

Human Trajectory Prediction via Counterfactual Analysis

Guangyi Chen^{1,2,3}, Junlong Li^{1,2,3}, Jiwen Lu^{1,2,3,*}, Jie Zhou^{1,2,3}

¹Department of Automation, Tsinghua University, China

²State Key Lab of Intelligent Technologies and Systems, China

³Beijing National Research Center for Information Science and Technology, China

chen-gyl6@mails.tsinghua.edu.cn; ljl原因.leo@gmail.com; {lujiwen, jzhou}@tsinghua.edu.cn

Abstract

Forecasting human trajectories in complex dynamic environments plays a critical role in autonomous vehicles and intelligent robots. Most existing methods learn to predict future trajectories by behavior clues from history trajectories and interaction clues from environments. However, the inherent bias between training and deployment environments is ignored. Hence, we propose a **counterfactual analysis** method for human trajectory prediction to investigate the causality between the predicted trajectories and input clues and alleviate the negative effects brought by the environment bias. We first build a causal graph for trajectory forecasting with history trajectory, future trajectory, and the environment interactions. Then, we cut off the inference from environment to trajectory by constructing the counterfactual intervention on the trajectory itself. Finally, we compare the factual and counterfactual trajectory clues to alleviate the effects of environment bias and highlight the trajectory clues. Our counterfactual analysis is a plug-and-play module that can be applied to any baseline prediction methods including RNN- and CNN-based ones. We show that our method achieves consistent improvement for different baselines and obtains the state-of-the-art results on public pedestrian trajectory forecasting benchmarks.¹

1. Introduction

Human trajectory prediction aims at forecasting the future trajectory of pedestrians based on their past positions in complex and crowd environments. It is a critical and fundamental task for many applications, including the planning and controlling of the autonomous vehicles, the robot navigation, and the tracking and re-identification in the crowd surveillance [39, 7, 8, 3, 9]. Thanks to these significances,

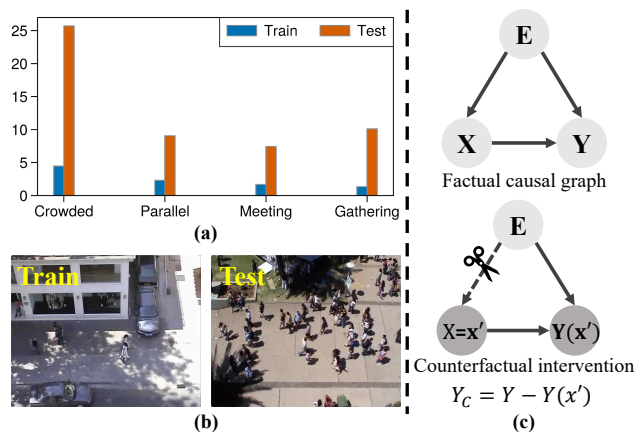


Figure 1. The motivation and key idea of our approach. In part (a), we show the bias between training and testing environments of the UNIV scene. We respectively compare the average numbers of neighbors, parallel trajectories, trajectories that pedestrians meet, and trajectories that pedestrians gather. In part (b), we visualize the training and testing environments. We can observe obvious difference between two environments. In part (c), we first build a causal graph with nodes E (environment interaction), X (history trajectory), and Y (future trajectory). Then, at the bottom of (c), we conduct the counterfactual intervention on the X history trajectory to cut off the dependence between history trajectory and environment interaction. The causal prediction Y_C is obtained by computing the difference between the original and counterfactual predictions to alleviate the negative effects of environment bias.

human trajectory prediction task has attracted much attention over the past few years [1, 53, 56, 15, 25].

Despite the recent progress, trajectory prediction is still a challenging problem due to complex social or physical environment interactions. In crowded environment, the pedestrians are always interacted with the social behavior from other pedestrians and the common scene context. For example, pedestrians may walk in parallel for social talking or stop for seconds to avoid collisions. Besides, the behaviors of pedestrians may be interacted by the traffic light, crosswalk, or just an obstacle like a tree.

Most existing methods concentrate on modeling the en-

*Corresponding author

¹Code and a video demo are available at <https://github.com/CHENGYI12/CausalHTP>

environment interactions and aggregate these interaction clues with history behavior clues for trajectory prediction. For example, Social LSTM [1] applies a social pooling module to extract clues from environment interaction. While Social-STGCNN [29] utilizes the spatial-temporal graph CNN to encode the environment interactions. However, these methods always ignore the inherent bias between training and deployment environments. As shown in part (a) of Figure 1, we statistically compare the interactions on training and testing environments of the UNIV scene. We observe an obvious gap between the interactions of different environments. We also visualize the environment difference in part (b) of Figure 1, where the training environment is a street scene while the testing environment is a more crowded public square. This environment bias causes heavy overfitting of complex environment interaction modules in most of trajectory prediction methods. When the training data contains many examples who turns left in the crossroads, the prediction of left trajectory will be wrongly attributed to crossroads, which misleads the prediction in the new environment where pedestrians turn right in the crossroads. Besides, deployment environments are always unpredictable in the applications of autonomous vehicles and intelligent robots. Thus, it is difficult to apply transfer learning methods to reduce environment bias.

To address this problem, in this paper, we propose a counterfactual analysis method to alleviate the overdependence of environment bias and highlight the trajectory clues itself. Inspired by the causal inference methods [37, 36, 12], we propose to use the counterfactual intervention to investigate the causality between the observed clues and predicted trajectories. Different from conventional causal inference methods [34, 10], we furthermore extend these to the training process for the optimization of the prediction model. Specifically, we first construct a factual causal graph by the human prior knowledge, whose nodes include past trajectory, environment interaction, and future trajectory. As shown in the left part of Figure 1 (b), environment interaction may be a negative confounder due to the bias between training and deployment environments. Then, we conduct the counterfactual intervention on the history trajectory, which cuts off the dependence between the environment and trajectory. Motivated by [48, 45], we replace the history trajectory feature with the counterfactual one, such as uniform rectilinear motion, mean trajectory, or random trajectory. This counterfactual prediction indicates the effect of biased environment clues. Finally, we subtract the counterfactual prediction from original prediction as the causality-aware prediction since the negative effect of confounder is alleviated. We highlight that the proposed counterfactual analysis method is a plug-and-play module which can be applied to any baseline prediction method including RNN- and CNN-based ones. In the experiments, we ap-

ply our method for two baseline methods including RNN-based STGAT [18] and CNN-based Social-STGCNN [29]. We show that our counterfactual analysis method achieves consistent improvement for both two baseline methods and obtains the state-of-the-art results on public pedestrian trajectory prediction datasets. We summarize three advantages of our method as follows:

- We show that the environment bias may cause the heavy overfitting of the complex environment interaction modules.
- We propose a counterfactual analysis method to alleviate the overdependence of environment bias and highlight the trajectory clues itself.
- Our counterfactual analysis is a simply plug-and-play module which can be easily applied to any baseline predictor, and consistently improves the performance on many human trajectory prediction benchmarks.

2. Related Work

Environment Interaction: The environment interaction consists of social environment interaction and physical environment interaction. To model the complex social environment interaction, previous methods always model human crowd motion with handcrafted rules or energy parameters, e.g. Social Force Model [17], continuum dynamics [47], Discrete Choice framework [2], Gaussian processes [46, 51] and crowd analysis [40, 54, 23]. In deep learning methods, early studies [1, 15] employ the pooling module to capture social interactions. While some methods [49, 11, 41, 55, 6] apply the attention model to distinguish the importance of different neighbors. Many recent methods employ the graph model to extract the clues from interactions, e.g. the States Refinement module for message passing in SR-LSTM [56]; the Graph Attention model for modeling interactions in [18, 21]; STR-GGRNN [16] uses data-driven adaptive online neighborhood recommendation. In addition, some methods [22, 41, 53, 21, 4] also incorporate physical environment interactions of the scene context for more reliable prediction. However, most of these methods ignore the environment bias between training and deployment environments, which causes heavy overfitting of the environment interaction modules. To address this, we propose a counterfactual analysis method to focus more on trajectory itself, instead of the biased interaction.

Predictor: Deep learning based methods always regard the trajectory forecasting as the sequence prediction problem and apply the RNN model [1, 15, 21, 56, 41] or temporal CNN model [31, 29] as the predictor. For example, Social LSTM [1] employs the LSTM to encode human motion, then fuses the environment interaction clues extracted by a pooling module and feeds them into a LSTM

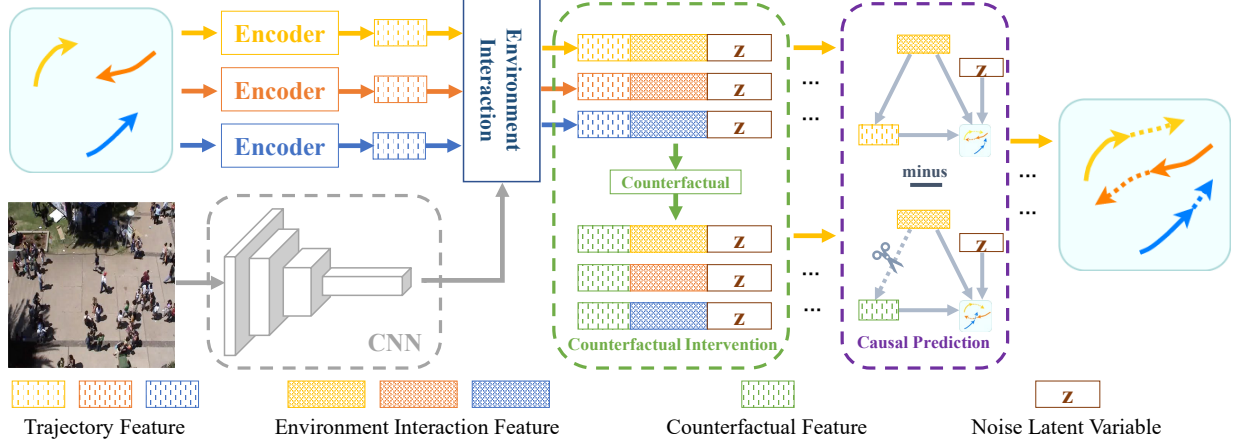


Figure 2. Training process of our counterfactual analysis method. We apply the counterfactual intervention by replacing the features of past trajectory with the counterfactual features such as uniform rectilinear motion (zero), mean trajectory, or random trajectory. The counterfactual prediction denotes the biased effect from environment confounder. To alleviate the negative effect of environment bias, we subtract the counterfactual prediction from original prediction as the final causal prediction. (Best viewed in color.)

decoder to sequentially predict the trajectory, while Nikhil *et al.* [31] first proposed to use the temporal CNN model to replace the RNN-based predictors. Besides, [13] proposes to use Transformers for predicting the future trajectories. However, both RNN-based and CNN-based network are likelihood-based predictor whose causalities between observed clues and predicted results are still non-transparent. In our counterfactual analysis, we can investigate the effects of different clues including history trajectory and environment interactions by the intervention.

Generative Modeling: Another popular research direction of human trajectory prediction is exploring the indeterminacy of the human behavior. Many methods [15, 19, 21, 57, 18, 43] employ generative model to predict multiple plausible paths, instead of the deterministic one. These methods employ Generative Adversarial Network (GAN) [14] to implicitly incorporate a latent variable into the encoded embedding [15, 21, 41, 57], or Variational Auto-Encoder (VAE) [20] to explicitly model the multi-modal distribution with a latent variable [22, 19, 42, 28]. While STGCNN [29] directly predicts a bi-variate Gaussian distribution to replace some deterministic trajectories with multi-modal predictions.

Causal Inference: Different from conventional associated inference, the causal inference aims to investigate the effect of variables when some cause is changed [36, 35]. Recently, many methods take effort to apply causal inference for the effect analysis of deep neural network in different fields including reinforcement learning [30, 12], natural language processing [50, 32], and visual representation learning [5, 27, 52, 45]. The clear causalities can improve the transparency of the deep models as most of them are treated as black-box nowadays. In this paper, we propose to apply the counterfactual analysis to the trajectory prediction

system to analyze the effects of different clues and enhance the reliability, which attempts to explain how neural networks make the prediction and optimize the prediction by removing the affects from biased dataset. It is helpful to understand the causal relation between the input clues and output trajectory predictions.

