



# ZeroEGGS: Zero-shot Example-based Gesture Generation from Speech

Saeed Ghorbani,<sup>1,2</sup>  Ylva Ferstl,<sup>1</sup>  Daniel Holden,<sup>1</sup> Nikolaus F. Troje<sup>2</sup>  and Marc-André Carbonneau<sup>3</sup> 

<sup>1</sup>Ubisoft, Toronto, Canada

{saeed.ghorbani, ylva.ferstl}@ubisoft.com

<sup>2</sup>York University, Toronto, Canada

<sup>3</sup>Ubisoft, Montreal, Canada

## Abstract

We present ZeroEGGS, a neural network framework for speech-driven gesture generation with zero-shot style control by example. This means style can be controlled via only a short example motion clip, even for motion styles unseen during training. Our model uses a Variational framework to learn a style embedding, making it easy to modify style through latent space manipulation or blending and scaling of style embeddings. The probabilistic nature of our framework further enables the generation of a variety of outputs given the input, addressing the stochastic nature of gesture motion. In a series of experiments, we first demonstrate the flexibility and generalizability of our model to new speakers and styles. In a user study, we then show that our model outperforms previous state-of-the-art techniques in naturalness of motion, appropriateness for speech, and style portrayal. Finally, we release a high-quality dataset of full-body gesture motion including fingers, with speech, spanning across 19 different styles. Our code and data are publicly available at <https://github.com/ubisoft/ubisoft-laforge-ZeroEGGS>.

**Keywords:** animation, gestures, character control, motion capture

**CCS Concepts:** • Computing methodologies → Machine learning approaches; Animation

## 1. Introduction

In human communication, gestures complement speech by providing additional information about thoughts, feelings, emotions, and intentions [ML04, DRBD12]. Efficiently animating realistic gesture behaviour is a key concern in many applications based on virtual humans, such as characters in video games and extended realities or customer service agents.

A common method for the automation of gesture animation is to trigger pre-recorded animations from a database based on tags. One can manually mark-up a dialog or program systemic events. While effective, this method requires significant amounts of time and manpower in order to produce a variety of engaging and realistic animations. Therefore this method does not scale up well with increasingly large open world video games and cannot be applied to avatars in virtual worlds where dialog content is not scripted. Moreover, pre-recorded animation may not be synchronized with speech rhythm.

These problems of scaling and synchronicity have motivated research into methods for automatic generation of gestures with

data-driven methods [AHKB20, YCL\*20]. Nonetheless, despite recent research efforts, generating realistic gesture motion remains a difficult problem, and addressing expressivity of speaker state and identity by providing control over style is even more challenging.

Speech is commonly used as a control input for gesture generation systems. Yet it is a weak control signal because a speech line may be associated with many different motions. This means that there is not only one true gesture match, but many equally appropriate gestures. In other words this is a one-to-many mapping problem. Moreover, every person exhibits their own style of movement, and therefore, producing relatable and engaging stories with variety of characters requires the ability to generate such distinct and appropriate styles. This implies that the model must capture this wide range motion variation, preserving characters' idiosyncrasies.

Previous works have sought to address the problem of creating distinct styles by modelling and generating gestures for specific speakers [NKAS08, GBK\*19, YCL\*20, ALNM20] and by



**Figure 1:** Samples generated by our model in different styles. Left to right: Neutral, Happy, Angry, Relaxed, Old, Sad, Oration.

modifying gesture motion through general statistics such as hand height and velocity [AHKB20, YPJ\*21]. These approaches lack flexibility because they are limited by the content of the training data. They require examples of every target speakers and every style prior to training the model, and cannot generalize outside of this range. This means a specific dataset needs to be captured for each individual and each style, which leads to a prohibitive amount of work in larger scale applications. To scale up, a gesture generation method must be able to capture an individual style with a very limited amount of data, ideally only one example. Secondly, motion style can be difficult to capture in words, rather, it can be easier to describe style by providing an example of the desired motion style.

In this work, we propose ZeroEGGS, a system that generates stylized full body gestures from speech. Unlike existing solutions, ZeroEGGS encodes gesture style from short example motions, which allows for zero-shot style transfer. This means that it can generalize to styles that were not covered by the training data, and does not require style labels. Moreover, the model learns a meaningful representation that enables style manipulation directly in the latent space. Our framework is probabilistic, allowing for repeated sampling to obtain a variety of output motion given the same input speech and style. Finally, we release a high-quality multi-modal dataset of synchronized speech and full-body gesture in 19 different styles. We also make our code publicly available for reproducibility<sup>1</sup>.

## 2. Related Work

Research addressing automatic production of gesture motion from speech can be divided into two main approaches: rule-based, and data-driven methods.

Rule-based methods use explicit, often manually created, mappings of speech markers to gestures. Some methods rely on text analysis [CVB01, LM06, KJLW03], while others use acoustic features [MXL\*13]. Rule-based approaches provide author-control, high motion definition and facilitate the generation of semantic gestures. However, motion diversity is limited by the number of de-

signed rules and database content. Moreover, the required manual labor hinders scalability.

Data-driven methods model correlations between speech and motion features rather than relying on hand-crafted mappings. Some methods still require manual labour due to using hand-annotated features such as the gesture shape [NKAS08] or timing of stressed phrases [YYH20]. The mapping between the motion and speech features is learned automatically. At inference, motion is produced by concatenating snippets from a database, which retains captured motion quality. Other approaches generate motion instead of drawing from a database. A model must learn a mapping between speech and gesture, which is a one-to-many problem. In that situation, minimizing the error between predicted and target motion often leads to mean collapsing problems and lethargic motion with small ranges (e.g., [FM18] and [KJvW\*20]). Generative adversarial networks have been used to address this problem with varying degree of success depending on the datasets size and gesture complexity [FNM19, GBK\*19, RGP21].

Despite recent research efforts, the naturalness of generated gestures still falls significantly behind motion-captured gestures, and often fails to outperform even mismatched real motion w.r.t. appropriateness [KJY\*21, GBK\*19].

While the aforementioned methods condition generation only on speech, several recent efforts propose to also provide control over gesture style. In [AHKB20], style control is achieved via input of four desired motion statistics, specifically the gesture speed, height, spacial extent and lateral symmetry. The authors employ a probabilistic model that predicts the next pose distribution instead of predicting a fixed pose; gesture motion can then be re-sampled repeatedly to obtain a variety of sequences. Similarly, in [YPJ\*21] a gesture generation toolkit is presented with the control parameters speed, spacial extent, and handedness. The system in [SGD21] uses the Laban Effort and Shape qualities as animation modifiers to impart the intended personality to the character. All these methods rely on handcrafted control features which are often not descriptive enough to encode to wide range of distinct styles that can be learned by the model. More recently, several text-to-motion models were proposed [KKC22, PBV22, TRG\*22, TGH\*22, ZCP\*22]. In theory, these models can produce motions in a style specified by a text query. However, they are not designed to synthesize gestures that synchronize with speech.

