

TITAN: Future Forecast using Action Priors

Srikanth Malla

Behzad Dariush

Chiho Choi

Honda Research Institute USA

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

Abstract

We consider the problem of predicting the future trajectory of scene agents from egocentric views obtained from a moving platform. This problem is important in a variety of domains, particularly for autonomous systems making reactive or strategic decisions in navigation. In an attempt to address this problem, we introduce TITAN (Trajectory Inference using Targeted Action priors Network), a new model that incorporates prior positions, actions, and context to forecast future trajectory of agents and future ego-motion. In the absence of an appropriate dataset for this task, we created the TITAN dataset that consists of 700 labeled video-clips (with odometry) captured from a moving vehicle on highly interactive urban traffic scenes in Tokyo. Our dataset includes 50 labels including vehicle states and actions, pedestrian age groups, and targeted pedestrian action attributes that are organized hierarchically corresponding to atomic, simple/complex-contextual, transportive, and communicative actions. To evaluate our model, we conducted extensive experiments on the TITAN dataset, revealing significant performance improvement against baselines and state-of-the-art algorithms. We also report promising results from our Agent Importance Mechanism (AIM), a module which provides insight into assessment of perceived risk by calculating the relative influence of each agent on the future ego-trajectory. The dataset is available at <https://usa.honda-ri.com/titan>

1. Introduction

The ability to forecast future trajectory of agents (individuals, vehicles, cyclists, etc.) is paramount in developing navigation strategies in a range of applications including motion planning and decision making for autonomous and cooperative (shared autonomy) systems. We know from observation that the human visual system possesses an uncanny ability to forecast behavior using various cues such as experience, context, relations, and social norms. For example, when immersed in a crowded driving scene, we are able to reasonably estimate the intent, future actions, and

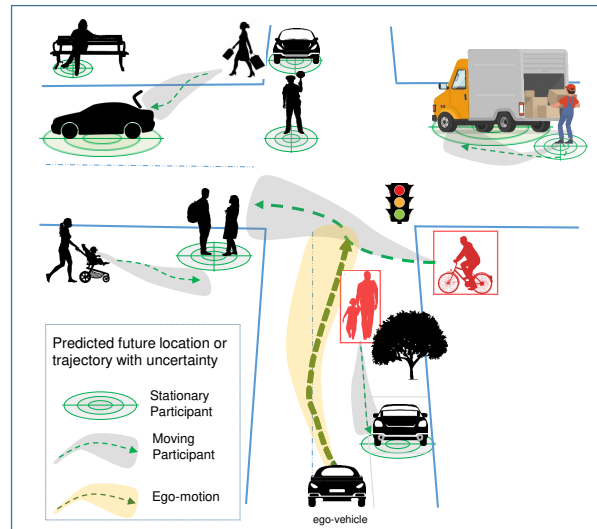


Figure 1. Our goal is to predict the future trajectory of agents from egocentric views obtained from a moving platform. We hypothesize that prior actions (and implicit intentions) play an important role in future trajectory forecast. To this end, we develop a model that incorporates prior positions, actions, and context to forecast future trajectory of agents and future ego-motion. This figure is a conceptual illustration that typifies navigation of ego-vehicle in an urban scene, and how prior actions/intentions and context play an important role in future trajectory forecast. We seek to also identify agents (depicted by the red bounding box) that influence future ego-motion through an Agent Importance Mechanism (AIM).

future location of the traffic participants in the next few seconds. This is undoubtedly attributed to years of prior experience and observations of interactions among humans and other participants in the scene. To reach such human level ability to forecast behavior is part of the quest for visual intelligence and the holy grail of autonomous navigation, requiring new algorithms, models, and datasets.

In the domain of behavior prediction, this paper considers the problem of future trajectory forecast from egocentric views obtained from a mobile platform such as a vehicle in a road scene. This problem is important for autonomous agents to assess risks or to plan ahead when making reactive or strategic decisions in navigation. Several recently reported models that predict trajectories incorporate social

norms, semantics, scene context, etc. The majority of these algorithms are developed from a stationary camera view in surveillance applications, or overhead views from a drone.

The specific objective of this paper is to develop a model that incorporates prior positions, actions, and context to simultaneously forecast future trajectory of agents and future ego-motion. In a related problem, the ability to predict future actions based on current observations has been previously studied in [25, 47, 46, 45, 50]. However, to the best of our knowledge, action priors have not been used in forecasting future trajectory, partly due to a lack of an appropriate dataset. A solution to this problem can help address the challenging and intricate scenarios that capture the interplay of observable actions and their role in future trajectory forecast. For example, when the egocentric view of a mobile agent in a road scene captures a delivery truck worker closing the tailgate of the truck, it is highly probable that the worker’s future behavior will be to walk toward the driver side door. Our aim is to develop a model that uses such action priors to forecast trajectory.

The algorithmic contributions of this paper are as follows. We introduce TITAN (Trajectory Inference using Targeted Action priors Network), a new model that incorporates prior positions, actions, and context to simultaneously forecast future trajectory of agents and future ego-motion. Our framework introduces a new interaction module to handle dynamic number of objects in the scene. While modeling pair-wise interactive behavior from all agents, the proposed interaction module incorporates actions of individuals in addition to their locations, which helps the system to understand the contextual meaning of motion behavior. In addition, we propose to use multi-task loss with aleatoric homoscedastic uncertainty [22] to improve the performance of multi-label action recognition. For ego-future, Agent Importance Mechanism (AIM) is presented to identify objects that are more relevant for ego-motion prediction.

