

# Safety-Compliant Generative Adversarial Networks for Human Trajectory Forecasting

Parth Kothari<sup>1</sup> and Alexandre Alahi<sup>1</sup>

**Abstract**—Human trajectory forecasting in crowds presents the challenges of modelling social interactions and outputting collision-free multimodal distribution. Following the success of Social Generative Adversarial Networks (SGAN), recent works propose various GAN-based designs to better model human motion in crowds. Despite superior performance in reducing distance-based metrics, current networks fail to output socially acceptable trajectories, as evidenced by high collisions in model predictions. To counter this, we introduce SGANv2: an improved safety-compliant SGAN architecture equipped with spatio-temporal interaction modelling and a transformer-based discriminator. The spatio-temporal modelling ability helps to learn the human social interactions better while the transformer-based discriminator design improves temporal sequence modelling. Additionally, SGANv2 utilizes the learned discriminator even at test-time via a collaborative sampling strategy that not only refines the colliding trajectories but also prevents mode collapse, a common phenomenon in GAN training. Through extensive experimentation on multiple real-world and synthetic datasets, we demonstrate the efficacy of SGANv2 to provide socially-compliant multimodal trajectories.

**Index Terms**—Trajectory forecasting, generative adversarial networks, transformers, multimodality.

## I. INTRODUCTION

**F**ORECASTING the motion of pedestrians in crowds is essential for autonomous systems like self-driving cars and social robots that will potentially co-exist with humans. To successfully predict how humans navigate in crowds, a forecasting model needs to tackle three crucial challenges:

- (1) **Modelling social interactions:** the model should learn how the trajectory of one person affects another person;
- (2) **Physically acceptable outputs:** the model predictions should be physically acceptable, *i.e.*, not undergo collisions;
- (3) **Multimodality:** given the history, the model needs to be able to output all futures without missing any mode.

The objective of multi-modal trajectory forecasting is to learn a generative model over future trajectories. Generative adversarial networks (GANs) [1] are a popular choice of generative models for trajectory forecasting as they can effectively capture all possible future modes by mapping samples from

Manuscript received 5 April 2022; revised 18 September 2022; accepted 29 December 2022. Date of publication 18 January 2023; date of current version 29 March 2023. This work was supported in part by the Honda Research and Development Company Ltd. and in part by the École Polytechnique Fédérale de Lausanne (EPFL). The Associate Editor for this article was S. Hamdar. (*Corresponding author: Parth Kothari.*)

The authors are with the Department of Civil Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland (e-mail: parth.kothari@epfl.ch).

Digital Object Identifier 10.1109/TITS.2022.3233906




Fig. 1. Given the history (solid), a forecasting model needs to account for social rules of human motion when predicting collision-free multimodal futures (dash). SGANv2 learns social interactions using spatio-temporal interaction modelling and refines unsafe outputs via collaborative sampling strategy.

a given noise distribution to samples in real data distribution. Gupta et al. [2] proposed Social GAN (SGAN), GANs with social mechanisms, to learn human interactions and output multimodal trajectories. Following the success of SGAN, recent works [3], [4], [5], [6] have proposed improved GAN architectures to better model human interactions in crowds. Indeed, these designs have been successful in reducing the distance-based metrics on real-world datasets [3]. However, we discover that they fail to model social interactions *i.e.*, the models output colliding trajectories.

The failure to output collision-free trajectories can be attributed to the fact that the current discriminator designs do not fully model human-human interactions; hence they are incapable of differentiating real trajectory data from fake data. *Only when the discriminator is capable of differentiating real data from fake data, can the supervised signal from it be meaningful to teach the generator.* To tackle this issue, we propose two architectural changes to the SGAN design: (1) Spatio-temporal interaction modelling to better discriminate between real and generated trajectories. (2) A transformer-based discriminator design to strengthen the sequence modelling capability and better guide the generator training. Equipped with these structural changes, our proposed architecture *SGANv2*, learns to better model the underlying etiquette of human motion as evidenced by reduced collisions.

To further reduce the prediction collisions, *SGANv2* leverages the trained discriminator even at test time. In particular,

TABLE I

HIGH-LEVEL COMPARISON OF PROPOSED ARCHITECTURE AGAINST SELECTED COMMON GENERATIVE MODEL-BASED FORECASTING MODELS

Method	Generative Model	Spatio-temporal Interaction Modelling in Generator	Multimodal	Spatio-temporal Interaction Modelling in Discriminator	Discriminator Design	Test-time Refinement
S-LSTM [7]	–	✓	✗	–	–	✗
DESIRE [8]	VAE	✓	✓	–	–	✓
Trajectron [9]	VAE	✓	✓	–	–	✗
SGAN [2]	GAN	✗	✓	✗	RNN	✗
S-BiGAT [3]	GAN	✓	✓	✗	RNN	✗
SGANv2 [Ours]	GAN	✓	✓	✓	Transformer	✓

we perform collaborative sampling [10] between the generator and discriminator at test-time to guide the unsafe trajectories sampled from the generator. Additionally, we empirically demonstrate that collaborative sampling not only helps to refine trajectories but also has the potential to prevent mode collapse, a phenomenon where the generator fails to capture all modes in the output distribution.

We empirically validate the efficacy of SGANv2 in outputting socially compliant predictions on both synthetic and real-world trajectory datasets. First, we shed light on the shortcomings of the current metric commonly used to measure the multimodal performance, namely Top-20 ADE/FDE [2]. Specifically, we demonstrate that a simple predictor that outputs uniformly spaced predictions performs at par with the state-of-the-art methods when evaluated using only Top-20 ADE/FDE. To counter this limitation, we propose an alternate evaluation scheme to better measure the socially-compliant multimodal performance of a model. We demonstrate that SGANv2 outperforms competitive baselines on both synthetic and real-world trajectory datasets under the new evaluation scheme. Finally, we demonstrate the ability of collaborative sampling to prevent mode collapse on the recently released Forging Paths [11] dataset. Our main contributions are:

- 1) We propose SGANv2, an improved SGAN architecture that incorporates spatio-temporal interaction modelling in both the generator and the discriminator. Moreover, our transformer-based discriminator better guides the learning process of the generator.
- 2) We demonstrate the efficacy of collaborative sampling between the generator and discriminator at test-time to reduce prediction collisions and prevent mode collapse in trajectory forecasting.

## II. RELATED WORK

Human trajectory forecasting in crowds has been an active area of research [7], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29] for various applications like autonomous systems [30], [31], [32], [33] and advanced surveillance [34]. In this section, we review model designs that learn social interactions and output socially compliant multimodal outputs. Table I provides a high-level overview of how SGANv2 architecture differs from selected generative model-based designs.

### A. Spatio-Temporal Interaction Modeling

The seminal work of Social LSTM [7] proposed to learn spatial interactions in a data-driven manner with a novel social pooling layer. Following the success of Social LSTM, various designs of data-driven interaction modules have been proposed [2], [9], [15], [16], [18], [20], [28], [35], [36], [37], [38], [39], [40], [41], [42] to effectively model interactions in crowds. For a detailed taxonomy on the designs of interaction modules, one can refer to Kothari et al. [43]. In this work, we highlight the importance of modelling both the spatial and temporal nature of social interactions.

Architectures that model dynamics of entities in spatio-temporal tasks have been well-studied. Structural-RNN [44], a specialized RNN design, proposed to model dynamics in spatio-temporal tasks like human-object interaction and driver maneuver anticipation. Specific to motion forecasting, several works consider the temporal evolution of spatial human interactions using recurrent mechanisms [14], [42], [45], graph convolutional networks [15], [46] as well as transformers [18]. However, many recent works advocated performing spatial interaction modelling only at the end of observation [2], [3], as this strategy did not impact the distance-based metrics and saved computational time. In this work, we study the importance of spatio-temporal interaction modelling from the perspective of reducing the collisions in model outputs.

### B. Multimodal Forecasting

Neural networks trained using  $L_2$  loss are condemned to output the average of all possible outcomes. To tackle this, one line of work proposes  $L_2$  loss variants [14], [47], [48], [49] capable of handling multiple hypotheses. However, these variants fail to penalize low quality predictions, *e.g.*, samples that are far away from the ground truth and undergo collisions. Thus, training using these variants can result in high diversity but low quality predictions.

