

It is not the Journey but the Destination: Endpoint Conditioned Trajectory Prediction

Karttikeya Mangalam¹, Harshayu Girase¹, Shreyas Agarwal¹, Kuan-Hui Lee²,
Ehsan Adeli³, Jitendra Malik¹, Adrien Gaidon²

¹ University of California, Berkeley

² Toyota Research Institute

³ Stanford University

mangalam@cs.berkeley.edu

Abstract. Human trajectory forecasting with multiple socially interacting agents is of critical importance for autonomous navigation in human environments, e.g., for self-driving cars and social robots. In this work, we present Predicted Endpoint Conditioned Network (PECNet) for flexible human trajectory prediction. PECNet infers distant trajectory endpoints to assist in long-range multi-modal trajectory prediction. A novel non-local social pooling layer enables PECNet to infer diverse yet socially compliant trajectories. Additionally, we present a simple “truncation-trick” for improving diversity and multi-modal trajectory prediction performance. We show that PECNet improves state-of-the-art performance on the Stanford Drone trajectory prediction benchmark by $\sim 20.9\%$ and on the ETH/UCY benchmark by $\sim 40.8\%$ ⁴.

Keywords: Multimodal Trajectory Prediction, Social Pooling

1 Introduction

Predicting the movement of dynamic objects is a central problem for autonomous agents, be it humans, social robots [1], or self-driving cars [2]. Anticipation by prediction is indeed required for smooth and safe path planning in a changing environment. One of the most frequently encountered dynamic objects are humans. Hence, predicting human motion is of paramount importance for navigation, planning, human-robot interaction, and other critical robotic tasks. However, predicting human motion is nuanced, because humans are not inanimate entities evolving under Newtonian laws [3]. Rather, humans have the will to exert causal forces to change their motion and constantly adjust their paths as they navigate around obstacles to achieve their goals [4]. This complicated planning process is partially internal, and thus makes predicting human trajectories from observations challenging. Hence, a multitude of aspects should be taken into account beyond just past movement history, for instance latent predetermined goals, other moving agents in the scene, and social behavioral patterns.

⁴ Code available at project homepage: <https://karttikeya.github.io/publication/htf/>



Fig. 1. Imitating the Human Path Planning Process. Our proposed approach to model pedestrian trajectory prediction (top left) breaks down the task in two steps: (a) inferring the local endpoint distribution (top right), and then (b) conditioning on sampled future endpoints (bottom left) for jointly planning socially compliant trajectories for all the agents in the scene (bottom right).

In this work, we propose to address **human trajectory prediction by modeling intermediate stochastic goals we call endpoints**. We hypothesize that three separate factors interact to shape the trajectory of a pedestrian. First, we posit that pedestrians have some understanding of their long-term desired destination. We extend this hypothesis to sub-trajectories, i.e. the pedestrian has one or multiple intermediate destinations, which we define as potential endpoints of the local trajectory. These sub-goals can be more easily correlated with past observations to predict likely next steps and disentangle potential future trajectories.

Second, the pedestrian plans a trajectory to reach one of these sub-goals, taking into account the present scene elements. Finally, as the agent goes about executing a plan, the trajectory gets modified to account for other moving agents, respecting social norms of interaction.

Following the aforementioned intuition, we propose to decompose the trajectory prediction problem into two sub-problems that also motivate our proposed architecture (Figure 1). First, given the previous trajectories of the humans in the scene, we propose to estimate a latent belief distribution modeling the pedestrians’ possible endpoints. Using this estimated latent distribution, we sample plausible endpoints for each pedestrian based on their observed trajectory. A socially-compliant future trajectory is then predicted, conditioned not only on the pedestrian and their immediate neighbors’ histories (observed trajectories) but also everybody’s estimated endpoints.

In conclusion, our contribution in this work is threefold. **First**, we propose a socially compliant, endpoint conditioned variational auto-encoder that closely imitates the multi-modal human trajectory planning process. **Second**, we pro-

pose a novel self-attention based social pooling layer that generalizes previously proposed social pooling mechanisms. **Third**, we show that our model can predict stable and plausible intermediate goals that enable setting a new state-of-the-art on several trajectory prediction benchmarks, improving by **20.9%** on SDD [5] & **40.8%** on ETH [6] & UCY [7].

2 Related work

There have been many previous studies [8] on how to forecast pedestrians’ trajectories and predict their behaviors. Several previous works propose to learn statistical behavioral patterns from the observed motion trajectories for future trajectory prediction [9,10,11,12,13,14,15,16,17,18]. Since then, many studies have developed models to account for agent interactions that may affect the trajectory — specifically, through scene and/or social information. Recently, there has been a significant focus on multi-modal trajectory prediction to capture different possible future trajectories given the past. There has also been some research on goal-directed path planning, which consider pedestrians’ goals while predicting a path.

2.1 Context-Based Prediction

Many previous studies have imported environment semantics, such as crosswalks, road, or traffic lights, to their proposed trajectory prediction scheme. Kitani et al. [19] encode agent-space interactions by a Markov Decision Process (MDP) to predict potential trajectories for an agent. Ballan et al. [20] leverage a dynamic Bayesian network to construct motion dependencies and patterns from training data and transferred the trained knowledge to testing data. With the great success of the deep neural network, the Recurrent Neural Network (RNN) has become a popular modeling approach for sequence learning. Kim et al. [21] train a RNN combining multiple Long Short-term Memory (LSTM) units to predict the location of nearby cars. These approaches incorporate rich environment cues from the RGB image of the scene for pedestrians’ trajectory forecasting.

