




# Take a NAP: Non-Autoregressive Prediction for Pedestrian Trajectories

Hao Xue<sup>1,2</sup>(✉) , Du Q. Huynh<sup>1</sup> , and Mark Reynolds<sup>1</sup> 

<sup>1</sup> The University of Western Australia, Perth, Australia  
{[du.huynh](mailto:du.huynh@uwa.edu.au), [mark.reynolds](mailto:mark.reynolds@uwa.edu.au)}@uwa.edu.au

<sup>2</sup> RMIT University, Melbourne, Australia  
[hao.xue@rmit.edu.au](mailto:hao.xue@rmit.edu.au)

**Abstract.** Pedestrian trajectory prediction is a challenging task as there are three properties of human movement behaviors which need to be addressed, namely, the social influence from other pedestrians, the scene constraints, and the multimodal (multi-route) nature of predictions. Although existing methods have explored these key properties, the prediction process of these methods is autoregressive. This means they can only predict future locations sequentially. In this paper, we present NAP, a non-autoregressive method for trajectory prediction. Our method comprises specifically designed feature encoders and a latent variable generator to handle the three properties above. It also has a future-time-agnostic context generator and a future-time-oriented context generator for non-autoregressive prediction. Through extensive experiments that compare NAP against eleven recent methods, we show that NAP achieves state-of-the-art trajectory prediction performance.

**Keywords:** Trajectory prediction · Non-autoregressive · Motion analysis

## 1 Introduction

Pedestrian trajectory prediction is an important component in a range of applications such as social robots and self-driving vehicles, and plays a key role in understanding human movement behaviors. This task is not trivial due to three key properties in pedestrian trajectory prediction: (i) **social interaction:** People do not always walk alone in public places. Pedestrians often socially interact with others to avoid collisions, walk with friends and keep a certain distance from strangers; (ii) **environmental scene constraints:** Besides social interaction, pedestrians' routes also need to obey scene constraints such as obstacles and building layouts; (iii) **multimodal nature of future prediction:** People can follow different routes as long as these routes are both socially and environmentally acceptable. For example, a person can choose to turn right or turn left to bypass an obstacle.

Recently, researchers have made progress in incorporating these properties into the trajectory prediction process. For example, the Social LSTM model [1]

is one of the methods that can capture social influence information in a local neighborhood around each pedestrian. Based on the generative model GAN [4], the SGAN model proposed by Gupta *et al.* [7] can handle multimodality in the prediction process while also capturing the social influence from other pedestrians in the scene. To deal with the scene constraints, Convolutional Neural Networks (CNNs) are often used to extract scene information in the trajectory prediction network, such as SS-LSTM [22], SoPhie [18], and Social-BiGAT [11].

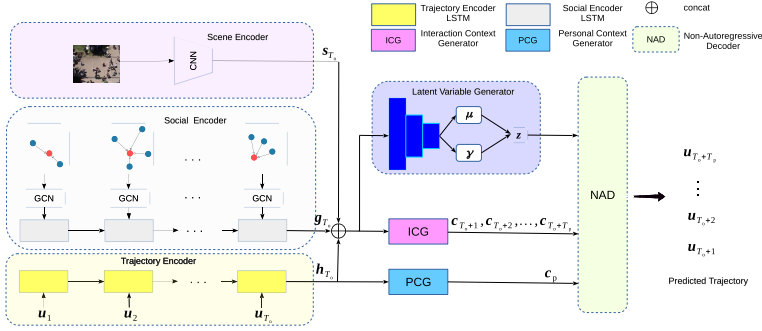
While other methods like [8, 14, 19, 20, 23, 24, 26] miss one or two aforementioned key properties, SoPhie [18], Liang *et al.* [15], and Social-BiGAT [11] are three typical papers that have taken all three properties into consideration. However, these methods predict the future locations recurrently. There are two main limitations in using autoregression to generate trajectory prediction: (i) the autoregressive prediction process works in a recursive manner and so the prediction errors accumulated from previous time steps are passed to the prediction for the next time step; (ii) the process cannot be parallelized, *i.e.*, predictions must be generated sequentially.

To overcome the above limitations and inspired by the application of non-autoregressive decoder in other areas such as machine translation [5, 6] and time series forecasting [21], we propose a novel trajectory prediction method that can predict future trajectories non-autoregressively. We name our method *NAP* (short for *Non-Autoregressive Prediction*). Our research contributions are threefold: (a) To the best of our knowledge, we are the first to explore non-autoregressive trajectory prediction. The network architecture of NAP includes trainable context generators to ensure that context vectors are available for the non-autoregressive decoder to forecast good quality predictions. The state-of-the-art performance of NAP is demonstrated through the extensive experiments and ablation study conducted. (b) Both the social and scene influences are handled by NAP through specially designed feature encoders. The social influence is captured by social graph features propagated through an LSTM; the scene influence is modeled by a CNN. The effectiveness of these encoders is confirmed from the performance of NAP. (c) Unlike previous work in the literature, NAP tackles multimodal predictions by training a latent variable generator to learn the latent variables of the sampling distribution for each pedestrian’s trajectory. The generator is shown to give NAP top performance in multimodal predictions.