Another line of work utilizes generative models [2], [3], [6], [8], [9], [50], with Variational Autoencoders (VAEs) and Generative Adversarial networks (GANs) being the most popular, to model future trajectory distribution. VAE models in trajectory forecasting [8], [9] employ a loss objective based on different variants of the euclidean distance. Such a formulation leads to low quality samples especially when the predictions are uncertain [51]. In contrast, the discriminator of the GAN framework acts as a learned loss function that

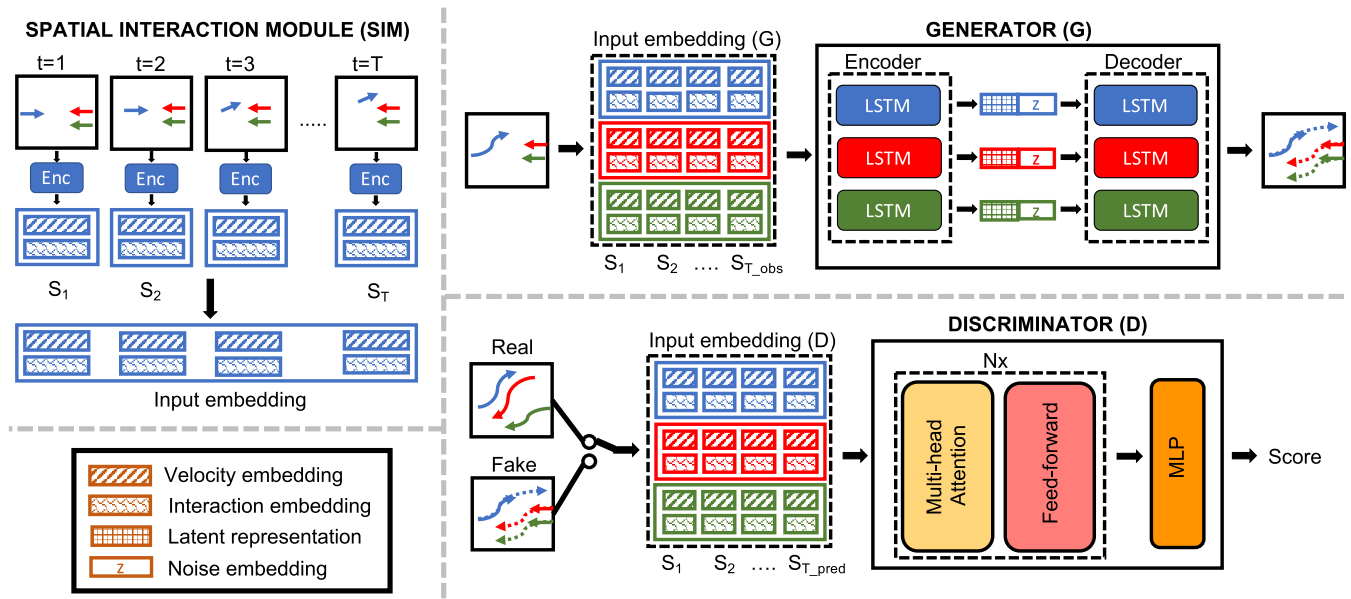


Fig. 2. Our proposed SGANv2 model: Our model consists of three main parts: the spatial interaction embedding module (SIM), the generator (G) and the discriminator (D). At each time-step, for each pedestrian, the SIM outputs the motion embedding and the spatial interaction embedding. G encodes the input embedding sequence using the encoder LSTM to obtain the latent representation. The latent representation along with the sampled noise vector  $z$  is used to generate multimodal predictions using the decoder LSTM. D inputs the stacked embedding sequence of real (or fake) trajectories, encodes it using the transformer architecture to obtain the real (or fake) score.

naturally penalizes the low quality samples under the adversarial training objective *i.e.*, penalty is incurred on the generator if a sample does not look real [1]. Thus, we choose GANs as our generative model as they can effectively produce diverse and high-quality modes by transforming samples from a noise distribution to samples in the real data.

### C. GANs in Trajectory Forecasting

SGAN [2] used an LSTM encoder-decoder with social mechanisms within the GAN framework [52] to perform multimodal forecasting. Following the success of SGAN, various GAN-based architectures have been proposed to better model multimodality in crowds [3], [6], [53] as well as on roads [54], [55]. Li [53] proposed to infer the latent decisions of the agents to model multimodality. Kosaraju et. al. [3] proposed to introduce two discriminators: a local discriminator for the local pedestrian trajectories, similar to [6] and [2], and a global discriminator that accounted for the spatial interactions. All these works exhibit two common design choices: (1) they do not perform spatio-temporal interaction modelling within the discriminator, (2) they utilize a recurrent LSTM-based discriminator.

It is crucial to equip the discriminator with the ability to model spatio-temporal interactions. Therefore, SGANv2 performs *spatio-temporal interaction modelling* within the discriminator, along with the generator. Transformers [56] have been shown to outperform RNNs in almost all sequence modelling tasks, including trajectory forecasting [17], [57]. Therefore, we design our discriminator using the transformer and demonstrate that it better guides the generator training. Giuliari et al. [17] do not take into account social interactions leading to high collisions in the outputs. The spatio-temporal transformer design of STAR [18] is most closely related to

the design of our discriminator. However, as discussed above, their  $L_2$  loss training objective can fail to effectively model multimodality. Further, in contrast to previous transformer and GAN-based works, SGANv2 performs test-time refinement that leads to further collision reduction, discussed next.

### D. Test-Time Refinement

This refers to the task of refining model predictions at test-time. Lee et al. [8] propose an inverse optimal control based module to refine the predicted trajectories. Sun et al. [58] refine trajectories using a reciprocal network that reconstructs input trajectories given the predictions. However, they rely on the strong assumption that both forward and backward trajectories follow identical rules of human motion. We propose to refine trajectories by performing collaborative sampling between the trained generator and discriminator [10]. This technique provides theoretical guarantees with respect to moving the generator distribution closer to real distribution.

### E. Mode Collapse

This is the phenomenon where the generator distribution fails to capture all modes of target distribution. SGAN collapses to a single mode of behavior. Social Ways [59] utilizes InfoGAN that overcomes this issue albeit on a toy dataset. We empirically show that the collaborative sampling technique in SGANv2 overcomes mode collapse on the more-diverse Forking Path dataset [11].

## III. METHOD

Modelling human trajectories using generative adversarial networks (GANs) has the potential to learn the underlying etiquette of human motion and output realistic multimodal

predictions. Indeed, recent GAN-based trajectory forecasting models have been successful in reducing distance-based metrics, however they suffer from high prediction collisions. In this section, we present SGANv2, an improvement over the SGAN architecture to output safety-compliant predictions. On a high level, we propose three structural changes: (1) *Spatio-temporal interaction modelling* within the discriminator and generator to better understand social interactions, (2) *Transformer-based discriminator* to better guide the generator, (3) *Collaborative sampling mechanism* between the generator and discriminator to refine the colliding trajectories at test-time. Our proposed changes are generic and can be employed on top of any existing GAN-based architecture.

### A. Problem Definition

Given a scene, we receive as input the trajectories of all people within the scene denoted by  $\mathbf{X} = \{X_1, X_2, \dots, X_n\}$ , where  $n$  is the number of people in the scene. The trajectory of a person  $i$ , is defined as  $X_i = (x_i^t, y_i^t)$ , for time  $t = 1, 2, \dots, T_{obs}$  and the future ground-truth trajectory is defined as  $Y_i = (x_i^t, y_i^t)$  for time  $t = T_{obs} + 1, \dots, T_{pred}$ . The objective is to accurately and simultaneously forecast the future trajectories of all people  $\hat{\mathbf{Y}} = \hat{Y}_1, \hat{Y}_2, \dots, \hat{Y}_n$ , where  $\hat{Y}_i$  is used to denote the predicted trajectory of person  $i$ . The velocity of a pedestrian  $i$  at time-step  $t$  is denoted by  $v_i^t$ .

