

Speech2AffectiveGestures: Synthesizing Co-Speech Gestures with Generative Adversarial Affective Expression Learning

Uttaran Bhattacharya
University of Maryland
College Park, MD, USA
uttaranb@umd.edu

Elizabeth Childs
University of Maryland
College Park, MD, USA
ehchilds@terpmail.umd.edu

Nicholas Rewkowski
University of Maryland
College Park, MD, USA
nick1@umd.edu

Dinesh Manocha
University of Maryland
College Park, MD, USA
dmanocha@umd.edu

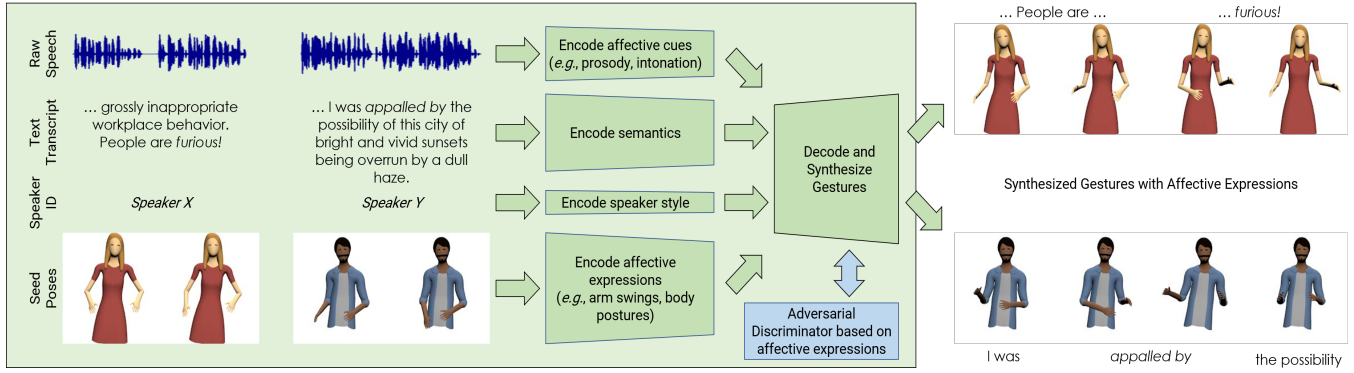


Figure 1: We synthesize 3D pose sequences of co-speech upper-body gestures with appropriate affective expressions. We extract the affective cues from the speech, the sentiments from the corresponding text transcripts, the individual speaker styles, and the joint-based affective expressions from the seed poses (shown on the left). We train a generative adversarial network to synthesize gestures aligned with the speech by leveraging the affective information in both the generation and the discrimination phases. We show two such affective gestures on the right, with the affects *furious* and *appalled* denoted in italics.

Abstract—We present a generative adversarial network to synthesize 3D pose sequences of co-speech upper-body gestures with appropriate affective expressions. Our network consists of two components: a generator to synthesize gestures from a joint embedding space of features encoded from the input speech and the seed poses, and a discriminator to distinguish between the synthesized pose sequences and real 3D pose sequences. We leverage the Mel-frequency cepstral coefficients and the text transcript computed from the input speech in separate encoders in our generator to learn the desired sentiments and the associated affective cues. We design an affective encoder using multi-scale spatial-temporal graph convolutions to transform 3D pose sequences into latent, pose-based affective features. We use our affective encoder in both our generator, where it learns affective features from the seed poses to guide the gesture synthesis, and our discriminator, where it enforces the synthesized gestures to contain the appropriate affective expressions. We perform extensive evaluations on two benchmark datasets for gesture synthesis from the speech, the TED Gesture Dataset and the GENE Challenge 2020 Dataset. Compared to the best baselines, we improve the mean absolute joint error by 10–33%, the mean acceleration difference by

8–58%, and the Fréchet Gesture Distance by 21–34%. We also conduct a user study and observe that compared to the best current baselines, around 15.28% of participants indicated our synthesized gestures appear more plausible, and around 16.32% of participants felt the gestures had more appropriate affective expressions aligned with the speech.

1. Introduction

Co-speech gestures are bodily expressions associated with a person’s speech [1]. They help underline the subject matter and the context of the speech, particularly in the form of beat, deictic, iconic, or metaphoric expressions [2]. Beat gestures are rhythmic movements following the speech, and deictic gestures point to an entity. Iconic gestures describe physical concepts, *e.g.*, spreading and contracting the arms to denote “large” and “small”, and metaphoric gestures describe abstract concepts, *e.g.*, putting a hand to the heart to denote “love”. Synthesizing co-speech gestures is an important task in creating socially engaging characters and virtual agents. These are useful in a variety of multimedia application such as online learning [3], [4], [5], interviewing and counseling [6], [7], robot assistants [1], character de-

signs and game development [8], [9], and visualizing stories and scripts [10].

In our work, we focus on synthesizing the upper-body gestures associated with speech. We consider the joints at the root, spine, head, and the two arms as part of the upper body, which are the joints most commonly used in co-speech gestures [11], [9]. Current state-of-the-art methods for co-speech upper-body gesture synthesis are based on an end-to-end learning approach [12], [11], [9]. These methods train deep neural networks using gestures (available as videos or motion-captured datasets), raw speech waveforms and the corresponding text transcripts, and individual speaker styles. While these methods can generate different beat, deictic, iconic, and metaphoric co-speech gestures and adapt to speaker-specific styles, they do not have any mechanism to reliably incorporate affective expressions in the gestures.

Affective expressions are the modulations in gestures resulting from the emotions experienced by the speakers [13], [14]. Even for a given speaker, the style of gesture expressions can change depending on the emotional context, and human observers are keenly alert to these changes [14]. The combined understanding of the content of the speech and the speaker’s gesture-based affective expressions are crucial to human-human interactions [15], [16]. Therefore, it is essential to incorporate affective expressions in co-speech gestures of animated characters and virtual agents to improve their plausibility in human-machine interactions.

In human-human interactions, we can break the gesture-based affective expressions down into a set of biomechanical features known as *affective features*, such as body postures, head positions, and arm motions [13]. Each affective expression is a combination of one or more affective features, e.g., rapid arm swings and head jerks are often used as expressions of anger or excitement [14]. A multitude of macroscopic and microscopic factors influence the affective features in a given context, including the social setting and the speaker’s idiosyncrasies, making an exhaustive enumeration of affective features tedious and challenging [17]. Nevertheless, it is essential to learn these affective features to understand and synthesize the desired affective expressions.

Moreover, co-speech affective gesture synthesis also requires aligning the gestures with the affective cues obtained from the speech. To this end, prior methods have either learned to map the raw speech waveforms to gestures via latent embeddings [11] or utilized the log-Mel spectrograms to obtain a richer understanding of the affective cues, including the prosody and the intonations in the speech [12], [9]. However, these are high-dimensional representations of speech that require significant computation overhead to be downscaled into convenient latent embedding spaces.