<sup>1</sup><https://github.com/ubisoft/ubisoft-laforge-ZeroEGGS>

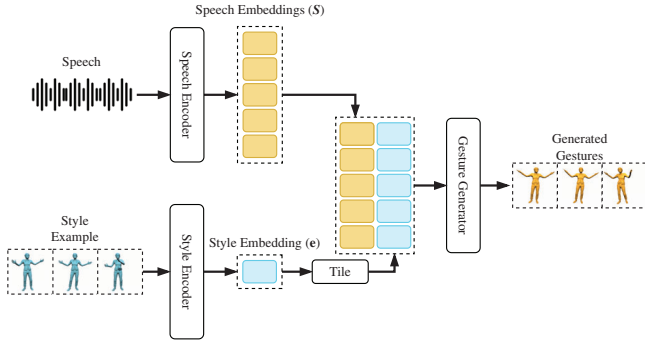


Figure 2: An overview of proposed model.

It was proposed to use style examples to address the shortcomings of hand-crafted features for other animation applications. For instance, Aberman et al. [AWL\*20] transferred style in human locomotion using example clips. While the method can generalize to unseen styles, providing a large enough dataset, it requires a set of labelled styles to train the discriminator in the model. Similarly, Valle-PÃ©rez et al. [VPHB\*21] introduced a model for dance motion generation conditioned on music and a short style example motion. However, the model was not shown to generalize beyond motion styles contained in the training data.

The inherent difficulty of accurately naming styles and collecting meaningful style datasets is not unique to the animation domain. For instance, recent speech synthesis efforts focused on speech stylization by example and enable style generalization beyond the training data. Some works [HZW\*19, WSZ\*18, ZSvNC21] augment a traditional text-to-speech model with a style encoder capturing the general style and prosody of a line to condition the generation. These models can generalize to new styles, but also allow for interpolation in the style latent space. Other work has even made style generalization for unseen speakers possible [ZSvNC21]. We take inspiration from these methods and adapt their ideas to gesture generation to inherit their advantages.

### 3. System Overview

Figure 2 shows an overview of the ZeroEGGS architecture which can be divided into three components: (1) Speech Encoder, (2) Style Encoder and (3) Gesture Generator. The Speech Encoder translates a sequence of raw speech data into a speech embedding sequence  $S$ . The Style Encoder summarizes the style of a reference animation sequence into a fixed-size style embedding vector  $e$ . Finally, an autoregressive decoder called Gesture Generator uses the style embedding vector concatenated to the speech embedding sequence to generate the corresponding gesture animation  $Y$ . This model is trained by maximizing the likelihood of the gesture animation  $p(Y | S, e) = \prod_t p(\mathbf{y}_t | \mathbf{y}_{<t}, S, e)$ .

#### 3.1. Speech encoder

The Speech Encoder, illustrated in Figure 3a, converts raw audio input into a sequence of speech embedding vectors. First, we convert raw audio samples to spectrograms. We use the log-amplitude of the

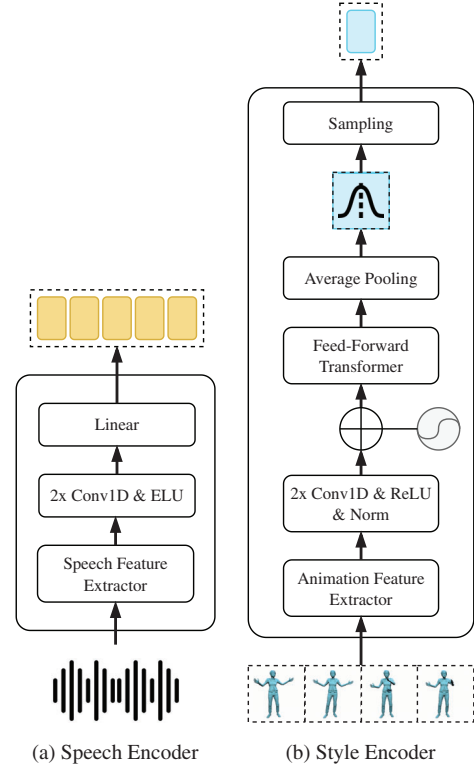


Figure 3: Architectures of the Speech Encoder and the Style Encoder.

spectrogram and the mel-frequency scale as done in many speech applications [ZSvNC21, HCC\*14]. Spectrograms capture how frequency components of a signal vary across time, and the mel scale for frequency better approximates how humans perceive sounds. We also extract the energy per frame as a supplementary feature. Then, extracted features are re-sampled and passed through 1D convolution layers followed by non-linear operators, and finally a frame-wise linear layer. This results in a sequence of embedding vectors  $S = [s_0, s_1, \dots, s_{T-1}] \in \mathbb{R}^{T \times D_S}$  where  $T$  is the number of frames in the sequence and  $D_S$  is the size of the speech embedding vector for each frame.

#### 3.2. Style encoder

The Style Encoder summarizes the reference style animation clip into a low dimensional, fixed size, embedding vector which captures general attributes of the reference style. The reference style clip can have an arbitrary length but should contain enough frames to be representative of the desired style. The style embedding is sampled from a multivariate Gaussian distribution as described in the Variational Auto-Encoder (VAE) framework [KW13]. VAE models are known to learn latent spaces that can be disentangled [RZM19] and interpolated [BVV\*16]. Moreover, in our application they naturally allow for the sampling of variations at inference time.

Each frame of the reference animation clip is represented by a feature vector  $\mathbf{a} = [\rho_p, \rho_r, \hat{\rho}_p, \hat{\rho}_r, \hat{\mathbf{r}}_p, \hat{\mathbf{r}}_r]$  where  $\rho_p \in \mathbb{R}^{3j}$ ,  $\rho_r \in \mathbb{R}^{6j}$  are the joint local translations and rotations,  $\hat{\rho}_p \in \mathbb{R}^{3j}$  and  $\hat{\rho}_r \in \mathbb{R}^{3j}$ ,

are the joint local translational and rotational velocities, and  $\dot{\mathbf{r}}_p \in \mathbb{R}^3$  and  $\dot{\mathbf{r}}_r \in \mathbb{R}^3$  are the character root translational and rotational velocity local to the character root transform.  $j$  corresponds to the number of joints in the kinematic tree. Joint rotations are represented by 2-axis rotation matrix while joint and root rotational velocities are specified using the scaled angle axis representation as done in [ZSKS18]. Unlike representations with less than five dimensions, such as quaternions, this representation is continuous and avoids interpolation problems, making it more suitable for neural network training [ZBL\*19]. We compute the root position of the character by projecting the position of the second spine joint on the ground. We project the  $z$ -axis of the hip joint onto the ground to obtain the root orientation.