Behaviour of surrounding dynamic agents is also a crucial cue for contextual trajectory prediction. Human behavior modeling studied from a crowd perspective, *i.e.*, how a pedestrian interacts with other pedestrians, has also been studied widely in human trajectory prediction literature. Traditional approaches use *social forces* [22,23,24,25] to capture pedestrians’ trajectories towards their goals with attractive forces, while avoiding collisions in the path with repulsive forces. These approaches require hand-crafted rules and features, which are usually complicated and insufficiently robust for complicated high-level behavior modeling. Recently, many studies applied Long Short Term Memory (LSTM [26]) networks to model trajectory prediction with the social cues. Alahi et al. [27] propose a Social LSTM which learns to predict a trajectory with joint interactions. Each pedestrian is modeled by an individual LSTM, and LSTMs are connected with their nearby individual LSTMs to share information from the hidden state.

2.2 Multimodal Trajectory Prediction

In [28,29], the authors raise the importance of accounting for the inherent multi-modal nature of human paths *i.e.*, given pedestrians’ past history, there are many plausible future paths they can take. This shift of emphasis to plan for multiple future paths has led many recent works to incorporate multi-modality in their trajectory prediction models. Lee et al. [28] propose a conditional variational autoencoder (CVAE), named DESIRE, to generate multiple future trajectories based on agent interactions, scene semantics and expected reward function, within a sampling-based inverse optimal control (IOC) scheme. In [29], Gupta et al. propose a Generative Adversarial Network (GAN) [30] based framework with a novel social pooling mechanism to generate multiple future trajectories in accordance to social norms. In [31], Sadeghian et al. also propose a GAN based framework named SoPhie, which utilizes path history of all the agents in a scene and the scene context information. SoPhie employs a social attention mechanism with physical attention, which helps in learning social information across the agent interactions. However, these socially-aware approaches do not take into account the pedestrians’ ultimate goals, which play a key role in shaping their movement in the scene. A few works also approach trajectory prediction via an inverse reinforcement learning (IRL) setup. Zou et al. [32] applies Generative Adversarial Imitation Learning (GAIL) [33] for trajectory prediction, named Social-Aware GAIL (SA-GAIL). With IRL, the authors model the human decision-making process more closely through modeling humans as agents with states (past trajectory history) and actions (future position). SA-GAIL generates socially acceptable trajectories via a learned reward function.

2.3 Conditioned-on-Goal

Goal-conditioned approaches are regarded as inverse planning or *prediction by planning* where the approach learns the final intent or goal of the agent before predicting the full trajectory. In [34], Rehder *et al.* propose a particle filtering based approach for modeling destination conditioned trajectory prediction and use explicit Von-Mises distribution based probabilistic framework for prediction. Later in a follow-up work, [35] Rehder *et al.* further propose a deep learning based destination estimation approach to tackle intention recognition and trajectory prediction simultaneously. The approach uses fully Convolutional Neural Networks (CNN) to construct the path planning towards some potential destinations which are provided by a recurrent Mixture Density Network (RMDN). While both the approaches make an attempt for destination conditioned prediction, a fully probabilistic approach trains poorly due to unstable training and updates. Further, they ignore the presence of other pedestrians in the scene which is key for predicting shorter term motions which are missed by just considering the environment. Rhinehart et al. [36] propose a goal-conditioned multi-agent forecasting approach named PRECOG, which learns a probabilistic forecasting model conditioned on drivers’ actions intents such as ahead, stop, etc. However, their approach is designed for vehicle trajectory prediction, and thus conditions

on semantic goal states. In our work, we instead propose to utilize destination position for pedestrian trajectory prediction.

In [37], Li *et al.* posit a Conditional Generative Neural System (CGNS), the previous established state-of-the-art result on the ETH/UCY dataset. They propose to use variational divergence minimization with soft-attention to predict feasible multi-modal trajectory distributions. Even more recently, Bhattacharyya *et al.* [38] propose a conditional flow VAE that proposed a general normalizing flow for structured sequence prediction and applies it to the problem of trajectory prediction. Concurrent to our work, Deo *et al.* [39] propose P2TIRL, a Maximum Entropy Reinforcement Learning based trajectory prediction module over a discrete grid. The work [38] shares state-of-the-art with [39] on the Stanford Drone Dataset (SDD) with the TrajNet [40] split. However, these works fail to consider the human aspect of the problem, such as interaction with other agents. We compare our proposed PECNet with all three of the above works on both the SDD & ETH/UCY datasets.

3 Proposed Method

In this work, we aim to tackle the task of human trajectory prediction by reasoning about all the humans in the scene jointly while also respecting social norms. Suppose a pedestrian p^k enters a scene \mathcal{I} . Given the previous trajectory of p for past t_p steps, as a sequence of coordinates $\mathcal{T}_p^k := \{\mathbf{u}^k\}_{i=1}^{t_p} = \{(x^k, y^k)\}_{i=1}^{t_p}$, the problem requires predicting the future position of p^k on \mathcal{I} for next t_f steps, $\mathcal{T}_f^k := \{\mathbf{u}^k\}_{i=t_p+1}^{t_p+t_f} = \{(x, y)\}_{i=t_p+1}^{t_p+t_f}$.