Main Contributions. We present an end-to-end learning approach for generating 3D pose sequences of co-speech gestures with appropriate affective expressions while maintaining the speakers’ individual styles and following a short sequence of seed poses.¹ We leverage the Mel-frequency

cepstral coefficients (MFCCs) from the speeches obtained by performing DCT on the log-Mel spectrograms. MFCCs are highly compressible representations containing sufficient information for speaker identification and also encode affective cues such as prosody and intonation for speech-based emotion recognition. We use separate encoders to encode the MFCCs from the raw speeches, the text transcripts obtained from the speeches, the speakers’ styles, and the seed poses. We use available text- and speaker-encoders proposed by Yoon et al. [11] to learn latent features from the text transcript and a latent style embedding space using a variational encoding of the speaker styles. We propose an encoder for the MFCCs that captures the affective cues in the speech. We also develop an “affective encoder” that transforms the 3D pose sequences to latent affective features using multi-scale spatial-temporal graph convolutions (STGCNs). We design our multi-scale STGCNs to expand attention from the local joints to the macroscopic body parts in a bottom-up manner. We use our affective encoder both in the generator to learn affective features from the seed poses to guide the gesture synthesis and in our discriminator to differentiate between the real and the synthesized gestures based on the affective expressions. To the best of our knowledge, we are the first to learn affective features directly from the gesture data to synthesize gestures with affective expressions. Our main contributions include:

- **Synthesizing co-speech affective gestures.** We synthesize 3D pose sequences of gestures with appropriate affective expressions given a speaker’s speech, maintaining the speakers’ individual styles of gesticulation and following a short sequence of seed poses.
- **Affective encoder for learning latent affective features.** Our affective encoder leverages the localized joint movements and the macroscopic body movements in the 3D pose sequences to learn latent affective features that are used for synthesizing the future poses from the seed poses and adversarially guiding the synthesis as per affective expressions.
- **MFCC encoder for leveraging the affective cues from the speech.** Our MFCC encoder takes in low-dimensional MFCCs containing information on the affective cues from the speech, including prosody and intonations, and transforms them into latent embeddings for affective gesture synthesis.

We evaluate the quantitative performance of our network on two benchmark datasets, the TED Gesture Dataset [1] and the GENE Challenge 2020 Dataset [19]. We observe an improvement of 10–33% on the mean absolute joint error, 8–58% on the mean acceleration difference, and 21–34% on the Fréchet Gesture Distance (FGD) [11] for our network compared to the current state-of-the-art baselines. We also conduct a user study to evaluate the plausibility of our synthesized gestures and the consistency between the affective expressions in the gestures and the speech. Around 15.28% participants indicated that our synthesized gestures are more plausible than the best current baseline of Yoon et al. [11], and around 16.32% participants felt the gestures had more appropriate affective expressions aligned with the

¹ Code and additional materials available at <https://gamma.umd.edu/s2ag>.

speech compared to the same baseline.

2. Related Work

We briefly summarize related prior work on how humans perceive affective body expressions and how these studies were leveraged to synthesize emotionally expressive characters. We also summarize works on synthesizing body motions, especially those aligned with a speech, a text transcript, or both.

2.1. Perceiving Affective Body Expressions

Affect is traditionally expressed in psychology in terms of its valence, arousal, and dominance (VAD) [20]. Valence measures the level of pleasantness (*e.g.*, happy vs. sad), arousal measures how animated the person is (*e.g.*, angry vs. bored), and dominance measures the level of control over the affect (*e.g.*, admiration vs. fear). Studies in both psychology and affective computing indicate the existence of biomechanical *affective* features that provide cues to a person’s perceived affect to human observers [13], [14], [21], [22], [23]. These affective features can be observed at different scales: they can be localized joint movements such as rapid arm swings and head jerks, indicating excitement or anger, as well as macroscopic body movements such as the upper body being expanded, indicating pride or confidence, or collapsed, indicating shame or nervousness. Subsequently, there has been work on detecting perceived emotions by leveraging known affective features either as input to a neural network [24] or to constrain the embedding space [17]. In contrast, we design our neural network to explicitly attend to the body movements at these multiple scales to learn latent affective features directly from the input gesture samples.

2.2. Synthesizing Affective Body Expressions

There has been substantial work on synthesizing affective expressions for embodied conversation agents [25], [26] and other social virtual agents to interact via facial expressions [27], [28] or gaits [29]. Furthermore, the synthesis of affective facial expressions has been aligned with a character’s speech using data-driven techniques [30]. While synthesizing speech-aligned affective facial expressions has been relatively well-studied, aligning the speech with affective body expressions has been more challenging. Some of the widely used approaches are rule-based systems such as that of DeVault et al. [7], which has a virtual human counselor expressing appropriate affective hand and body gestures following known mappings between the emotional states and the stored animations. Recent methods utilize gait datasets annotated with categorical emotions such as happy, sad, and angry to generate emotive gaits [31], [32]. Other techniques have extended to the VAD space of affect, where body gestures are generated given the text transcripts of speech and the corresponding intended emotion as a point

in the VAD space [33]. Our approach is based on designing an end-to-end system that can synthesize body expressions by automatically understanding the affective content in the input speech.

2.3. Synthesizing Gestures

There is a rich body of work gesture synthesis using rule-based systems, as surveyed comprehensively by Wagner et al. [34]. However, scalability to novel scenarios remains a challenge for rule-based systems on account of manually designing new rules. Instead, we focus on a summary of the recent data-driven approaches of automated gesture synthesis in novel scenarios [35], which are in line with our learning-based approach. Existing techniques have utilized hidden Markov models [36], recurrent neural network variants [37], [1], and autoencoders [38] to learn robust latent features that encode the input speech, available as either an audio or a text transcript, and can be used to decode the output gestures. Other approaches have opted to learn stochastic generation processes using tools such as invertible sub-transformations [39] to map between the speech and the gesture spaces. To improve the realism of the generated speech-driven gestures, more recent works incorporate the speech semantics into the training process [9], and even combined the synthesized gestures with rule-based head nods and hand waves for embodied conversation agents [40].

Our approach is complementary to these approaches in that we learn mappings from the text transcripts of speech to gestures. It eliminates the noise in speech signals and helps us focus only on the relevant content and context. Learning from the text also enables us to focus on a broader range of gestures, including iconic, deictic, and metaphoric gestures [2]. Our work is most closely related to that of Yoon et al. [1]. They learn upper body gestures as PCA-based, low-dimensional pose features, corresponding to text transcripts from a dataset of TED-talk videos, then map these 3D gestures to an NAO robot. They have also followed up this work by generating upper-body gestures aligned with the three modalities of speech, text transcripts, and person identity [11]. On the other hand, we learn to map text transcripts to 3D pose sequences corresponding to semantic-aware, full-body gestures of more human-like virtual agents using an end-to-end trainable transformer network and blend in emotional expressiveness.