### B. Generative Adversarial Networks

GANs consist of two neural networks, namely the generator  $G$  and the discriminator  $D$ , which are trained together in tandem. The objective of  $D$  is to correctly identify whether a sample belongs to the real data distribution or is generated by the generator. The objective of  $G$  is to produce realistic samples which can fool the discriminator.  $G$  takes as input a noise vector  $z$  sampled from a given noise distribution  $p_z$  and transforms it into a real looking sample  $G(z)$ .  $D$  outputs a probability score indicating whether a sample comes from the generator distribution  $p_g$  or the real data distribution  $p_r$ . Training GANs is essentially a minimax game between the generator and the discriminator:

$$\min_G \max_D \mathbb{E}_{x \sim p_r} [\log(D(x))] + \mathbb{E}_{z \sim p_z} [1 - \log(D(G(z)))]. \quad (1)$$

### C. Interaction Modeling Designs

Modelling social interactions is the key to outputting safe and accurate future trajectories. In this work, we argue that current works do not model interactions between agents sufficiently within both the generator and discriminator leading to large number of prediction collisions. Here, we differentiate between the notion of performing *spatial interaction modelling* and performing *spatio-temporal interactions modelling*. On one hand, an architectural design is said to perform *spatial interaction modelling* if it models the interaction between pedestrians at a **single time-step only**. For instance, SGAN performs spatial interaction modelling within the generator as it encodes the neighbourhood information only once, at the end of the observation. On the other hand, an architectural design is said to perform *spatio-temporal interaction modelling* if

it performs the spatial interaction modelling at **every** time-step (from  $t = 1$  to  $t = T_{pred}$ ) and the temporal evolution of the interactions are captured using any sequence encoding mechanism, *e.g.*, an LSTM or a Transformer. We empirically demonstrate that spatio-temporal interactions modelling within both the generator and the discriminator are essential to output safer trajectories.

### D. SGANv2

We now describe our proposed model design in detail (see Fig. 2). Our architecture consists of three key components: the Spatial Interaction embedding Module (SIM), the Generator (G), and the Discriminator (D). SIM is responsible for spatial interaction modelling while the G and D perform temporal modelling. Thus, *G and D in congregation with SIM perform spatio-temporal interaction modelling (STIM)*. In particular, SIM performs motion embedding and spatial interaction embedding for each pedestrian at each time-step. G encodes the embedded sequence through time and outputs multimodal predictions using an LSTM encoder-decoder framework. D, modelled using a transformer [56], inputs the entire sequence comprising the observed trajectory  $\mathbf{X}$  and the future prediction  $\hat{\mathbf{Y}}$  (or ground-truth  $\mathbf{Y}$ ), and classifies it as real/fake.

1) *Spatial Interaction Embedding Module*: One important characteristic that differentiates human motion forecasting from other sequence prediction tasks is the presence of social interactions: the trajectory of a person is affected by other people in their vicinity. SIM performs the task of encoding human motion and human-human interactions in the spatial domain at a particular time-step. We embed the velocity  $v^t$  of pedestrian  $i$  at time  $t$  using a single layer MLP to get the motion embedding vector  $e_i^t$  given as:

$$e_i^t = \phi(v_i^t; W_{emb}), \quad (2)$$

where  $\phi$  is the embedding function with weights  $W_{emb}$ .

The design of SIM is flexible and it can utilize any spatial interaction module proposed in literature [3], [43]. It embeds the spatial configuration of the scene and outputs the interaction embedding  $p_i^t$  for pedestrian  $i$  at time-step  $t$ . We then concatenate the motion embedding with the spatial interaction embedding, *i.e.*,  $s_i^t = [e_i^t; p_i^t]$ , and provide the concatenated embedding  $s_i^t$  to the G (or the D). The input embedding is constructed using the ground-truth observations from  $[1, T_{obs}]$ , and generator predictions from  $[T_{obs} + 1, T_{pred}]$ .

2) *Generator*: Within the generator, the encoder LSTM encodes the input embedding sequence provided by the SIM. The encoder LSTM helps to model the temporal evolution of spatial interactions in the form of the following recurrence:

$$h_i^t = LSTM_{enc}(h_i^{t-1}, s_i^t; W_{encoder}), \quad (3)$$

where  $h_i^t$  denotes the hidden state of pedestrian  $i$  at time  $t$ ,  $W_{encoder}$  are the weights of encoder LSTM that are learned.

The output of the LSTM encoder for each pedestrian at the end of the observation period represents his/her *observed scene representation*. Similar to SGAN, we utilize this representation to condition our GAN for prediction. In other words, SGANv2 take as input noise  $z$  and the observed scene representation to produce future trajectories that are conditioned on the past



observations. The decoder hidden-state of each pedestrian is initialized with the final hidden-state of the encoder LSTM. The input noise  $z$  is concatenated with the inputs of the decoder LSTM, resulting in the following recurrence for the decoder LSTM:

$$h_i^t = LSTM_{dec}(h_i^{t-1}, [s_i^t; z_i]; W_{decoder}), \quad (4)$$

where  $W_{decoder}$  are the weights of decoder LSTM.

The decoder hidden-state at time-step  $t$  of pedestrian  $i$  is then used to predict the next velocity at time-step  $t + 1$ . Similar to Alahi et al. [7], we model the next velocity as a bivariate Gaussian distribution parametrized by the mean  $\mu^{t+1} = (\mu_x, \mu_y)^{t+1}$ , standard deviation  $\sigma^{t+1} = (\sigma_x, \sigma_y)^{t+1}$  and correlation coefficient  $\rho^{t+1}$ :

$$[\mu^t, \sigma^t, \rho^t] = \phi_{dec}(h_i^{t-1}, W_{norm}), \quad (5)$$

where  $\phi_{dec}$  is an MLP and  $W_{norm}$  is learned.

3) *Discriminator*: The social interactions between humans evolve with time. Therefore, we design our discriminator to perform spatio-temporal interaction modelling. Also, in recent times, transformers [56] have become the de-facto model for modelling temporal sequences, replacing recurrent architectures [17], [18]. Therefore, we design the discriminator as a transformer to perform the temporal sequence modelling of the output provided by SIM.

The discriminator takes as input  $Traj_{real} = [\mathbf{X}, \mathbf{Y}]$  or  $Traj_{fake} = [\mathbf{X}, \hat{\mathbf{Y}}]$  and classifies them as real/fake. The discriminator has its own SIM, which provides the spatial interaction embedding  $s_i^t$  for each pedestrian  $i$  at each time-step  $t$  in the input sequence. Instead of passing  $s_i^t$  through an LSTM (similar to the generator), we stack these embedded vectors together to form an embedded sequence  $S_i$  for each pedestrian  $i$  (similar to an embedded sequence obtained after embedding word tokens in the field of natural language [56]):

$$S_i = [s_i^1; s_i^2; \dots; s_i^{T_{pred}}]. \quad (6)$$

This sequence  $S_i$  is given as input to the encoder of the transformer proposed in [56]. The ability of transformers to capture the temporal correlations within the spatial interaction embedding lies mainly in its self-attention module. Within the attention module, each element of the sequence  $S_i$  is decomposed into query (Q), key (K) and value (V). The matrix of outputs is computed using the following equation [56]:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (7)$$

where  $d_k$  is the dimension of the SIM embedding  $s_i^t$ . The output of the attention layer is normalized and passed through a feedforward layer to obtain the latent representation of the input sequence, denoted by  $R_i$ :

$$R_i = \max(0, A_i * W1 + b1) * W2 + b2, \quad (8)$$

where the weights  $W1, W2, b1, b2$  are learned,  $*$  represents matrix multiplication and  $A_i$  denotes the normalized representation of the output of the attention module. We utilize the last element of  $R_i$ , as the representation of the input sequence. This embedding gets scored using an MLP  $\phi_d$  to determine if the sequence is real or fake.

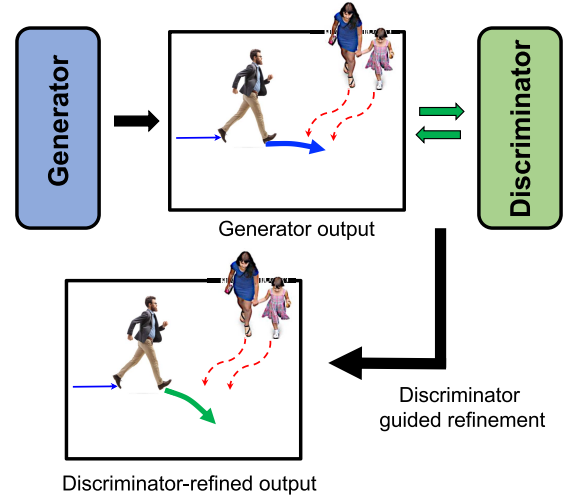


Fig. 3. Illustration of trajectory refinement using collaborative sampling. The trained discriminator provides feedback to improve the generated samples during test-time.