Apart from algorithmic contributions, we introduce a novel dataset, referred to as TITAN dataset, that consists of 700 video clips captured from a moving vehicle on highly interactive urban traffic scenes in Tokyo. The pedestrians in each clip were labeled with various action attributes that are organized hierarchically corresponding to atomic, simple/complex contextual, transportive, and communicative actions. The action attributes were selected based on commonly observed actions in driving scenes, or those which are important for inferring intent (*e.g.*, waiting to cross). We also labeled other participant categories, including vehicle category (4 wheel, 2 wheel), age-groups, and vehicle state. The dataset contains synchronized ego-motion information from an IMU sensor. To our knowledge, this is the only comprehensive and large scale dataset suitable for studying action priors for forecasting the future trajectory of agents from ego-centric views obtained from a moving platform.

Furthermore, we believe our dataset will contribute to advancing research for action recognition in driving scenes.

2. Related Work

2.1. Future Trajectory Forecast

Human Trajectory Forecast Encoding interactions between humans based on their motion history has been widely studied in the literature. Focusing on input-output time sequential processing of data, recurrent neural network (RNN)-based architectures have been applied to the future forecast problem in the last few years [2, 26, 17, 56, 60]. More recently, RNNs are used to formulate a connection between agents with their interactions using graph structures [54, 30]. However, these methods suffer from understanding of environmental context with no or minimal considerations of scene information. To incorporate models of human interaction with the environment, [57] takes local to global scale image features into account. More recently, [10] visually extracts relational behavior of humans interacting with other agents as well as environments.

Vehicle Trajectory Forecast Approaches for vehicle motion prediction have developed following the success of interaction modeling using RNNs. Similar to human trajectory forecast, [13, 35, 30, 29] only consider the past motion history. These methods perform poorly in complex road environments without the guidance of structured layouts. Although the subsequent approaches [40, 28, 11] partially overcome these issues by using 3D LIDAR information as inputs to predict future trajectories, their applicability to current production vehicles is limited due to the higher cost. Recent methods [3, 58, 31] generate trajectories of agents from an egocentric view. However, they do not consider interactions between road agents in the scene and the potential influence to the ego-future. In this work, we explicitly model pair-wise interactive behavior from all agents to identify objects that are more relevant for the target agent.

2.2. Action Recognition

With the success of 2D convolutions in image classification, frame-level action recognition has been presented in [20]. Subsequently, [44] separates their framework into two streams: one to encode spatial features from RGB images and the other to encode temporal features from corresponding optical flow. Their work motivated studies that model temporal motion features together with spatial image features from videos. A straightforward extension has been shown in [51, 52], replacing 2D convolutions by 3D convolutions. To further improve the performance of these models, several research efforts have been provided such as I3D [7] that inflates a 2D convolutional network into 3D to benefit from the use of pre-trained models and 3D ResNet [18] that adds residual connections to build a very

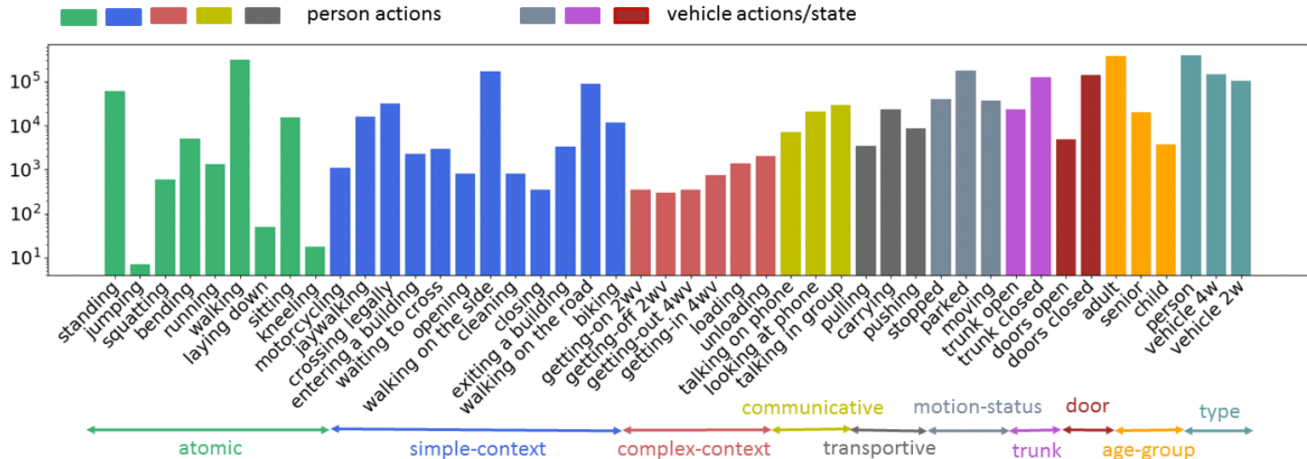


Figure 2. Distribution of labels sorted according to person actions, vehicle actions/state, and other labels such as age groups and types.

deep 3D network. Apart from them, other approaches capture pair-wise relations between actor and contextual features [49] or those between pixels in space and in time [55]. More recently, Timeception [19] models long range temporal dependencies, particularly focusing on complex actions.

2.3. Datasets

Future Trajectory Several influential RGB-based datasets for pedestrian trajectory prediction have been reported in the literature. These datasets are typically created from a stationary surveillance camera [27, 37, 34], or from aerial views obtained from a static drone-mounted camera [41]. In driving scenes, the 3D point cloud-based datasets [15, 36, 23, 5, 1, 9] were originally introduced for detection, tracking, etc., but recently used for vehicle trajectory prediction as well. Also, [58, 8] provide RGB images captured from an egocentric view of a moving vehicle and applied to future trajectory forecast problem. The JAAD [39], CMU-UAH [33], and PIE [38] datasets are most similar to our TITAN dataset in the sense that they are designed to study the intentions and actions of objects from on-board vehicles. However, their labels are limited to simple actions such as walking, standing, looking, and crossing. These datasets, therefore, do not provide an adequate number of actions to use as priors in order to discover contextual meaning of agents’ motion behavior. To address these limitations, our TITAN dataset provides 50 labels including vehicle states and actions, pedestrian age groups, and targeted pedestrian action attributes that are hierarchically organized as illustrated in the supplementary material. **Action Recognition** A variety of datasets have been introduced for action recognition with a single action label [24, 48, 20, 32, 21] and multiple action labels [43, 59, 4] in videos. Recently released datasets such as AVA [16], READ [14], and EPIC-KITCHENS [12] contain actions with corresponding localization around a person or object. Our TITAN dataset is similar to AVA in the sense that it

provides spatio-temporal localization for each agent with multiple action labels. However, the labels of TITAN are organized hierarchically from primitive atomic actions to complicated contextual activities that are typically observed from on-board vehicles in driving scenes.