2.4. Incorporating Speaker Styles

Co-speech gesture generation is intrinsically related to stylized gesture generation. There has been considerable progress on stylized generation of head motions [41], [42], facial motions [43], [28] as well as locomotions [44], [45], [46]. At the same time, many techniques have been proposed to generate appropriately styled body motions from textual descriptions of the actions [47], [48]. Other approaches have developed separate gesture generation networks for individual speakers to adapt to their individual styles [12], together with adversarial losses to improve the fidelity of the

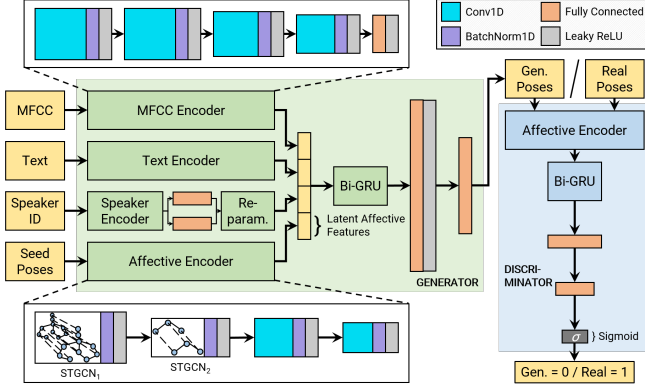


Figure 2: Our network consists of a generator (pale-green box) and a discriminator (pale-blue box). Our generator takes in the MFCC from the speech, the text transcript, the speaker ID, and a sequence of 3D seed poses. We use four encoders: the MFCC encoder (Sec. 3.1.1), the text encoder (Sec. 3.1.2), the speaker encoder (Sec. 3.1.3), and the affective encoder (Sec. 3.1.4). We feed the concatenation of these latent features into our Bi-GRU followed by a set of FC layers to synthesize the gestures aligned with the speech. Our discriminator learns to discriminate between the real and the synthesized gestures based on the latent affective features from the affective encoder, constraining the generator to synthesize appropriate affective expressions.

generation [49]. Recently, Yoon et al. [11] proposed a unified architecture that considers the speech, its text transcript, and the speaker identity to generate co-speech gestures with continuously varying speaker styles. We extend such speaker-aware gesture synthesis to further incorporate the appropriate affective body expressions that align with the affective content in the speech.

3. Approach

Our goal is to generate 3D pose sequences of co-speech upper-body gestures with appropriate affective expressions and speaker styles, given the raw speech waveform, the speaker identity, and a short sequence of seed poses. We consider affective expressions to be specific sequences of joint movements, generally as a combination of the affective features [33]. We learn these affective expressions both at the localized joint neighborhoods and the macroscopic body movements, and use them to condition the training of a generative adversarial network. We show our overall network architecture in Fig. 2.

3.1. Synthesizing Co-Speech Gestures

Our generative network takes in the raw speech waveform as a 1D array, the corresponding text transcript as a sequence of words, the speaker identity as a unique number, and the seed poses as a 3D pose sequence. Similar to Yoon et al. [11], we encode the speech waveform, the text transcript, and the speaker identities using separate encoders. However,

unlike Yoon et al., we convert the speech waveform to Mel-Frequency Cepstral Coefficients (MFCCs) to guide the encoding process based on the affective cues from speech. We also propose an affective encoder to encode the pose-based affective expressions into latent features for both gesture generation and discrimination. In the generation process, we combine the latent embeddings learned from the four encoders, speech, text, speaker, and affective, into a joint embedding for learning the upper-body gestures.

3.1.1. MFCC Encoder. MFCCs are known to encode signal frequencies consistent with how humans perceive sound, and are therefore particularly useful for tasks such as speech recognition [50], speaker identification [51] and speech-based emotion recognition [52]. In our case, we design our MFCC encoder to embed the speech-based affective cues such as prosody and intonations captured by the MFCCs and incorporate them in gesture synthesis. Given a raw waveform as a 1D array, we transform it to its top 14 MFCCs. These include the log-energy spectrum and 13 coefficients containing sufficient information on the speaker’s pitch, intonation, prosody, and other relevant parameters [53], [43]. We also append the first- and the second-order discrete forward differences of the 13 coefficients, obtaining a total of 37 values. Using a window size W on a input waveform of length L , we obtain individual MFCCs of shape $\lceil L/W \rceil$, leading to a combined feature tensor $f_m \in \mathbb{R}^{37 \times \lceil L/W \rceil}$. We pass these features through a series of 1D temporal convolutions, followed by a single fully-connected (FC) layer, to obtain a latent feature sequence $\hat{f}_m \in \mathbb{R}^{D_m \times T}$ of sequence length T equal to 3D pose sequence length of the seed poses, as

$$\hat{f}_m = \text{Conv} \circ \text{FC}_{mfcc}(f_m; W_{mfcc}), \quad (1)$$

where D_m is the dimension of the latent features, $\text{Conv} \circ \text{FC}_{mfcc}$ denotes the series of 1D convolutions followed by the FC layer, and W_{mfcc} its set of trainable parameters.

3.1.2. Text Encoder. Given the text transcript corresponding to the speech, we first pad the transcript with padding tokens following the approach of Yoon et al., to ensure that the text transcript has the same sequence length T as the seed poses. We then use the pre-trained FastText [54] word embedding model to transform the word sequence into 300-dimensional features, leading to a feature tensor $f_x \in \mathbb{R}^{300 \times T}$. We use FastText for its memory efficiency and its usefulness in sentiment analysis [55], which is important in understanding text-based affect. We pass the FastText features through a series of temporal 1D convolutions to obtain a latent feature sequence $\hat{f}_x \in \mathbb{R}^{D_x \times T}$, as

$$\hat{f}_x = \text{Conv}_{text}(f_x; W_{text}), \quad (2)$$

where D_x is the dimension of the latent features, Conv_{text} denotes the series of 1D convolutions with trainable parameters W_{text} .

3.1.3. Speaker Encoder. For the speaker IDs, we use one-hot vectors $f_s \in \{0, 1\}^S$, assuming S is the number of

available speakers. Following Yoon et al. [11], we use two sets of FC layers to learn an embedding space capturing the mean $\mu_s \in \mathbb{R}^{D_s}$ and the variance $\Sigma_s \in \mathbb{R}_+^{D_s \times D_s}$ of the latent distribution of the speaker styles as

$$\mu_s = \text{FC}_\mu(f_s; w_\mu) \quad (3)$$

$$\log \Sigma_s = \text{FC}_\Sigma(f_s; w_\Sigma), \quad (4)$$

where D_s is the dimension of the latent distribution space, FC_μ and FC_Σ denote the two sets of FC layers, and W_μ and W_Σ denote the corresponding sets of trainable parameters. Intuitively, this latent distribution space consists of all the available speakers, plus speakers that can be “constructed” by linear combinations of those speakers in the latent space. As a result, we can pick a random point from the latent space to use in the synthesis, resulting in some variability in the synthesized gestures even when the speech remains the same. We term this variability as having “speaker-aware” styles. Given the parameters μ_s and Σ_s of latent distribution space, we use the re-parametrization trick [56] to generate a random speaker-aware style sample $\hat{f}_s \in \mathbb{R}^{D_s}$ and repeat it for all the T time steps of the input pose sequence.

3.1.4. Affective Encoder. We propose an encoding mechanism that transforms the pose-based affective expressions into a latent embedding. Since gestures typically consist of movements in the trunk, arms, and head, we only consider ten joints corresponding to these parts of the body: root, spine, neck, head, left and right shoulders, left and right elbows, and left and right wrists. We consider a directed graph for the pose, where the joints are the vertices, and the edges are directed from the root towards the extremities. We assume the edge lengths are known for each input and train our encoder only on the directions of the edges. We consider nine unit-vector sequences $U = [u_1; \dots; u_9]$, each of sequence length T , to denote the edge directions at the corresponding T time steps of the input pose sequence.