The sequence of  $M$  feature frames,  $A = [\mathbf{a}_0, \mathbf{a}_1, \dots, \mathbf{a}_{M-1}]$ , is normalized and fed to a neural network to obtain the style embedding vector  $\mathbf{e} \in \mathbb{R}^{D_e}$ , where  $D_e$  is the dimensionality of this final conditioning style embedding vector. The architecture, shown in Figure 3b, is inspired from [ZSvNC21] which was originally proposed for speech prosody encoding. First, the sequence of animation features,  $A$ , are passed through two 1D convolution layers each followed by a ReLU and a layer normalization layer. Positional encoding encourages the model to encode the sequence ordering [VSP\*17]. Then, similarly to [ZSvNC21], we apply a Feed-Forward Transformer block that implements a multi-head self-attention layer [VSP\*17] and two 1D convolution layers each followed by a residual connection and layer normalization. This results in a sequence of shape  $M \times 2D_e$ . We average over the whole sequence to obtain the parameters  $\mu$  and  $\sigma$  of the  $D_e$ -dimensional multivariate Gaussian distribution from which we sample the final style embedding vector  $\mathbf{e}$ .

### 3.3. Gesture generator

The Gesture Generator, shown in Figure 4, is a conditional auto-regressive model which produces the final animated gesture sequence  $Y = \{\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{T-1}\}$  from the speech embedding sequence  $S$  and the reference style embedding vector  $\mathbf{e}$ . The full parametrization of the output pose state for each frame is given by  $\mathbf{y} = [\rho_p, \rho_r, \dot{\rho}_p, \dot{\rho}_r, \mathbf{r}_p, \mathbf{r}_r, \dot{\mathbf{r}}_p, \dot{\mathbf{r}}_r]$ . Similar to style feature representation (see Section 3.2),  $\rho_p, \rho_r, \dot{\rho}_p, \dot{\rho}_r$  are joint local translations and

rotations along with their velocities, and  $\dot{\mathbf{r}}_p$  and  $\dot{\mathbf{r}}_r$  are the character root translational and rotational velocity local to the character root transform.  $\mathbf{r}_p \in \mathbb{R}^3$  and  $\mathbf{r}_r \in \mathbb{R}^4$  are the position and orientation of the root (represented as quaternions), respectively, which are updated using the root translational and rotational velocities at each frame.

The core of the gesture generator is the Recurrent Decoder. It is an auto-regressive neural network built by two layers of Gated Recurrent Units (GRU). It produces the pose encoding for a new frame  $i$  from the corresponding speech frame  $s_i$ , the reference style embedding vector  $\mathbf{e}$ , and the previous pose state vector  $\mathbf{y}_{i-1}$ .

The Update Character State block in Figure 4 formats the Recurrent Decoder output, computes the pose state, and updates the character facing direction. To compute the pose state at each frame, we denormalize the output of the Recurrent Decoder and use the predicted root translational and rotational velocities to update the root transform. In addition to the last pose encoding, we condition the Recurrent Decoder to a fixed target facing direction to avoid rotational drifting over time. We start by converting the target facing direction from world space to the character root transform. Then, we concatenate it with the previously generated pose state, normalize the resulting vector and provide it to the Recurrent Decoder.

The Hidden State Initializer is a separate neural network that provides hidden states for the GRU layers based on the initial pose, the character facing direction and the style embedding. As in [HP18], we found that using a separate initializing network improved the quality of our results. We implemented the Hidden State Initializer using three linear layers followed by ELU activation functions.

## 4. Implementation and Training

### 4.1. Dataset and data preparation

We recorded a high quality dataset of synchronized motion capture and audio. It contains 67 sequences of monologue performed by a female actor speaking in English and covers 19 different motion styles. The styles were chosen to cover a variety of motions (e.g., *Tired* vs. *Oration* in Figure 5) as well as hand and head movement (e.g., *Oration*, *Agreement*). The total length of the dataset is 135 minutes. The style label of each sequence is the instruction given to the actor, thus might not always correspond to how an external annotator would perceive the motion. Table 1 summarizes the information about the captured styles.

We recorded full-body motion at 60 frames per second (fps) and represented animation data via a skeleton of  $j = 75$  joints, including hands and fingers. For training, we added the mirrored version of all animation data to double the amount of data and extracted feature vector  $\mathbf{a}$  and pose state vector  $\mathbf{y}$  for each frame as explained in Sections 3.2 and 3.3, respectively. In addition, we extracted the head  $z$ -axis direction for all frames in the sequence and projected it onto the ground. Then we used the median of the extracted head direction across all frames as the global target facing direction. During runtime, we set the target facing direction to the global  $z$ -axis direction.

Audio data was recorded at a sampling rate of 48 kHz. For the speech encoding, we extract spectrograms using an FFT Hanning-window of 50 ms and a hop length of 12.5 ms. We project the

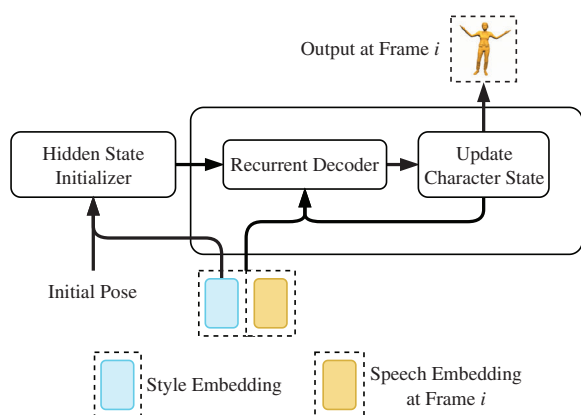


Figure 4: The architecture of Gesture Generator block.



**Table 1:** Details of the recorded motion and audio dataset in minutes.

Style	Length (min)	Style	Length (min)
Agreement	5.25	Pensive	6.21
Angry	7.95	Relaxed	10.81
Disagreement	5.33	Sad	11.80
Distracted	5.29	Sarcastic	6.52
Flirty	3.27	Scared	5.58
Happy	10.08	Sneaky	6.27
Laughing	3.85	Still	5.33
Oration	3.98	Threatening	5.84
Neutral	11.13	Tired	7.13
Old	11.37	<b>Total</b>	<b>134.65</b>

spectrogram into the mel frequency scale, and use the log amplitude of each of the 80 channels as well as the total frame energy as final speech features. We then re-sampled the sequence of speech features to 60 fps and normalized it.