3. TITAN Dataset

In the absence of an appropriate dataset suitable for our task, we introduce the TITAN dataset for training and evaluation of our models as well as to accelerate research on trajectory forecast. Our dataset is sourced from 10 hours of video recorded at 60 FPS in central Tokyo. All videos are captured using a GoPro Hero 7 Camera with embedded IMU sensor which records synchronized odometry data at 100 HZ for ego-motion estimation. To create the final annotated dataset, we extracted 700 short video clips from the original (raw) recordings. Each clip is between 10-20 seconds in duration, image size width:1920px, height:1200px and annotated at 10 HZ sampling frequency. The characteristics of the selected video clips include scenes that exhibit a variety of participant actions and interactions.

The taxonomy and distribution of all labels in the dataset are depicted in Figure 2. The total number of frames annotated is approximately 75,262 with 395,770 persons, 146,840 4-wheeled vehicles and 102,774 2-wheeled vehicles. This includes 8,592 unique persons and 5,504 unique vehicles. For our experiments, we use 400 clips for training, 200 clips for validation and 100 clips for testing. As mentioned in Section 2.3, there are many publicly available datasets related to mobility and driving, many of which include ego-centric views. However, since those datasets do not provide action labels, a meaningful quantitative comparison of the TITAN dataset with respect to existing mobility datasets is not possible. Furthermore, a quantitative comparison with respect to action localization datasets such as AVA is not warranted since AVA does not include ego-




Figure 3. Example scenarios of the TITAN Dataset: a pedestrian bounding box with tracking ID is shown in , vehicle bounding box with ID is shown in , future locations are displayed in ○. Action labels are shown in different colors following Figure 2.

centric views captured from a mobile platform.

In the TITAN dataset, every participant (individuals, vehicles, cyclists, etc.) in each frame is localized using a bounding box. We annotated 3 labels (person, 4-wheeled vehicle, 2-wheeled vehicle), 3 age groups for person (child, adult, senior), 3 motion-status labels for both 2 and 4-wheeled vehicles, and door/trunk status labels for 4-wheeled vehicles. For action labels, we created 5 mutually exclusive person action sets organized hierarchically (Figure 2). In the first action set in the hierarchy, the annotator is instructed to assign exactly one class label among 9 atomic whole body actions/postures that describe primitive action poses such as sitting, standing, standing, bending, etc. The second action set includes 13 actions that involve single atomic actions with simple scene context such as jaywalking, waiting to cross, etc. The third action set includes 7 complex contextual actions that involve a sequence of atomic actions with higher contextual understanding, such as getting in/out of a 4-wheel vehicle, loading/unloading, etc. The fourth action set includes 4 transportive actions that describe the act of manually transporting an object by carrying, pulling or pushing. Finally, the fifth action set includes 4 communicative actions observed in traffic scenes such as talking on the phone, looking at phone, or talking in groups. In each action sets 2-5, the annotators were instructed to assign ‘None’ if there is no label. This hierarchical strategy was designed to produce unique (unambiguous) action labels while reducing the annotators’ cognitive workload and thereby improving annotation quality. The tracking ID’s of all localized objects are associated within each video clip. Example scenarios are displayed in Figure 3.

4. Methodology

Figure 4 shows the block diagram of the proposed TITAN framework. A sequence of image patches $\mathcal{I}_{t=1:T_{obs}}^i$ is obtained from the bounding box¹ $x^i = \{c_u, c_v, l_u, l_v\}$ of

¹We assume that the bounding box detection using past images is provided by the external module since detection is not the scope of this paper.

agent i at each past time step from 1 to T_{obs} , where (c_u, c_v) and (l_u, l_v) represent the center and the dimension of the bounding box, respectively. The proposed TITAN framework requires three inputs as follows: $\mathcal{I}_{t=1:T_{obs}}^i$ for the action detector, x_t^i for both the interaction encoder and past object location encoder, and $e_t = \{\alpha_t, \omega_t\}$ for the ego-motion encoder where α_t and ω_t correspond to the acceleration and yaw rate of the ego-vehicle at time t , respectively. During inference, the multiple modes of future bounding box locations are sampled from a bi-variate Gaussian generated by the noise parameters, and the future ego-motions \hat{e}_t are accordingly predicted, considering the multi-modal nature of the future prediction problem.

Henceforth, the notation of the feature embedding function using multi-layer perceptron (MLP) is as follows: Φ is without any activation, and Φ_r , Φ_t , and Φ_s are associated with ReLU, tanh, and a sigmoid function, respectively.

