our network end-to-end by weighting each of the loss terms:

$$\mathcal{L}_{Final} = \lambda_1 \mathcal{L}_{GPN} + \lambda_2 \mathcal{L}_{int} + \lambda_3 \mathcal{L}_{traj}$$

4.2.4 Evaluation Metrics

For trajectory prediction evaluation, we use the standard Average Displacement Error (ADE) and Final Displacement Error (FDE) metrics:

$$ADE = \frac{\sum_{j=t_{ob}+1}^{t_f} \|\hat{u}_j - u_j\|_2}{(t_f - t_{ob})} \quad FDE = \|\hat{u}_{t_f} - u_{t_f}\|_2$$

where \hat{u} and u are the estimated and ground truth positions respectively. Furthermore, we use the *minADE-N* and *minFDE-N* error metric introduced in [23] for multimodal evaluation. The metric is simply the minimum ADE and FDE out of N future trajectories predicted at test-time.

For intention prediction, we evaluate frame-wise classification accuracy of intents and visualize the confusion matrix to analyze classification performance.

5. Experiments

In this section, we present results of our model on trajectory & intent prediction tasks and demonstrate a superior performance against prior state-of-the-art baselines (with publicly available code) across a variety of settings. We benchmark against PECNet [3], a strong scene agnostic trajectory prediction method with state-of-the-art performance on standard intention agnostic prediction datasets. S-STGCNN [46] and S-GAN [23] are strong socially-aware models that achieved prior state-of-the-art on various benchmarks. We also report an interesting ablation on the effect of annotation frequency on the final performance, which confirms our hypothesis for the effectiveness of detailed intent annotations in trajectory prediction.

Trajectory Prediction Performance. We report our model’s performance and benchmark it against prior state-of-the-art models for unimodal (single shot, $N = 1$) prediction in Table 2 and for multimodal predictions ($N = 20$ shots) in Table 3. Our ablations are with Ours (without action/intention labels), IC (with action/intention labels for intention conditioning), SG (with scene graph for social reasoning and environmental cues).

Several interesting trends emerge. First, we observe that in the single shot setting, our intention conditioned model outperforms prior state-of-the-art method by a significant margin of 12% in ADE, 9% in FDE. Second, we see a similar trend in multi-shot prediction setting as well with our model outperforming PECNet by 33% in ADE and 9% in FDE for pedestrians and a delta of 26% in ADE and 13% in FDE for moving vehicles. Third, notice that the performance gap is significant in hard non-linear cases such as lane changes and turns, where our model achieves 30% and 16% better performance in ADE and FDE respectively.

Also noteworthy is the crucial effect of conditioning predictions on intentions and incorporating social and environmental cues through the scene graph, which is also shown in Table 2 and Table 3. We note that both intention cues and scene graph information are critical to overall performance, with intention improving ADE performance by up to 7% and 8% across all agent types (especially nonlinear trajectories such as lane changes and turns) for the unimodal and multimodal settings. We notice that the scene graph boosts performance by 3% in ADE for the multimodal setting across all agent types.

We notice an interesting behavior with pedestrians. Conditioning on pedestrian intent such as crossing vs. waiting to cross helps for single-shot prediction as shown in Table 2. However, we do not see a benefit for multimodal prediction.

	S-STGCNN		EvolveGraph		PECNet		Ours		Ours + IC		Ours + IC + SG	
	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE
Pedestrians	0.96	1.98	0.83	1.49	0.79	1.31	0.61	1.38	0.56	1.24	0.55	1.21
Vehicles	3.03	7.01	3.03	2.58	6.63	6.34	2.37	6.20	2.23	5.80	2.24	5.82
Lane Change	4.41	10.17	2.96	7.92	2.78	7.60	2.93	7.88	2.47	6.78	2.52	6.71
Turn	3.48	8.15	3.13	7.85	2.97	7.44	2.76	7.26	2.69	7.03	2.69	7.02

Table 2: **Trajectory error metrics for N=1 samples:** ADE and FDE of various state-of-the-art baselines and our method using unimodal (single-shot) evaluation. Reported errors are in meters. Lower is better. We show results evaluated on separate classes to gain more insight on prediction performance. We report errors on 1) pedestrians, 2) vehicles (non-static), 3) agents that change lanes, and 4) agents that turn.