### E. Training

As mentioned earlier, SGANv2 is a conditional GAN model. It takes as input noise vector  $z$ , sampled from  $\mathcal{N}(0, 1)$ , and outputs future trajectories  $\hat{\mathbf{Y}}$  conditioned on the past observations  $\mathbf{X}$ . We found the least-square training objective [60] to be effective in training SGANv2:

$$\min_G \mathcal{L}(G) = \frac{1}{2} \mathbb{E}_{z \sim p_z} [(D(X, G(X, z)) - 1)^2], \quad (9)$$

$$\min_D \mathcal{L}(D) = \frac{1}{2} \mathbb{E}_{x \sim p_r} [(D(X, Y) - 1)^2] + \frac{1}{2} \mathbb{E}_{z \sim p_z} [(D(X, G(X, z)))^2]. \quad (10)$$

Additionally, we utilize the variety loss [2] to further encourage the network to produce diverse samples. For each scene, we generate  $k$  output predictions by randomly sampling  $z$  and penalize the prediction closest to the ground-truth based on L2 distance.

$$\mathcal{L}_{variety} = \min_k \|Y - G(X, z)^{(k)}\|_2^2. \quad (11)$$

Following the strategy in [43], the generator predicts only the trajectory of the pedestrian of interest in each scene and uses the ground-truth future of neighbours during training. During test time, we predict the trajectories of *all* the pedestrians simultaneously in the scene. All the learnable weights are shared between all pedestrians in the scene.

### F. Collaborative Sampling in GANs

The common practice in GANs is to sample from the generator and discard the discriminator during test time. However, our trained discriminator has knowledge regarding the social etiquette of human motion. We can utilize this knowledge to refine the *bad* predictions proposed by the generator. We define a prediction as *bad* if the pedestrian of interest undergoes collision in the model prediction. We propose to refine such trajectories by performing collaborative sampling [10] between the generator and discriminator, as demonstrated in Fig. 3.

TABLE II

QUANTITATIVE EVALUATION OF VARIOUS METHODS ON ETH-UCY. ERRORS REPORTED ARE TOP-K ADE/FDE (IN M) AND COLLISION (IN %). ONLY OBSERVING THE TOP-20 METRIC (AS DONE BY PREVIOUS WORK) CAN LEAD TO INCORRECT CONCLUSIONS. WE SHOW THAT A HIGH-ENTROPY UNIFORM PREDICTOR IS HIGHLY COMPETITIVE WITH RESPECT TO STATE-OF-THE-ART METHODS IN MULTIMODAL FORECASTING USING TOP 20 METRIC. SEE TABLE IV FOR A MORE INFORMATIVE EVALUATION

Model	ETH			HOTEL			UNIV			ZARA1			ZARA2		
	Top-3	Top-20	Col	Top-3	Top-20	Col	Top-3	Top-20	Col	Top-3	Top-20	Col	Top-3	Top-20	Col
Transformer <sup>†</sup> [17]	1.0/1.9	0.6/0.9	5.8	0.5/0.9	0.3/0.5	8.2	2.3/4.2	0.8/1.3	10.9	0.5/1.0	0.3/0.4	7.1	0.4/0.8	<b>0.2/0.3</b>	11.3
STGAT <sup>†</sup> [14]	<b>0.9/1.8</b>	0.7/1.2	1.7	0.7/1.4	0.5/1.0	4.2	0.6/1.2	0.3/0.7	13.9	0.4/0.9	<b>0.2/0.4</b>	3.9	0.4/0.7	0.2/0.4	6.9
Social-STGCNN <sup>†</sup> [15]	1.0/1.8	0.7/1.2	6.7	0.4/0.8	0.3/0.6	10.4	0.7/1.3	0.5/0.8	25.0	0.5/0.9	0.3/0.5	12.1	0.4/0.8	0.3/0.5	19.4
Uniform Predictor (UP)	1.1/2.2	<b>0.6/0.9</b>	3.3	0.5/0.9	<b>0.2/0.4</b>	5.1	0.6/1.3	<b>0.3/0.6</b>	15.7	0.5/1.0	0.3/0.6	4.7	0.4/0.8	0.2/0.4	7.5
SGANv2 [Ours]	1.0/1.9	0.7/1.2	<b>1.0</b>	<b>0.4/0.7</b>	0.3/0.5	<b>1.2</b>	<b>0.6/1.3</b>	0.5/0.8	<b>8.3</b>	<b>0.4/0.8</b>	0.3/0.6	<b>1.3</b>	<b>0.3/0.7</b>	0.3/0.5	<b>2.2</b>

To summarize collaborative sampling for the case of trajectory forecasting, our goal is to refine the generator prediction using gradients from the discriminator *without* updating the parameters of the generator. We leverage the gradient information provided by the discriminator to continuously refine the generator predictions of the pedestrian of interest  $i$  through the following iterative update:

$$\hat{Y}_i^{m+1} = \hat{Y}_i^m - \lambda \nabla \mathcal{L}_G(\hat{Y}_i^m), \quad (12)$$

where  $m$  is the iteration number,  $\lambda$  is the stepsize,  $\mathcal{L}_G$  is the loss of the generator in Eq. 9. The authors demonstrate that the above iteration process theoretically, under mild assumptions, shifts the learned generator distribution towards the real distribution [10]. The trajectories are updated till either the discriminator score goes above a defined threshold or the maximum number of iterations is reached.

#### IV. EXPERIMENTS

In this section, we highlight the ability of SGANv2 to output socially-compliant multimodal futures. We evaluate the performance of our architecture against several state-of-the-art methods on the ETH/UCY datasets [61], [62] and on the interaction-centric TrajNet++ benchmark [43]. Additionally, we highlight the potential of collaborative sampling to prevent mode collapse on the Forking Paths [11] dataset. We evaluate two variants of our model against various baselines:

- **SGANv2 w/o CS:** Our GAN architecture comprising of a transformer-based discriminator that performs spatio-temporal interaction modelling.
- **SGANv2:** Our complete GAN architecture in combination with collaborative sampling at test-time.

##### A. Evaluation Metrics

- 1) **Top-K Average Displacement Error (ADE):** Average  $l_2$  distance between ground truth and closest prediction (out of  $k$  samples) over all predicted time steps.
- 2) **Top-K Final Displacement Error (FDE):** The distance between the final destination of closest prediction (out of  $k$  samples) and the ground truth final destination at the end of the prediction period  $T_{pred}$ .
- 3) **Prediction collision (Col)** [43]: The percentage of collision between the primary pedestrian and the neighbors in the *forecasted future* scene.




Fig. 4. 20 uniformly spread predictions (solid) of a handcrafted predictor conditioned on the last observed velocity (dotted).

##### B. Limitations of Current Multimodal Evaluation Scheme

Current multimodal forecasting works utilize metrics that measure model performance at the *individual level* such as the top- $k$  ADE/FDE [2], [14]. This metric evaluates the quality of the predicted distribution per pedestrian; and does not measure the interaction between different pedestrians. Further, the value of  $k$  is very high ( $k = 20$  being most common). Almost all the recent works [2], [3], [14], [17], [24] in human trajectory forecasting utilize the Top-20 ADE/FDE metric [2] to quantify multimodal performance. We argue that measuring multimodal performance based solely on this metric can be misleading.

The Top-20 ADE/FDE metric can be easily cheated by predicting a high entropy distribution that covers all the space but is not precise [63]. We empirically validate this claim by comparing state-of-the-art baselines against a simple hand-crafted uniform predictor (UP). UP takes as input the last observed velocity of each pedestrian and outputs 20 uniformly spread trajectories (see Fig 4). UP outputs 20 predictions using the combination of 5 different relative direction profiles  $[0, 25, 50, -25, -50]$  (in degrees relative to current direction of motion) and 4 different relative speed profiles  $[1, 0.75, 1.25, 0.25]$  (factors of the current speed).

Table II compares the performance of recent state-of-the-art methods [14], [15], [17] and UP on ETH-UCY datasets. It is apparent that by observing the *Top-20 metric only*, UP seems to perform better (or at par) against the state-of-the-art baselines. If we note the prediction collisions, it is apparent that UP is not a good multimodal predictor. This corroborates our conjecture that a high entropy distribution can easily cheat the Top-20 metric leading to incorrect conclusions.