4.1. Action Recognition

We use the existing state-of-the-art method as backbone for the action detector. We finetune single-stream I3D [7] and 3D ResNet [18] architecture pre-trained on Kinetics-600 [6]. The original head of the architecture is replaced by a set of new heads (8 action sets of TITAN except age group and type) for multi-label action outputs. The action detector takes $\mathcal{I}_{t=1:T_{obs}}^i$ as input, which is cropped around the agent i . Then, each head outputs an action label including a ‘None’ class if no action is shown. From our experiments, we observed that certain action sets converge faster than others. This is due in part because some tasks are relatively easier to learn, given the shared representations. Instead of tuning the weight of each task by hand, we adopt the multi-task loss in [22] to further boost performance of our action detector. Note that each action set of the TITAN dataset is mutually exclusive, thus we consider the outputs are independent to each other as follows:

$$p(y_m, \dots, y_n | f(\mathcal{I})) = \prod_{i=m}^n p(y_i | f(\mathcal{I})), \quad (1)$$

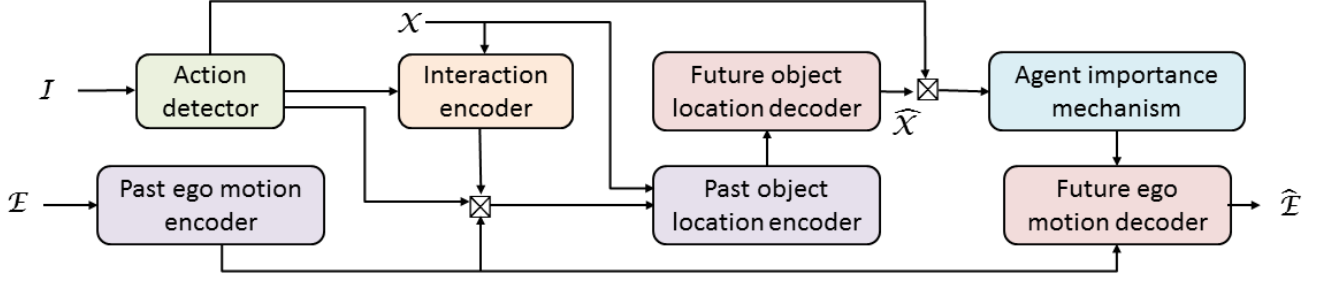


Figure 4. The proposed approach predicts the future motion of road agents and ego-vehicle in egocentric view by using actions as a prior. The notation \mathcal{I} represents input images, \mathcal{X} is the input trajectory of other agents, E is the input ego-motion, $\hat{\mathcal{X}}$ is the predicted future trajectory of other agents, and \hat{E} is the predicted future ego-motion.

where y_i is the output label of i^{th} action set and f is the action detection model. Then, multi-task loss is defined as:

$$\mathcal{L}_a = \sum_{i=m}^n \frac{ce(\widehat{cls}_i, cls_i)}{\sigma_i^2} + \log \sigma_i, \quad (2)$$

where ce is the cross entropy loss between predicted actions \widehat{cls}_i and ground truth cls_i for each label $i = m : n$. Also, σ_i is the task dependent uncertainty (aleatoric homoscedastic). In practice, the supervision is done separately for vehicles and pedestrians as they have different action sets. The efficacy of the multi-task loss is detailed in the supplementary material, and the performance of the action detector with different backbone is compared in Table 1.

4.2. Future Object Localization

Unlike existing methods, we model the interactions using the past locations of agents conditioned on their actions, which enables the system to explicitly understand the contextual meaning of motion behavior. At each past time step t , the given bounding box $x_t^i = \{c_u, c_v, l_u, l_v\}_t$ is concatenated with the multi-label action vector a_t^i . We model the pair-wise interactions between the target agent i and all other agents j through MLP, $v_t^{ij} = \Phi_r(x_t^i \boxtimes a_t^i \boxtimes x_t^j \boxtimes a_t^j)$ where \boxtimes is a concatenation operator. The resulting interactions v_t^{ij} are evaluated through the dynamic RNN with GRUs to leave more important information with respect to the target agent, $h_t^{i(j+1)} = GRU(v_t^{ij}, h_t^{ij}; W_{\text{INT}})$, where W_{INT} are the weight parameters. Note that we pass the messages of instant interaction with each agent at time t , which enables us to find their potential influence at that moment. Then, we aggregate the hidden states to generate interaction features $\psi_t^i = \frac{1}{n} \sum_i h_t^{ij}$ for the target agent i , computed from all other agents in the scene at time t as in Figure 5.

The past ego motion encoder takes $e_t = (\alpha_t, \omega_t)$ as input and embeds the motion history of ego-vehicle using the GRU. We use each hidden state output h_t^e to compute future locations of other agents. The past object location encoder uses the GRU to embed the history of past motion into a feature space. The input to this module is a bounding box




Figure 5. Interaction encoding for agent i against others at time t .

x_t^i of the target agent i at each past time step, and we use the embedding $\Phi(x_t^i)$ for the GRU. The output hidden state h_t^p of the encoder is updated by $\hat{h}_t^p = \Phi(H_t^{x^i} \boxtimes h_t^p)$, where $H_t^{x^i} = \Phi_r(a_t^i \boxtimes \psi_t^i \boxtimes \Phi_r(h_t^e))$ is the concatenated information. Then, \hat{h}_t^p is used as a hidden state input to the GRU by $h_{t+1}^p = GRU(\hat{h}_t^p, \Phi(x_t^i); W_{\text{POL}})$, where W_{POL} are the weight parameters. We use its final hidden state as an initial hidden state input of the future object location decoder.

The future bounding boxes of the target agent i are decoded using the GRU-based future object location decoder from time step $T_{\text{obs}} + 1$ to T_{pred} . At each time step, we output a 10-dimensional vector where the first 5 values are the center $\mu_c = (c_u, c_v)$, variance $\sigma_c = (\sigma_{c_u}, \sigma_{c_v})$, and its correlation ρ_c and the rest 5 values are the dimension $\mu_l = (l_u, l_v)$, variance $\sigma_l = (\sigma_{l_u}, \sigma_{l_v})$, and its correlation ρ_l . We use two bi-variate Gaussians for bounding box centers and dimensions, so that they can be independently sampled. We use the negative log-likelihood loss function as:

$$\mathcal{L}_O = -\frac{1}{T} \sum_{t=T_{\text{obs}}+1}^{T_{\text{pred}}} \log p(c|\mu_c^t, \sigma_c^t, \rho_c) p(l|\mu_l^t, \sigma_l^t, \rho_l). \quad (3)$$