#### 4.2. Model implementation

We adjust the kernels of 1D convolution layers for the speech encoder network to obtain a final receptive field covering roughly 1 second of speech. The first convolutional layer has 64 channels and a kernel size of 3, the second has 64 channels and a kernel size of 31. Both convolution layers are followed by a dropout layer with a rate of 0.2 and ELU activations. Our final encoder dimensionality ( $D_S$ ) is 64.

For the style encoder, we set the dimensionality of the style embedding vector  $\mathbf{e}$  to 64. All convolutional layers have a kernel size of 3 with 512 output channels, followed by dropout layers with a drop rate of 0.2. We constructed our Feed-Forward Transformer block with a 4-head self-attention layer, followed by a dropout at a rate of 0.1 and two 1D convolution layers with a kernel size of 3 and 64 channels.

The Recurrent Decoder in our gesture generator is constructed by two GRU cells with a hidden state size of 1024. The cell state encoder is a 3-layer fully connected feed-forward network with ELU activation and a hidden size of 1024.

#### 4.3. Training and losses

We train the network end-to-end using the Rectified Adam optimizer [LJH\*19] with a learning rate of  $1e-4$ , a decay factor of 0.995 is applied at every 1k iterations. We used a batch size of 32 and stopped training after 120k iterations based on visual result quality. In a batch, the length of the sequences  $T$  is set to 256 frames (4.26 s). We do not use teacher forcing during training, but instead train the model on its own predictions. Although this decelerates convergence, it ensures that the model learns to recuperate from its own errors which leads to more robustness. The style example sequence  $A$  is sampled from the same animation clip as the target sequence. Its length  $M$  randomly spans from 256 to 512 frames, but always encompasses the target sequence. This random sam-

pling scheme ensures de-correlation between styles and clip length. We also randomly alter the animation and speech speed by 10% in whole batches as a data augmentation strategy.

We can consider our model as a conditional VAE where the objective is to maximize the evidence lower bound (ELBO) of the marginal log likelihood of gesture motion given a speech sequence. Thus we can formulate the training loss as the negative ELBO expressed as

$$\begin{aligned} \mathcal{L} &= \mathbb{E}_{q(\mathbf{z}|\mathbf{e})}[-\log p(Y | S, \mathbf{z})] + D_{KL}(q(\mathbf{z} | \mathbf{e})\|p(\mathbf{z})) \\ &= \mathcal{L}_{recon} + D_{KL}(q(\mathbf{z} | \mathbf{e})\|p(\mathbf{z})) \end{aligned} \quad (1)$$

The first term is hereby the expected negative log-likelihood of the gesture motion, or reconstruction loss. The second term is the regularization term expressed as the Kullback–Leibler divergence between the posterior distribution  $q(\mathbf{z} | \mathbf{e})$  predicted by the style encoder and the prior distribution  $p(\mathbf{z})$ , which is a standard multivariate Gaussian distribution. In our implementation, we used the cost annealing strategy proposed by [BVV\*16] for weighting the regularization term.

Our reconstruction loss divides into several terms:

$$\begin{aligned} \mathcal{L}_{recon} &= \lambda_p \mathcal{L}_p + \lambda_r \mathcal{L}_r + \lambda_{vp} \mathcal{L}_{vp} + \lambda_{vr} \mathcal{L}_{vr} + \\ &\quad \lambda_{dp} \mathcal{L}_{dp} + \lambda_{dr} \mathcal{L}_{dr} + \lambda_f \mathcal{L}_f \end{aligned} \quad (2)$$

where  $\mathcal{L}_p$ ,  $\mathcal{L}_r$ ,  $\mathcal{L}_{vp}$ , and  $\mathcal{L}_{vr}$ , are the mean absolute error (MAE) between predicted and target joint positions, rotations, translational velocities, and rotational velocities, respectively, in both local and world spaces. In addition to direct velocity predictions,  $\mathcal{L}_{dp}$  and  $\mathcal{L}_{dr}$  penalize the velocity MAE by computing the translational and rotational velocities in the local and world spaces on-the-fly via finite-difference. This is inspired by the reconstruction loss proposed in [HKPP20] for Learned Motion Matching. Finally,  $\mathcal{L}_f$  penalizes the MAE for the facing direction in the world space to prevent any character rotational drift. Loss terms are empirically weighted to be on a similar scale.

## 5. Experiments and Results

We perform a number of experiments to assess ZeroEGGS’s performance regarding generalization, style control flexibility, and subjective quality. Still images of motion generated for various input styles are shown in Figure 1; please also refer to the supplemental video for additional visual results.

### 5.1. Unseen styles and speakers

With our Style Encoder, we can extract style control features from arbitrary, previously unseen gesture motion samples. To test this, we removed all *Oration* style samples from the training set, a style with noticeably larger average hand height than all other styles. We visualize the model’s generalizability to this unseen style in Figure 5a. We also retain one recording sample for each style and show that at inference time, the model can produce appropriate style for these samples (see Figure 5b).



(a) Oration: This style was not part of the training data in any form.



(b) Old: This style was part of the training data, but not the specific style example clip used for conditioning.

**Figure 5:** Samples of generated gestures given unseen style examples. Left column: Given example. Right column: Generated gesture.



**Figure 6:** Samples of generated gestures given unseen style examples and new speaker. Left column: given example. Right column: generated gesture.

Next, we assessed our model's generalization capabilities to new speakers. We generated results for a male speaker with a deep voice plus the unseen *Oration* style. The output is visualized in Figure 6.

Because ZeroEGGS relies solely on the speech spectrum amplitude, it can be used with languages unseen in the training set. We test and confirm this language agnosticism for French, German, and gibberish, with results illustrated in the supplementary material.

## 5.2. Manipulating style in the style embedding space

### 5.2.1. Blending styles

The variational framework used in our Style Encoder provides a morphable and continuous style embedding space. This allows to mix the styles of multiple samples via linear interpolation. We illustrate interpolations of the *Old* and *Oration* styles in Figure 7. As can be seen, the character's posture and the hand position gradually changes as we interpolate.



**Figure 7:** Screenshot of generated gestures by interpolating between *Old* and *Oration*.

### 5.2.2. Control via PCA components

We can control some of the low level style characteristics by projecting the style embedding vector onto a PCA space and manipulating the components. Figure 8a shows a scatter plot of the first two principal components for non-overlapping samples in different styles. We observe that the first principal component roughly corresponds to body sway. The more static styles such as *Still* and *Sad* are located on the left side of the plot, while more dynamic styles, such as *Happy* and *Angry*, are located on the right side of the plot.

Similarly, the second principal component is associated with hand motion height and radius. For example *Oration* samples, for which hands are usually above the shoulders, are located on the upper parts of the plot, whereas styles such as *Tired*, during which the actor put her hands on her knees, are on the lower parts of the plot.