##### C. Multimodal Evaluation Scheme

To counter the above issues with current multimodal evaluation strategy, we propose to set  $k$  to a lower value in our

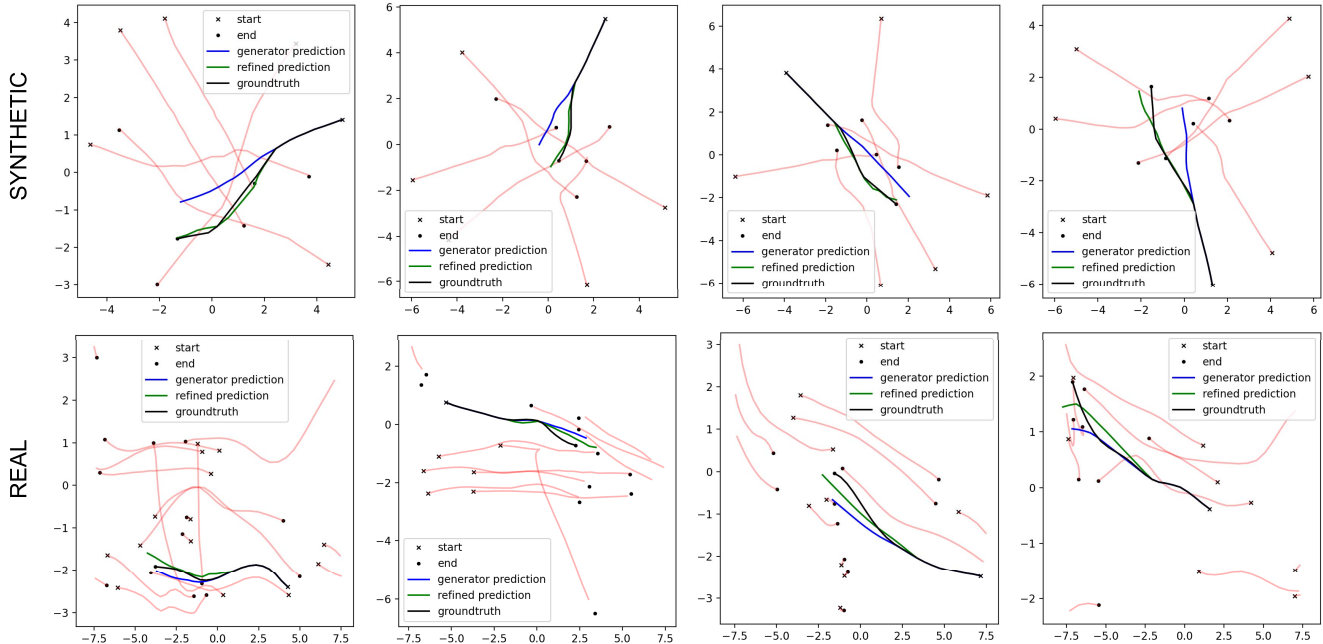


Fig. 5. Illustration of collaborative sampling at test-time to reduce model collisions in both TrajNet++ synthetic and real-world datasets. Given a generator prediction of the pedestrian of interest (blue) that undergoes collision with the neighbours (red), our discriminator, equipped with spatio-temporal interaction modelling, provides feedback based on its learned understanding human-human interactions. Consequently, the resulting refined prediction (green) does not undergo collision and in some cases, is closer to the ground-truth (black).

experiments; as a lower  $k$  is a better proxy for likelihood estimation for implicit generative models [63]. Specific to our problem, we will demonstrate that when  $k$  is low ( $k = 3$ ), the uniform predictor due to a lack of modeling social interactions performs poorly compared to interaction-based baselines [14], [15]. Further, to measure the interaction-modelling capability, we focus on the percentage of collisions between the primary pedestrian and the neighbors in the *forecasted future scene*.

#### D. Synthetic Experiments

We first demonstrate the efficacy of our proposed architectural changes in SGANv2 compared to other generative model designs in the TrajNet++ synthetic setup. We observe that SGANv2 greatly improves upon the Top-3 ADE/FDE metric with a lower collision metric compared to training a model using only variety loss (see Table III).

Next, we utilize collaborative sampling technique to refine trajectories that undergo collision at test-time. The trained discriminator provides feedback to the colliding samples which helps to reduce the collisions. For each colliding prediction, we perform 5 refinement iterations with stepsize 0.01. We observe that this scheme greatly reduces the collision rate by  $\sim 70\%$ . The first row of Fig 5 illustrates the ability of collaborative sampling to refine predictions in the synthetic scenario.

#### E. Real-World Experiments

Next, we evaluate the performance of our SGANv2 architecture in real-world datasets of ETH/UCY and the TrajNet++ benchmark. For ETH/UCY, we observe the trajectories for

TABLE III  
QUANTITATIVE EVALUATION ON TRAJNET++ SYNTHETIC DATASET. ERRORS REPORTED ARE TOP-3 ADE/FDE (IN M) AND COLLISION (IN %). SGANv2 WITH COLLABORATIVE SAMPLING GREATLY REDUCES THE MODEL PREDICTION COLLISIONS WITHOUT COMPROMISING ON THE DISTANCE-BASED METRICS. \*UNIMODAL METHODS

Method	Top-3	Col
CV* [64]	0.4/1.0	21.1
LSTM* [65]	0.3/0.6	19.0
S-LSTM* [43]	0.2/0.5	2.2
D-LSTM* [43]	0.2/0.5	2.2
CVAE [8]	0.2/0.5	4.6
WTA [48]	0.2/0.4	2.4
SGAN [2]	0.2/0.4	2.8
SGANv2 w/o CS [Ours]	0.2/0.4	1.9
SGANv2 [Ours]	<b>0.2/0.4</b>	<b>0.6</b>

8 times steps (3.2 seconds) and show prediction results for 12 (4.8 seconds) time steps. For TrajNet++, we observe the trajectories for 9 times steps (3.6 seconds) and show prediction results for 12 (4.8 seconds) time steps.

Table IV provides the quantitative evaluation of various baselines and state-of-the-art forecasting methods on the ETH/UCY dataset. We observe that SGANv2 outputs safer predictions in comparison to competitive baselines without compromising on the prediction accuracy. Our Top-3 ADE/FDE are on par with (if not better than) state-of-the-art methods while our collision rate is significantly reduced thanks to spatio-temporal interaction modelling. It is further interesting to note that Trajectory Transformer [17] and the

TABLE IV

QUANTITATIVE EVALUATION OF OUR PROPOSED METHOD ON ETH/UCY DATASETS. WE OBSERVE THE TRAJECTORIES FOR 8 TIMES STEPS (3.2 SECONDS) AND SHOW PREDICTION RESULTS FOR THE NEXT 12 TIME STEPS (4.8 SECONDS). ERRORS REPORTED ARE TOP-3 ADE / FDE (IN M), COL (IN %). SGANv2 IMPROVES IN COLLISION METRIC WITHOUT COMPROMISING ON THE DISTANCE-BASED METRICS. \*UNIMODAL

Model	ETH		HOTEL		UNIV		ZARA1		ZARA2	
	Top-3	Col	Top-3	Col	Top-3	Col	Top-3	Col	Top-3	Col
CV* [64]	1.1/2.3	5.3	0.4/0.8	7.2	0.6/1.4	20.3	0.4/1.0	6.0	0.3/0.7	9.6
LSTM* [65]	1.0/2.1	5.8	0.5/0.9	6.7	0.6/1.3	20.2	0.5/1.0	5.2	0.4/0.8	9.5
Uniform Predictor	1.1/2.2	3.3	0.5/0.9	5.1	0.6/1.3	15.7	0.5/1.0	4.7	0.4/0.8	7.5
Transformer <sup>†</sup> [17]	1.0/1.9	5.8	0.5/0.9	8.2	2.3/4.2	10.9	0.5/1.0	7.1	0.4/0.8	11.3
S-LSTM* [7]	1.1/2.1	2.2	0.5/0.9	2.5	0.7/1.5	11.8	0.4/0.9	2.7	0.4/0.8	3.7
CVAE [8]	1.1/2.2	2.8	0.4/0.8	1.5	0.7/1.5	12.6	0.4/0.9	2.6	0.4/0.8	3.5
WTA [48]	1.0/1.9	2.5	0.4/0.7	2.3	0.6/1.3	12.7	0.4/0.8	2.2	0.3/0.7	4.1
SGAN [2]	1.0/2.0	2.2	0.4/0.7	1.7	0.6/1.3	11.8	0.4/0.8	2.3	0.3/0.7	3.2
STGAT <sup>†</sup> [14]	<b>0.9/1.8</b>	1.7	0.7/1.4	4.2	0.6/1.2	13.9	0.4/0.9	3.9	0.4/0.7	6.9
Social-STGCNN <sup>†</sup> [15]	1.0/1.8	6.7	0.4/0.8	10.4	0.7/1.3	25.0	0.5/0.9	12.1	0.4/0.8	19.4
S-BiGAT [3]	1.0/1.9	3.3	0.4/0.7	1.7	0.6/1.3	11.5	0.4/0.8	2.2	0.3/0.7	3.3
SGANv2 w/o CS [Ours]	1.0/1.9	1.7	0.4/0.7	1.4	0.6/1.3	11.5	0.4/0.8	2.1	0.3/0.7	3.6
SGANv2 [Ours]	1.0/1.9	<b>1.0</b>	<b>0.4/0.7</b>	<b>1.2</b>	<b>0.6/1.3</b>	<b>8.3</b>	<b>0.4/0.8</b>	<b>1.3</b>	<b>0.3/0.7</b>	<b>2.2</b>