4.3. Future Ego-motion prediction

We first embed the predicted future bounding box of all agents $\hat{X} = \{\hat{x}^1, \dots, \hat{x}^N\}$ through MLP at each future time step $T_{\text{obs}} + 1$ to T_{pred} . We further condition it on the previously computed action labels in a feature space through $H_t^{e_i} = \Phi(r_{T_{\text{obs}}}^i \boxtimes \Phi_r(\hat{x}_t^i))$, where $r_{T_{\text{obs}}}^i = \Phi_r(a_{T_{\text{obs}}}^i)$. By using the action labels as a prior constraint, we explicitly




Figure 6. Agent Importance Mechanism (AIM) module.

lead the model to understand about the contextual meaning of locations. The resulting features of each agent i are weighted using the AIM module $\hat{H}_t^{ei} = w_t^i * H_t^{ei}$, where the weights $w_t^i = \Phi_t(H_t^{ei})$, similar to self-attention [53]. Then, we sum all features $H_t^e = \sum_i \hat{H}_t^{ei}$ for each future time step. This procedure is detailed in Figure 6. Note that our AIM module is simultaneously learned with the future ego-motion prediction, which results in weighting other agents more or less based on their influence/importance to the ego-vehicle. It thus provides insight into assessment of perceived risk while predicting the future motion. We qualitatively evaluate it in Sec. 5.

The last hidden state h_T^e of the past ego motion encoder is concatenated with H_t^e through $\hat{h}_T^e = \Phi(H_t^e \boxtimes h_T^e)$ and fed into the future ego motion decoder. The intermediate hidden state h_t^f is accordingly updated by H_t^e at each future time step for recurrent update of the GRU. We output the ego-future using each hidden state h_t^f through $\hat{e}_t^i = \Phi(h_t^f)$ at each future time $T_{obs} + 1$ to T_{pred} . For training, we use task dependent uncertainty with L2 loss for regressing both acceleration and angular velocity as shown below:

$$\mathcal{L}_E = \frac{\|\alpha_t - \hat{\alpha}_t\|^2}{\sigma_1^2} + \frac{\|\omega_t - \hat{\omega}_t\|^2}{\sigma_2^2} + \log \sigma_1 \sigma_2. \quad (4)$$

Note that the predicted future ego-motion is deterministic in its process. However, its multi-modality comes from sampling of the predicted future bounding boxes of other agents. In this way, we capture their influence with respect to the ego-vehicle, and AIM outputs the importance weights consistent with the agents' action and future motion.

5. Experiments

In all experiments performed in this work, we predict up to 2 seconds into the future while observing 1 second of past observations as proposed in [31]. We use average distance error (ADE), final distance error (FDE), and final intersection over union (FIOU) metrics for evaluation of future object localization. We include FIOU in our evaluation since ADE/FDE only capture the localization error of the final bounding box without considering its dimensions. For action recognition, we use per frame mean average precision (mAP). Finally, for ego-motion prediction, we use root mean square error (RMSE) as an evaluation metric.

	Method	I3D [7]	3D ResNet [18]
	Backbone	InceptionV1	ResNet50
person	atomic	0.9219	0.7552
	simple	0.5318	0.3173
	complex	0.9881	0.9880
	communicative	0.8649	0.8648
	transportive	0.9080	0.9081
	overall	0.8429	0.7667
vehicle	motion	0.9918	0.7132
	trunk	1.00	1.00
	doors	1.00	1.00
	overall	0.9921	0.9044
overall \uparrow	0.8946	0.8128	

Table 1. Action recognition results (mAP) on TITAN.

5.1. Action Recognition

We evaluate two state-of-the-art 3D convolution-based architectures, I3D with InceptionV1 and 3D ResNet with ResNet50 as backbone. Both models are pre-trained on Kinetics-600 and finetuned using TITAN with the multi-task loss in Eqn. 2. As detailed in Sec. 4.1, we modify the original structure using new heads that corresponds to the 8 action sets of the TITAN dataset. Their per frame mAP results are compared in Table 1 for each action set. We refer to the supplementary material for the detailed comparison on individual action categories. Note that we use the I3D-based action detector for the rest of our experiments.

5.2. Future Object Localization

The results of future object localization performance is shown in Table 2. The constant velocity (Const-Vel [42]) baseline is computed using the last two observations for linearly interpolating future positions. Since the bounding box dimensions error is not captured by ADE or FDE, we evaluate on FIOU using two baselines: 1) without scaling the box dimensions, and 2) with scaling linearly the box dimensions. Titan_vanilla is an encoder and decoder RNN without any priors or interactions. It shows better performance than linear models. Both Social-GAN [17] and Social-LSTM [2] improve the performance in ADE and FDE compared to the simple recurrent model (Titan_vanilla) or linear approaches. Note that we do not evaluate FIOU for Social-GAN and Social-LSTM since their original method is not designed to predict dimensions. Titan_AP adds action priors to the past positions and performs better than Titan_vanilla, which shows that the model better understands contextual meaning of the past motion. However, its performance is worse than Titan_EP that includes ego-motion as priors. This is because Titan_AP does not consider the motion behavior of other agents in egocentric view. Titan_IP includes interaction priors as shown in Figure 5 without concatenating actions. Interestingly, its performance is better than Titan_AP (action priors) and Titan_EP (ego priors) as well as Titan_EP+AP (both ego and action priors). It validates the efficacy of our interaction encoder that aims to pass the interactions over