## 2 Proposed Method

### 2.1 Problem Definition and System Overview

Pedestrian trajectory prediction is defined as forecasting the future trajectory of the person  $i$  given his/her observed trajectory. We assume that trajectories have already been obtained in the format of time sequences of coordinates (*i.e.*,  $\mathbf{u}_t^i = (x_t^i, y_t^i) \in \mathbb{R}^2, \forall i$ ). The lengths of the observed trajectories and predicted trajectories are represented by  $T_o$  and  $T_p$ . Thus, considering an observed trajectory  $X^i = \{\mathbf{u}_t^i | t = 1, \dots, T_o\}$ , our target is to generate the prediction  $\hat{Y}^i = \{\hat{\mathbf{u}}_t^i | t = T_o + 1, \dots, T_o + T_p\}$ .



**Fig. 1.** The network architecture of NAP. There are three encoders to extract features, two context generators to generate context vectors for the non-autoregressive decoder (NAD), and a latent variable generator to handle multimodal predictions. The embedding layer and the superscript  $i$  are not shown to simplify visualization.

Our proposed NAP comprises four major parts (Fig. 1): (i) three feature encoders for encoding the input information such as observed trajectories and scene images (Sect. 2.2); (ii) two context generators to yield context vectors for prediction (Sect. 2.3); (iii) a latent variable generator for multimodality (Sect. 2.4); (iv) a non-autoregressive decoder for predicting future trajectories (Sect. 2.5). Details of these parts are described in the following subsections.

### 2.2 Feature Encoders

In NAP, there are three feature encoders: a trajectory encoder, to learn the representation of the observed history movement of each pedestrian; a social encoder, to learn the representation of the influence from other pedestrians; and a scene encoder, to learn the representation of the scene features.

**Trajectory Encoder.** The coordinates of the  $i^{\text{th}}$  pedestrian in the observation phase ( $t = 1, \dots, T_o$ ) are firstly embedded into a higher dimensional space through an embedding layer  $\phi(\cdot)$ . Then, across different time steps, the embedded features are used as inputs of an LSTM layer (denoted by  $\text{LSTM}_{\text{ENC}}(\cdot)$ ) to get the encoded hidden state  $\mathbf{h}_t^i$  which captures the observed path information. This trajectory encoding is given by:

$$\mathbf{e}_t^i = \phi(x_t^i, y_t^i; \mathbf{W}_{\text{EMB}}), \tag{1}$$

$$\mathbf{h}_t^i = \text{LSTM}_{\text{ENC}}(\mathbf{h}_{t-1}^i, \mathbf{e}_t^i; \mathbf{W}_{\text{ENC}}), \tag{2}$$

where  $\mathbf{W}_{\text{EMB}}$  and  $\mathbf{W}_{\text{ENC}}$  are trainable weights of the corresponding layers.

**Social Encoder.** At each time step  $t$ , NAP captures the social influence on the  $i^{\text{th}}$  pedestrian through a graph  $\mathcal{G}_t^i = (V_t^i, E_t^i)$ . The  $i^{\text{th}}$  pedestrian and all other

pedestrians  $\mathcal{N}_t^{(i)}$  at the same time step are considered as nodes in the set  $V_t^i$ . Edges linking the  $i^{\text{th}}$  pedestrian and pedestrians in  $\mathcal{N}_t^{(i)}$  form the edge set  $E_t^i$ .

We then use a graph convolutional network (GCN) to process these graphs. In the  $\ell^{\text{th}}$  graph convolutional layer, the node feature of pedestrian  $i$  is aggregated as follows:

$$\mathbf{a}_t^{i,\ell} = \text{ReLU}\left(\mathbf{b}^\ell + \frac{1}{|\mathcal{N}_t^{(i)}|} \sum_{j \in \mathcal{N}_t^{(i)}} \mathbf{W}^\ell \mathbf{a}_t^{j,\ell-1}\right), \quad (3)$$

where  $\mathbf{W}^\ell$  and  $\mathbf{b}^\ell$  are the weight matrix and bias term. At the first layer, we initialize the node feature  $\mathbf{a}_t^{i,0}$  as the location coordinates of the  $i^{\text{th}}$  pedestrian, *i.e.*,  $\mathbf{a}_t^{i,0} = (x_i^t, y_i^t)$ .