3. Approach

3.1. Problem Definition

Pedestrian trajectory prediction task can be defined as a sequential prediction problem, whose inputs include the history trajectory and environment interaction. Many method designs complex model to learn the clues from environment interactions. However, the inherent bias between training and deployment environments may cause the overfitting of interaction model. In the conventional trajectory prediction framework, given N pedestrians in a scene, we can define the history trajectory of the i th pedestrian as $X_i = \{(x_i^t, y_i^t) \in R^2 | t = 1, 2, \dots, t_{obs}\}$, where the (x_i^t, y_i^t) is the 2D location at time t . The ground-truth future trajectory of the i th pedestrian can be defined as $Y_i = \{(x_i^t, y_i^t) \in R^2 | t = t_{obs} + 1, t_{obs} + 2, \dots, t_{pred}\}$. For the normalization of scale, many methods employ the relative locations even the relative speeds to replace the absolute locations. The environment interactions can be extracted from the trajectories including other pedestrians and the physical scene context as $E_i = \mathcal{E}_\phi(\{X_i\}_{i=1}^N, C)$, where C denotes the physical scene context. The sequential trajectory prediction process is modeled as

$$\hat{Y}_i = \mathcal{F}_\theta(X_i, E_i), \quad (1)$$

where $\hat{Y}_i = \{(\hat{x}_i^t, \hat{y}_i^t) \in R^2 | t = t_{obs} + 1, t_{obs} + 2, \dots, t_{pred}\}$ denotes the predicted trajectory.

Given a set of training pedestrian trajectory data $\{(X_i, Y_i)\} \in \Omega$, the predictor is optimized by the L2 loss:

$$\mathcal{L}_{tra}(\theta, \phi | \Omega) = L_{L2}(Y_i, \hat{Y}_i), \quad (2)$$

where θ and ϕ denote the parameters of the pedestrian trajectory prediction system.

3.2. Causal Graph

In this subsection, we reformulate above trajectory prediction framework with the causal graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ with the prior knowledge. As shown in part (c) of Figure 1, the nodes \mathcal{V} in the graph denote the variables including history trajectory X , environment interactions E (including social and physical environment interactions), and future trajectory Y . While the causal links \mathcal{E} indicate the hidden causal relations and how these variables interact with each other. In a link $E \rightarrow X \rightarrow Y$, we can call the node E as the parent node of X , while Y is the child node of X . The link from the parent node to the child node denotes a causality. $E \rightarrow (X, Y)$ indicates all the human trajectories are influenced by the environment. In the causal inference theory [33], when a variable (E) simultaneously affects two variables, it becomes a confounder in the causal analysis of these two variables. For example, if turning left is often observed with the environment interaction of crossroads, the prediction of left trajectory will be wrongly attributed to crossroads. As a result, the real causality underlying the history trajectory may be ignored.

3.3. Counterfactual Analysis

For conventional likelihood-based predictors, the causalities between different observed clues and predicted results are non-transparent. The prediction model is easy to be “cheated” by the short cut between the biased environment interaction and final future trajectory. Inspired by the causal inference methods [33, 34] which attempt to analyze the causalities among different clues, we propose to apply the counterfactual analysis to trajectory prediction system to mitigate the negative effect of biased environment and encourage the model to focus more on trajectory itself. Counterfactual intervention means to imagine a non-existent observation to replace the original factual clues in the trajectory prediction system. In the causal inference methods [37], the intervention is formulated as $do(\cdot)$. Once one variable is intervened, its all in-coming links in the causal graph are cut off and its value is independently given, while other variables who are not affected still maintain the original value. In our method, as shown in part (c) of Figure 1, we replace the history trajectory features in the system with the imagined ones. We can obtain the prediction results under this intervention as:

$$\hat{Y}_{X_i=x'} = \mathcal{F}_\theta(do(X_i = x'), E_i), \quad (3)$$

where $X = x'$ indicates the counterfactual value. Specifically, we can apply the different counterfactual interventions such as the zero vector (which denotes uniform rectilinear motion when the relative speed is used as input), the mean vector of all history trajectories or the random trajectories as the counterfactual intervention.

The original factual prediction \hat{Y}_i is dependent on both trajectories and environment interactions, while the counterfactual predicted result $\hat{Y}_{X_i=x'}$ is only dependent with environment interactions since the trajectory is replaced by counterfactual intervention. To investigate the real effect of the trajectory itself, we compute the difference between the counterfactual prediction and factual prediction and name it as the causal prediction:

$$\hat{Y}_{causal} = \hat{Y}_i - \hat{Y}_{X_i=x'}. \quad (4)$$

Compared with the original likelihood prediction, the causal prediction is more reliable by avoiding the biased effects from environment confounders. In the training process, we optimize the network to predict the causal prediction:

$$\mathcal{L}_{Causal}(\theta, \phi | \Omega) = L_{L2}(Y_i, \hat{Y}_{causal}). \quad (5)$$

3.4. Generative Model

It is worth noting that our counterfactual analysis is robust to the slight change of trajectory prediction systems, such as different predictors, whether the physical environment content is used, or whether the generative model is used. These changes cannot affect the structure of the causal graph. When using the generative methods, we will add the noise latent variable in the trajectory prediction system and predict the stochastic future trajectory. Taking the GAN as the example of generative model, the loss function in (2) is reformulated as

$$\begin{aligned} \mathcal{L}_{GAN}(\theta, \phi, \psi | \Omega) &= L_{L2}(Y_i, \hat{Y}_i) + \log D_\psi(Y_i) \\ &\quad + \log(1 - D_\psi(\hat{Y}_i)), \end{aligned} \quad (6)$$

where $\hat{Y}_i = \mathcal{F}_\theta(X_i, E_i, Z)$ indicates the trajectory is generated with a noise latent variable Z for multi-modal distributions. While D_ψ is the discriminator in the GAN to judge whether the trajectory is generated with the noise latent variable. Our counterfactual analysis can work together with generative methods, which only needs to add the noise latent variable Z in the side of predicted trajectory Y with $Z \rightarrow Y$. The causal prediction with the latent variable is formulated as:

$$\begin{aligned} \hat{Y}_{causal} &= \hat{Y}_i(X_i, E_i, Z) \\ &\quad - \hat{Y}_{X_i=x'}(do(X_i = x'), E_i, Z). \end{aligned} \quad (7)$$