We employ a hierarchical encoding strategy using spatial-temporal graph convolutions (STGCNs) [57]. STGCNs are adapted to leverage localized dependencies in generalized graph-structured data, and are therefore suitable for our pose graph sequences. We use two levels of hierarchy, the first at the level of individual bones and the second at the level of the three body parts, the trunk and the two arms. At the first level, our unweighted adjacency matrix $A_1 \in \{0, 1\}^{9 \times 9 \times T}$ captures the temporal counterparts of each edge at the four nearest time steps (past two and future two), and spatially adjacent edges with a maximum hop of two, *i.e.*, we consider two edges to be spatially adjacent if they either share a vertex or are connected to the two ends of a third edge. This size of the adjacent neighborhood sufficiently groups the edges influenced by typical affective expressions such as arm swings, head jerks, and upper-body collapse. Consequently, the convolution filters can learn a latent feature sequence $\hat{f}_{a_1} \in \mathbb{R}^{D_{a_1} \times 9 \times T}$ from the edges based on the variations in the affective expressions, obtained as

$$\hat{f}_{a_1} = \text{STGCN}_1(U, A_1; W_{a_1}), \quad (5)$$

where D_{a_1} is the dimension of the per-edge latent features, STGCN_1 denotes the first-level STGCN with trainable parameters W_{a_1} . At the second level, the three body parts, the trunk and the two arms, capture the macroscopic body movements such as raising or crossing the arms, and bending or straightening the trunk. In the second-level adjacency matrix $A_2 \in \{0, 1\}^{3 \times 3 \times T}$, we assume both the arms to be adjacent to the torso but not to each other, since the movements on one arm need not influence the other. We again consider the temporal counterparts of each body part in the four nearest time steps in the temporal adjacency. We reshape the latent features \hat{f}_{a_1} to $3D_{a_1} \times 3 \times T$, to collect the per-edge features corresponding to the three body parts in the feature dimension. Our second-level STGCN then operates on these reshaped features to produce the second-level latent features $\hat{f}_{a_2} \in \mathbb{R}^{D_{a_2} \times 3 \times T}$ as

$$\hat{f}_{a_2} = \text{STGCN}_2(\hat{f}_{a_1}, A_2; W_{a_2}), \quad (6)$$

where D_{a_2} is the dimension of the per-edge latent features, STGCN_2 denotes the second-level STGCN with trainable parameters W_{a_2} . We then apply a series of 1D convolutions on the reshaped second-level features $\hat{f}_{a_2} \in \mathbb{R}^{3D_{a_2} \times T}$ to obtain the latent affective feature sequence $\hat{f}_a \in \mathbb{R}^{D_a \times T}$, as

$$\hat{f}_a = \text{Conv}_{aff}(\hat{f}_{a_2}; W_a), \quad (7)$$

where D_a is the dimension of the latent affective features, and Conv_{aff} denotes the series of 1D convolutions with trainable parameters W_a .

3.1.5. Gesture Generator. Given the latent feature sequences \hat{f}_m , \hat{f}_x , \hat{f}_s , and \hat{f}_a , we concatenate them, pass them through a bidirectional gated recurrent unit (Bi-GRU), and sum the bidirectional outputs to obtain the predicted edge embeddings sequence $\hat{u}_e \in \mathbb{R}^{D_e \times T}$, as

$$\text{out}_{frw}, \text{out}_{bkw} = \text{GRU}_e\left(\left[\hat{f}_m; \hat{f}_x; \hat{f}_s; \hat{f}_a\right]; W_e\right), \quad (8)$$

$$\hat{u}_e = \text{out}_{frw} + \text{out}_{bkw}, \quad (9)$$

where D_e is the dimension of the predicted edge embeddings, GRU_e denotes the bidirectional GRU with the corresponding set of trainable parameters W_e , and out_{frw} and out_{bkw} respectively denote the outputs of the forward and the backward channels of the GRU. As in Yoon et al., we then transform the predicted edge embeddings to predicted edge vector sequences $\hat{U} = [\hat{u}_1; \dots; \hat{u}_9]$, each of sequence length T , using a set of FC layers as

$$\hat{U} = \text{FC}_{gen}(\hat{u}_e; W_{gen}), \quad (10)$$

where FC_{gen} denotes the set of FC layers with the trainable parameters W_{gen} . Thus, our generator is designed to take in a sequence of seed poses of length T and predicts a pose sequence of gestures for the next T time steps. Finally, we scale each predicted edge vector \hat{u}_i to have the corresponding bone length b_i , $i = 1, \dots, 9$. We add it to the 3D position $\text{pos}_{s(i)}$ of the source joint $s(i)$ of that edge vector to obtain 3D position $\text{pos}_{d(i)}$ of the destination joint

TABLE 1: Hyperparameters (HPs) for our network. We chose all the values via empirical search.

HP	Description	Value
D_m	Latent feature from the MFCC encoder	32
D_x	Latent feature from the text encoder	32
D_s	Latent distribution space of speaker styles	16
D_{a_1}	Per-edge latent features after STGCN ₁ in the affective encoder	16
D_{a_2}	Per-edge latent features after STGCN ₂ in the affective encoder	16
D_a	Latent affective features from the affective encoder	16
D_e	Predicted edge embeddings from the GRU in the generator	150
D_d	Predicted embeddings from the GRU in the discriminator	150

$d(i)$ of the same edge vector, as

$$pos_{d(i)} = pos_{s(i)} + b_i \cdot \frac{\hat{u}_i}{\|\hat{u}_i\|}. \quad (11)$$

3.2. Discriminating Gestures

Our discriminator takes in a gesture of sequence length T and computes its latent affective feature sequence $\hat{f}_a \in \mathbb{R}^{D_a \times T}$ using our affective encoder (Sec. 3.1.4). We pass this feature sequence through another bidirectional GRU, and sum the bidirectional outputs to obtain the discriminator embeddings sequence $\hat{d} \in \mathbb{R}^{h \times T}$, as

$$out_{frw}, out_{bkw} = \text{GRU}_{disc}(\hat{f}_a; W_{GRU, disc}), \quad (12)$$

$$\hat{d} = out_{frw} + out_{bkw}, \quad (13)$$

where D_d is the dimension of the predicted discriminator embeddings, GRU_{disc} denotes the bidirectional GRU with trainable parameters $W_{GRU, disc}$, and out_{frw} and out_{bkw} respectively denote the outputs of the forward and the backward channels of the GRU. We then transform the discriminator embeddings to a probability vector $c \in [0, 1]$ using a set of FC layers as

$$c = \text{FC}_{disc}(\hat{d}; W_{FC, disc}), \quad (14)$$

where FC_{disc} denotes the set of FC layers with trainable parameters $W_{FC, disc}$, and c is such that $c \geq 0.5$ implies the discriminator predicts the input gesture to be real, and generated otherwise.