The social graph feature  $\mathbf{g}_t^i$  (Eq. (4)) is designed to model the surrounding (or *social*) information of pedestrian  $i$  at each time step  $t$ . To compute this feature, the node features, denoted by  $\{\mathbf{a}_t^i \mid t = 1, \dots, T_o\}$ , from the final GCN layer across all the time steps in the observation phase are passed through an LSTM layer with trainable weights  $\mathbf{W}_{\text{SG}}$ , *i.e.*,

$$\mathbf{g}_t^i = \text{LSTM}_{\text{SG}}(\mathbf{g}_{t-1}^i, \mathbf{a}_t^i; \mathbf{W}_{\text{SG}}). \quad (4)$$

**Scene Encoder.** Different from other methods (such as [18, 22]) that process each image frame in the observation phase, the scene encoder of NAP takes only image  $I_{T_o}^i$  as input, since the scene encoder focuses on the static information like scene layouts and obstacles. Not only does this save the computation time, but it also supplies the most up-to-date and sufficient scene context before prediction kicks in at  $t = T_o + 1$ . We use a CNN to model the scene feature  $\mathbf{s}_{T_o}^i$  as follows:

$$\mathbf{s}_{T_o}^i = \text{CNN}(I_{T_o}^i; \mathbf{W}_{\text{CNN}}). \quad (5)$$

In the  $\text{CNN}(\cdot)$ , we take the merit of the state-of-the-art semantic segmentation architecture DeepLabv3+ [2] and adopt their encoder module to extract semantic features. Similar to SoPhie [18], the encoded semantic features is then passed through a convolutional layer and embedded to the scene feature  $\mathbf{s}_{T_o}^i$  through a single fully connected layer. The encoder module of DeepLabv3+ used in our  $\text{CNN}(\cdot)$  is initialized with the weight matrix that has been pre-trained on the Cityscapes dataset [3]. The scene encoder ( $\text{CNN}(\cdot)$ ) is then trained together with other parts of NAP.

### 2.3 Context Generators

The role of the context generators is to aggregate the outputs of the feature encoders for the downstream decoder for trajectory forecasting. We use two context generators in NAP: (i) a *personal context generator* that is *future-time-agnostic*, as its input is the hidden state  $\mathbf{h}_t^i$  computed from the  $i^{\text{th}}$  pedestrian’s own trajectory only; (ii) an *interaction context generator* that is *future-time-oriented* as its input includes both social graph and scene interaction features also.

**Personal Context Generator (PCG).** We use a Multi-Layer Perceptron (MLP) to model this context generator. The output context vector  $\mathbf{c}_p^i$  is computed as

$$\mathbf{c}_p^i = \text{MLP}_A(\mathbf{h}_{T_o}^i; \mathbf{W}_A), \quad (6)$$

where  $\mathbf{W}_A$  is the corresponding weight matrix. The context  $\mathbf{c}_p^i$  captures the “personal” cues such as the pedestrian’s preferred walking speed and direction in the observation phase, oblivious of his/her surrounding. This context is *future-time-agnostic* because, without considering the social and scene influences, such personal cues can remain the same for the entire trajectory, *e.g.*, the pedestrian can continue to walk in a straight line or at a fast pace with no penalty when bumping into obstacles or other pedestrians since both the social graph and scene features are not present in the equation.

**Interaction Context Generator (ICG).** This context generator incorporates both the social graph and scene features. These two types of influences allow the context generator to be future-time-oriented, *e.g.*, while a pedestrian can walk at a fast pace in the initial part of his/her trajectory, he/she would need to slow down or detour at a later part of the trajectory in order to avoid other pedestrians or scene obstacles. Similar to PCG, we use an MLP to model ICG but its input, being the concatenation of  $\mathbf{h}_{T_o}^i$ ,  $\mathbf{g}_{T_o}^i$ , and  $\mathbf{s}_{T_o}^i$ , contains richer information. The output of ICG comprises different context vectors for different time steps in the prediction phase, as given below:

$$(\mathbf{c}_{T_o+1}^i, \mathbf{c}_{T_o+2}^i, \dots, \mathbf{c}_{T_o+T_p}^i) = \text{MLP}_B(\mathbf{h}_{T_o}^i \oplus \mathbf{g}_{T_o}^i \oplus \mathbf{s}_{T_o}^i; \mathbf{W}_B), \quad (7)$$

where  $\mathbf{W}_B$  is the corresponding weight matrix.