As mentioned in Section 1, we break the problem into two daisy chained steps. First, we model the sub-goal of p^k , i.e. the last observed trajectory points of p^k say, $\mathcal{G}^k = \mathbf{u}^k|_{t_p+t_f}$ as a representation of the predilection of p^k to go its pre-determined route. This sub-goal, also referred to as the endpoint of the trajectory, the pedestrian’s desired end destination for the current sequence. Then in the second step, we jointly consider the past histories $\{\mathcal{T}_p^k\}_{k=1}^\alpha$ of all the pedestrians $\{p^k\}_{k=1}^\alpha$ present in the scene and their estimated endpoints $\{\mathcal{G}^k\}_{k=1}^\alpha$ for predicting socially compliant future trajectories \mathcal{T}_f^k . In the rest of this section we describe in detail, our approach to achieve this, using the endpoint estimation VAE for sampling the future endpoints \mathcal{G} and a trajectory prediction module to use the sampled endpoints $\hat{\mathcal{G}}^k$ to predict \mathcal{T}_f .

3.1 Endpoint VAE

We propose to model the predilection of the pedestrian as a sub-goal endpoint $\mathcal{G} := \mathbf{u}_{t_f} = (x_{t_f}, y_{t_f})$ which is the last observed trajectory point for pedestrian p^k . First, we infer a distribution on \mathcal{G} based on the previous location history \mathcal{T}_i of p^k using the Endpoint VAE.

As illustrated in Figure 2, we extract the previous history \mathcal{T}_i^k and the ground truth endpoint \mathcal{G}^k for all pedestrian p_k in the scene. We encode the past trajectory \mathcal{T}_i^k of all p_k independently using a past trajectory encoder \mathbf{E}_{past} . This

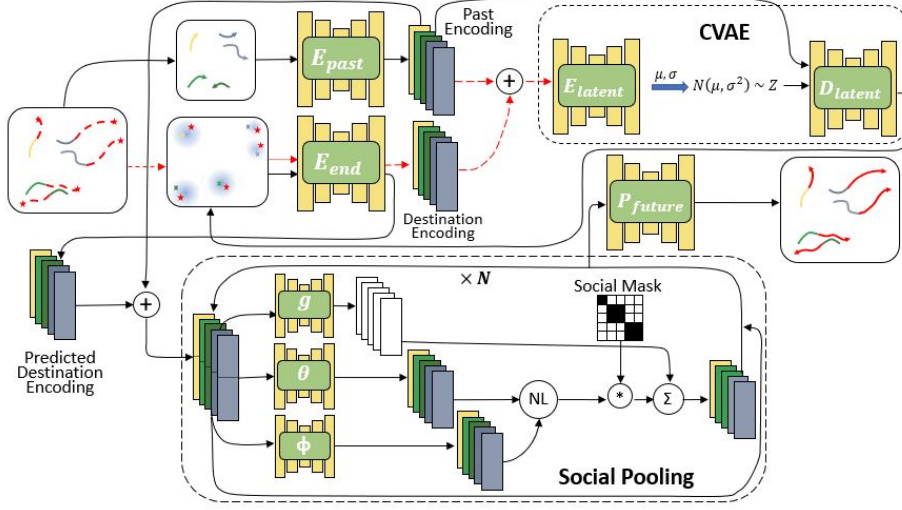


Fig. 2. Architecture of PECNet: PECNet uses past history, \mathcal{T}_i along with ground truth endpoint \mathcal{G}_c to train a VAE for multi-modal endpoint inference. Ground-truth endpoints are denoted by \star whereas \mathbf{x} denote the sampled endpoints $\hat{\mathcal{G}}_c$. The sampled endpoints condition the social-pooling & predictor networks for multi-agent multi-modal trajectory forecasting. Red connections denote the parts utilized only during training. Shades of the same color denote spatio-temporal neighbours encoded with the block diagonal social mask in social pooling module. Further Details in Section 3.1.

yields us $\mathbf{E}_{past}(\mathcal{T}_i)$, a representation of the motion history. Similarly, the future endpoint \mathcal{G}^k is encoded with an Endpoint encoder \mathbf{E}_{end} to produce $\mathbf{E}_{end}(\mathcal{G}^k)$ independently for all k . These representations are concatenated together and passed into the latent encoder \mathbf{E}_{latent} which produces parameter (μ, σ) for encoding the latent variable $z = \mathcal{N}(\mu, \sigma)$ of the VAE. Finally, we sample possible latent future endpoints from $\mathcal{N}(\mu, \sigma)$, concatenate it with $\mathbf{E}_{past}(\mathcal{T}_i)$ for past context and decode using the latent decoder \mathbf{D}_{latent} to yield our guesses for $\hat{\mathcal{G}}^k$. Since the ground truth \mathcal{G}^k belongs to the future, and is unavailable at test time, during evaluation we sample z from $\mathcal{N}(0, \sigma_T \mathbf{I})$, concatenate with $\mathbf{E}_{past}(\mathcal{T}_i)$ (as done in training) and then use the learned \mathbf{D}_{latent} to estimate the future $\hat{\mathcal{G}}^k$. This is illustrated in Figure 2 where the red connections are only used in the training and not in the evaluation phase.

Truncation trick: In [41], Brock *et al.* introduce the ‘Truncation Trick’ as a method of trade-off between the fidelity and variety of samples produced by the generator in BigGAN. In this work, we propose an analogous trick for evaluation phase in multi-modal trajectory forecasting where the variance of the latent endpoint sampling distribution is changed according to the number of samples (K) allowed for multi-modal prediction. In a situation requiring few shot multi-modal prediction, such as under computation constraints, where only a few samples ($K = 1, 2$ or 3) are permissible, we propose to use $\sigma_T = 1$ and truncate the sampling distribution at $\pm c\sqrt{K-1}$. In contrast, in situations where

a high number of predictions are to be generated (such as $K = 20$, a standard setting on benchmarks) we propose to use $\sigma_T > 1$ with no truncation. We posit that this procedure allows simple adjustment of prediction diversity in favor of overall performance for different K , thereby providing a simple method of achieving good performance in all settings without requiring any retraining.