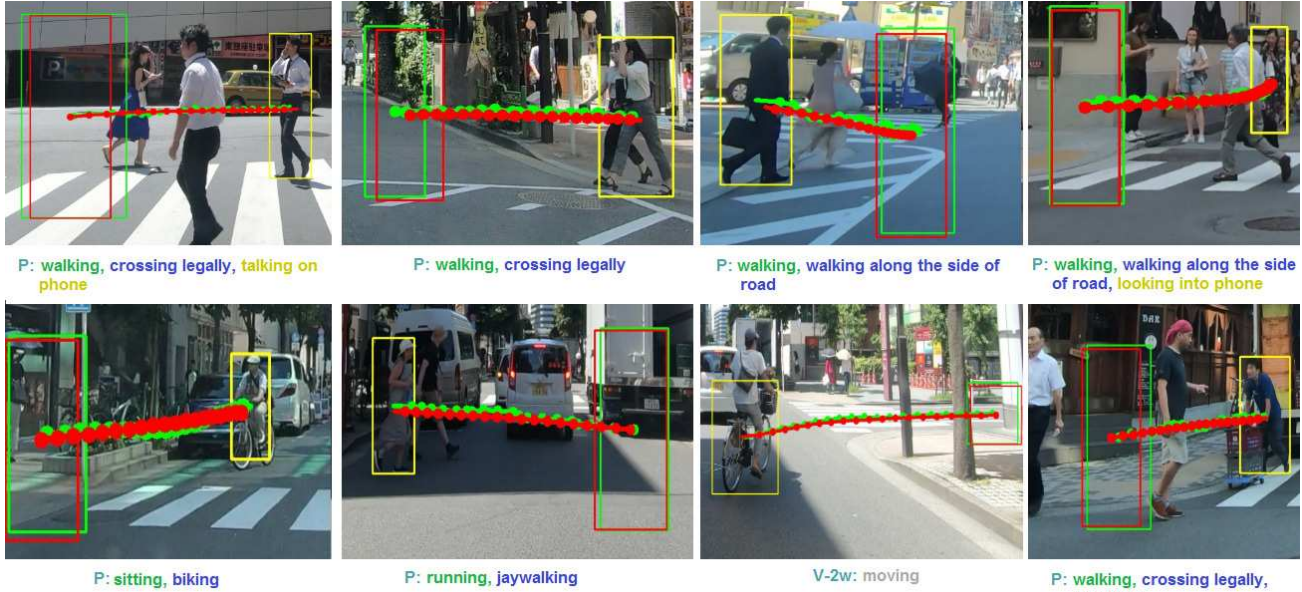


Figure 7. Qualitative evaluation on the TITAN dataset: ground truth future trajectory $\text{---}\bullet\text{---}$, TITAN prediction $\text{---}\bullet\text{---}$, last observation bounding box $\text{---}\square\text{---}$. The color of detected action labels indicates each action set described in Figure 2. Images are cropped for better visibility.




Figure 8. Comparison with others: ground truth $\text{---}\bullet\text{---}$, Titan_EP+IP+AP (ours) $\text{---}\bullet\text{---}$, Titan_EP+IP (w/o action) $\text{---}\bullet\text{---}$, Social-LSTM [2] $\text{---}\bullet\text{---}$, Social-GAN [17] $\text{---}\square\text{---}$, Const-Vel [42] $\text{---}\bullet\text{---}$, bounding box at T_{obs} $\text{---}\square\text{---}$. Images are cropped for better visibility.

all agents. This is also demonstrated by comparing Titan_IP with two state-of-the-art methods. With ego priors as default input, interaction priors (Titan_EP+IP) finally perform better than Titan_IP. Interactions with action information (Titan_EP+IP+AP) significantly outperforms all other baselines, suggesting that interactions are important and can be more meaningful with the information of actions².

The qualitative results are shown in Figure 7. The proposed method predicts natural motion for the target with respect to their detected actions (listed below each example). In Figure 8, we compare ours with the baseline models. The performance improvement against Titan_EP+IP further validates our use of action priors for future prediction. Additional results can be found in the supplementary material.

5.3. Future Ego-Motion Prediction

The quantitative results for future ego-motion prediction are shown in Table 3. Between Const-Vel [42] and Const-

²Using ground-truth actions as a prior, we observed further improvement in overall ADE by 2 pixels and overall FDE by 3.5 pixels.

Method	ADE ↓	FDE ↓	FIOU ↑
Const-Vel (w/o scaling) [42]	44.39	102.47	0.1567
Const-Vel (w/ scaling) [42]	44.39	102.47	0.1692
Social-LSTM [2]	37.01	66.78	-
Social-GAN [17]	35.41	69.41	-
Titan_vanilla	38.56	72.42	0.3233
Titan_AP	33.54	55.80	0.3670
Titan_EP	29.42	41.21	0.4010
Titan_IP	22.53	32.80	0.5589
Titan_EP+AP	26.03	38.78	0.5360
Titan_EP+IP	17.79	27.69	0.5650
Titan_EP+IP+AP (ours)	11.32	19.53	0.6559

Table 2. Quantitative evaluation for future object localization. ADE are FDE in pixels on the original size 1920x1200.

Acc (acceleration), the Const-Vel baseline performs better in predicting angular velocity (yaw-rate) and Const-Acc performs better for predicting acceleration. Titan_vanilla only takes the past ego-motion as input, performing better than Const-Vel and Const-Acc for acceleration prediction. Although incorporating information of other agents' future predictions (Titan_FP) does not improve the performance over Titan_vanilla, the addition of their action priors