## 2.4 Latent Variable Generator

For multimodal prediction, NAP is designed to generate multiple trajectories through the latent variables  $\boldsymbol{\mu}$  and  $\boldsymbol{\gamma}$  (see Fig. 1). Although several existing trajectory prediction methods such as [7, 9, 12, 14] also use latent variables to handle multimodality, the latent variables in these methods are either directly sampled from the normal distribution or a multivariate normal distribution conditioned on the observed trajectories. To make our latent variables more aware of the social and scene cues, we design NAP to learn the parameters ( $\boldsymbol{\mu}_i$  and  $\boldsymbol{\gamma}_i$ ) of the sampling distribution from the observed trajectories, the social influence, and the scene influence features. To this end, the concatenated feature  $\mathbf{h}_{T_o}^i \oplus \mathbf{g}_{T_o}^i \oplus \mathbf{s}_{T_o}^i$  is passed to two different MLPs (Eqs. (8)–(9)) to yield the mean vector  $\boldsymbol{\mu}_i$  and logarithmic variance  $\boldsymbol{\gamma}_i$  and finally  $\mathbf{z}_i$  for the downstream non-autoregressive decoder:

$$\boldsymbol{\mu}_i = \text{MLP}_\mu(\mathbf{h}_{T_o}^i \oplus \mathbf{g}_{T_o}^i \oplus \mathbf{s}_{T_o}^i; \mathbf{W}_\mu), \quad (8)$$

$$\log \boldsymbol{\sigma}_i^2 \triangleq \boldsymbol{\gamma}_i = \text{MLP}_\sigma(\mathbf{h}_{T_o}^i \oplus \mathbf{g}_{T_o}^i \oplus \mathbf{s}_{T_o}^i; \mathbf{W}_\sigma), \quad (9)$$

$$\mathbf{z}_i \sim \mathcal{N}(\boldsymbol{\mu}_i, \text{diag}(\boldsymbol{\sigma}_i^2)), \quad (10)$$

where  $\mathbf{W}_\mu$  and  $\mathbf{W}_\sigma$  are trainable weights of  $\text{MLP}_\mu(\cdot)$  and  $\text{MLP}_\sigma(\cdot)$ . The reparameterization trick [10] is applied to sample the latent variable  $\mathbf{z}_i$ .

## 2.5 Non-Autoregressive Decoder (NAD)

In the work of Gu *et al.* [5], the authors introduce in their model a component that enhances the inputs passed to the decoder for their machine translation problem. The idea is to help the model learn the internal dependencies (which are absent in their non-autoregressive translator) within a sentence. In the work of Guo *et al.* [6], the authors use a positional module to improve the decoder’s ability to perform local reordering. The context generators in NAP play a similar role as these two approaches. In the testing stage, the trained PCG and ICG are able to generate context vectors for new (unseen) observed trajectories to help the decoder improve its time awareness for trajectory prediction. Both the ICG, which generates future-time-oriented contexts  $\{\mathbf{c}_t \mid T_o+1 \leq t \leq T_o+T_p\}$ , and the future-time-agnostic PCG are needed in the NAD for generating trajectories.

To make multimodal predictions at time step  $t$ , the NAD therefore takes the concatenation of the two contexts  $\mathbf{c}_p^i$  and  $\mathbf{c}_t^i$  and the latent variable  $\mathbf{z}_i$  as inputs, modeled by the MLP below:

$$(\hat{x}_t^i, \hat{y}_t^i) = \text{MLP}_{\text{out}}(\mathbf{c}_t^i \oplus \mathbf{c}_p^i \oplus \mathbf{z}_i; \mathbf{W}_{\text{out}}), \quad (11)$$

where  $\text{MLP}_{\text{out}}(\cdot)$  is the MLP used for predicting the location coordinates. Its parameter  $\mathbf{W}_{\text{out}}$  is shared across all the time steps in the prediction phase. Note that the input passed to the NAD is different for each time step  $t$  in the prediction phase as  $\mathbf{c}_t^i$  depends on  $t$ . We can consider that the contexts  $\{\mathbf{c}_t^i\}$  function is like the hidden states in the decoder of an LSTM except that they are not recursively defined.

## 2.6 Implementation Details

The embedding layer  $\phi$  in Eq. (1) is modeled as a perceptron that outputs 32-dimensional embedded vectors for the input location coordinates. The dimensions of the hidden states of the LSTM layers for both the Trajectory and Social Encoders are 32. The GCN in the Social Encoder is a single graph convolutional layer (*i.e.*,  $\ell = 1$  in Eq. (3)). For the ICG,  $\text{MLP}_B$  is a three-layer MLP with ReLU activations. All the other MLPs used in Eqs. (6), (8), (9), and (11) are single-layer MLPs. Except for in Sect. 3.4 where we explore the prediction performance for different prediction lengths, the observed length of input trajectories is 8 time steps ( $T_o = 8$ ) and the prediction length is 12 time steps ( $T_p = 12$ ) for all other experiments. We implemented NAP and its variants (Sect. 3.3) using the PyTorch framework in Python. The Adam optimizer was used to train our models with the learning rate set to 0.001 and the mini-batch size to 128.