TABLE V

QUANTITATIVE EVALUATION OF OUR PROPOSED METHOD ON TRAJNET++ REAL-WORLD DATASET. ERRORS REPORTED ARE TOP-3 ADE/FDE (IN M) AND COLLISION (IN %). SGANv2 IN COMBINATION WITH COLLABORATIVE SAMPLING (CS) IMPROVES IN COLLISION METRIC WITHOUT COMPROMISING ON THE DISTANCE-BASED METRICS. \*UNIMODAL

Method	Top-3	Col
CV* [64]	0.6/1.3	10.9
LSTM* [65]	0.5/1.2	9.3
S-LSTM* [7]	0.5/1.0	4.9
D-LSTM* [43]	0.5/1.1	3.9
CVAE [8]	0.5/1.1	3.9
WTA [48]	0.5/1.0	3.5
SGAN [2]	0.5/1.0	3.5
S-NCE [23]	0.5/1.1	4.0
PECNet [24]	<b>0.4/0.9</b>	10.7
Uniform Predictor	0.6/1.2	8.4
Transformer <sup>†</sup> [17].	0.7/1.3	9.4
STGCNN <sup>†</sup> [15]	0.6/1.1	12.6
STGAT <sup>†</sup> [14]	0.5/1.1	5.6
S-BiGAT [3]	0.5/1.0	3.3
SGANv2 w/o CS [Ours]	0.5/1.0	3.1
SGANv2 [Ours]	0.5/1.0	<b>2.3</b>

simple uniform predictor (UP) that performed the best on Top-20 ADE/FDE in Table II are not among the top performing methods when evaluated on the more-strict Top-3 ADE/FDE. Next, we benchmark on the TrajNet++ with interaction-centric scenes with a standardized evaluator that provides a more objective comparison [43].

Table V compares SGANv2 against other competitive baselines on TrajNet++ real-world benchmark. The first part of Table V reports simple baselines and the top-3 official submissions on AICrowd made by different works literature [23], [24], [43]. SGANv2 performs at par with the top-ranked PECNet [24] on the Top-3 evaluation while having **3x** lower

collisions demonstrating that spatio-temporal interaction modelling is key to outputting safer trajectories.<sup>1</sup> Additionally, we utilize the open-source implementation of three additional state-of-the-art methods (denoted by <sup>†</sup>) and evaluate them on the TrajNet++ benchmark. Compared to these competing baselines, SGANv2 improves upon the Top-3 ADE/FDE metric by  $\sim 10\%$  and the collision metric by  $\sim 40\%$ .

We perform collaborative sampling to refine trajectories that undergo collision in real world datasets. For each colliding prediction, we perform 5 refinement iterations with stepsize 0.01. We observe that this procedure reduces the collision rate by  $\sim 30\%$  on both ETH/UCY and TrajNet++. The trained discriminator understands human social interactions, and provides feedback to the bad samples, and consequently helps to reduce collisions. The second row of Fig 5 illustrates a few real-world scenarios where collaborative sampling demonstrates the ability to refine generator predictions that undergo collisions. In conclusion, we observe that SGANv2 beats competitive baselines in generating socially-compliant trajectories without compromising on the distance-based metrics.

#### F. Ablation: Interaction Modeling

In Table VI, we empirically demonstrate that modelling interactions is the key to reducing prediction collisions. We consider the performance of different variants of our proposed SGANv2 architecture based on the interaction modelling schemes within the generator and discriminator. It is apparent that modelling interaction within both the generator and discriminator is necessary to output safe multimodal trajectories.

#### G. Multimodal Analysis

In this final experiment, we demonstrate the potential of collaborative sampling to prevent mode collapse in trajectory

<sup>1</sup>PECNet performs spatial interaction modelling once at end of observation.



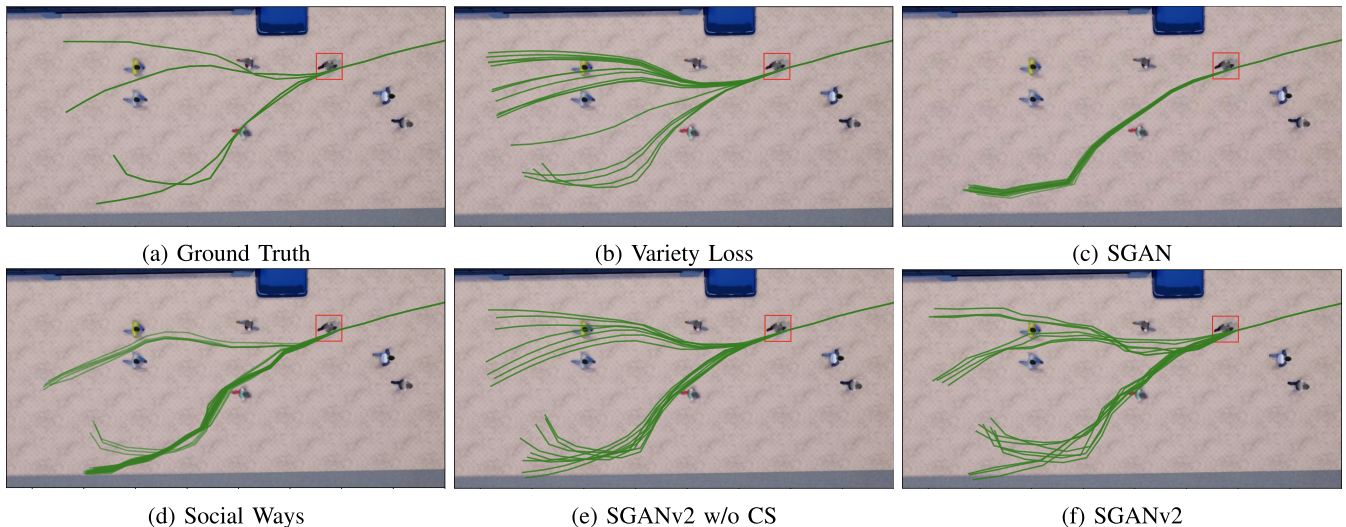


Fig. 6. Qualitative illustration of effectiveness of collaborative sampling on Forking Paths [11]. (b) Training models using variety loss [48] leads to uniform output distribution. (c) Training using SGAN objective [2] leads to mode collapse while (d) InfoGAN [6] helps to mitigate the mode collapse issue. (e) SGANv2 helps to cover all the modes, and in combination with (f) collaborative sampling, we can successfully recover all modes with high accuracy.

TABLE VI

INTERACTION MODULES OF SGANv2. ERRORS REPORTED ARE TOP-3 ADE/FDE (IN M) AND COLLISION (IN %). MODELLING INTERACTIONS GREATLY REDUCES COLLISIONS ON TRAJNET++

$G_{Pool}$	$D_{Pool}$	TrajNet++ Synth		TrajNet++ Real	
		Top-3	Col	Top-3	Col
✗	✗	0.3 / 0.5	18.3	0.5 / 1.1	9.6
✓	✗	0.2 / 0.4	4.1	0.5 / 1.0	3.9
✓	✓	0.2 / 0.4	<b>2.9</b>	0.5 / 1.0	<b>3.1</b>

generation. We utilize the sample scene ‘Zara01’ from the Forking Paths dataset. We choose this scene as the multimodal futures of the ‘Zara01’ scene is only affected by social interactions, and not physical obstacles. It forms the ideal test ground to check the multimodal performance of forecasting models. In this experiment, we observe the trajectories for 8 times steps (3.2 seconds) and show prediction results for 13 (5.2 seconds) time steps.