	S-GAN		S-STGCNN		EvolveGraph		PECNet		Ours		Ours + IC		Ours + IC + SG	
	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE
Pedestrians	1.04	2.18	0.63	1.01	0.55	0.79	0.51	0.70	0.36	0.70	0.37	0.71	0.34	0.64
Vehicles	3.57	8.05	2.28	4.46	1.72	3.41	1.59	3.05	1.33	3.09	1.20	2.63	1.18	2.64
Lane Change	3.50	8.41	3.00	6.09	1.86	3.39	1.62	2.85	1.42	3.30	1.26	2.70	1.22	2.71
Turn	3.75	9.01	2.68	5.71	2.25	4.32	1.96	4.07	1.54	3.59	1.45	3.24	1.40	3.13

Table 3: **Trajectory error metrics for N=20 samples:** ADE and FDE of various state-of-the-art baselines and our method using multi-modal evaluation. Reported errors are in meters. Lower is better. We report errors on the same classes described in Table 2




Figure 5: Visualization of top-1 trajectory prediction result (green: past observation, blue: ground truth, red: prediction) and frame-wise intention of a particular agent in dark green circle at the start of the observation time step (GI: Ground truth Intention, PI: Predicted Intention) is shown at the bottom of each scenario. More detailed visualizations and comparisons are provided in supplementary material.

We hypothesize that this is because the type of intent we label for pedestrian is not as granular as for vehicles in that it does not change drastically frame-by-frame. This is validated in Figure 8 which shows experiments with downsampled intention annotations. We observe that for pedestrians, lower frequency annotations does not diminish performance




Figure 6: Intention prediction confusion matrices. (a-b) results for vehicles under both unimodal and multimodal sampling, (c-d) those for pedestrians.

as compared to vehicles due to more unconstrained behavior, we cannot have as detailed intent labels that are used for vehicles such as turn or lane change. This may explain the behavior of why intention conditioning only helps for the single-shot case for pedestrians.

In Figure 5, we visualize our model’s best-of-20 performance. We observe that predicted trajectories are fairly accurate and with underlying turning intentions. While there are limitations in exact frame-wise intention predictions, we notice it can capture key future actions of turning and can help guide predictions.

Intention Prediction: In addition to trajectory prediction, our dataset enables for a more high level understanding of agent intent to mimic how they plan their trajectory. Figure 7 illustrates the performance of intention prediction over a 25 frame (5s) prediction horizon. Our work is the first to




Figure 7: **Accuracy vs. Future Horizon (in frames).** The x axis of each figure is time and the y axis of each figure is accuracy (from 0 to 1). The change of intention prediction accuracy over a time horizon for both unimodal and multimodal predictions. In (a-b) we plot intention accuracy over time for vehicles for N=1 and N=20 samples respectively. In (c-d) we plot intention accuracy over time for pedestrians with N=1 and N=20 samples.

baseline both pedestrian and vehicle intent on a frame-wise level. We notice that prediction performance monotonically worsens over the horizon. However, we notice that for vehicles the intention accuracy in the multimodal setting are significantly improved from the unimodal case. This explains why intention conditioning helps more in the multimodal case, as agent intents are much more accurately understood. In contrast, only a slight improvement in intention performance for pedestrians. We posit this is because the intents for pedestrians do not change as frequently and are not as granular capturing direction such as "turn left"; thus, having more samples does not necessarily increase performance.

To better understand intention estimation, we visualize the confusion matrices as shown in Figure 6. For vehicles, we use the following set of discrete actions: *moving*, *stopped*, *parking*, *lane change*, *turn left*, and *turn right*. We observed improved performance for vehicle intention prediction with multimodal goal destination sampling, indicating that our model can correlate long-term goals with short-term intent. For pedestrians, we use *moving*, *waiting to cross*, *crossing*, and *stopped*. The intents for pedestrians do not rapidly change unlike those for vehicles. Thus, we see that multimodal predictions do not actually improve pedestrian intention estimation. These results corroborate the results in Table 3 where multimodal predictions with intention fail to outperform predictions without intentions. This is further examined in the next section.