3.2 Endpoint conditioned Trajectory Prediction

Using the sampled estimate of the *endpoints* $\hat{\mathcal{G}}$ from Endpoint VAE, we employ the endpoint encoder \mathbf{E}_{end} once again (within the same forward pass) to obtain encodings for the sampled endpoints $\mathbf{E}_{end}(\hat{\mathcal{G}}^k)$. This is used along with prediction network to *plan* the path \mathcal{T}_f starting to \mathcal{G} thereby predicting the future path.

Note that, another design choice could have been that even during training, the ground truth $\mathbf{E}_{end}(\mathcal{G}^k)$ are used to predict the future \mathcal{T}_f . This seems reasonable as well since it provides cleaner, less noisy signals for the downstream social pooling & prediction networks while still training the overall module end to end (because of coupling through \mathbf{E}_{past}). However, such a choice will decouple training of the Endpoint VAE (which would then train only with KL Divergence and AWL loss, refer Section 3.3) and social pooling network (which would then train only with ATL loss, refer 3.3) leading to inferior performance empirically.

The sampled endpoints' representations $\mathbf{E}_{end}(\hat{\mathcal{G}}^k)$ are then concatenated with corresponding $\mathbf{E}_{past}(\mathcal{T}_i)$ (as in Section 3.1) and passed through N rounds of social pooling using a social pooling mask \mathbf{M} for all the pedestrians in the scene jointly. The social pooling mask \mathbf{M} is $\alpha \times \alpha$ block diagonal matrix denoting the social neighbours for all $\{p_i\}_{i=1}^\alpha$ pedestrians in the scene. Mathematically,

$$\mathbf{M}[i, j] = \begin{cases} 0 & \text{if } \min_{1 \leq m, n \leq t_p} \|\mathbf{u}_m^i - \mathbf{u}_n^j\|_2 > t_{dist} \\ 0 & \text{if } \min_{1 \leq m \leq t_p} |\mathcal{F}(\mathbf{u}_m^i) - \mathcal{F}(\mathbf{u}_m^j)| * \min_{1 \leq m \leq t_p} |\mathcal{F}(\mathbf{u}_m^i) - \mathcal{F}(\mathbf{u}_0^j)| > 0 \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

where $\mathcal{F}(\cdot)$ denoted the actual frame number the trajectory was observed at. Intuitively, \mathbf{M} defines the spatio-temporal neighbours of each pedestrian p_i using proximity threshold t_{dist} for distance in space and ensure temporal overlap. Thus, the matrix \mathbf{M} encodes crucial information regarding social locality of different trajectories which gets utilized in attention based pooling as described below.

Social Pooling: Given the concatenated past history and sampled way-point representations $X_k^{(1)} = (\mathbf{E}_{past}(\mathcal{T}_p^k), \mathbf{E}_{end}(\hat{\mathcal{G}}^k))$ we do N rounds of social pooling where the $(i+1)$ th round of pooling recursively updates the representations $X_k^{(i)}$ from the last round according to the non-local attention mechanism [42]:

$$X_k^{(i+1)} = X_k^{(i)} + \frac{1}{\sum_{j=1}^\alpha \mathbf{M}_{ij} \cdot e^{\phi(X_k^{(i)})^T \theta(X_j^{(i)})}} \sum_{j=1}^\alpha \mathbf{M}_{ij} \cdot e^{\phi(X_k^{(i)})^T \theta(X_j^{(i)})} \mathbf{g}(X_k^{(i)}) \quad (2)$$

where $\{\theta, \phi\}$ are encoders of X_k to map to a learnt latent space where the representation similarity between p_i and p_j trajectories is calculated using the embedded gaussian $\exp(\phi(X_k)^T \theta(X_j))$ for each round of pooling. The social mask, \mathbf{M} is used point-wise to allow pooling only on the spatio-temporal neighbours masking away other pedestrians in the scene. Finally, \mathbf{g} is a transformation encoder for X_k used for the weighted sum with all other neighbours. The whole procedure, after being repeated N times yields $X_k^{(N)}$, the pooled prediction features for each pedestrian with information about the past positions and future destinations of all other neighbours in the scene.

Our proposed social pooling is a novel method for extracting relevant information from the neighbours using non-local attention. The proposed social non local pooling (S-NL) method is permutation *invariant* to pedestrian indices as a useful inductive bias for tackling the social pooling task. Further, we argue that this method of learnt social pooling is more robust to social neighbour mis-identification such as say, mis-specified distance (t_{dist}) threshold compared to previously proposed method such as max-pooling [29], sorting based pooling [31] or rigid grid-based pooling [27] since a learning based method can ignore spurious signals in the social mask \mathbf{M} .

The pooled features $X_k^{(N)}$ are then passed through the prediction network \mathbf{P}_{future} to yield our estimate of rest of trajectory $\{\mathbf{u}^k\}_{k=t_p+1}^{t_p+t_f}$ which are concatenated with sampled endpoint $\hat{\mathcal{G}}$ yields $\hat{\mathcal{T}}_f$. The complete network is trained end to end with the losses described in the next subsection.

3.3 Loss Functions

For training the entire network end to end we use the loss function,

$$\mathcal{L}_{PECNet} = \lambda_1 \underbrace{D_{KL}(\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}) \| \mathcal{N}(0, \mathbf{I}))}_{KL \text{ Div in latent space}} + \lambda_2 \underbrace{\|\hat{\mathcal{G}}_c - \mathcal{G}_c\|_2^2}_{AEL} + \underbrace{\|\hat{\mathcal{T}}_f - \mathcal{T}_f\|^2}_{ATL} \quad (3)$$

where the KL divergence term is used for training the Variational Autoencoder, the Average endpoint Loss (AEL) trains \mathbf{E}_{end} , \mathbf{E}_{past} , \mathbf{E}_{latent} and \mathbf{D}_{latent} and the Average Trajectory Loss (ATL) trains the entire module together.