Fig. 6 qualitatively illustrates the performance of a GAN model trained using variety loss [2], [48] and other GAN objectives on the chosen scene. As there are 4 dominant modes in the scene, we chose  $k = 4$  for the variety loss. The model trained using variety loss (Fig. 6b) ends up learning a uniform distribution, *i.e.*, high diversity and low quality, as there is no penalty on the bad samples during training. Variety loss only penalizes the sample closest to the ground-truth. SGAN training [2] (Fig. 6c) results in mode collapse, *i.e.*, low diversity and high quality as standard GAN training is highly unstable. Social Ways [6] proposed infoGAN objective [66] to mitigate the mode collapse issue. The InfoGAN improves upon SGAN, however, it still fails to cover all the modes (Fig. 6d).

Empirically, we found that training SGANv2 with the gradient penalty objective (Fig. 6e), proposed in [67], provides a better mode coverage compared to InfoGAN, but the resulting distribution is still not accurate. As shown in Fig. 6f, our proposed collaborative sampling at test-time helps to improve the accuracy of the SGANv2 predictions, recovering modes

TABLE VII

COMPUTATIONAL TIME COMPARISON AT INFERENCE PER SCENE FOR VARIOUS FORECASTING DESIGNS. THE ADDITIONAL COMPUTATIONAL TIME FOR SGANv2 CORRESPONDS TO THE SAMPLE REFINEMENT PROCESS THAT OCCURS FOR FIVE ITERATIONS

	LSTM	D-LSTM	SGANv2 w/o CS	SGANv2
Time	10ms	22ms	22ms	77ms

with low coverage. The trained discriminator guides the generated samples to these modes. Thus, we see that collaborative sampling is not only effective in refining trajectories at test time, but also can help to prevent mode collapse.

#### H. Computational Time

Speed is crucial for a method to be used in a real world setting like autonomous vehicles where you need accurate predictions about pedestrian behavior. We provide the computational time at inference for our method against baseline unimodal LSTMs with and without interaction modelling. All the run times have been benchmarked on a single NVIDIA 2080 Ti GPU. We provide the run time per scene (averaged over all the scenes in the TrajNet++ real world benchmark).

The runtimes of D-LSTM and SGANv2 without collaborative sampling are similar as the multiple future predictions in the latter case can be generated in parallel, albeit at the cost of additional memory complexity. The relatively higher computational time of collaborative sampling corresponds to the sample refinement process based on the gradients from the discriminator. Nevertheless, the absolute computational time of collaborative sampling (77ms per scene) is suitable for real-time applications like autonomous systems.

#### V. CONCLUSION

We presented SGANv2, an improved SGAN architecture equipped with two crucial architectural changes in order to output safety-compliant trajectories. First, SGANv2 incorporates spatio-temporal interaction modelling that can help to

understand the subtle nuances of human interactions. Secondly, the transformer-based discriminator better guides the generator learning process. Furthermore, the collaborative sampling strategy helps leverage the trained discriminator during test-time to identify and refine the socially-unacceptable trajectories output by the generator. We empirically demonstrated the strength of SGANv2 to reduce the model collisions without comprising the distance-based metrics. We additionally highlighted the potential of collaborative sampling to overcome mode collapse in a challenging multimodal scenario.

Our work aims at expanding the current horizon of trajectory forecasting models for real-world applications where humans' lives are at risk, such as social robots or autonomous vehicles. Accuracy, safety, and robustness are all mandatory keywords. Over the past years, researchers have focused their evaluation on distance-based metrics. Yet, if we compare the methods on the safety-critical "collision" metric, we observe a difference in performance above 50%. Hence, we believe that one should focus more on this metric and develop methods that aim for zero collisions.

#### ACKNOWLEDGMENT

The authors would like to thank VITA members and reviewers for their valuable comments.

#### REFERENCES

- [1] I. J. Goodfellow et al., "Generative adversarial nets," in *Proc. NIPS*, 2014, pp. 1–9.
- [2] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social GAN: Socially acceptable trajectories with generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2255–2264.
- [3] V. Kosaraju, A. Sadeghian, R. Martín-Martín, I. D. Reid, S. H. Rezafooghi, and S. Savarese, "Social-BiGAT: Multimodal trajectory forecasting using bicycle-GAN and graph attention networks," in *Proc. NeurIPS*, 2019, pp. 1–10.
- [4] A. Sadeghian, V. Kosaraju, A. Sadeghian, N. Hirose, and S. Savarese, "Sophie: An attentive GAN for predicting paths compliant to social and physical constraints," *CoRR*, vol. abs/1806.01482, pp. 1–9, Jun. 2018.
- [5] T. Zhao et al., "Multi-agent tensor fusion for contextual trajectory prediction," in *Proc. CVPR*, Jun. 2019, pp. 12126–12134.
- [6] J. Amirian, J.-B. Hayet, and J. Pettré, "Social ways: Learning multimodal distributions of pedestrian trajectories with GANs," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2019, pp. 2964–2972.
- [7] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social LSTM: Human trajectory prediction in crowded spaces," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 961–971.
- [8] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. S. Torr, and M. Chandraker, "DESIRE: Distant future prediction in dynamic scenes with interacting agents," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2165–2174.
- [9] B. Ivanovic and M. Pavone, "The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2018, pp. 2375–2384.
- [10] Y. Liu, P. Kothari, and A. Alahi, "Collaborative sampling in generative adversarial networks," in *Proc. AAAI Conf. Artif. Intell. (AAAI)*, 2020, pp. 4948–4956.
- [11] J. Liang, L. Jiang, K. Murphy, T. Yu, and A. Hauptmann, "The garden of forking paths: Towards multi-future trajectory prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10508–10518.
- [12] D. Helbing and P. Molnár, "Social force model for pedestrian dynamics," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 51, no. 5, pp. 4282–4286, May 1995.
- [13] J. Li, H. Ma, Z. Zhang, and M. Tomizuka, "Social-WaGDAT: Interaction-aware trajectory prediction via Wasserstein graph double-attention network," 2020, *arXiv:2002.06241*.
- [14] Y. Huang, H. Bi, Z. Li, T. Mao, and Z. Wang, "STGAT: Modeling spatial-temporal interactions for human trajectory prediction," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6271–6280.
- [15] A. Mohamed, K. Qian, M. Elhoseiny, and C. Claudel, "Social-STGCNN: A social spatio-temporal graph convolutional neural network for human trajectory prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 14412–14420.
- [16] Y. Zhu, D. Qian, D. Ren, and H. Xia, "StarNet: Pedestrian trajectory prediction using deep neural network in star topology," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Nov. 2019, pp. 8075–8080.
- [17] F. Giuliari, I. Hasan, M. Cristani, and F. Galasso, "Transformer networks for trajectory forecasting," in *Proc. 25th Int. Conf. Pattern Recognit. (ICPR)*, Jan. 2021, pp. 10335–10342.
- [18] C. Yu, X. Ma, J. Ren, H. Zhao, and S. Yi, "Spatio-temporal graph transformer networks for pedestrian trajectory prediction," in *Proc. ECCV*, 2020, pp. 507–523.
- [19] Y. Su et al., "Trajectory forecasting based on prior-aware directed graph convolutional neural network," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 9, pp. 16773–16785, Sep. 2022.
- [20] P. Zhang, W. Ouyang, P. Zhang, J. Xue, and N. Zheng, "SR-LSTM: State refinement for LSTM towards pedestrian trajectory prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 12077–12086.
- [21] P. Kothari, B. Siffringer, and A. Alahi, "Interpretable social anchors for human trajectory forecasting in crowds," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 15551–15561.
- [22] P. Kothari and A. Alahi, "Adversarial loss for human trajectory prediction," in *Proc. 8th Symp. Eur. Assoc. Res. Transp. (hEART)*, 2019, pp. 1–9.
- [23] Y. Liu, Q. Yan, and A. Alahi, "Social NCE: Contrastive learning of socially-aware motion representations," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 15098–15109.
- [24] N. Daniel et al., "PECNet: A deep multi-label segmentation network for eosinophilic esophagitis biopsy diagnostics," 2021, *arXiv:2103.02015*.
- [25] S. Saadatnejad, M. Bahari, P. Khorsandi, M. Saneian, S.-M. Moosavi-Dezfooli, and A. Alahi, "Are socially-aware trajectory prediction models really socially-aware?" *Transp. Res. C, Emerg. Technol.*, vol. 141, Aug. 2022, Art. no. 103705.
- [26] Y. Liu, R. Cadei, J. Schweizer, S. Bahmani, and A. Alahi, "Towards robust and adaptive motion forecasting: A causal representation perspective," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 17081–17092.
- [27] P. Kothari, D. Li, Y. Liu, and A. Alahi, "Motion style transfer: Modular low-rank adaptation for deep motion forecasting," 2022, *arXiv:2211.03165*.
- [28] C. Xu, M. Li, Z. Ni, Y. Zhang, and S. Chen, "GroupNet: Multiscale hypergraph neural networks for trajectory prediction with relational reasoning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 6488–6497.
- [29] Y. Xu, L. Wang, Y. Wang, and Y. Fu, "Adaptive trajectory prediction via transferable GNN," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 6510–6521.
- [30] Waymo Report. Accessed: Apr. 1, 2022. [Online]. Available: <https://storage.googleapis.com/sdc-prod/v1/safety-report/safety%20report%202018.pdf>
- [31] Uber Report. Accessed: Apr. 1, 2022. [Online]. Available: <https://uber.app.box.com/v/uberatsafetyreport>
- [32] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2019, pp. 6015–6022.
- [33] A. Rasouli and J. K. Tsotsos, "Autonomous vehicles that interact with pedestrians: A survey of theory and practice," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 3, pp. 900–918, Mar. 2020.
- [34] R. Mehran, A. Oyama, and M. Shah, "Abnormal crowd behavior detection using social force model," in *Proc. CVPR*, Jun. 2009, pp. 935–942.
- [35] M. Pfeiffer, G. Paolo, H. Sommer, J. Nieto, R. Siegwart, and C. Cadena, "A data-driven model for interaction-aware pedestrian motion prediction in object cluttered environments," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 1–8.
- [36] X. Shi, X. Shao, Z. Guo, G. Wu, H. Zhang, and R. Shibasaki, "Pedestrian trajectory prediction in extremely crowded scenarios," *Sensors*, vol. 19, no. 5, p. 1223, Mar. 2019.