Effect of Annotation Frequency: Our dataset provides very detailed frame-wise intention labels at 5FPS for all




Figure 8: ADE Performance based on varying ground-truth intention annotation frequency.

agents. To examine the importance of having a dataset with such detailed annotations, we experiment with how changing annotation frequency can affect performance. We provide our model with oracle intentions available at varying frequencies. As shown in Figure 8 trajectory prediction performance worsens roughly linearly as the frequency of intention labels reduces. This highlights the importance of our highly detailed annotations, as a choice to annotate every other frame (2.5FPS) clearly affects performance. Note that this effect is witnessed for primarily for vehicles, especially those that change lanes or turn. Pedestrian performance is not affected much, as the intention labels used for pedestrians do not change drastically for each frame. This also explains why intention conditioning did not help for multimodal evaluation for pedestrians as seen in Table 3.

6. Conclusion

In this work, we presented a large-scale heterogeneous dataset with detailed, frame-wise intention annotations. This dataset allows for both traditional trajectory prediction as well as understanding how intent changes over a long time horizon. In doing so, this dataset is the first that can be used as a benchmark for intention understanding for both vehicles and pedestrians. Furthermore, we formulate a joint trajectory and intention prediction framework which outperforms state-of-the-art on trajectory prediction metrics and offers a strong baseline for intention prediction. We bridge the gap between trajectory prediction and intention prediction and show that combining the two can better model agents' decision-making process, assisting in trajectory prediction. We believe our dataset can inspire future works that consider intention prediction in addition to traditional trajectory forecasting. Doing so can give more insight into models' decisions and will be critical in designing and maintaining a safe forecasting system.

Acknowledgement

We thank our Honda Research Institute USA colleagues – Behzad Dariush for his advice and support, Jiawei Huang for sensor calibration, and Huan Doung Nugen for data inspection and quality control.

References

- [1] Susan Carey and Elizabeth Spelke. Domain-specific knowledge and conceptual change. *Mapping the mind: Domain specificity in cognition and culture*, 169:200, 1994.
- [2] Vivian V Valentin, Anthony Dickinson, and John P O’Doherty. Determining the neural substrates of goal-directed learning in the human brain. *Journal of Neuroscience*, 27(15):4019–4026, 2007.
- [3] Karttikeya Mangalam, Yang An, Harshayu Girase, and Jitendra Malik. From goals, waypoints & paths to long term human trajectory forecasting. *arXiv preprint arXiv:2012.01526*, 2020.
- [4] Chiho Choi, Srikanth Malla, Abhishek Patil, and Joon Hee Choi. Drogon: A trajectory prediction model based on intention-conditioned behavior reasoning. In *Proceedings of the Conference on Robot Learning*, 2020.
- [5] Hang Zhao, Jiyang Gao, Tian Lan, Chen Sun, Benjamin Sapp, Balakrishnan Varadarajan, Yue Shen, Yi Shen, Yuning Chai, Cordelia Schmid, et al. Tnt: Target-driven trajectory prediction. *arXiv preprint arXiv:2008.08294*, 2020.
- [6] Bingbin Liu, Ehsan Adeli, Zhangjie Cao, Kuan-Hui Lee, Abhijeet Sheno, Adrien Gaidon, and Juan Carlos Niebles. Spatiotemporal relationship reasoning for pedestrian intent prediction, 2020.
- [7] Amir Rasouli, Iuliia Kotseruba, Toni Kunic, and John K Tsotsos. Pie: A large-scale dataset and models for pedestrian intention estimation and trajectory prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6262–6271, 2019.
- [8] Sergio Casas, Wenjie Luo, and Raquel Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *Conference on Robot Learning*, pages 947–956. PMLR, 2018.
- [9] Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, Kuan-Hui Lee, Ehsan Adeli, Jitendra Malik, and Adrien Gaidon. It is not the journey but the destination: Endpoint conditioned trajectory prediction. In *European Conference on Computer Vision*, pages 759–776. Springer, 2020.
- [10] Amir Rasouli. Pedestrian simulation: A review. *arXiv preprint arXiv:2102.03289*, 2021.
- [11] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>.
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- [13] Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. Are they going to cross? a benchmark dataset and baseline for pedestrian crosswalk behavior. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 206–213, 2017.
- [14] Srikanth Malla, Behzad Dariush, and Chiho Choi. Titan: Future forecast using action priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11186–11196, 2020.
- [15] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrila, and Kai O Arras. Human motion trajectory prediction: A survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020.
- [16] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019.
- [17] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [18] Parth Kothari, Sven Kreiss, and Alexandre Alahi. Human trajectory forecasting in crowds: A deep learning perspective. *arXiv preprint arXiv:2007.03639*, 2020.
- [19] A Robicquet, A Sadeghian, A Alahi, and S Savarese. Learning social etiquette: Human trajectory prediction in crowded scenes. In *European Conference on Computer Vision (ECCV)*, 2020.
- [20] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M. Kitani, Dariu M. Gavrila, and Kai O. Arras. Human Motion Trajectory Prediction: A Survey. *arXiv e-prints*, 2019.
- [21] Jiachen Li, Fan Yang, Masayoshi Tomizuka, and Chiho Choi. Evolvegraph: Multi-agent trajectory prediction with dynamic relational reasoning. In *Advances in neural information processing systems*, 2020.
- [22] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezafofighi, and Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1349–1358, 2019.
- [23] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2255–2264, 2018.
- [24] Hengbo Ma, Yaofeng Sun, Jiachen Li, Masayoshi Tomizuka, and Chiho Choi. Continual multi-agent interaction behavior prediction with conditional generative memory. *IEEE Robotics and Automation Letters*, 2021.
- [25] Matteo Lisotto, Pasquale Coscia, and Lamberto Ballan. Social and scene-aware trajectory prediction in crowded spaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [26] Sumit Kumar, Yiming Gu, Jerrick Hoang, Galen Clark Haynes, and Micol Marchetti-Bowick. Interaction-based trajectory prediction over a hybrid traffic graph. *arXiv preprint arXiv:2009.12916*, 2020.