4 Experiments

4.1 Datasets

Stanford Drone Dataset: Stanford Drone Dataset [5] is a well established benchmark for human trajectory prediction in bird’s eye view. The dataset consists of 20 scenes captured using a drone in top down view around the university campus containing several moving agents like humans and vehicles. It consists of over 11,000 unique pedestrians capturing over 185,000 interactions between agents and over 40,000 interactions between the agent and scene [5]. We use the standard test train split as used in [31,29,39] and other previous works.

ETH/UCY: Second is the ETH [6] and UCY [7] dataset group, which consists of five different scenes – ETH & HOTEL (from ETH) and UNIV, ZARA1,

	SoPhie	S-GAN	DESIRE	CF-VAE*	P2TIRL [†]	SimAug [†]	O-S-TT	O-TT	Ours	PECNet (Ours)
K	20	20	5	20	20.	20	20	20	5	20
ADE	16.27	27.23	19.25	12.60	12.58	10.27	10.56	10.23	12.79	9.96
FDE	29.38	41.44	34.05	22.30	22.07	19.71	16.72	16.29	25.98	15.88

Table 1. Comparison of our method against several recently published multi-modal baselines and previous state-of-the-art method (denoted by *) on the Stanford Drone Dataset [5]. ‘-S’ & ‘-TT’ represents ablations of our method without social pooling & truncation trick. We report results for in pixels for both $K = 5$ & 20 and for several other K in Figure 5. [†] denotes concurrent work. Lower is better.

& ZARA2 (from UCY). All the scenes report the position of pedestrians in world-coordinates and hence the results we report are in metres. The scenes are captured in unconstrained environments with few objects blocking pedestrian paths. Hence, scene constraints from other physical non-animate entities is minimal. For bench-marking, we follow the commonly used leave one set out strategy i.e., training on four scenes and testing on the fifth scene [31,29,37].

4.2 Implementation Details

All the sub-networks used in proposed module are Multi-Layer perceptrons with ReLU non-linearity. Network architecture for each of the sub-networks are mentioned in Table 4.1. The entire network is trained end to end with the \mathcal{L}_{E-VAE} loss using an ADAM optimizer with a batch size of 512 and learning rate of 3×10^{-4} for all experiments. For the loss coefficient weights, we set $\lambda_1 = \lambda_2 = 1$. We use $N = 3$ rounds of social pooling for Stanford Drone Dataset and $N = 1$ for ETH & UCY scenes. Using social masking, we perform the forward pass in mini-batches instead of processing all the pedestrians in the scene in a single forward pass (to avoid memory overflow) constraining all the neighbours of a pedestrian to be in the same mini-batch.

Metrics: For prediction evaluation, we use the Average Displacement Error (ADE) and the Final Displacement Error (FDE) metrics which are commonly used in literature [29,27,25,37]. ADE is the average ℓ_2 distance between the predictions and the ground truth future and FDE is the ℓ_2 distance between the predicted and ground truth at the last observed point. Mathematically,

$$ADE = \frac{\sum_{j=t_i+1}^{t_p+t_f+1} \|\hat{\mathbf{u}}_j - \mathbf{u}_j\|_2}{t_f} \quad FDE = \|\hat{\mathbf{u}}_{t_p+t_f+1} - \mathbf{u}_{t_p+t_f+1}\|_2 \quad (4)$$

where \mathbf{u}_j , $\hat{\mathbf{u}}_j$ are the ground truth and our estimated position of the pedestrian at future time step j respectively.

Baselines: We compare our PECNet against several published baselines including previous state-of-the-art methods briefly described below.

	Network Architecture
\mathbf{E}_{way}	$2 \rightarrow 8 \rightarrow 16 \rightarrow 16$
\mathbf{E}_{past}	$16 \rightarrow 512 \rightarrow 256 \rightarrow 16$
\mathbf{E}_{latent}	$32 \rightarrow 8 \rightarrow 50 \rightarrow 32$
\mathbf{D}_{latent}	$32 \rightarrow 1024 \rightarrow 512 \rightarrow 1024 \rightarrow 2$
ϕ, θ	$32 \rightarrow 512 \rightarrow 64 \rightarrow 128$
\mathbf{g}	$32 \rightarrow 512 \rightarrow 64 \rightarrow 32$
$\mathbf{P}_{predict}$	$32 \rightarrow 1024 \rightarrow 512 \rightarrow 256 \rightarrow 22$

Fig. 3. Network architecture details for all the sub-networks used in the module.