### 3 Experiments

#### 3.1 Datasets and Metrics

We use the popular ETH [17] and UCY [13] datasets, which, altogether, include 5 scenes: ETH, HOTEL, UNIV, ZARA1, and ZARA2. Similar to [18, 24], we normalize each pedestrian’s coordinates and augment the training data by rotating trajectories. As raw scene images are used as inputs for extracting scene features, we also rotate the input images when input trajectories are rotated. Same as previous work in the literature [1, 7, 9, 18], the leave-one-out strategy is adopted for training and testing. All methods are evaluated based on two standard metrics: the Average Displacement Error (ADE) [17] and the Final Displacement Error (FDE) [1]. Smaller errors indicate better prediction performance.

**Table 1.** The ADEs/FDEs (in meters) of various methods. The settings are:  $T_o = 8$  and  $T_p = 12$ . The results with a † are taken from the authors’ papers. The result with ‡ is taken from [7].

Method	#	ETH & UCY scenes					
		ETH	HOTEL	UNIV	ZARA1	ZARA2	Average
Social-LSTM [1]‡		1.09/2.35	0.79/1.76	0.67/1.40	0.47/1.00	0.56/1.17	0.72/1.54
SGAN 1V-1 [7]†		1.13/2.21	1.01/2.18	0.60/1.28	0.42/0.91	0.52/1.11	0.74/1.54
MX-LSTM [8]†		–	–	<b>0.49/1.12</b>	0.59/1.31	0.35/0.79	–
Nikhil & Morris [16]†		1.04/2.07	0.59/1.17	0.57/1.21	0.43/0.90	0.34/0.75	0.59/1.22
Liang <i>et al.</i> [15]†		0.88/1.98	0.36/0.74	0.62/1.32	0.42/0.90	0.34/0.75	0.52/1.14
MATF [25]†		1.33/2.49	0.51/0.95	0.56/1.19	0.44/0.93	0.34/0.73	0.64/1.26
SR-LSTM [24]†		0.63/1.25	0.37/0.74	0.51/1.10	<b>0.41/0.90</b>	0.32/0.70	<b>0.45/0.94</b>
STGAT 1V-1 [9]†		0.88/1.66	0.56/1.15	0.52/1.13	<b>0.41/0.91</b>	<b>0.31/0.68</b>	0.54/1.11
NAP (ours)		<b>0.59/1.13</b>	<b>0.30/0.51</b>	0.59/1.23	<b>0.41/0.86</b>	0.36/0.72	<b>0.45/0.89</b>
SGAN 20V-20 [7]†	✓	0.81/1.52	0.72/1.61	0.60/1.26	0.34/0.69	0.42/0.84	0.58/1.18
SoPhie [18]†	✓	0.70/1.43	0.76/1.67	0.54/1.24	0.30/0.63	0.38/0.78	0.54/1.15
Liang <i>et al.</i> [15]†	✓	0.73/1.65	0.30/0.59	0.60/1.27	0.38/0.81	0.31/0.68	0.46/1.00
MATF GAN [25]†	✓	1.01/1.75	0.43/0.80	<b>0.44/0.91</b>	0.26/0.45	0.26/0.57	0.48/0.90
IDL [14]†	✓	0.59/1.30	0.46/0.83	0.51/1.27	<b>0.22/0.49</b>	<b>0.23/0.55</b>	0.40/0.89
STGAT 20V-20 [9]†	✓	0.65/1.12	0.35/0.66	0.52/1.10	0.34/0.69	0.29/0.60	0.43/0.83
Social-BiGAT [11]†	✓	0.69/1.29	0.49/1.01	0.55/1.32	0.30/0.62	0.36/0.75	0.48/1.00
NAP (ours)	✓	<b>0.53/1.08</b>	<b>0.26/0.46</b>	0.58/1.22	0.30/0.65	0.28/0.60	<b>0.39/0.80</b>

#### 3.2 Comparison with Other Methods

We compare NAP against the following state-of-the-art trajectory prediction methods: Social-LSTM [1], SGAN [7], MX-LSTM [8], Nikhil & Morris [16], Liang *et al.* [15], MATF [25], SR-LSTM [24], SoPhie [18], IDL [14], STGAT [9], and Social-BiGAT [11].

In Table 1, all the compared methods are put into two groups depending on whether they generate only one prediction (top half of the table) or multiple predictions (bottom half and indicated by a tick under the # column) for each

input observed trajectory. The multimodal predictions being considered is 20. The reported ADEs and FDEs are computed from the best predictions out of the 20. Five methods report both single and multimodal prediction results and so they appear in both halves of the table: SGAN, MATF, STGAT, Liang *et al.*, and NAP.