- [37] N. Bisagno, B. O. Zhang, and N. Conci, "Group LSTM: Group trajectory prediction in crowded scenarios," in *Proc. ECCV Workshops*, 2018, pp. 1–12.
- [38] J. Liang, L. Jiang, J. C. Niebles, A. Hauptmann, and L. Fei-Fei, "Peeking into the future: Predicting future person activities and locations in videos," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2019, pp. 5718–5727.
- [39] A. Tordeux, M. Chraïbi, A. Seyfried, and A. Schadschneider, "Prediction of pedestrian dynamics in complex architectures with artificial neural networks," *J. Intell. Transp. Syst.*, vol. 24, no. 6, pp. 556–568, Nov. 2020.
- [40] Y. Ma, E. W. M. Lee, and R. K. K. Yuen, "An artificial intelligence-based approach for simulating pedestrian movement," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 11, pp. 3159–3170, Nov. 2016.
- [41] I. Hasan, F. Setti, T. Tsesmelis, A. D. Bue, F. Galasso, and M. Cristani, "MX-LSTM: Mixing tracklets and vislets to jointly forecast trajectories and head poses," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 6067–6076.
- [42] J. Li, F. Yang, M. Tomizuka, and C. Choi, "EvolveGraph: Heterogeneous multi-agent multi-modal trajectory prediction with evolving interaction graphs," 2020, *arXiv:2003.13924*.
- [43] P. Kothari, S. Kreiss, and A. Alahi, "Human trajectory forecasting in crowds: A deep learning perspective," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 7, pp. 7386–7400, Jul. 2022.
- [44] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena, "Structural-RNN: Deep learning on spatio-temporal graphs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 5308–5317.
- [45] A. Vemula, K. Muelling, and J. Oh, "Social attention: Modeling attention in human crowds," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 1–7.
- [46] J. Sun, Q. Jiang, and C. Lu, "Recursive social behavior graph for trajectory prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 657–666.
- [47] A. Guzmán-Rivera, D. Batra, and P. Kohli, "Multiple choice learning: Learning to produce multiple structured outputs," in *Proc. NIPS*, 2012, pp. 1–12.
- [48] C. Rupprecht et al., "Learning in an uncertain world: Representing ambiguity through multiple hypotheses," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 3611–3620.
- [49] O. Makansi, E. Ilg, Ö. Çiçek, and T. Brox, "Overcoming limitations of mixture density networks: A sampling and fitting framework for multimodal future prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 7137–7146.
- [50] L. Huang, J. Zhuang, X. Cheng, R. Xu, and H. Ma, "STI-GAN: Multimodal pedestrian trajectory prediction using spatiotemporal interactions and a generative adversarial network," *IEEE Access*, vol. 9, pp. 50846–50856, 2021.
- [51] A. Dosovitskiy and T. Brox, "Generating images with perceptual similarity metrics based on deep networks," in *Proc. NIPS*, 2016, pp. 1–9.
- [52] I. Goodfellow et al., "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, vol. 27, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Red Hook, NY, USA: Curran Associates, 2014, pp. 2672–2680. [Online]. Available: <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
- [53] Y. Li, "Which way are you going? Imitative decision learning for path forecasting in dynamic scenes," in *Proc. CVPR*, Jun. 2019, pp. 294–303.
- [54] D. Roy, T. Ishizaka, C. K. Mohan, and A. Fukuda, "Vehicle trajectory prediction at intersections using interaction based generative adversarial networks," in *Proc. IEEE Intell. Transp. Syst. Conf. (ITSC)*, Oct. 2019, pp. 2318–2323.
- [55] J. Jin et al., "A GAN-based short-term link traffic prediction approach for urban road networks under a parallel learning framework," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 9, pp. 16185–16196, Sep. 2022.
- [56] A. Vaswani et al., "Attention is all you need," in *Proc. NIPS*, 2017, pp. 1–11.
- [57] B. Yu, K. Zhu, K. Wu, and M. Zhang, "Improved OpenCL-based implementation of social field pedestrian model," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 7, pp. 2828–2839, Jul. 2020.
- [58] H. Sun, Z. Zhao, and Z. He, "Reciprocal learning networks for human trajectory prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 7414–7423.
- [59] J. Amirian, W. van Toll, J.-B. Hayet, and J. Pettré, "Data-driven crowd simulation with generative adversarial networks," in *Proc. CASA*, Jul. 2019, pp. 7–10.
- [60] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2813–2821.
- [61] S. Pellegrini, A. Ess, K. Schindler, and L. van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, Sep. 2009, pp. 261–268.
- [62] A. Lerner, Y. Chrysanthou, and D. Lischinski, "Crowds by example," *Comput. Graph. Forum*, vol. 26, no. 3, pp. 655–664, 2007.
- [63] H. Eghbal-zadeh and G. Widmer, "Likelihood estimation for generative adversarial networks," 2017, *arXiv:1707.07530*.
- [64] C. Schöller, V. Aravantinos, F. Lay, and A. Knoll, "What the constant velocity model can teach us about pedestrian motion prediction," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 1696–1703, Apr. 2020.
- [65] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [66] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel, "InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets," in *Advances in Neural Information Processing Systems*, vol. 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. Red Hook, NY, USA: Curran Associates, 2016, pp. 2172–2180. [Online]. Available: <http://papers.nips.cc/paper/6399-infoGAN-interpretability-representation-learning-by-information-maximizing-generative-adversarial-nets.pdf>
- [67] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 214–223.



**Parth Kothari** received the bachelor's degree (Hons.) in electrical engineering from IIT Bombay in 2018. He is currently pursuing the doctoral degree with the Visual Intelligence for Transportation (VITA) Laboratory, EPFL, Switzerland. His research interests include trajectory forecasting, modeling crowd behaviors, and deep learning.



**Alexandre Alahi** received the Ph.D. degree from EPFL. He was a Post-Doctoral Researcher and a Research Scientist at Stanford University for five years. He has worked on the theoretical challenges and practical applications of socially-aware artificial intelligence, i.e., systems equipped with perception and social intelligence. His research interests include enabling machines to perceive the world and make decisions in the context of transportation problems and smart environments. He was awarded the Swiss NSF Early and Advanced Researcher Grants for his work on predicting human social behavior. He has also co-founded multiple startups, such as Visiosafe and won several startup competitions. He was elected as one of the Top 20 Swiss Venture leaders in 2010.