- Social GAN (S-GAN) [29]: Gupta *et al.* propose a multi-modal human trajectory prediction GAN trained with a variety loss to encourage diversity.
- SoPhie [31]: Sadeghian *et al.* propose a GAN employing attention on social and physical constraints from the scene to produce human-like motion.
- CGNS [37]: Li *et al.* posit a Conditional Generative Neural System (CGNS) that uses conditional latent space learning with variational divergence minimization to learn feasible regions to produce trajectories. They also established the previous state-of-the-art result on the ETH/UCY datasets.
- DESIRE [28]: Lee *et al.* propose an Inverse optimal control based trajectory planning method that uses a refinement structure for predicting trajectories.
- CF-VAE [38]: Recently, a conditional normalizing flow based VAE proposed by Bhattacharyya *et al.* pushes the state-of-the-art on SDD further. Notably, their method also does not also rely on the RGB scene image.
- P2TIRL [39]: A concurrent work by Deo *et al.* proposes a method for trajectory forecasting using a grid based policy learned with maximum entropy inverse reinforcement learning policy. They closely tie with the previous state-of-the-art [38] in ADE/FDE performance.
- SimAug [43]: More recently, a concurrent work by Liang *et al.* proposes to use additional 3D multi-view simulation data adversarially, for novel camera view adaptation. [43] improves upon the P2TIRL as well, with performance close to PECNet’s base model. However our best model (with pooling and truncation) still achieves a better ADE/FDE performance.
- Ours-TT: This represents an ablation of our method without using the truncation trick. In other words, we set σ_T to be identically 1 for all K settings. Truncation trick ablations with different K are shown in Fig 5 & Table 1.
- Ours-S-TT: This represents an ablation of our method without using both the social pooling module and the truncation trick *i.e.* the base PECNet. We set $\sigma_T = 1$ and $N = 0$ for the number of rounds of social pooling and directly transmit the representations to \mathbf{P}_{future} , the prediction sub-network.

4.3 Quantitative Results

In this section, we compare and discuss our method’s performance against above mentioned baselines on the ADE & FDE metrics.

Stanford Drone Dataset: Table 1 shows the results of our proposed method against the previous baselines & state-of-the-art methods. Our proposed method achieves a superior performance compared to the previous state-of-the-art [38,39] on both ADE & FDE metrics by a significant margin of **20.9%**. Even without using the proposed social pooling module & truncation trick (OUR-S-TT), we achieve a very good performance (10.56 ADE), underlining the importance of future endpoint conditioning in trajectory prediction. As observed by the difference in performance between Ours-S-TT and Our-TT, the social pooling module also plays a crucial role, boosting performance by 0.33 ADE ($\sim 2.1\%$). Note that, while both P2TIRL [39] & SimAug [43] are concurrent works, we compare with

	S-GAN		SoPhie		CGNS*		S-LSTM		Ours - S - TT		PECNet (Ours)	
	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE
ETH	0.81	1.52	0.70	1.43	0.62	1.40	1.09	2.35	0.58	0.96	0.54	0.87
HOTEL	0.72	1.61	0.76	1.67	0.70	0.93	0.79	1.76	0.19	0.34	0.18	0.24
UNIV	0.60	1.26	0.54	1.24	0.48	1.22	0.67	1.40	0.39	0.67	0.35	0.60
ZARA1	0.34	0.69	0.30	0.63	0.32	0.59	0.47	1.00	0.23	0.39	0.22	0.39
ZARA2	0.42	0.84	0.38	0.78	0.35	0.71	0.56	1.17	0.24	0.35	0.17	0.30
AVG	0.58	1.18	0.54	1.15	0.49	0.97	0.72	1.54	0.32	0.54	0.29	0.48

Table 2. Quantitative results for various previously published methods and state-of-the-art method (denoted by *) on commonly used trajectory prediction datasets. Both ADE and FDE are reported in metres in world coordinates. ‘Our-S-TT’ represents ablation of our method without social pooling & truncation trick.

their methods’ performance as well in Table 1 for experimental comprehensiveness. All reported results averaged for 100 separate trials.

ETH/UCY: Table 2 shows the results for evaluation of our proposed method on the ETH/UCY scenes. We follow the leave-one-out evaluation protocol with $K = 20$ as in CGNS [37]/Social-GAN [29]. All reported numbers are *without* the truncation trick. In this setting too, we observe that our method outperforms previously proposed methods, including the previous state-of-the-art [37]. We push the state-of-the-art on average by $\sim 40.8\%$ with the effect being the most on HOTEL (74.2%) and least on ETH (12.9%). Also, without the social pooling & truncation trick (OUR-S-TT) the performance is still superior to the state-of-the-art by 34.6%, underlining the usefulness of conditioning on the endpoint in PECNet.

Conditioned Way-point positions

& Oracles: For further evaluation of our model, we condition on future trajectory points other than the last observed point which we refer to as *way-points*. Further, to decouple the errors in inferring the conditioned position from errors in predicting a path to that position, we use a destination (endpoint) oracle. The destination oracle provides ground truth information of the conditioned position to the model, which uses it to predict the rest of the trajectory. All of the models, with and without the destination oracle are trained from scratch for each of the conditioning positions.

Referring to Figure 4, we observe several interesting and informative trends that support our earlier hypotheses. (A) As a sanity check, we observe that as we condition on positions further




Fig. 4. Conditioned Way-point positions & Oracles: We evaluate the performance of the proposed method against the choice of future conditioning position on ADE & FDE metrics. Further, we evaluate the performance of a destination oracle version of the model that receives perfect information on conditioned position for predicting rest of the trajectory.

into the future, the FDE for both the Oracle model & the proposed model decrease with a sharp trend after the 7th future position. This is expected since points further into the future provide more information for the final observed point. (B) The ADE error curves for both the oracle and the proposed model have the same decreasing trend albeit with a gentler slope than FDE because the error in predicting the other points (particularly the noisy points in the middle of the trajectory) decreases the gradient. (C) Interestingly, our model’s ADE and FDE is not significantly different from that of the Oracle model for points close in the future and the error in the two models are approximately the same until about the 7th future position. This suggests that till around the middle of the future, the conditioned way-points do not hold significant predictive power on the endpoint and hence using our noisy guesses vs. the oracle’s ground truth for their position does not make a difference.