Then, we only need to replace the original prediction with causal prediction as:

$$\mathcal{L}_{CausalGAN}(\theta, \phi, \psi | \Omega) = L_{L2}(Y_i, \hat{Y}_{causal}) + \log D_{\psi}(Y_i) + \log(1 - D_{\psi}(\hat{Y}_{causal})). \quad (8)$$

3.5. Implementation Details

As shown in Figure 2, the framework of trajectory prediction system always contains 4 modules including history path encoder, scene perceptron, interactions analysis module, and trajectory predictor. Instead of the conventional likelihood-based trajectory predictor, in this paper, we conduct the counterfactual intervention for some clues and apply a causal predictor by computing the difference between the prediction results with factual and counterfactual clues. Our causal prediction is simple yet effective for the reliability of trajectory prediction system. To evaluate the generality of our counterfactual analysis method, we incorporate it as an plug-and-play module into two different baseline methods including RNN-based STGAT [18] and CNN-based Social-STGCNN [29]. We simply introduce these two implementations in the following. Note that though the whole counterfactual analysis method contains the features of environment scene, here we **do not** use the scene features in the implementations since our baseline models do not use scene images.

Causal-STGAT: The STGAT method employs two LSTM (M-LSTM motion encoding and G-LSTM for graph) and a graph attention network (GAT) to encode the history trajectory and social interaction clues. In Causal-STGAT, we replace the feature vector from M-LSTM with the counterfactual intervention (the mean vector of all trajectory features, or the zero vector). Then we respectively conduct predictions with original connected feature and counterfactual connected feature whose history trajectory part is changed. Finally, we adopt the difference as the prediction.

Causal-STGCNN: In Social-STGCNN, the trajectories in a scene are represented as a graph $G = \{N, E\}$, where nodes N denote the past trajectories, while the adjacency matrix A in the GCN denotes the social interactions. Then a time-extrapolator (TE) CNN is applied as the trajectory predictor. In Causal-STGCNN, we replace the original nodes N by counterfactual intervention and maintain the values of interaction clues A . Then we respectively feed factual and counterfactual outputs to the TE-CNN predictor and compute the difference to get the causal prediction.

4. Experiments

4.1. Datasets and Experimental Settings

Datasets: We conduct experiments on two publicly available human trajectory prediction datasets: ETH [38]

and UCY [24]. The human trajectories in both two datasets are captured in the real world scenes with rich social interactions. These datasets contain five unique scenes: Zara1, Zara2, Univ, ETH, and Hotel with 1536 detected pedestrians. All trajectories in datasets are sampled every 0.4 seconds. The experimental settings follow the previous methods like Social-GAN [15], Social BiGAT [21], and STGAT [18]. We also use the leave-one-out approach to train and validate our model on 4 sets and test on the remaining one. During evaluation, the first 3.2 seconds (8 frames) are observed and the next 4.8 seconds (12 frames) are to be predicted.

Evaluation Metrics: We adopt the same evaluation metrics as prior methods [15, 21, 18], including Average Displacement Error (ADE) and Final Displacement Error (FDE). ADE computes the mean square error (MSE) of predicted trajectory and ground-truth trajectory, while FDE computes the L2 distance between the final locations of predicted and ground-truth trajectories. Since both baselines Social-STGCNN [29] and STGAT [18] are generative methods, we also follow the evaluation method in Social-GAN [15]. We generate 20 samples based on the predicted distribution and select the closest sample to the ground truth for ADE and FDE metrics.

4.2. Quantitative Evaluation

Evaluation of Counterfactual Analysis: To verify the effectiveness of our method, we compare the performance of our Causal-STGAT, Causal-STGCNN with the RNN-based baseline STGAT [18] and CNN-based Social-STGCNN [29] methods. The performance comparison about our causal-based methods, baseline methods, and other SOTA methods is summarized as Table 1. Note that the * in Table 1 means our reproduced results trained with the official released code.² Due to different implementation environments, the performance of our reproduced baselines are slightly lower than the performance reported in the paper. For the fair comparison, all hyperparameters and environments of our Causal-STGAT and Causal-STGCNN methods are same with the reproduced baselines. As shown in Table 1, our causal prediction method can consistently improve the performance than the original likelihood-based prediction with different baselines on most of datasets. Specifically, Causal-STGAT obtains +0.07/+0.16 ADE/FDE improvement than the baseline STGAT method on the average of 5 scenes, while Causal-STGCNN outperforms the reproduced Social-STGCNN method by over +0.02/+0.09 ADE/FDE results.

Besides, we also compare the proposed causal based

²We reproduced the STGAT [18] methods with the released code from <https://github.com/huang-xx/STGAT>, while the code of Social-STGCNN [29] comes from <https://github.com/abdullahmohamed/Social-STGCNN>.

Table 1. Comparison with several state-of-the-art models. We notice that our causal-based methods have the best average error for RNN-based and CNN-based methods. The lower is the better.