Our proposed method is able to achieve results on par with the state-of-the-art methods for the single prediction setting. Specifically, NAP has the same smallest average ADE (0.45 m) as SR-LSTM while outperforming all methods on the average FDE (0.89 m). In addition to NAP, SR-LSTM, MX-LSTM, and STGAT 1V-1 also have the best performance on one or more scenes. In the lower half of Table 1, results of multimodal predictions are given and compared. On average, our NAP achieves the smallest ADE of 0.39 m and the smallest FDE of 0.80 m. For each scene, the best performers that achieve the smallest ADE/FDE in the lower half of the table include NAP, IDL, MATF, and GAN. Taken together, these results demonstrate the efficacy of our proposed method in both single and multimodal prediction settings.

### 3.3 Ablation Study

To explore the effectiveness of different contexts working together in our proposed method, we consider four variants of NAP listed below:

- NAP-P: This variant only uses the Personal Context Generator (future-time-agnostic context, the light blue PCG box in Fig. 1), *i.e.*, the interaction context  $\mathbf{c}_t^i$  in Eq. (11) is removed. Accordingly, a different MLP taking only  $\mathbf{c}_p^i$  as input is used in the NAD so that it is still capable of returning different outputs at different time steps.
- NAP-ISS: In contrast to the NAP-P above, NAP-ISS disables the personal context and forecasts predictions based on the future-time-oriented interaction context (the pink box in Fig. 1). The personal context  $\mathbf{c}_p^i$  in Eq. (11) is removed so the  $\text{MLP}_{\text{out}}$  in that equation has fewer input neurons. The rest of NAP-ISS is the same as NAP.
- NAP-ISg: In order to further investigate the impact of removing the scene influence, we drop the scene feature  $\mathbf{s}_{T_o}^i$  from the Interaction Context Generator in Eq. (7) to form this variant. That is, the interaction context  $\mathbf{c}_t^i$  in NAP-ISS is computed using both the social graph and scene features, whereas the  $\mathbf{c}_t^i$  in NAP-ISg is computed using the social graph feature only. The  $\text{MLP}_B$  (Eq. (7)) used in this variant has only input neurons for  $\mathbf{h}_{T_o}^i \oplus \mathbf{g}_{T_o}^i$ .
- NAP-ISc: Similar to NAP-ISg, this variant is designed to investigate the impact of removing the social influence. We drop the social graph feature  $\mathbf{g}_{T_o}^i$  but keep the scene feature  $\mathbf{s}_{T_o}^i$  in Eq. (7) so the context  $\mathbf{c}_t^i$  is computed from the scene feature only.

In our ablation study, we compare only the single prediction performance (see Table 2) of these four variants, *i.e.*, the latent variable  $\mathbf{z}_i$  for multimodality is removed from Eq. (11) in the experiments. In general, NAP-P, which uses only

**Table 2.** The ADEs / FDEs (in meters) of the four variants for single predictions in the ablation study, with  $T_o = 8$  and  $T_p = 12$ .

	NAP-P	NAP-ISS	NAP-ISg	NAP-ISc
ETH	0.87/1.63	0.66/1.22	0.69/1.31	0.74/1.52
HOTEL	0.43/0.77	0.34/0.61	0.37/0.70	0.38/0.73
UNIV	0.71/1.42	0.68/1.37	0.68/1.35	0.70/1.39
ZARA1	0.46/0.95	0.45/0.94	0.47/0.96	0.45/0.95
ZARA2	0.44/0.88	0.42/0.83	0.44/0.84	0.44/0.86
Average	0.58/1.13	0.51/0.99	0.53/1.03	0.54/1.09

the personal context (future-time-agnostic), has a poorer performance than the other three variants. This is not unexpected as, without the future-time-oriented context, the NAD is not able to forecast good predictions for different time steps in the prediction phase. Comparing the three interaction context based variants against each other, it is not surprising to see that NAP-ISS outperforms the other two variants due to the presence of both the social graph and scene features. As for NAP-ISg versus NAP-ISc, we observe that NAP-ISg slightly outperforms NAP-ISc. This demonstrates that the social influence is more important than the scene influence. However, it should be noted that the five scenes in the ETH/UCY datasets do not have many obstacles scattered in the pedestrian pathways. The slightly better performance of NAP-ISg confirms that there are more social (pedestrian) interactions than scene interactions in these datasets.

Comparing the results from the four variants in Table 2 and from NAP in Table 1, we observe that NAP outperforms all the four variants. Our ablation study justifies the need for all the contexts to be present in NAP.

**Table 3.** The ADEs/FDEs (in meters) of various methods for different prediction lengths ( $T_o = 8$ ). A method with \* indicates that it generates 20 predictions for each input observed trajectory.