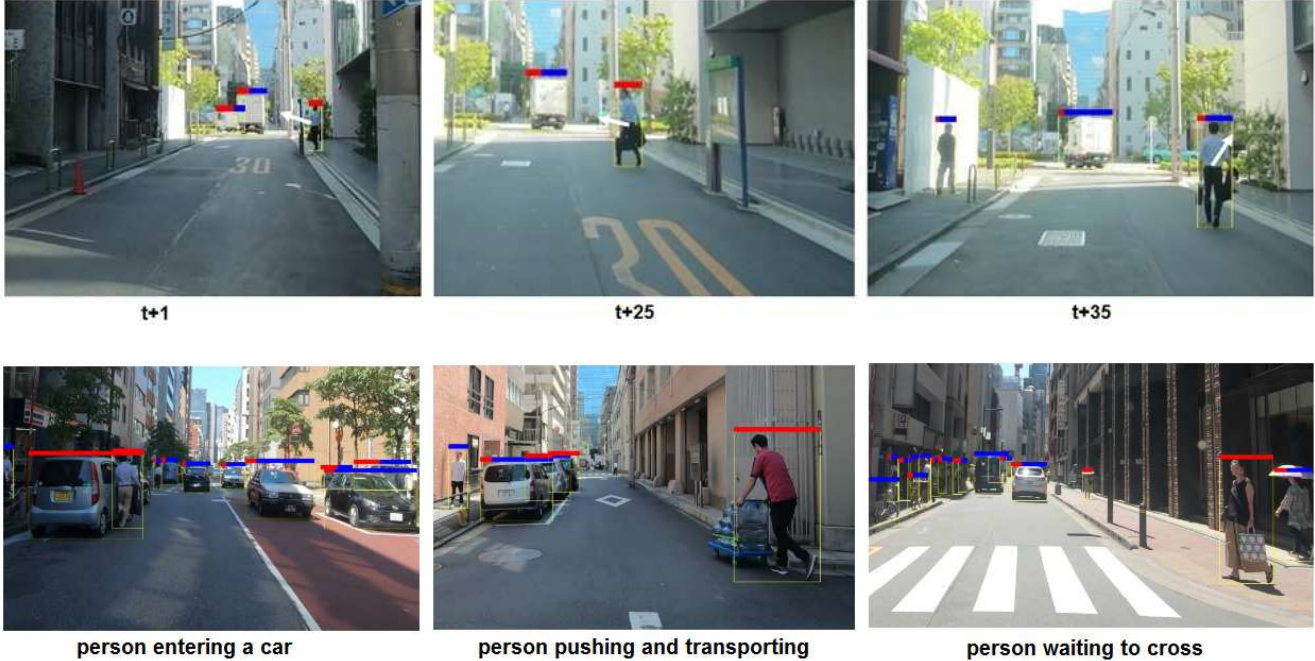


Figure 9. The importance (or degree of influence) of each agent toward the ego-vehicle’s future trajectory is illustrated by the proportion of red bar relative to the blue bar displayed across the top width of the agent’s bounding box. A red bar spanning across the top width represents the maximum importance derived from the AIM module, while a blue bar spanning across the top width represents minimum importance. (top row) images from same sequence. (bottom row) images from different sequences.

Method	acc RMSE ↓	yaw rate RMSE ↓
Const-Vel [42]	1.745	0.1249
Const-Acc	1.569	0.1549
Titan_vanilla	1.201	0.1416
Titan_FP	1.236	0.1438
Titan_FP+AP	1.182	0.1061
Titan_AIM_FP	1.134	0.0921
Titan_AIM (ours)	1.081	0.0824

Table 3. Comparison of Future ego motion prediction. acceleration error in m/s^2 and yaw rate error in rad/s .

(Titan_FP+AP) shows better performance for both acceleration and yaw rate prediction. By adding just future position in the AIM module (Titan_AIM_FP), the system can weight the importance of other agents’ behavior with respect to the ego-future, resulting in decreased error rates. Finally, by incorporating future position and action in the AIM module as a prior yields the best performance, Titan_AIM.

To show the interpretability of which participant is more important for ego-future, we visualize the importance weights in Figure 9. In particular, the top row illustrates that the importance weight of the pedestrian increases as the future motion direction (in white arrow) is towards the ego-vehicle’s future motion. Although the agent is closer to the ego-vehicle at a later time step, the importance decreases as the future motion changes. This mechanism provides insight into assessment of perceived risk for other agents from the perspective of the ego-vehicle.

6. Conclusion

We presented a model that can reason about the future trajectory of scene agents from egocentric views obtained from a mobile platform. Our hypothesis was that action priors provide meaningful interactions and also important cues for making future trajectory predictions. To validate this hypothesis, we developed a model that incorporates prior positions, actions, and context to simultaneously forecast future trajectory of agents and future ego-motion. For evaluation, we created a novel dataset with over 700 video clips containing labels of a diverse set of actions in urban traffic scenes from a moving vehicle. Many of those actions implicitly capture the agent’s intentions. Comparative experiments against baselines and state-of-art prediction algorithms showed significant performance improvement when incorporating action and interaction priors. Importantly, our framework introduces an Agent Importance Mechanism (AIM) module to identify agents that are influential in predicting the future ego-motion, providing insight into assessment of perceived risk in navigation. For future work, we plan to incorporate additional scene context to capture participant interactions with the scene or infrastructure.

Acknowledgement We thank Akira Kanehara for supporting our data collection and Yuji Yasui, Rei Sakai, and Isht Dwivedi for insightful discussions.