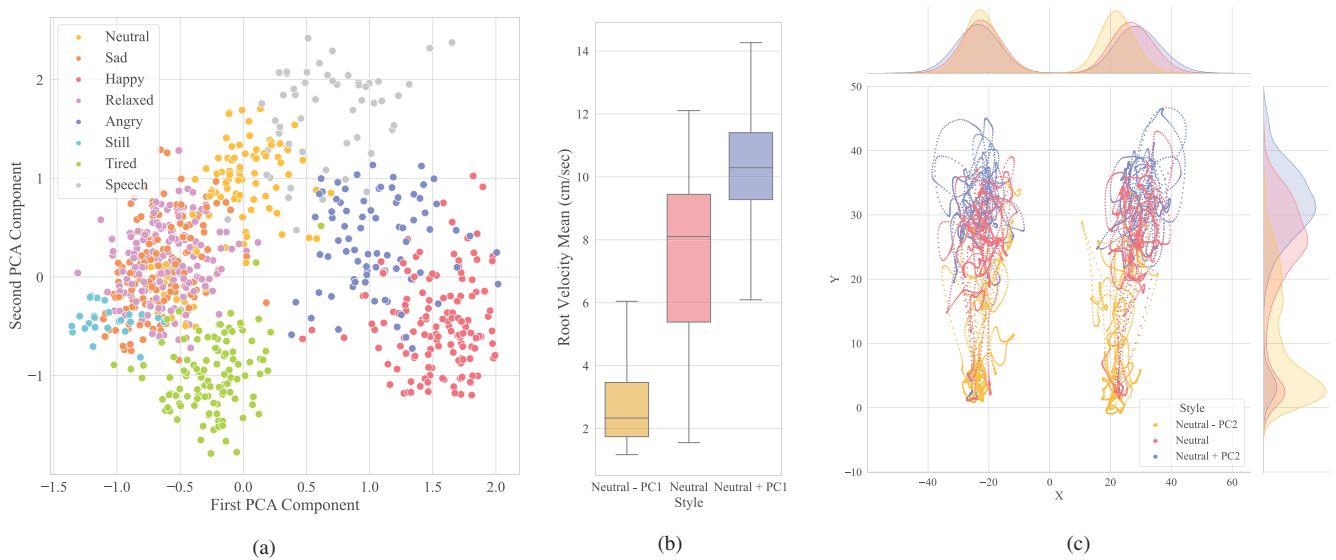
We can modify these gestures characteristics in the PCA space and project them back into the original style embedding space. Figure 8b shows the distribution of the root velocity as an indicator of body sway for three variation a generated gesture: A *Neutral* style example and its two versions obtained by changing its first principal component. We can see that modifying the first principal component affects root velocity. Figure 8c shows the hand movement trajectories relative to the hip joint in a frontal view for the same *Neutral* sample, after manipulation of the second principal component. The distribution plot on the Y axis of the figure highlights the effect of this component on hand height.

## 5.3. User study

We evaluated subjective result quality in three separate user studies, assessing the motion output with regard to (1) naturalness, (2) appropriateness w.r.t. the speech, and (3) recognizability of style. We compared performance of our model to the ground truth reference motion as well as three baselines, as described below.

### 5.3.1. Study design

We used a MUSHRA-like (Multiple Stimuli with Hidden Reference and Anchor) [ITU15] interface based on [JYW\*21]. We had a total of 131 participants, with a minimum of 20 per study (ages 20-55 years  $\mu = 33.6$ ,  $\sigma = 8.1$ ). In each study, participants performed 12 ratings comparing the motion of rendered animation clips with the female character shown in Figure 5, for the four systems listed



**Figure 8:** We control low level gesture characteristics by projecting style embedding vector onto the PCA space. (a) Shows the first two principal components of the style embedding vector in the PCA space for a subset of styles. The first dimension corresponds to the body sway, while the second dimension corresponds to the hand height and radius. (b) We can control the root velocity of the character by changing the first principal component of the style embedding vector in the PCA space. We modify the first principal component by one standard deviation of the Neutral style. (c) Shows the hand trajectories for the neutral style after modifying the second principal component of the style embedding vector in the PCA space.

**Table 2:** Comparison of our model (ZEGGS) with the baseline model (MG).

System	#Parameters	Inference time (per frame)	Rating on our dataset			Rating on TSG		
			Nat.	Appr.	Style	Nat.	Appr.	Style
MG	86M (88M with style)	29 ms	11.8 ± 12.2	7.8 ± 9.6	12.9 ± 16.1	32.7 ± 21.0	32.2 ± 20.6	19.8 ± 18.0
ZEGGS	25M	4 ms	56.3 ± 22.2	48.9 ± 20.3	58.1 ± 25.9	36.7 ± 18.4	38.3 ± 23.3	61.9 ± 23.3

further below. The order of systems and rating pages was randomized. For *naturalness*, participants were asked how natural the character’s motion looks to them. For *appropriateness*, participants rated how appropriate the gesture motion was for the speech. For *style portrayal*, videos played without sound and participants judged how well the motion represented a given target style. For this, a fixed set of 6 representative styles (*Happy, Sad, Relaxed, Old, Angry, Orat-ion*), and participants rated two samples of each.

Participants compared and rated the following 4 systems:

1. Ground truth motion (GT)
2. Mismatched motion (MM)
3. Our model (ZEGGS)
4. MoGlow (MG)

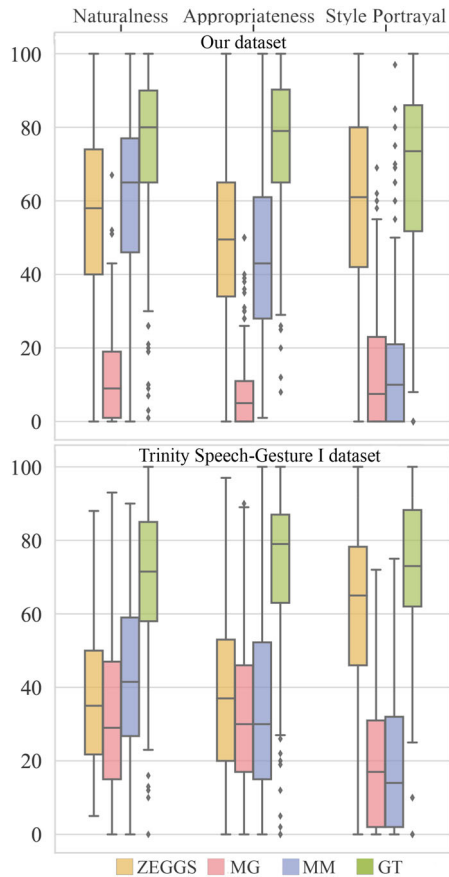
GT is the original motion-captured animation associated with the given speech segment. MM is also motion-captured animation but used in the wrong context. For the naturalness and appropriateness studies, MM is a motion originally associated with a *different* speech segment of the same style, while in the style portrayal study, MM is a motion associated with a different style.