	$T_p = 8$	$T_p = 12$	Error increment
Social-LSTM	0.45/0.91	0.72/1.54	60.00%/69.23%
SGAN 1V-1	0.49/1.00	0.74/1.54	51.02%/54.00%
STGAT 1V-1	0.39/0.81	0.54/1.11	38.46%/37.03%
NAP (ours)	0.35/0.67	0.45/0.89	<b>28.57%/32.84%</b>
SGAN 20V-20*	0.39/0.78	0.58/1.18	48.72%/51.28%
STGAT 20V-20*	0.31/0.62	0.43/0.83	38.71%/33.87%
NAP (ours)*	0.31/0.61	0.39/0.80	<b>25.81%/31.15%</b>

### 3.4 Different Prediction Lengths

In addition to the prediction length setting ( $T_p = 12$  frames, corresponding to 4.8 seconds) used in Tables 1 and 2 and similar to previous work such as SGAN [7] and STGAT [9], we conduct experiments for the prediction length  $T_p = 8$  frames (or 3.2 seconds) to further evaluate the performance of NAP. Table 3 shows the average ADE/FDE results for this prediction length setting. The figures under the ‘ $T_p = 12$ ’ column are copied from the *Average* column of Table 1. Each error increment (last column) due to the increase of  $T_p$  is calculated as:  $(e_{p12} - e_{p8})/e_{p8} \times 100\%$ , where  $e_{p12}$  and  $e_{p8}$  are errors (ADEs or FDEs) for  $T_p = 12$  and  $T_p = 8$  of the same method.

As expected, all methods shown in Table 3 have better performance for the shorter prediction length. In the top half of the table, when generating a predicted trajectory for each input, the error increments of Social-LSTM and SGAN 1V-1 are over 50%. Compared to these two methods, STGAT 1V-1 has smaller error increments for both ADE and FDE. For the multimodal predictions (bottom half of the table), STGAT 20V-20 again outperforms SGAN 20V-20.

We observe from Table 3 that NAP consistently outperforms all other methods for both prediction length settings and for both single and multimodal predictions. Furthermore, NAP also has the smallest error increments for both ADE and FDE when  $T_p$  increases. This demonstrates that NAP is more robust in generating long trajectories. The reason is due to the non-autoregressive nature of the decoder, which not only allows the location coordinates at different time steps to be independently forecast but also helps minimize the accumulation of prediction errors when the prediction length increases.

### 3.5 Qualitative Results

Figure 2 illustrates some prediction examples generated by NAP. The observed and ground truth trajectories are shown in yellow and green; the best trajectory of the 20 predictions of each pedestrian is shown in pink. For better visualization, the video frames have been darkened and blurred. These examples cover different movement behaviors of pedestrians. For example, Fig. 2(a) shows two simple straight path scenarios, Fig. 2(b) and (c) show a gentle turning scenario, and Fig. 2(d) shows a more difficult scenario in which an abrupt turn occurs near the end of the observation phase. Although the predicted trajectory (in pink) in Fig. 2(d) does not perfectly overlap with the ground truth trajectory, NAP is still able to correctly predict the trajectory from the late turning cue.



**Fig. 2.** Examples of predicted trajectories (shown in pink) generated by NAP. The observed trajectories and ground truth trajectories are shown in yellow and green. (Color figure online)



**Fig. 3.** Examples of multiple predicted trajectories shown as heatmaps. The observed trajectories and ground truth trajectories are shown in yellow and green. (Color figure online)

All the 20 predicted trajectories of each pedestrian are displayed as a heatmap in Fig. 3. The generated heatmaps in Fig. 3(a) for these pedestrians show a well coverage of the ground truth trajectories. Two difficult scenarios are shown in Fig. 3(b) and (c). For the pedestrian in Fig. 3(b) and the right pedestrian in Fig. 3(c), each made an abrupt turn at almost the last frame of the observation phase. However, NAP is still able to give good predicted trajectories, as all the plausible paths (including the ground truth trajectories (green)) are well covered by the heatmaps. The left pedestrian in Fig. 3(c) is a stopping scenario. After stopping, the pedestrian can remain still or resume walking in any direction. The generated heatmap shows a good coverage of possible paths; however, it has a small dent in the bottom left hand region due to the presence of a bench there, showing that the pedestrian must bypass the obstacle. This example shows the importance of including scene influence in the method.

## 4 Conclusion

We have presented a novel method called NAP which can handle both social influence and scene influence in the pedestrian trajectory prediction process. NAP captures these influences using the trainable feature encoders in the network. In addition, NAP handles multimodal predictions via a latent variable generator which models the sampling distribution that describes the multiple plausible paths of each pedestrian. Unlike existing trajectory prediction methods, the decoder of NAP is non-autoregressive. NAP is therefore able to forecast predictions for different time steps simultaneously or to forecast only for those time steps that are of interest. From our extensive experiments and ablation study, the context encoders used in NAP have been demonstrated to be effective. Not only does NAP achieve state-of-the-art performance, it also has smaller error increments when the prediction length increases.