RNN-based Method	Performance (ADE/FDE)					
	ETH	HOTEL	ZARA1	ZARA2	UNIV	AVG
LSTM	1.09/2.41	0.86/1.91	0.41/0.88	0.52/1.11	0.61/1.31	0.70/1.52
S-LSTM [1]	1.09/2.35	0.79/1.76	0.47/1.00	0.56/1.17	0.67/1.40	0.72/1.54
SGAN [15]	0.81/1.52	0.72/1.61	0.34/0.69	0.42/0.84	0.60/1.26	0.58/1.18
Sophie [41]	0.70/1.43	0.76/1.67	0.30/0.63	0.38/0.78	0.54/1.24	0.54/1.15
SR-LSTM [56]	0.63/1.25	0.37/0.74	0.41/0.90	0.32/0.70	0.51/1.10	0.45/0.94
Social-BiGAT [21]	0.69/1.29	0.49/1.01	0.30/0.62	0.36/0.75	0.55/1.32	0.48/1.00
MATF [57]	1.33/2.49	0.51/0.95	0.44/0.93	0.34/0.73	0.56/1.19	0.64/1.26
MATF GAN [57]	1.01/1.75	0.43/0.80	0.26/ 0.45	0.26/0.57	0.44/0.91	0.48/0.90
IDL [54]	0.59 /1.30	0.46/0.83	0.22 /0.49	0.23 / 0.55	0.51/1.27	0.40 /0.89
PIF [26]	0.73/1.65	0.30 /0.59	0.60/1.27	0.38/0.81	0.31 / 0.68	0.46/1.00
STGAT [18]	0.65/1.12	0.35/0.66	0.34/0.69	0.29/0.60	0.52/1.10	0.43/0.83
STGAT*	0.73/1.39	0.38/0.72	0.35/0.69	0.32/0.64	0.57/1.22	0.47/0.93
Causal-STGAT	0.60/ 0.98	0.30 / 0.54	0.32/0.64	0.28/0.58	0.52/1.10	0.40 / 0.77

CNN-based Method	Performance (ADE/FDE)					
	ETH	HOTEL	ZARA1	ZARA2	UNIV	AVG
CNN [31]	1.04/2.07	0.59/1.17	0.43/0.90	0.34/0.75	0.57/2.32	0.59/1.22
Social-STGCNN [29]	0.64 /1.11	0.49/0.85	0.34 / 0.53	0.30 /0.48	0.44 / 0.79	0.44/0.75
Social-STGCNN*	0.66/1.12	0.40/0.65	0.35/0.56	0.31/ 0.47	0.51/0.93	0.45/0.75
Causal-STGCNN	0.64 / 1.00	0.38 / 0.45	0.34 / 0.53	0.32/0.49	0.49/0.81	0.43 / 0.66

Table 2. Evaluation of Causal-STGAT for conducting different counterfactual interventions.

Method	Performance (ADE/FDE)					
	ETH	HOTEL	ZARA1	ZARA2	UNIV	AVG
STGAT* (baseline)	0.73/1.39	0.38/0.72	0.35/0.69	0.32/0.64	0.57/1.22	0.47/0.93
Causal-STGAT-Zero	0.60/0.98	0.30/0.54	0.32/0.64	0.28/0.58	0.52/1.10	0.40/0.77
Causal-STGAT-Mean	0.65/1.17	0.33/0.60	0.35/0.71	0.29/0.60	0.54/1.15	0.44/0.84
Causal-STGAT-Random	0.59/1.05	0.31/0.55	0.34/0.70	0.28/0.59	0.53/1.13	0.42/0.80

Table 3. Parameters size and inference time of baseline models compared to ours. The lower is better.

Method	Social-STGCNN	Causal-STGCNN	STGAT	Causal-STGAT
Inference Speed	0.0116	0.0124	0.3343	0.3418
Parameters Count	7.6k	7.6k	56k	56k

method with other SOTA methods, such as IDL [54], Social-BiGAT [21] and MATF [57]. As shown in Table 1, we achieves the SOTA performance with slight improvement over the RNN-based IDL [54] method and CNN-based Social-STGCNN [29] method. It also demonstrates the effectiveness of the proposed counterfactual analysis method. For both our methods in Table 1, we conducted the counterfactual intervention by replacing the features with zero vector. Other attempts are introduced below.

Evaluation of Counterfactual Implementations: We have attempted different implementations of counterfactual intervention, including using the zero vector, the mean of all feature vector or a random vector sampled from a uniform distribution with $[-0.1, 0.1]$. These intervention implemen-

tations (“Zero”, “Mean”, “Random”) with limited computing cost are usually adopted [44] for causal model. For the zero vector and mean vector, the counterfactual intervention in the training and testing stage is invariant. While for random vector, we sampled the vector from a uniform distribution in the training stage and using the expectation of uniform distribution (zero vector) in the testing stage to avoid introducing the bias for testing data. As shown in bottom part of Table 2, all implementations obtain obvious improvements for the baselines, which demonstrates the generality of the proposed counterfactual analysis method. Furthermore, 1) The “Zero” obtains slightly better performance than the others. It might be because the “Zero” is a harder counterfactual intervention, which better emphasizes

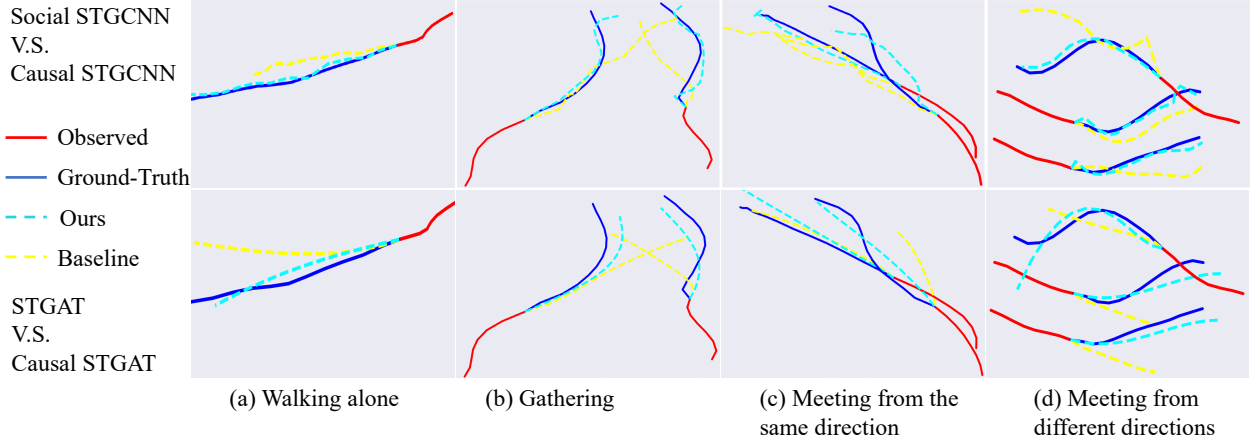


Figure 3. Qualitative analysis for our causal prediction methods. We compare our causal prediction methods (Causal-STGCNN and Causal-STGAT) with corresponding baselines in four scenarios: (a) a person walking alone; (b) two pedestrians gathering; (c) two persons meeting from the same direction; (d) two persons meeting from different directions.