As a strong baseline model, we chose MoGlow (MG) by Alexander et al. [AHKB20] which represents style via their three

proposed gesture features: hand height, velocity and radius. We chose this model based on its style-control capabilities, competitive performance [KJY\*21], and code availability. We compare model statistics of ZEGGS and MG in Table 2. We first retrained the model using our dataset, however, this led to severely jittery, low quality output. We therefore decided to train MoGlow using its original training dataset, the Trinity Speech-Gesture dataset I (TSG) [FM18]. TSG contains 4 h of a single male speaker producing spontaneous speech in monologue format in an overall neutral style. To level the playing field, we therefore ran two versions of each of the three studies of naturalness, appropriateness, and style: The first version used our dataset, the second version used the TSG dataset. For the TSG version of the style study, the style input motion features were taken from our own dataset since TSG does not contain any style data. A number of video samples from our user study are included in the supplemental material.

### 5.3.2. Results

We analyzed the collected perceptual data using Analysis of Variance (ANOVA) when the normality assumption was not violated (Shapiro–Wilk test) and corrected degrees of freedom using the



**Figure 9:** Results from our three user studies. The x-axis represents the system and the y-axis the average participant ratings.

Greenhouse–Geisser when sphericity was violated. When normality was violated, we used Aligned Rank Transform (ART) instead of ANOVA. We performed post-hoc comparisons using Estimated Marginal Means. We had one within-subject factor for naturalness and appropriateness measures, the gesture generation system, with four levels: GT, MM, ZEGGS, MG. For style portrayal, we additionally had the factor style. Results are visualized in Figure 9 and summarized in Table 2.

The results of our user studies show that our model outperforms the comparison model [AHKB20] for all three collected measures, motion naturalness, appropriateness for speech, and recognizability of portrayed style for our own dataset, and for appropriateness and style recognizability for the TSG data. There was a main effect of system for each of the three measures, and GT performed best across each measure, as expected ( $p < 0.001$  for all pairwise comparisons with GT).

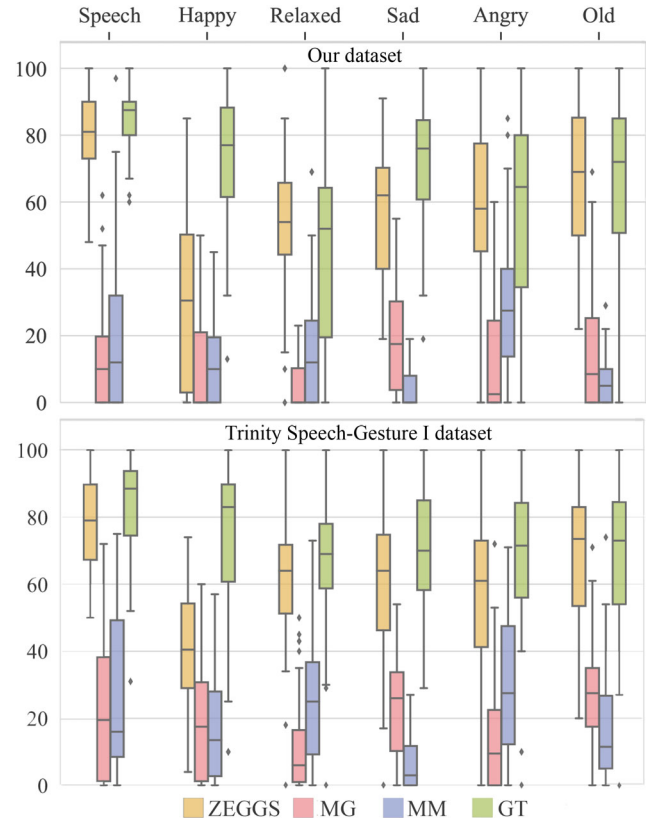
For **naturalness**, we would expect a similar level of naturalness for MM as for GT because they are both motion-capture clip. However, the mismatch with the speech audio resulted in lower ratings which highlights the importance of synchronicity for the perception of naturalness. MM was rated significantly lower than GT and higher than MG (for TSG,  $p < 0.05$  for ZEGGS vs. MM, all other

differences  $p < 0.001$ ). For our dataset, MM was not rated significantly different to ZEGGS.

For **appropriateness**, for our own dataset, ZEGGS scored higher than MM ( $p < .01$ ), indicating that our model was indeed able to capture some relation between the speech and motion modalities and produced output that was more appropriate for the given speech than the baseline systems. This reinforces the idea that synchronicity plays an important role in the perception of gesture quality. Here, even if MM is a motion-captured clip in the same style as the target, it seems inappropriate for the speech segment, because gesture are out of sync with voice inflections. However, for the TSG dataset, MM was rated on-par with ZEGGS and MG.

When measuring **style portrayal** specifically, there is also a significant interaction between the factors system and style. This time, MM is a motion-captured animation from another style. Therefore it performs significantly worse than GT as it accurately portrays a different style ( $p < 0.001$ ). MM was also rated significantly worse than ZEGGS ( $p < 0.001$ ) but not MG, indicating that MG was unable to portray the desired styles whereas our method produced good results. Visually inspecting the interaction of system and style, this appears mainly stem from a lower than average performance of ZEGGS for the style *Happy* (see Figure 10).

We believe there are several reasons explaining why ZeroEGGS outperformed MoGlow in most subjective evaluations. Firstly, as



**Figure 10:** Rating scores per style for the style portrayal study.



described in Section 5.3.1, training MoGlow on our own dataset yielded poor results. MoGlow captures style information in three statistical measures extracted from hand movement. While reasonable for styles like public speaking or neutral conversational bearing as in the TSG, there is a lot more information to capture when the styles are more diverse. For instance, MoGlow cannot accurately capture style differences of gestures with the general same movement intensity, but different trajectories. Having widely different targets for the same input impedes the model to converge during training. In addition to the more consistent style within the TSG dataset, this dataset is also four times the size of our dataset. These factors are why we resorted to training MoGlow on the single-style TSG dataset with which it was originally developed. This training explains the relatively lower performance of MoGlow on our dataset when compared to the performance on the TSG dataset as seen in Figure 9. While the difference in performance of ZeroEGGS relative to the ground truth is larger on TSG compared to our dataset, it still outperforms MoGlow that was trained on TSG.