4. Dataset and Training

We train our network on the TED Gesture Dataset [1], which consists of videos of English-language speakers at TED Talks. It provides 3D pose sequences of the upper-body gestures of the speakers, their speech audio, and the associated text transcripts. Each data sample has a sequence length of $T = 34$ time steps at a rate of 15 fps. There are 200,038 training samples in total, constituting around 80% of the dataset. The evaluation set consists of 26,903 samples or around 10% of the dataset. The test set consists of 26,245 samples, making up the remaining 10% of the dataset.

TABLE 2: Evaluation of our method with baselines and ablated versions of our method on two benchmark dataset, using the objective metrics of mean absolute joint error (MAJE), mean acceleration difference (MAD), and the Fréchet Gesture Distance (FGD). Bold indicates best.

Dataset	Method	MAJE (mm)	MAD (mm/s ²)	FGD
TED Gesture [1]	Seq2Seq [1]	45.62	6.33	6.62
	S2G-IS [12]	45.11	7.22	6.73
	JEM [60]	48.56	4.31	5.88
	GTC [11]	27.30	3.20	4.49
	Ours w/o MFCC Enc.	27.84	3.02	4.21
	Ours w/o Aff. Enc.	25.38	3.51	4.84
	Ours	24.49	2.93	3.54
GENEA Challenge 2020 [19]	Gesticulator [9]	82.41	3.62	31.04
	Ours w/o MFCC Enc.	105.71	1.57	23.03
	Ours w/o Aff. Enc.	92.90	2.81	24.28
	Ours	54.93	1.49	20.36

We use loss functions \mathcal{L}_G and \mathcal{L}_D identical to Yoon et al. [11] to train our generator and discriminator respectively:

$$\mathcal{L}_G = \lambda_{Hub} \mathcal{L}_{Hub} + \lambda_{gen} \mathcal{L}_{gen} + \lambda_{stl} \mathcal{L}_{stl} + \lambda_{KLD} \mathcal{L}_{KLD}, \quad (15)$$

$$\mathcal{L}_D = -\mathbb{E}[\log(\text{Disc}(U))] - \mathbb{E}[\log(1 - \text{Disc}(\hat{U}))], \quad (16)$$

where Disc denotes the discriminator network (Sec. 3.2), λ_* are the weights of the corresponding loss terms with the same values as in Yoon et al. [11], and the individual loss terms of the generator are:

- **Huber loss** [58] between the ground truth and predicted edge vectors,
- **generative adversarial loss** on the output of the discriminator,

$$\mathcal{L}_{gen} = -\mathbb{E}[\log(\text{Disc}(\hat{U}))], \quad (17)$$

- **diversity regularization** between the synthesized gestures and other gestures in the dataset to ensure that the styles of different speakers appear visually different,
- **Kullback-Leibler (KL) divergence** between the latent distribution space of the styles defined by μ_s and Σ_s , and the normal distribution $\mathcal{N}(0, I)$.

Table 1 lists the latent dimensions we use for training our network. We use the Adam optimizer [59] with $\beta_1 = 0.5$, $\beta_2 = 0.999$, batch size of 512, and learning rate of 5E^{-4} for the generator and 1E^{-4} for the discriminator with no warm-up epochs (*i.e.*, $\lambda_{gen} > 0$ starting from the first epoch). We train our network for 300 epochs, which took close to 45 hours on an NVIDIA GeForce GTX 1080 Ti GPU.

5. Experiments

We describe the objective evaluation of our method compared to current baseline methods. We highlight the benefit of our proposed components via ablation studies on the objective evaluation metrics. We also show the qualitative performance of our method on selected samples from the TED Gesture Dataset [1] and the perceived quality of our synthesized gestures through a user study.





























































Input Speech	... process of research, I was <i>very excited</i> to find I was, I believe, <i>ecologically bored</i> .					
Original Speaker												
GTC [57]												
Ours w/o MFCC Encoder												
Ours w/o Affective Encoder												
Ours												

Figure 3: Qualitative results on the gestures synthesized by our method for two sample speech excerpts from the TED Gesture Dataset [1]. The italicized words *very excited* and *bored* indicate the primary affect in the corresponding speeches. We compare with the corresponding gestures of the original speakers, the output of GTC [11], and that of the two ablated versions of our network (Sec. 5.3). See Sec. 5.4 for a detailed discussion of the results.

5.1. Baseline Methods

We compare our method with the baseline methods on two benchmark datasets, the TED Gesture Dataset [1], and the GENE Challenge 2020 Dataset [18], [19].

On the TED Gesture Dataset, we compare with the methods of Seq2Seq [1], Speech to Gestures with Individual Styles (S2G-IS) [12], Joint Embedding Model (JEM) [60], and Gestures from Trimodal Context (GTC) [11]. Seq2Seq and JEM generate gestures based only on the text transcript of the speech, whereas S2G-IS uses only the speech to generate the gestures. GTC uses the speech, the corresponding text transcript, and the speaker styles to generate gestures. Seq2Seq follows an encoder-decoder architecture, where the authors transform the text to latent features and predict gestures based on both the latent features and a short gesture history. The authors of S2G-IS employ a generative adversarial network that generates gestures from a latent space obtained from the input log-Mel spectrograms. JEM maps both the text and the target gesture into a common latent embedding space and uses a decoder to reconstruct the gestures from the embedding space. The authors train the model to learn to align the text-based and the gesture-based embeddings for the same input and decode gestures from only the text-based embeddings. For Seq2Seq, S2G-IS, and JEM, we follow the training routine and the hyperparameters used by Yoon et al. [11]. For GTC, we directly use the pre-trained model provided by Yoon et al. [11].

The GENE Challenge 2020 Dataset is the publicly available version of the Trinity Gesture Dataset [18], [19]. It consists of the speech and the full-body motion capture of

a male actor talking unrestrained about various topics over multiple recording sessions. The full dataset is about 242 minutes long, of which 221 minutes are used as training data, and the remaining 21 minutes are kept for testing. We do not fine-tune our network on this dataset and evaluate our network on the test partition. Since we consider only upper-body gestures, we consider the ten relevant upper-body joints at the root, the spine, the head, and the two arms for evaluating our performance. On this dataset, we compare with the method of Gesticulator [9], which leverages the acoustics and the semantics of the speech to generate semantically consistent beat, deictic, metaphoric, and iconic gestures. For a fair comparison, we use the pre-trained model provided by the authors and compare the performance on the same ten joints that we use for our method.

5.2. Objective Evaluation

While evaluation metrics for gesture synthesis are not standardized, we evaluate on the commonly used metrics of mean absolute joint error (MAJE), mean acceleration difference (MAD), and the Fréchet Gesture Distance (FGD) proposed by Yoon et al. [11]. MAJE measures the mean of the absolute differences between the ground truth and the predicted joint positions over all the time steps, joints, and samples. MAD measures the mean ℓ_2 -norm error between the ground truth and predicted joint accelerations over all the time steps, joints, and samples. FGD measures the difference between the distributions of the latent features of the ground truth and the predicted gestures. The latent features are computed from an autoencoder network trained on the

well-known Human 3.6M dataset [61] of human motions using the ten joints in the TED Gesture Dataset [1]. MAJE indicates how closely the predicted joint positions follow the ground truth joint positions. MAD indicates how closely the ground truth and predicted joint movements match. Since affective expressions are based on joint movements, a lower MAD is especially desirable for our stated aim of generating gestures with appropriate affective expressions. FGD is shown to align well with the perceived plausibility of the synthesized gestures to human users [11]; therefore, a lower FGD is equally desirable to gauge the quality of our synthesized gestures.