- [27] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.
- [28] Nachiket Deo and Mohan M Trivedi. Trajectory forecasts in unknown environments conditioned on grid-based plans. *arXiv preprint arXiv:2001.00735*, 2020.
- [29] Jiachen Li, Fan Yang, Hengbo Ma, Masayoshi Tomizuka, and Chiho Choi. Rain: Reinforced hybrid attention inference network for motion forecasting. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.
- [30] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, S Hamid Rezatofighi, and Silvio Savarese. Socialbigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. *arXiv preprint arXiv:1907.03395*, 2019.
- [31] Chiho Choi, Joon Hee Choi, Jiachen Li, and Srikanth Malla. Shared cross-modal trajectory prediction for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 244–253, 2021.
- [32] Alexandre Alahi, Vignesh Ramanathan, and Li Fei-Fei. Socially-aware large-scale crowd forecasting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2203–2210, 2014.
- [33] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.
- [34] Takuma Yagi, Karttikeya Mangalam, Ryo Yonetani, and Yoichi Sato. Future person localization in first-person videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7593–7602, 2018.
- [35] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip HS Torr, and Manmohan Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 336–345, 2017.
- [36] Jiachen Li, Hengbo Ma, and Masayoshi Tomizuka. Conditional generative neural system for probabilistic trajectory prediction. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6150–6156. IEEE, 2019.
- [37] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. *arXiv preprint arXiv:2001.03093*, 2020.
- [38] Nicholas Rhinehart, Rowan McAllister, Kris Kitani, and Sergey Levine. Precog: Prediction conditioned on goals in visual multi-agent settings. *arXiv preprint arXiv:1905.01296*, 2019.
- [39] Yu Yao, Ella Atkins, Matthew Johnson-Roberson, Ram Vasudevan, and Xiaoxiao Du. Bitrap: Bi-directional pedestrian trajectory prediction with multi-modal goal estimation. *IEEE Robotics and Automation Letters*, 2021.
- [40] Lingyao Zhang, Po-Hsun Su, Jerrick Hoang, Galen Clark Haynes, and Micol Marchetti-Bowick. Map-adaptive goal-based trajectory prediction. *arXiv preprint arXiv:2009.04450*, 2020.
- [41] Patrick Dendorfer, Aljosa Osep, and Laura Leal-Taixé. Goalgan: Multimodal trajectory prediction based on goal position estimation. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [42] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Etinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [43] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [44] Wei Zhan, Liting Sun, Di Wang, Haojie Shi, Aubrey Claude, Maximilian Naumann, Julius Kümmerle, Hendrik Königshof, Christoph Stiller, Arnaud de La Fortelle, and Masayoshi Tomizuka. INTERACTION Dataset: An INTERNATIONAL, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps. *arXiv:1910.03088 [cs, eess]*, 2019.
- [45] Yunsheng Shi, Zhengjie Huang, Shikun Feng, and Yu Sun. Masked label prediction: Unified message passing model for semi-supervised classification. *arXiv preprint arXiv:2009.03509*, 2020.
- [46] Abdullah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14424–14432, 2020.