Regarding naturalness, MoGlow produced more jittery animations and generated at a lower frame rate (20 fps instead of 60 fps) which may have also played a role in the evaluation.

By capturing style information in three statistical measures extracted only from hand movement, MoGlow discards a lot of useful information when compared to ZeroEGGS which is free to encode any relevant information in a larger encoding space. By focusing solely on hand movements, MoGlow ignores information about posture, stance and leg movement. This negatively impacts style portrayal accuracy compared to ZeroEGGS in styles such as *Old* where the posture of the character is important, or *Happy* where the actor shifts her weight frequently. The nature of MUSHRA-like tests magnifies these differences between models because scores are not an absolute quality rating but rather an indication of the relative performance of each models.

There appears to be a larger performance gap for naturalness and appropriateness between the ground truth and the other methods when comparing results for TSG versus our dataset. We hypothesize that the nature of the style and content in TSG may play a role. When compared to our dataset, the speaker in TSG performs semantic gestures much more frequently, which when present and well timed systematically lead to a higher naturalness and appropriateness. ZeroEGGS and MoGlow do not model such gestures which would explain lower comparative ratings. The fact that mismatched motion (MM) also underperforms on TSG is an indication that semantic gestures are performed at the wrong moment, leading to lower performance ratings.

## 6. Discussion

Our experiments show that our model generalizes to new voices, languages and to new styles that were not part of the training data. This zero-shot style transfer capability constitutes a major advantage when compared to existing methods. In practice, it means that a single model can be used for different characters and that it does not have to be retrained each time a new style is needed. Moreover, specifying style directly using examples enables the model to capture smaller details that could be otherwise discarded by the strict

information bottleneck imposed by handcrafted features and statistics. This results in more realistic and compelling motions. Finally, by using examples instead of labels, users do not have to elaborate a precise consentaneous style taxonomy.

We also show that our variational framework learns a meaningful style embedding that enables manipulation and interpolation within the latent space. This allows, for example, the weighted mixing of style reference samples. This proves useful in practice when transitioning from one style to another in the same animated sequence. In addition, the probabilistic nature of variational methods allows for variations of gesture motion for a given input by re-sampling.

In our subjective evaluations, ZeroEGGS outperformed the baseline model w.r.t. naturalness, appropriateness for speech, and style portrayal. This means that for ZeroEGGS, generalization does not come at the cost of visual quality and style diversity. Improvement on state-of-the-art mainly comes from the superior descriptive power of our example-based conditioning over handcrafted features and statistics. As explained above, capturing more detail in the style embedding improves naturalness, but it also helps generating subtleties that helps differentiate between styles. Moreover, our loss ensures smooth and stable motion, which also plays a role in the perception of the naturalness of the animation.

While our model outperforms our baseline models, it does fall significantly behind ground truth motion, motivating further work to improve the naturalness of the generated output. Moreover, some more subtle motions that are specific to certain styles are often over-smoothed or even not produced. For instance, in the *Agreement* style, head nods are very infrequently generated, while abundantly present in the training data. Enforcing explicit disentanglement in the style latent space could help capture motion localized in specific parts of the body. Our model showed a slight decrease in performance on the Trinity Speech-Gesture dataset, on which it was not trained, likely due to the frequent displacement of the speaker, which is not captured in our dataset. Our method performs well on more static style, but we would like to address character displacement as a next step.

Finally, ZeroEGGS generates small amounts of foot sliding, a common problem in non physically-based animation generation that can be addressed with IK-based post-processing.

ZeroEGGS uses raw spectral audio features as input which provides robustness to language and speaker voice. However, this means that the generation is limited to so-called “beat gestures” that do not carry semantic information about the speech content. This explains in part why the ground truth motion scored significantly higher w.r.t. appropriateness for the speech in the user study. Future work could include semantic markers in the speech as input to the model.

## 7. Conclusion

We proposed ZeroEGGS, a model that generates stylized gesture animations from speech. The desired style is specified to the model as a short example clip. The model generalizes beyond training data allowing the generation of unseen styles, for different voices and languages. Moreover, the latent style representation allows for

control over generation. Our experiments show that ZeroEGGS generates state-of-the-art gesture animations.

There are many directions for future work, including enforcing user-specified disentanglement in the learned style latent space, providing support for semantic gestures as well as exploring new architectures based on affine transform coupling layers.

### Acknowledgments

This work was started as part of S. G's internship at Ubisoft Inc. and was supported by the Mitacs Accelerate Program.