## References

1. Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social LSTM: human trajectory prediction in crowded spaces. In: CVPR, June 2016
2. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: ECCV, September 2018
3. Cordts, M., et al.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016)
4. Goodfellow, I., et al.: Generative adversarial networks. In: NeurIPS, pp. 2672–2680 (2014)
5. Gu, J., Bradbury, J., Xiong, C., Li, V.O., Socher, R.: Non-autoregressive neural machine translation. In: ICLR (2018). <https://openreview.net/forum?id=B118BtlCb>
6. Guo, J., Tan, X., He, D., Qin, T., Xu, L., Liu, T.Y.: Non-autoregressive neural machine translation with enhanced decoder input. In: AAAI, vol. 33, pp. 3723–3730 (2019)
7. Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., Alahi, A.: Social GAN: socially acceptable trajectories with generative adversarial networks. In: CVPR, June 2018
8. Hasan, I., Setti, F., Tsemelis, T., Del Bue, A., Galasso, F., Cristani, M.: MX-LSTM: mixing tracklets and vislets to jointly forecast trajectories and head poses. In: CVPR, June 2018
9. Huang, Y., Bi, H., Li, Z., Mao, T., Wang, Z.: STGAT: modeling spatial-temporal interactions for human trajectory prediction. In: ICCV, October 2019
10. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013)
11. Kosaraju, V., Sadeghian, A., Martín-Martín, R., Reid, I., Rezatofighi, S.H., Savarese, S.: Social-BiGAT: multimodal trajectory forecasting using bicycle-GAN and graph attention networks. In: NeurIPS (2019)
12. Lee, N., Choi, W., Vernaza, P., Choy, C.B., Torr, P.H.S., Chandraker, M.: DESIRE: distant future prediction in dynamic scenes with interacting agents. In: CVPR (2017)
13. Lerner, A., Chrysanthou, Y., Lischinski, D.: Crowds by example. In: Computer Graphics Forum, vol. 26, pp. 655–664. Wiley Online Library (2007)
14. Li, Y.: Which way are you going? Imitative decision learning for path forecasting in dynamic scenes. In: CVPR, June 2019
15. Liang, J., Jiang, L., Niebles, J.C., Hauptmann, A.G., Fei-Fei, L.: Peeking into the future: predicting future person activities and locations in videos. In: CVPR, June 2019
16. Nikhil, N., Morris, B.T.: Convolutional neural network for trajectory prediction. In: ECCV Workshop, September 2018
17. Pellegrini, S., Ess, A., Schindler, K., Van Gool, L.: You’ll never walk alone: modeling social behavior for multi-target tracking. In: ICCV, pp. 261–268 (2009)
18. Sadeghian, A., Kosaraju, V., Sadeghian, A., Hirose, N., Rezatofighi, H., Savarese, S.: SoPhie: an attentive GAN for predicting paths compliant to social and physical constraints. In: CVPR, June 2019
19. Su, H., Zhu, J., Dong, Y., Zhang, B.: Forecast the plausible paths in crowd scenes. In: IJCAI, pp. 2772–2778 (2017)
20. Vemula, A., Muelling, K., Oh, J.: Social attention: modeling attention in human crowds. In: ICRA, pp. 1–7, May 2018

21. Wen, R., Torkkola, K., Narayanaswamy, B., Madeka, D.: A multi-horizon quantile recurrent forecaster. In: NIPS Workshop (2017)
22. Xue, H., Huynh, D.Q., Reynolds, M.: SS-LSTM: a hierarchical LSTM model for pedestrian trajectory prediction. In: WACV, pp. 1186–1194. IEEE (2018)
23. Xue, H., Huynh, D.Q., Reynolds, M.: Pedestrian trajectory prediction using a social pyramid. In: Nayak, A.C., Sharma, A. (eds.) PRICAI 2019. LNCS (LNAI), vol. 11671, pp. 439–453. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-29911-8\\_34](https://doi.org/10.1007/978-3-030-29911-8_34)
24. Zhang, P., Ouyang, W., Zhang, P., Xue, J., Zheng, N.: SR-LSTM: state refinement for LSTM towards pedestrian trajectory prediction. In: CVPR, June 2019
25. Zhao, T., et al.: Multi-agent tensor fusion for contextual trajectory prediction. In: CVPR, June 2019
26. Zou, H., Su, H., Song, S., Zhu, J.: Understanding human behaviors in crowds by imitating the decision-making process. In: AAAI (2018)