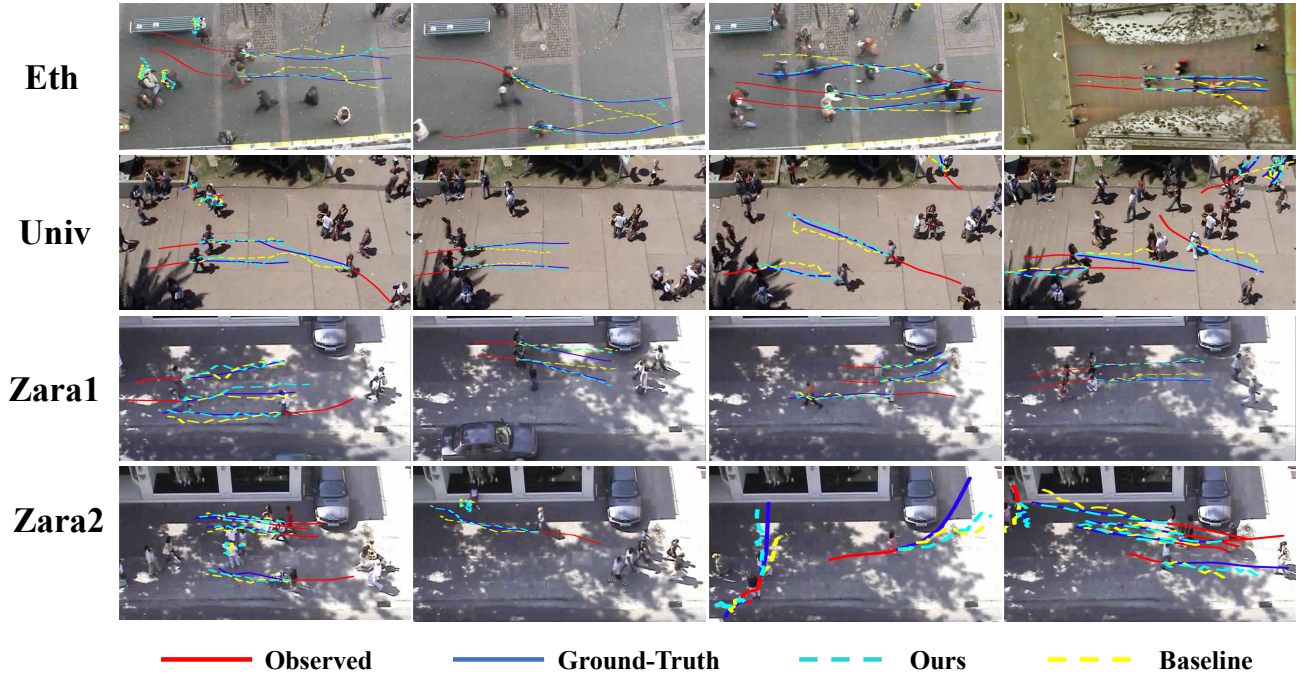


Figure 4. Visualization examples of our Causal-STGCNN method in different scenes. We compare the trajectories generated by our Causal-STGCNN method and baseline Social-STGCNN method in the ETH and UCY datasets. Our causal-based method can solve complex scenes and social interactions with multiple humans by the counterfactual analysis. (Best viewed in color.)

the causal effect from history behaviors to future trajectories. 2) The performance of these implementations is close, which demonstrates the robustness of the implementation of counterfactual intervention.

Evaluation of Inference Speed and Model Size: Model size and inference speed are also critical for the deployment of the method in the real-world environment. We evaluated the speed of all models with all data in five scenes and compute the average inference time for one trajectory on one GTX 2080Ti GPU. For fairness, we fixed the batch-

size as 1 for all methods and repeat the evaluation 3 times for average. The parameters cost and inference speed of our methods and baselines are summarized in Table 3. First, our counterfactual analysis method does not need any extra parameters since all parameters are shared with factual and counterfactual parts. Second, there is no such thing as a free lunch. The speed of our methods are lower than the baselines because of the extra computations for counterfactual analysis. However the extra speed cost is not heavy since the counterfactual analysis only uses part of network.