### References

- [AHKB20] ALEXANDERSON S., HENTER G. E., KUCHERENKO T., BESKOW J.: Style-controllable speech-driven gesture synthesis using normalising flows. In *Computer Graphics Forum*. (2020), vol. 39, Wiley Online Library, pp. 487–496.
- [ALNM20] AHUJA C., LEE D. W., NAKANO Y. I., MORENCY L.-P.: Style transfer for co-speech gesture animation: A multi-speaker conditional-mixture approach. In *European Conference on Computer Vision*. (2020), Springer, pp. 248–265.
- [AWL\*20] ABERMAN K., WENG Y., LISCHINSKI D., COHEN-OR D., CHEN B.: Unpaired motion style transfer from video to animation. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 64–1.
- [BVV\*16] BOWMAN S. R., VILNIS L., VINYALS O., DAI A. M., JÓZEFOWICZ R., BENGIO S.: Generating sentences from a continuous space. In *SIGLL Conference on Computational Natural Language Learning (CONLL)* (2016). <http://arxiv.org/abs/1511.06349>
- [CVB01] CASSELL J., VILHJÁLMSOHN H. H., BICKMORE T.: beat: the behavior expression animation toolkit. *ACM Transactions on Graphics (TOG)* (2001), 477–486.
- [DRBD12] DE RUITER J. P., BANGERTER A., DINGS P.: The interplay between gesture and speech in the production of referring expressions: Investigating the tradeoff hypothesis. *Topics in Cognitive Science* 4, 2 (2012), 232–248.
- [FM18] FERSTL Y., MCDONNELL R.: Investigating the use of recurrent motion modelling for speech gesture generation. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents* (2018), pp. 93–98.
- [FNM19] FERSTL Y., NEFF M., MCDONNELL R.: Multi-objective adversarial gesture generation. In *Motion, Interaction and Games*. 2019, pp. 1–10.
- [GBK\*19] GINOSAR S., BAR A., KOHAVI G., CHAN C., OWENS A., MALIK J.: Learning individual styles of conversational gesture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 3497–3506.
- [HCC\*14] HANNUN A., CASE C., CASPER J., CATANZARO B., DIAMOS G., ELSÉN E., PRENGER R., SATHEESH S., SENGUPTA S., COATES A., ET AL.: Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567* (2014).
- [HKPP20] HOLDEN D., KANOUN O., PEREPICHKA M., POPA T.: Learned motion matching. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 53–1.
- [HP18] HARVEY F. G., PAL C.: Recurrent transition networks for character locomotion. In *SIGGRAPH Asia 2018 Technical Briefs* (New York, NY, USA, 2018), SA '18, Association for Computing Machinery. <https://doi.org/10.1145/3283254.3283277>
- [HZW\*19] HSU W.-N., ZHANG Y., WEISS R. J., ZEN H., WU Y., WANG Y., CAO Y., JIA Y., CHEN Z., SHEN J., NGUYEN P., PANG R.: Hierarchical Generative Modeling for Controllable Speech Synthesis. In *ICLR* (2019). <http://arxiv.org/abs/1810.07217>
- [ITU15] ITU.: *Method for the subjective assessment of intermediate quality level of audio systems (mushra)*. Tech. rep, 2015.
- [JYW\*21] JONELL P., YOON Y., WOLFERT P., KUCHERENKO T., HENTER G. E.: Hemvip: Human evaluation of multiple videos in parallel. In *Proceedings of the 2021 International Conference on Multimodal Interaction* (New York, NY, USA, 2021), ICMI '21, Association for Computing Machinery, p. 707–711. <https://doi.org/10.1145/3462244.3479957>
- [KJLW03] KOPP S., JUNG B., LESSMANN N., WACHSMUTH I.: Max-a multimodal assistant in virtual reality construction. *KI* 17, 4 (2003), 11.
- [KJvW\*20] KUCHERENKO T., JONELL P., VAN WAVEREN S., HENTER G. E., ALEXANDERSON S., LEITE I., KJELLSTRÖM H.: Gesticulator: A framework for semantically-aware speech-driven gesture generation. In *Proceedings of the 2020 International Conference on Multimodal Interaction* (2020), pp. 242–250.
- [KJY\*21] KUCHERENKO T., JONELL P., YOON Y., WOLFERT P., HENTER G. E.: A large, crowdsourced evaluation of gesture generation systems on common data: The genea challenge 2020. In *26th International Conference on Intelligent User Interfaces* (2021), pp. 11–21.
- [KKC22] KIM J., KIM J., CHOI S.: Flame: Free-form language-based motion synthesis & editing. *arXiv preprint arXiv:2209.00349* (2022).
- [KW13] KINGMA D. P., WELING M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [LJH\*19] LIU L., JIANG H., HE P., CHEN W., LIU X., GAO J., HAN J.: On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265* (2019).
- [LM06] LEE J., MARSELLA S.: Nonverbal behavior generator for embodied conversational agents. In *International Workshop on Intelligent Virtual Agents* (2006), pp. 243–255.

- [ML04] MELINGER A., LEVELT W. J.: Gesture and the communicative intention of the speaker. *Gesture* 4, 2 (2004), 119–141.
- [MXL\*13] MARSELLA S., XU Y., LHOMMET M., FENG A., SCHERER S., SHAPIRO A.: Virtual character performance from speech. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2013), pp. 25–35.
- [NKAS08] NEFF M., KIPP M., ALBRECHT I., SEIDEL H.-P.: Gesture modeling and animation based on a probabilistic re-creation of speaker style. *ACM Transactions on Graphics (TOG)* 27, 1 (2008), 1–24.
- [PBV22] PETROVICH M., BLACK M. J., VAROL G.: TEMOS: Generating diverse human motions from textual descriptions. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2022). <http://arxiv.org/abs/2204.14109>
- [RGP21] REBOL M., GÜTI C., PIETROSZEK K.: Passing a non-verbal turing test: Evaluating gesture animations generated from speech. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. (2021), IEEE, pp. 573–581.
- [RZM19] ROLINEK M., ZIETLOW D., MARTIUS G.: Variational autoencoders pursue pca directions (by accident). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019).
- [SGD21] SONLU S., GÜDÜKBAY U., DURUPINAR F.: A conversational agent framework with multi-modal personality expression. *ACM Transactions on Graphics (TOG)* 40, 1 (2021), 1–16.
- [TGH\*22] TEVET G., GORDON B., HERTZ A., BERMANO A. H., COHEN-OR D.: Motionclip: Exposing human motion generation to clip space. *arXiv preprint arXiv:2203.08063* (2022).
- [TRG\*22] TEVET G., RAAB S., GORDON B., SHAFIR Y., COHEN-OR D., BERMANO A. H.: Human motion diffusion model. *arXiv preprint arXiv:2209.14916* (2022).
- [VPHB\*21] VALLE-PÉREZ G., HENTER G. E., BESKOW J., HOLZAPFEL A., OUDEYER P.-Y., ALEXANDERSON S.: Transflower: probabilistic autoregressive dance generation with multimodal attention. *arXiv preprint arXiv:2106.13871* (2021).
- [VSP\*17] VASWANI A., SHAZEER N., PARMAR N., USZKOREIT J., JONES L., GOMEZ A. N., KAISER Ł., POLOSUKHIN I.: Attention is all you need. In *Advances in neural information processing systems*. (2017), pp. 5998–6008.
- [WSZ\*18] WANG Y., STANTON D., ZHANG Y., SKERRY-RYAN R., BATTENBERG E., SHOR J., XIAO Y., REN F., JIA Y., SAUROUS R. A.: Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis. In *International Conference on Machine Learning* (2018), pp. 5180–5189. PMLR.
- [YCL\*20] YOON Y., CHA B., LEE J.-H., JANG M., LEE J., KIM J., LEE G.: Speech gesture generation from the trimodal context of text, audio, and speaker identity. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.
- [YPJ\*21] YOON Y., PARK K., JANG M., KIM J., LEE G.: Sgtoolkit: An interactive gesture authoring toolkit for embodied conversational agents. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (2021), pp. 826–840.
- [YYH20] YANG Y., YANG J., HODGINS J.: Statistics-based motion synthesis for social conversations. *Computer Graphics Forum* (2020), 39, 201–212.
- [ZBL\*19] ZHOU Y., BARNES C., LU J., YANG J., LI H.: On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 5745–5753.
- [ZCP\*22] ZHANG M., CAI Z., PAN L., HONG F., GUO X., YANG L., LIU Z.: Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001* (2022).
- [ZSKS18] ZHANG H., STARKE S., KOMURA T., SAITO J.: Mode-adaptive neural networks for quadruped motion control. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–11.
- [ZSvNC21] ZAÏDI J., SEUTÉ H., VAN NIEKERK B., CARBONNEAU M.-A.: Daft-exprt: Robust prosody transfer across speakers for expressive speech synthesis. *arXiv preprint arXiv:2108.02271* (2021).

### Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Video S1