Fig. 5. Performance across K : ADE & FDE performance of our method against number of samples used for evaluation. Several previous baselines are mentioned as well with their number of samples used. Our method significantly outperforms the state-of-the-art reaching their performance with much lesser number of samples & performing much better with same number of samples as theirs ($K = 20$).

Way-point Prediction Error: The way-point position error is the ℓ_2 distance between the prediction of location of the conditioned position and its ground truth location (in the future). Referring to Figure 4, we observe an interesting trend in the way-point error as we condition on points further into the future. The way-point prediction error increases at the start which is expected since points further into the future have a higher variance. However, after around the middle (7th point) the error plateaus and then even slightly decreases. This lends support to our hypothesis that pedestrians, having predilection towards their destination, exert their will towards it. Hence, *predicting the last observed way-point allows for lower prediction error* than way-points in the middle! This in a nutshell, confirms the motivation of this work.

Effect of Number of samples (K): All the previous works use $K = 20$ samples (except DESIRE which uses $K = 5$) to evaluate the multi-modal predictions for metrics ADE & FDE. Referring to Figure 5, we see the expected decreasing trend in ADE & FDE with time as K increases. Further, we observe that our proposed method achieves the same error as the previous works with much smaller K . Previous state-of-the-art achieves 12.58 [39] ADE using $K = 20$ samples which is matched by PECNet at half the number of samples, $K = 10$. This further lends support to our hypothesis that conditioning on the inferred way-

point significantly reduces the modeling complexity for multi-modal trajectory forecasting, providing a better estimate of the ground truth.

Lastly, as K grows large ($K \rightarrow \infty$) we observe that the FDE slowly gets closer to 0 with more number of samples, as the ground truth \mathcal{G}_c is eventually found. However, the ADE error is still large (6.49) because of the errors in the rest of the predicted trajectory. This is in accordance with the observed ADE (8.24) for the oracle conditioned on the last observed point (*i.e.* 0 FDE error) in Fig. 4.

Design choice for VAE: We also evaluate our design choice of using the inferred future way-points $\hat{\mathcal{G}}_c$ for training subsequent modules (social pooling & prediction) instead of using the ground truth \mathcal{G}_c . As mentioned in Section 3.2, this is also a valid choice for training PECNet end to end. Empirically, we find that such a design achieves 10.87 ADE and 17.03 FDE. This is worse ($\sim 8.8\%$) than using $\hat{\mathcal{G}}_c$ which motivates our design choice for using $\hat{\mathcal{G}}_c$ (Section 3.2).

Truncation Trick: Fig. 5 shows the improvements from the truncation trick for an empirically chosen hyperparameter $c \approx 1.2$. As expected, small values of K gain the most from truncation, with the performance boosting from 22.85 ADE (48.8 FDE) to 17.29 ADE (35.12 FDE) for $K = 1$ ($\sim 24.7\%$).

4.4 Qualitative Results

In Fig. 6, we present several visualizations of PECNet predictions. As shown, PECNet can produce diverse predictions taking into account the past motion history & inferred endpoints. In Fig. 7, we present animations of several socially




Fig. 6. Visualizing Multimodality: We show visualizations for some multi-modal and diverse predictions produced by PECNet. White represents the past 3.2 seconds while red & cyan represents predicted & ground truth future respectively over next 4.8 seconds. Predictions capture a wide-range of plausible trajectory behaviours while discarding improbable ones like, endpoints opposite to pedestrian’s direction of motion.

Fig. 7. Social Interaction Animation: Circles show the past 3.2 sec & stars show the future 19.2 sec (top) & 4.8 sec (bottom) for both ground truth (left) & predictions (right). On top, PECNet neatly captures the purple pedestrian’s overtake of the red pedestrian predicting a smooth cut-in trajectory while blue’s trajectory remains unaffected (a neighbour in social mask \mathbf{M}). At the bottom, the blue pedestrian avoids collision at the intersection by speeding up the trajectory that was originally linear which PECNet accurately captures. Animation best viewed in *Adobe Acrobat Reader*.

compliant predictions. The visualizations show that along with good metric performance PECNet also performs rich multi-modal multi-agent forecasting.

5 Conclusion

In this work we present PECNet, a Pedestrian endpoint conditioned trajectory prediction network. We show that PECNet predicts rich and diverse multi-modal socially compliant trajectories across a variety of scenes. Furthermore, we perform extensive ablations on several design choices such as endpoint conditioning position, number of samples, and choice of training signal to pinpoint the performance gains from PECNet. We also introduce the “truncation trick” [41] for trajectory prediction, a simple method for boosting trajectory prediction accuracy in the few-shots regime. Finally, we benchmark PECNet across multiple datasets including Stanford Drone Dataset [5], ETH [6], and UCY [7], in all of which PECNet achieved the state-of-the-art.