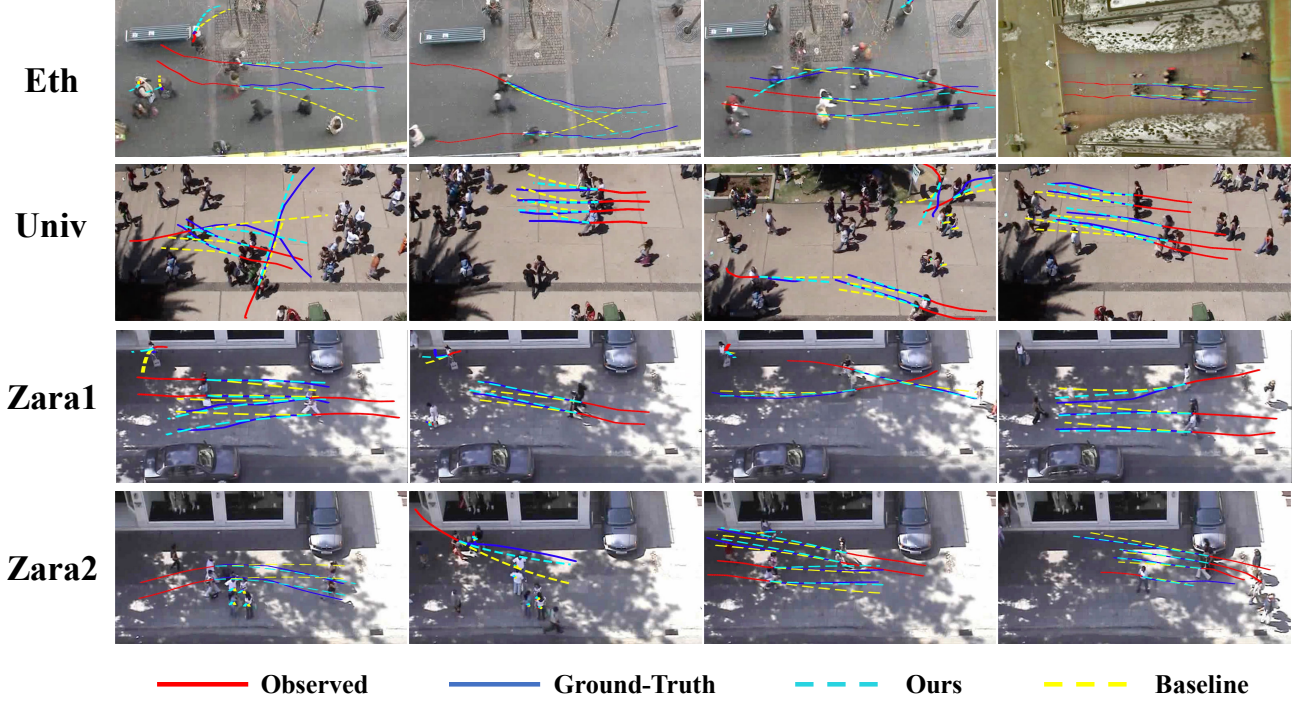


Figure 5. Visualization examples of our Causal-STGAT method and baseline Social-STGAT method in the different scenes in the ETH and UCY datasets. The comparisons quantitatively demonstrate the effectiveness of our method on the RNN-based baselines.

4.3. Qualitative Evaluation

We qualitatively analyze how our causal predictions are more reliable than conventional likelihood-based predictions. As shown in Figure 3, we analyze our methods in 4 different scenarios. Taking the (b) as an example, given two trajectories gathering, the conventional likelihood-based methods always predict they will meet each other since the negative effect brought by the training bias of the environment interaction, (like “most trajectories getting closer finally meet”). These training biases mislead the predictor to learn the spurious correlations of environments and ignore the real clues from history trajectory, (like “these persons tend to gather and move forward together”). By counterfactual analysis, we remove the negative effect contained in the counterfactual prediction from original predictions. It is effective to overcome the training bias and encourage the model to highlight the real causation itself.

Besides, we also provide the visualization examples with the environment scenes. As shown in Figure 4 and Figure 5, we respectively compare our Causal-STGCNN and Causal-STGAT methods with corresponding baselines in different scenes. We observe that our counterfactual analysis method can effectively captures the real causal relations instead of the biased environment interactions. Taking the second scene in ETH as an example, both CNN-based and RNN-based counterfactual analysis methods significantly outperform the baseline methods and generate more reliable predicted trajectories.

5. Conclusion

In this paper, we have presented a counterfactual analysis method to investigate the causalities between the observed clues and predicted trajectories. We apply the counterfactual intervention by replacing the features of trajectory clues with the counterfactual ones, and subtract this counterfactual prediction from original prediction. By this comparison, we encourage the model to learn the real causations of trajectory and alleviate the negative effects brought by the bias between training and deployment environments. Our counterfactual analysis is a plug-and-play module, which can be employed for off-the-shell trajectory prediction models. In the experiments, we demonstrate the effectiveness of our counterfactual analysis method for different scenes, analyze the effects of different counterfactual implementations, and evaluate the generalization ability for different baseline methods.

Acknowledgements

This work was supported in part by the National Key Research and Development Program of China under Grant 2018AAA0102803, in part by the National Natural Science Foundation of China under Grant 61822603, Grant U1813218, and Grant U1713214, in part by a grant from the Beijing Academy of Artificial Intelligence (BAAI), and in part by a grant from the Institute for Guo Qiang, Tsinghua University.