Table 2 summarizes the performance of all the methods on all these evaluation metrics. Our method consistently has the lowest MAJE, MAD, and FGD on both the benchmark datasets. On the TED Gesture Dataset, we observe improvements of 10.29%, 8.44%, and 21.16% on MAJE, MAD, and FGD, respectively, over the best current baseline of GTC. On the GENE Challenge 2020 Dataset, we observe improvements of 33.34%, 58.84%, and 34.41% on MAJE, MAD, and FGD, respectively, over the baseline of Gesticulator. We note that the absolute FGD values are significantly higher on the GENE Challenge 2020 Dataset than on the TED Gesture Dataset. We hypothesize that this is because the gestures in the GENE Challenge 2020 Dataset are more abstract and unscripted compared to the well-defined actions in the Human 3.6M Dataset or polished speeches in the TED Gesture Dataset. As a result, the pre-trained latent embeddings used for FGD are not as good at reconstructing the joints movements in the GENE Challenge 2020 Dataset.


5.3. Ablation Studies

We perform ablation studies on our two proposed components: the MFCC encoder (Sec. 3.1.1) and the affective encoder (Sec. 3.1.4). In one study, we replace only our MFCC encoder with an encoder for the raw audio waveform, identical to that of GTC [11], and train the resultant network. In the other study, we remove the affective encoder from our network. Our generator takes in the raw seed poses instead of the latent affective features. Our discriminator uses a convolution filter identical to GTC [11] to transform the input gestures to latent features for the bidirectional GRU.


Without the MFCC encoder, our generator cannot take the speech-based affective cues into account. It results in a degradation of the synthesis, leading to higher MAJE, MAD, and FGD. However, without the affective encoder, our network is severely limited in understanding both the affective expressions in the seed poses and the affective expressions of the synthesized poses. It results in more severe performance degradation on all the evaluation metrics, and the synthesized gestures appear less diverse and plausible.

5.4. Qualitative Results

We show qualitative results on two sample speech excerpts from the TED Gesture Dataset [1] in Fig. 3. It has five rows of gestures: those of the original speakers’, those



(a) Plausibility of the different types of gestures.



(b) Synchronization of the movements and the affective expressions of the different types of gestures with the speech.

Figure 4: Mean fraction of participant responses on each point of the Likert scales across the 12 speech excerpts from the TED Gesture Dataset [1] and the corresponding gestures in our user study. See Sec. 5.5 for details.

synthesized by GTC [11] (the current state-of-the-art), those by the two ablated versions of our network: one without our MFCC encoder (Sec. 3.1.1) and the other without our affective encoder (Sec. 3.1.4), and those by our proposed network with all the encoders. We observe a diversity of speaker styles in the synthesized gestures compared to the original speaker, which results from using a variational embedding of speaker styles using the speaker encoder (Sec. 3.1.3). GTC, however, cannot generate affective expressions except for a few words with strong intonations in the speech, such as “excited” (second row, left column). Without our MFCC encoder, our network can still match the speech content but cannot align the gestures with the affective cues from the speech. For example, it can match the words “I was, I believe” with a deictic gesture pointing to the speaker himself (third row, right column) but cannot generate any expressions for “bored”. Without our affective encoder, we observe only slight body movements but no appreciable affective expressions in the synthesized gestures. With all our encoders in place, we observe appropriate affective expressions that align well with the speech. For example, we observe rapid arm movements when saying “excited” (fifth row, left column) and dropping of the arms and shoulders when saying “bored” (fifth row, right column).

5.5. User Study

We conducted a user study to evaluate the perceptual quality of our synthesized gestures in terms of how plausible they appear and how well-aligned are their affective expressions with the corresponding speeches. 24 participants took part in our study, of which 20 were male, and 4 were female. 10 participants were between 18 and 24 years of age, 13 were between 25 and 34, and one was above 35. Each participant observed gestures corresponding to the same 12 speech excerpts, each taken from a different TED Talk in the TED Gesture Dataset [1]. For each speech excerpt, the participants observed three different types of gestures: that

of the original speaker as a 3D pose sequence (provided in the dataset), those synthesized by GTC [11], the current state-of-the-art, and those synthesized by our network. The order of the gestures was unknown to and randomized for each participant. We then asked the participants to answer two questions. The first question was how plausible the gestures appeared on a five-point Likert scale ranging from “very unnatural” (1) to “look like they could be from a real person” (5). The second question was how well the gestures synchronized with the corresponding speeches on a five-point Likert scale, ranging from “no or arbitrary gestures” (1) to “well-synchronized with the speech, and are appropriately emotionally expressive” (5). Intuitively, our Likert-scale points for both questions reflect the participants’ individual assessments of quality, with 1 being the worst, 3 being average, and 5 being the best. The entire study took around 20 minutes on average for each participant.

We summarize the participants’ responses in Fig. 4. When adjudging the plausibility of the gestures (Fig. 4a), we observe that 15.28% more participants marked our synthesized gestures either 4 or 5 compared to the gestures synthesized by GTC [11]. Further, 3.82% more participants marked our synthesized gestures 4 or 5 than the original speakers’ gestures, indicating that the participants found our synthesized gestures to have visual quality comparable to that of the original data. When adjudging the synchronization of the movements and the affective expressions of different types of gestures with speech (Fig. 4b), we observed that 16.32% more participants marked our synchronization quality either 4 or 5 compared to that of GTC [11]. Also, 4.86% more participants marked out synchronization quality 4 or 5 than that of the original speakers, indicating that the participants perceived our synthesized gestures to be as well-synchronized and expressive as the original data.

6. Conclusion, Limitations and Future Work

We have presented an end-to-end learning approach to generate 3D pose sequences of co-speech gestures with appropriate affective expressions. Our contributions include an MFCC encoder to guide the gesture synthesis based on the speech-based affective cues such as prosody and intonation, and an affective encoder to learn joint-based affective features from the gesture data. Using these encoders in a generative adversarial learning framework, we have synthesized affective gestures that advance the state-of-the-art on co-speech gesture synthesis on multiple evaluation metrics. Our synthesized gestures also appeared more plausible and well-synced with the corresponding speeches to participants in a user study.

Our work has some limitations. First, we do not build a mechanism to control the affective expressions; we compute them automatically from the input modalities. We plan to investigate the interface between affective expressions from the speech and the gestures, especially when expressing contradictory cues such as sarcasm and irony. Second, we plan to use a finer representation of poses in the future since affective expressions in the gestures are often associated

with subtle movements not captured by our current representation. Lastly, our network uses only upper-body gestures with a fixed root. We plan to expand to affect-aware gestures of the whole body, incorporating locomotion in the global coordinate space. We also plan to align the body gestures with the corresponding facial expressions, leading to fully emotive characters.

Acknowledgment

This work has been supported in part by ARO Grants W911NF1910069 and W911NF1910315, and Intel.