References

1. Maren Bennewitz, Wolfram Burgard, and Sebastian Thrun. Learning motion patterns of persons for mobile service robots. In *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, volume 4, pages 3601–3606. IEEE, 2002.
2. S Thrun, W Burgard, and D Fox. Probabilistic robotics (intelligent robotics and autonomous agents series), ser. intelligent robotics and autonomous agents, 2005.
3. Chris L Baker, Rebecca Saxe, and Joshua B Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.
4. Brian D Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, J Andrew Bagnell, Martial Hebert, Anind K Dey, and Siddhartha Srinivasa. Planning-based prediction for pedestrians. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3931–3936. IEEE, 2009.
5. 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.
6. 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.
7. Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer graphics forum*, volume 26, pages 655–664. Wiley Online Library, 2007.
8. Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M. Kitani, Darius M. Gavrilu, and Kai O. Arras. Human Motion Trajectory Prediction: A Survey. *arXiv e-prints*, 2019.
9. Eckhard Kruse and Friedrich M Wahl. Camera-based observation of obstacle motions to derive statistical data for mobile robot motion planning. In *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No. 98CH36146)*, volume 1, pages 662–667. IEEE, 1998.
10. Lin Liao, Dieter Fox, Jeffrey Hightower, Henry Kautz, and Dirk Schulz. Voronoi tracking: Location estimation using sparse and noisy sensor data. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, volume 1, pages 723–728. IEEE, 2003.
11. Maren Bennewitz, Wolfram Burgard, Grzegorz Cielniak, and Sebastian Thrun. Learning motion patterns of people for compliant robot motion. *The International Journal of Robotics Research*, 24(1):31–48, 2005.
12. Meng Keat Christopher Tay and Christian Laugier. Modelling smooth paths using gaussian processes. In *Field and Service Robotics*, pages 381–390. Springer, 2008.
13. Eugen Käfer, Christoph Hermes, Christian Wöhler, Helge Ritter, and Franz Kummert. Recognition of situation classes at road intersections. In *2010 IEEE International Conference on Robotics and Automation*, pages 3960–3965. IEEE, 2010.
14. Georges Aoude, Joshua Joseph, Nicholas Roy, and Jonathan How. Mobile agent trajectory prediction using bayesian nonparametric reachability trees. In *Infotech@Aerospace 2011*, page 1512. 2011.
15. Christoph G Keller and Darius M Gavrilu. Will the pedestrian cross? a study on pedestrian path prediction. *IEEE Transactions on Intelligent Transportation Systems*, 15(2):494–506, 2013.
16. Michael Goldhammer, Konrad Doll, Ulrich Brunsman, André Gensler, and Bernhard Sick. Pedestrian’s trajectory forecast in public traffic with artificial neural networks. In *2014 22nd International Conference on Pattern Recognition*, pages 4110–4115. IEEE, 2014.

17. Shuang Xiao, Zhan Wang, and John Folkesson. Unsupervised robot learning to predict person motion. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 691–696. IEEE, 2015.
18. Tomasz Piotr Kucner, Martin Magnusson, Erik Schaffernicht, Victor Hernandez Bennetts, and Achim J Lilienthal. Enabling flow awareness for mobile robots in partially observable environments. *IEEE Robotics and Automation Letters*, 2(2):1093–1100, 2017.
19. Kris M Kitani, Brian D Ziebart, James Andrew Bagnell, and Martial Hebert. Activity forecasting. In *European Conference on Computer Vision*, pages 201–214. Springer, 2012.
20. Lamberto Ballan, Francesco Castaldo, Alexandre Alahi, Francesco Palmieri, and Silvio Savarese. Knowledge transfer for scene-specific motion prediction. In *European Conference on Computer Vision*, pages 697–713. Springer, 2016.
21. ByeoungDo Kim, Chang Mook Kang, Jaekyum Kim, Seung Hi Lee, Chung Choo Chung, and Jun Won Choi. Probabilistic vehicle trajectory prediction over occupancy grid map via recurrent neural network. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 399–404. IEEE, 2017.
22. Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.
23. Ramin Mehran, Alexis Oyama, and Mubarak Shah. Abnormal crowd behavior detection using social force model. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 935–942. IEEE, 2009.
24. Kota Yamaguchi, Alexander C Berg, Luis E Ortiz, and Tamara L Berg. Who are you with and where are you going? In *CVPR 2011*, pages 1345–1352. IEEE, 2011.
25. Alexandre Alahi, Vignesh Ramanathan, and Li Fei-Fei. Socially-aware large-scale crowd forecasting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2203–2210, 2014.
26. Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
27. Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.
28. 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.
29. 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.
30. Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
31. Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezaatoughi, and Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1349–1358, 2019.

32. Haosheng Zou, Hang Su, Shihong Song, and Jun Zhu. Understanding human behaviors in crowds by imitating the decision-making process. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
33. Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pages 4565–4573, 2016.
34. Eike Rehder and Horst Kloeden. Goal-directed pedestrian prediction. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 50–58, 2015.
35. Eike Rehder, Florian Wirth, Martin Lauer, and Christoph Stiller. Pedestrian prediction by planning using deep neural networks. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–5. IEEE, 2018.
36. Nicholas Rhinehart, Rowan McAllister, Kris Kitani, and Sergey Levine. Precog: Prediction conditioned on goals in visual multi-agent settings. *arXiv preprint arXiv:1905.01296*, 2019.
37. Jiachen Li, Hengbo Ma, and Masayoshi Tomizuka. Conditional generative neural system for probabilistic trajectory prediction. *arXiv preprint arXiv:1905.01631*, 2019.
38. Apratim Bhattacharyya, Michael Hanselmann, Mario Fritz, Bernt Schiele, and Christoph-Nikolas Straehle. Conditional flow variational autoencoders for structured sequence prediction. *arXiv preprint arXiv:1908.09008*, 2019.
39. Nachiket Deo and Mohan M Trivedi. Trajectory forecasts in unknown environments conditioned on grid-based plans. *arXiv preprint arXiv:2001.00735*, 2020.
40. Amir Sadeghian, Vineet Kosaraju, Agrim Gupta, Silvio Savarese, and Alexandre Alahi. Trajnet: Towards a benchmark for human trajectory prediction. *arXiv preprint*, 2018.
41. Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
42. 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.
43. Junwei Liang, Lu Jiang, and Alexander Hauptmann. Simaug: Learning robust representations from 3d simulation for pedestrian trajectory prediction in unseen cameras. 2020.