References

- [1] Waymo open dataset: An autonomous driving dataset, 2019.
- [2] Alexandre Alahi, Krathar 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.
- [3] Apratim Bhattacharyya, Mario Fritz, and Bernt Schiele. Long-term on-board prediction of people in traffic scenes under uncertainty. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4194–4202, 2018.
- [4] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–970, 2015.
- [5] 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. *arXiv preprint arXiv:1903.11027*, 2019.
- [6] Joao Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. *arXiv preprint arXiv:1808.01340*, 2018.
- [7] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [8] Rohan Chandra, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. Taphic: Trajectory prediction in dense and heterogeneous traffic using weighted interactions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8483–8492, 2019.
- [9] 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 Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019.
- [10] Chiho Choi and Behzad Dariush. Looking to relations for future trajectory forecast. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [11] Chiho Choi, Abhishek Patil, and Srikanth Malla. Drogon: A causal reasoning framework for future trajectory forecast. *arXiv preprint arXiv:1908.00024*, 2019.
- [12] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, 2018.
- [13] Nachiket Deo and Mohan M Trivedi. Multi-modal trajectory prediction of surrounding vehicles with maneuver based lstms. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1179–1184. IEEE, 2018.
- [14] Valentina Fontana, Gurkirt Singh, Stephen Akrigg, Manuele Di Maio, Suman Saha, and Fabio Cuzzolin. Action detection from a robot-car perspective. *arXiv preprint arXiv:1807.11332*, 2018.
- [15] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [16] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6047–6056, 2018.
- [17] 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.
- [18] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6546–6555, 2018.
- [19] Noureldien Hussein, Efstratios Gavves, and Arnold WM Smeulders. Timeception for complex action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 254–263, 2019.
- [20] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014.
- [21] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [22] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7482–7491, 2018.
- [23] R. Kesten, M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinsky, W. Jiang, and V. Shet. Lyft level 5 av dataset 2019. <https://level5.lyft.com/dataset/>, 2019.
- [24] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: a large video database for human motion recognition. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2011.
- [25] Tian Lan, Tsung-Chuan Chen, and Silvio Savarese. A hierarchical representation for future action prediction. In *European Conference on Computer Vision*, pages 689–704. Springer, 2014.
- [26] 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.
- [27] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer graphics forum*, volume 26, pages 655–664. Wiley Online Library, 2007.
- [28] Jiachen Li, Hengbo Ma, and Masayoshi Tomizuka. Conditional generative neural system for probabilistic trajectory prediction. In *2019 IEEE Conference on Robotics and Systems (IROS)*, 2019.
- [29] Jiachen Li, Hengbo Ma, and Masayoshi Tomizuka. Interaction-aware multi-agent tracking and probabilistic behavior prediction via adversarial learning. In *2019 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2019.
- [30] Yuexin Ma, Xinge Zhu, Sibozhang, Ruigang Yang, Wenping Wang, and Dinesh Manocha. Trafficpredict: Trajectory prediction for heterogeneous traffic-agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6120–6127, 2019.
- [31] Srikanth Malla and Chiho Choi. Nemo: Future object localization using noisy ego priors. *arXiv preprint arXiv:1909.08150*, 2019.
- [32] Marcin Marszałek, Ivan Laptev, and Cordelia Schmid. Actions in context. In *CVPR 2009-IEEE Conference on Computer Vision & Pattern Recognition*, pages 2929–2936. IEEE Computer Society, 2009.
- [33] Raúl Quintero Mínguez, Ignacio Parra Alonso, David Fernández-Llorca, and Miguel Ángel Sotelo. Pedestrian path, pose, and intention prediction through gaussian process dynamical models and pedestrian activity recognition. *IEEE Transactions on Intelligent Transportation Systems*, 20(5):1803–1814, 2018.
- [34] Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, JK Aggarwal, Hyungtae Lee, Larry Davis, et al. A large-scale benchmark dataset for event recognition in surveillance video. In *CVPR 2011*, pages 3153–3160. IEEE, 2011.
- [35] Seong Hyeon Park, ByeongDo Kim, Chang Mook Kang, Chung Choo Chung, and Jun Won Choi. Sequence-to-sequence prediction of vehicle trajectory via lstm encoder-decoder architecture. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1672–1678. IEEE, 2018.
- [36] Abhishek Patil, Srikanth Malla, Haiming Gang, and Yi-Ting Chen. The h3d dataset for full-surround 3d multi-object detection and tracking in crowded urban scenes. *arXiv preprint arXiv:1903.01568*, 2019.
- [37] Stefano Pellegrini, Andreas Ess, and Luc Van Gool. Improving data association by joint modeling of pedestrian trajectories and groupings. In *European conference on computer vision*, pages 452–465. Springer, 2010.
- [38] 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 *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [39] 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*, pages 206–213, 2017.
- [40] Nicholas Rhinehart, Kris M Kitani, and Paul Vernaza. R2p2: A reparameterized pushforward policy for diverse, precise generative path forecasting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 772–788, 2018.
- [41] Alexandre Robicquet, Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *European conference on computer vision*, pages 549–565. Springer, 2016.
- [42] Christoph Schöller, Vincent Aravantinos, Florian Lay, and Alois Knoll. The simpler the better: Constant velocity for pedestrian motion prediction. *arXiv preprint arXiv:1903.07933*, 2019.
- [43] Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *European Conference on Computer Vision*, 2016.
- [44] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014.
- [45] Gurkirt Singh, Suman Saha, and Fabio Cuzzolin. Predicting action tubes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018.
- [46] Gurkirt Singh, Suman Saha, Michael Sapienza, Philip HS Torr, and Fabio Cuzzolin. Online real-time multiple spatiotemporal action localisation and prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3637–3646, 2017.
- [47] Khurram Soomro, Haroon Idrees, and Mubarak Shah. Predicting the where and what of actors and actions through online action localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2648–2657, 2016.
- [48] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [49] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-centric relation network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 318–334, 2018.
- [50] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Rahul Sukthankar, Kevin Murphy, and Cordelia Schmid. Relational action forecasting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 273–283, 2019.
- [51] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.
- [52] Gül Varol, Ivan Laptev, and Cordelia Schmid. Long-term temporal convolutions for action recognition. *IEEE*

transactions on pattern analysis and machine intelligence, 40(6):1510–1517, 2017.

- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [54] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–7. IEEE, 2018.
- [55] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018.
- [56] Yanyu Xu, Zhixin Piao, and Shenghua Gao. Encoding crowd interaction with deep neural network for pedestrian trajectory prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5275–5284, 2018.
- [57] Hao Xue, Du Q Huynh, and Mark Reynolds. Ss-lstm: A hierarchical lstm model for pedestrian trajectory prediction. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1186–1194. IEEE, 2018.
- [58] Yu Yao, Mingze Xu, Chiho Choi, David J Crandall, Ella M Atkins, and Behzad Dariush. Egocentric vision-based future vehicle localization for intelligent driving assistance systems. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9711–9717. IEEE, 2019.
- [59] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision*, 126(2-4):375–389, 2018.
- [60] Pu Zhang, Wanli Ouyang, Pengfei Zhang, Jianru Xue, and Nanning Zheng. Sr-lstm: State refinement for lstm towards pedestrian trajectory prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12085–12094, 2019.