References

- [1] Y. Yoon, W.-R. Ko, M. Jang, J. Lee, J. Kim, and G. Lee, “Robots learn social skills: End-to-end learning of co-speech gesture generation for humanoid robots,” in *Proc. of The International Conference in Robotics and Automation (ICRA)*, 2019. [Online]. Available: https://github.com/youngwoo-yoon/Co-Speech_Gesture_Generation
- [2] D. McNeill, *Hand and mind: What gestures reveal about thought*. University of Chicago press, 1992.
- [3] J. Li, R. Kizilcec, J. Bailenson, and W. Ju, “Social robots and virtual agents as lecturers for video instruction,” *Computers in Human Behavior*, vol. 55, pp. 1222 – 1230, 2016. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0747563215002885>
- [4] M. Liao, C. Sung, H. Wang, and W. Lin, “Virtual classmates: Embodying historical learners’ messages as learning companions in a vr classroom through comment mapping,” in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, 2019, pp. 163–171.
- [5] A. L. Simeone, M. Speicher, A. Molnar, A. Wilde, and F. Daiber, “Live: The human role in learning in immersive virtual environments,” in *Symposium on Spatial User Interaction*, ser. SUI ’19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: <https://doi.org/10.1145/3357251.3357590>
- [6] T. Baur, I. Damian, P. Gebhard, K. Porayska-Pomsta, and E. André, “A job interview simulation: Social cue-based interaction with a virtual character,” in *2013 International Conference on Social Computing*, 2013, pp. 220–227.
- [7] D. DeVault, R. Artstein, G. Benn, T. Dey, E. Fast, A. Gainer, K. Georgila, J. Gratch, A. Hartholt, M. Lhommet *et al.*, “Simsensei kiosk: A virtual human interviewer for healthcare decision support,” in *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 2014, pp. 1061–1068.
- [8] S. Mascarenhas, M. Guimarães, R. Prada, J. Dias, P. A. Santos, K. Star, B. Hirsh, E. Spice, and R. Kommeren, “A virtual agent toolkit for serious games developers,” in *2018 IEEE Conference on Computational Intelligence and Games (CIG)*, 2018, pp. 1–7.
- [9] T. Kucherenko, P. Jonell, S. van Waveren, G. E. Henter, S. Alexandersson, I. Leite, and H. Kjellström, “Gesticulator: A framework for semantically-aware speech-driven gesture generation,” ser. ICMI ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 242–250. [Online]. Available: <https://doi.org/10.1145/3382507.3418815>
- [10] K. Watson, S. S. Sohn, S. Schriber, M. Gross, C. M. Muniz, and M. Kapadia, “Storyprint: An interactive visualization of stories,” in *Proceedings of the 24th International Conference on Intelligent User Interfaces*, ser. IUI ’19. New York, NY, USA: Association for Computing Machinery, 2019, p. 303–311. [Online]. Available: <https://doi.org/10.1145/3301275.3302302>
- [11] Y. Yoon, B. Cha, J.-H. Lee, M. Jang, J. Lee, J. Kim, and G. Lee, “Speech gesture generation from the trimodal context of text, audio, and speaker identity,” *ACM Transactions on Graphics*, vol. 39, no. 6, 2020.

- [12] S. Ginosar, A. Bar, G. Kohavi, C. Chan, A. Owens, and J. Malik, "Learning individual styles of conversational gesture," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [13] A. Kleinsmith and N. Bianchi-Berthouze, "Affective body expression perception and recognition: A survey," *IEEE Transactions on Affective Computing*, vol. 4, no. 1, pp. 15–33, 2013.
- [14] M. Karg, A. Samadani, R. Gorbet, K. Kühnlenz, J. Hoey, and D. Kulić, "Body movements for affective expression: A survey of automatic recognition and generation," *IEEE Transactions on Affective Computing*, vol. 4, no. 4, pp. 341–359, 2013.
- [15] D. Matsumoto, M. G. Frank, and H. S. Hwang, *Nonverbal communication: Science and applications*. Sage Publications, 2012.
- [16] M. L. Knapp, J. A. Hall, and T. G. Horgan, *Nonverbal communication in human interaction*. Cengage Learning, 2013.
- [17] U. Bhattacharya, C. Roncal, T. Mittal, R. Chandra, K. Kapsaskis, K. Gray, A. Bera, and D. Manocha, "Take an emotion walk: Perceiving emotions from gaits using hierarchical attention pooling and affective mapping," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 145–163.
- [18] Y. Ferstl and R. McDonnell, "Iva: Investigating the use of recurrent motion modelling for speech gesture generation," in *IVA '18 Proceedings of the 18th International Conference on Intelligent Virtual Agents*, Nov 2018. [Online]. Available: <https://trinityspeechgesture.scss.tcd.ie>
- [19] T. Kucherenko, P. Jonell, Y. Yoon, P. Wolfert, and G. E. Henter, "A large, crowdsourced evaluation of gesture generation systems on common data: The GENE Challenge 2020," in *26th International Conference on Intelligent User Interfaces*, ser. IUI '21. New York, NY, USA: Association for Computing Machinery, 2021, pp. 11–21. [Online]. Available: <https://doi.org/10.1145/3397481.3450692>
- [20] A. Mehrabian and J. A. Russell, *An approach to environmental psychology*. the MIT Press, 1974.
- [21] G. Castillo and M. Neff, "What do we express without knowing? emotion in gesture," in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS '19. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2019, p. 702–710.
- [22] A. Banerjee, U. Bhattacharya, and A. Bera, "Learning unseen emotions from gestures via semantically-conditioned zero-shot perception with adversarial autoencoders," *arXiv preprint arXiv:2009.08906*, 2020.
- [23] S. Kim, S. J. Guy, W. Liu, D. Wilkie, R. W. Lau, M. C. Lin, and D. Manocha, "Brvo: Predicting pedestrian trajectories using velocity-space reasoning," *The International Journal of Robotics Research*, vol. 34, no. 2, pp. 201–217, 2015.
- [24] U. Bhattacharya, T. Mittal, R. Chandra, T. Randhavane, A. Bera, and D. Manocha, "Step: Spatial temporal graph convolutional networks for emotion perception from gaits," in *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*, ser. AAAI'20. AAAI Press, 2020, p. 1342–1350.
- [25] A. Chowanda, P. Blanchfield, M. Flintham, and M. Valstar, "Computational models of emotion, personality, and social relationships for interactions in games: (extended abstract)," in *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS '16. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2016, p. 1343–1344.
- [26] S. S. Sohn, X. Zhang, F. Geraci, and M. Kapadia, "An emotionally aware embodied conversational agent," in *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS '18. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2018, p. 2250–2252.
- [27] T. Karras, T. Aila, S. Laine, A. Herva, and J. Lehtinen, "Audio-driven facial animation by joint end-to-end learning of pose and emotion," *ACM Trans. Graph.*, vol. 36, no. 4, Jul. 2017. [Online]. Available: <https://doi.org/10.1145/3072959.3073658>
- [28] Y. Ferstl and R. McDonnell, "A perceptual study on the manipulation of facial features for trait portrayal in virtual agents," in *Proceedings of the 18th International Conference on Intelligent Virtual Agents*, ser. IVA '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 281–288. [Online]. Available: <https://doi.org/10.1145/3267851.3267891>
- [29] T. Randhavane, A. Bera, K. Kapsaskis, R. Sheth, K. Gray, and D. Manocha, "Eva: Generating emotional behavior of virtual agents using expressive features of gait and gaze," in *ACM Symposium on Applied Perception 2019*. ACM, 2019, p. 6.
- [30] J. H. Chuah, B. Rossen, and B. Lok, "Automated generation of emotive virtual humans," in *Intelligent Virtual Agents*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 490–491.
- [31] T. Randhavane, A. Bera, K. Kapsaskis, U. Bhattacharya, K. Gray, and D. Manocha, "Identifying emotions from walking using affective and deep features," *arXiv preprint arXiv:1906.11884*, 2019.
- [32] U. Bhattacharya, N. Rewkowski, P. Guhan, N. L. Williams, T. Mittal, A. Bera, and D. Manocha, "Generating emotive gaits for virtual agents using affect-based autoregression," in *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2020, pp. 24–35.
- [33] U. Bhattacharya, N. Rewkowski, A. Banerjee, P. Guhan, A. Bera, and D. Manocha, "Text2gestures: A transformer-based network for generating emotive body gestures for virtual agents," in *2021 IEEE Conference on Virtual Reality and 3D User Interfaces (IEEE VR)*. IEEE, 2021.
- [34] P. Wagner, Z. Malisz, and S. Kopp, "Gesture and speech in interaction: An overview," *Speech Communication*, vol. 57, pp. 209 – 232, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167639313001295>
- [35] C.-C. Chiu, L.-P. Morency, and S. Marsella, "Predicting co-verbal gestures: A deep and temporal modeling approach," in *Intelligent Virtual Agents*. Cham: Springer International Publishing, 2015, pp. 152–166.
- [36] S. Levine, P. Krähenbühl, S. Thrun, and V. Koltun, "Gesture controllers," in *ACM SIGGRAPH 2010 Papers*, ser. SIGGRAPH '10. New York, NY, USA: Association for Computing Machinery, 2010. [Online]. Available: <https://doi.org/10.1145/1833349.1778861>
- [37] D. Hasegawa, N. Kaneko, S. Shirakawa, H. Sakuta, and K. Sumi, "Evaluation of speech-to-gesture generation using bi-directional lstm network," in *Proceedings of the 18th International Conference on Intelligent Virtual Agents*, ser. IVA '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 79–86. [Online]. Available: <https://doi.org/10.1145/3267851.3267878>
- [38] T. Kucherenko, D. Hasegawa, G. E. Henter, N. Kaneko, and H. Kjellström, "Analyzing input and output representations for speech-driven gesture generation," in *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, ser. IVA '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 97–104. [Online]. Available: <https://doi.org/10.1145/3308532.3329472>
- [39] S. Alexanderson, G. E. Henter, T. Kucherenko, and J. Beskow, "Style-controllable speech-driven gesture synthesis using normalising flows," *Computer Graphics Forum*, vol. 39, no. 2, pp. 487–496, 2020. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13946>
- [40] N. Sadoughi and C. Busso, "Speech-driven animation with meaningful behaviors," *Speech Communication*, vol. 110, pp. 90 – 100, 2019. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167639318300013>