References

- [1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *CVPR*, pages 961–971, 2016.
- [2] Gianluca Antonini, Michel Bierlaire, and Mats Weber. Discrete choice models of pedestrian walking behavior. *Transportation Research Part B: Methodological*, 40(8):667–687, 2006.
- [3] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *ICCV*, pages 6182–6191, 2019.
- [4] Huikun Bi, Ruisi Zhang, Tianlu Mao, Zhigang Deng, and Zhaoqi Wang. How can i see my future? fvtraj: Using first-person view for pedestrian trajectory prediction. In *ECCV*, 2020.
- [5] Krzysztof Chalupka, Pietro Perona, and Frederick Eberhardt. Visual causal feature learning. In *UAI*, 2015.
- [6] Rohan Chandra, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. Taphic: Trajectory prediction in dense and heterogeneous traffic using weighted interactions. In *CVPR*, pages 8483–8492, 2019.
- [7] Changan Chen, Yuejiang Liu, Sven Kreiss, and Alexandre Alahi. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *ICRA*, pages 6015–6022, 2019.
- [8] Guangyi Chen, Chunze Lin, Liangliang Ren, Jiwen Lu, and Zhou Jie. Self-critical attention learning for person re-identification. In *ICCV*, 2019.
- [9] Guangyi Chen, Jiwen Lu, Ming Yang, and Jie Zhou. Spatial-temporal attention-aware learning for video-based person re-identification. *TIP*, 28(9):4192–4205, 2019.
- [10] Liangji Fang, Qinhong Jiang, Jianping Shi, and Bolei Zhou. Tpnnet: Trajectory proposal network for motion prediction. In *CVPR*, 2020.
- [11] Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Soft+ hardwired attention: An lstm framework for human trajectory prediction and abnormal event detection. *Neural Network*, 108:466–478, 2018.
- [12] Andrew Forney, Judea Pearl, and Elias Bareinboim. Counterfactual data-fusion for online reinforcement learners. In *ICML*, pages 1156–1164, 2017.
- [13] Francesco Giuliari, Irtiza Hasan, Marco Cristani, and Fabio Galasso. Transformer networks for trajectory forecasting. In *ICPR*, pages 10335–10342, 2021.
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014.
- [15] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *CVPR*, pages 2255–2264, 2018.
- [16] Sirin Haddad and Siew-Kei Lam. Self-growing spatial graph networks for pedestrian trajectory prediction. In *WACV*, pages 1151–1159, 2020.
- [17] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282, 1995.
- [18] Yingfan Huang, HuiKun Bi, Zhaoxin Li, Tianlu Mao, and Zhaoqi Wang. Stgat: Modeling spatial-temporal interactions for human trajectory prediction. In *ICCV*, pages 6272–6281, 2019.
- [19] Boris Ivanovic and Marco Pavone. The trajctron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *ICCV*, pages 2375–2384, 2019.
- [20] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- [21] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezaatofghi, and Silvio Savarese. Socialbigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. In *NeurIPS*, pages 137–146, 2019.
- [22] 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 *CVPR*, pages 336–345, 2017.
- [23] Namhoon Lee and Kris M Kitani. Predicting wide receiver trajectories in american football. In *WACV*, pages 1–9, 2016.
- [24] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer Graphics Forum*, volume 26, pages 655–664, 2007.
- [25] Junwei Liang, Lu Jiang, and Alexander Hauptmann. Simaug: Learning robust representations from simulation for trajectory prediction. In *ECCV*, 2020.
- [26] Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander G Hauptmann, and Li Fei-Fei. Peeking into the future: Predicting future person activities and locations in videos. In *CVPR*, pages 5725–5734, 2019.
- [27] David Lopez-Paz, Robert Nishihara, Soumith Chintala, Bernhard Scholkopf, and Léon Bottou. Discovering causal signals in images. In *CVPR*, pages 6979–6987, 2017.
- [28] 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 *ECCV*, 2020.
- [29] Abdullah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *CVPR*, 2020.
- [30] Suraj Nair, Yuke Zhu, Silvio Savarese, and Li Fei-Fei. Causal induction from visual observations for goal directed tasks. *arXiv preprint arXiv:1910.01751*, 2019.
- [31] Nishant Nikhil and Brendan Tran Morris. Convolutional neural network for trajectory prediction. In *ECCV*, pages 0–0, 2018.
- [32] Sunghyun Park, Seung-won Hwang, Fuxiang Chen, Jaegul Choo, Jung-Woo Ha, Sunghun Kim, and Jinyeong Yim. Paraphrase diversification using counterfactual debiasing. In *AAAI*, pages 6883–6891, 2019.
- [33] Judea Pearl. *Causality*. Cambridge University Press, 2009.
- [34] Judea Pearl. Direct and indirect effects. *arXiv preprint arXiv:1301.2300*, 2013.

- [35] Judea Pearl et al. Causal inference in statistics: An overview. *Statistics surveys*, 3:96–146, 2009.
- [36] Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.
- [37] Judea Pearl and Dana Mackenzie. *The book of why: the new science of cause and effect*. Basic Books, 2018.
- [38] Stefano Pellegrini, Andreas Ess, and Luc Van Gool. Improving data association by joint modeling of pedestrian trajectories and groupings. In *ECCV*, pages 452–465, 2010.
- [39] Scott Drew Pendleton, Hans Andersen, Xinxin Du, Xiaotong Shen, Malika Meghiani, You Hong Eng, Daniela Rus, and Marcelo H Ang. Perception, planning, control, and coordination for autonomous vehicles. *Machines*, 5(1):6, 2017.
- [40] Mikel Rodriguez, Josef Sivic, Ivan Laptev, and Jean-Yves Audibert. Data-driven crowd analysis in videos. In *ICCV*, pages 1235–1242, 2011.
- [41] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezaatofghi, and Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In *CVPR*, pages 1349–1358, 2019.
- [42] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *ECCV*, 2020.
- [43] Charlie Tang and Russ R Salakhutdinov. Multiple futures prediction. In *NeurIPS*, pages 15398–15408, 2019.
- [44] Kaihua Tang and et.al. Unbiased scene graph generation from biased training. In *CVPR*, pages 3716–3725, 2020.
- [45] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *CVPR*, 2020.
- [46] Meng Keat Christopher Tay and Christian Laugier. Modelling smooth paths using gaussian processes. In *Field and Service Robotics*, pages 381–390, 2008.
- [47] Adrien Treuille, Seth Cooper, and Zoran Popović. Continuum crowds. In *TOG*, volume 25, pages 1160–1168, 2006.
- [48] Tyler J VanderWeele. A three-way decomposition of a total effect into direct, indirect, and interactive effects. *Epidemiology (Cambridge, Mass.)*, 24(2):224, 2013.
- [49] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In *ICRA*, pages 1–7, 2018.
- [50] Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. Causal mediation analysis for interpreting neural nlp: The case of gender bias. *arXiv preprint arXiv:2004.12265*, 2020.
- [51] Jack M Wang, David J Fleet, and Aaron Hertzmann. Gaussian process dynamical models for human motion. *TPAMI*, 30(2):283–298, 2007.
- [52] Tan Wang, Jianqiang Huang, Hanwang Zhang, and Qianru Sun. Visual commonsense r-cnn. In *CVPR*, 2020.
- [53] Hao Xue, Du Q Huynh, and Mark Reynolds. Ss-lstm: A hierarchical lstm model for pedestrian trajectory prediction. In *WACV*, pages 1186–1194, 2018.
- [54] Kota Yamaguchi, Alexander C Berg, Luis E Ortiz, and Tamara L Berg. Who are you with and where are you going? In *CVPR*, pages 1345–1352, 2011.
- [55] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *ECCV*, 2020.
- [56] Pu Zhang, Wanli Ouyang, Pengfei Zhang, Jianru Xue, and Nanning Zheng. Sr-lstm: State refinement for lstm towards pedestrian trajectory prediction. In *CVPR*, pages 12085–12094, 2019.
- [57] Tianyang Zhao, Yifei Xu, Mathew Monfort, Wongun Choi, Chris Baker, Yibiao Zhao, Yizhou Wang, and Ying Nian Wu. Multi-agent tensor fusion for contextual trajectory prediction. In *CVPR*, pages 12126–12134, 2019.