- [41] D. Greenwood, S. Laycock, and I. Matthews, "Predicting head pose from speech with a conditional variational autoencoder." *ISCA*, 2017.
- [42] N. Sadoughi and C. Busso, "Novel realizations of speech-driven head movements with generative adversarial networks," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 6169–6173.
- [43] S. Suwajanakorn, S. M. Seitz, and I. Kemelmacher-Shlizerman, "Synthesizing obama: Learning lip sync from audio," *ACM Trans. Graph.*, vol. 36, no. 4, Jul. 2017. [Online]. Available: <https://doi.org/10.1145/3072959.3073640>
- [44] D. Holden, J. Saito, and T. Komura, "A deep learning framework for character motion synthesis and editing," *ACM Trans. Graph.*, vol. 35, no. 4, Jul. 2016. [Online]. Available: <https://doi.org/10.1145/2897824.2925975>
- [45] D. Holden, T. Komura, and J. Saito, "Phase-functioned neural networks for character control," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, p. 42, 2017.
- [46] S. Starke, H. Zhang, T. Komura, and J. Saito, "Neural state machine for character-scene interactions," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 6, p. 209, 2019.
- [47] Y. Li, M. R. Min, D. Shen, D. E. Carlson, and L. Carin, "Video generation from text," in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, ser. AAAI'18. AAAI Press, pp. 7065–7072.
- [48] X. Zhou, S. Huang, B. Li, Y. Li, J. Li, and Z. Zhang, "Text guided person image synthesis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [49] Y. Ferstl, M. Neff, and R. McDonnell, "Multi-objective adversarial gesture generation," in *Motion, Interaction and Games*, ser. MIG '19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: <https://doi.org/10.1145/3359566.3360053>
- [50] R. Vergin, D. O'Shaughnessy, and A. Farhat, "Generalized mel frequency cepstral coefficients for large-vocabulary speaker-independent continuous-speech recognition," *IEEE Transactions on Speech and Audio Processing*, vol. 7, no. 5, pp. 525–532, 1999.
- [51] K. S. R. Murty and B. Yegnanarayana, "Combining evidence from residual phase and mfcc features for speaker recognition," *IEEE Signal Processing Letters*, vol. 13, no. 1, pp. 52–55, 2006.
- [52] D. Neiberg, K. Elenius, and K. Laskowski, "Emotion recognition in spontaneous speech using gmms," in *Ninth international conference on spoken language processing*, 2006.
- [53] A. Bagher Zadeh, P. P. Liang, S. Poria, E. Cambria, and L.-P. Morency, "Multimodal language analysis in the wild: CMU-MOSEI dataset and interpretable dynamic fusion graph," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia: Association for Computational Linguistics, Jul. 2018, pp. 2236–2246. [Online]. Available: <https://www.aclweb.org/anthology/P18-1208>
- [54] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching word vectors with subword information," *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 135–146, 2017.
- [55] I. Santos, N. Nedjah, and L. de Macedo Mourelle, "Sentiment analysis using convolutional neural network with fasttext embeddings," in *2017 IEEE Latin American Conference on Computational Intelligence (LA-CCI)*, 2017, pp. 1–5.
- [56] D. P. Kingma and M. Welling, "An introduction to variational autoencoders," *arXiv preprint arXiv:1906.02691*, 2019.
- [57] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, April 2018. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/12328>
- [58] P. J. Huber, "A robust version of the probability ratio test," *The Annals of Mathematical Statistics*, pp. 1753–1758, 1965.
- [59] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014. [Online]. Available: <https://arxiv.org/pdf/1412.6980v9.pdf>
- [60] C. Ahuja and L. Morency, "Language2pose: Natural language grounded pose forecasting," in *2019 International Conference on 3D Vision (3DV)*, 2019, pp. 719–728.
- [61] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu, "Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 7, pp. 1325–1339, 2014.