

Retrieving Semantics from the Deep: an RAG Solution for Gesture Synthesis

M. Hamza Mughal¹ Rishabh Dabral¹ Merel C.J. Scholman³

Vera Demberg^{1,2} Christian Theobalt^{1,2}

¹Max Planck Institute for Informatics, SIC ²Saarland University ³Utrecht University

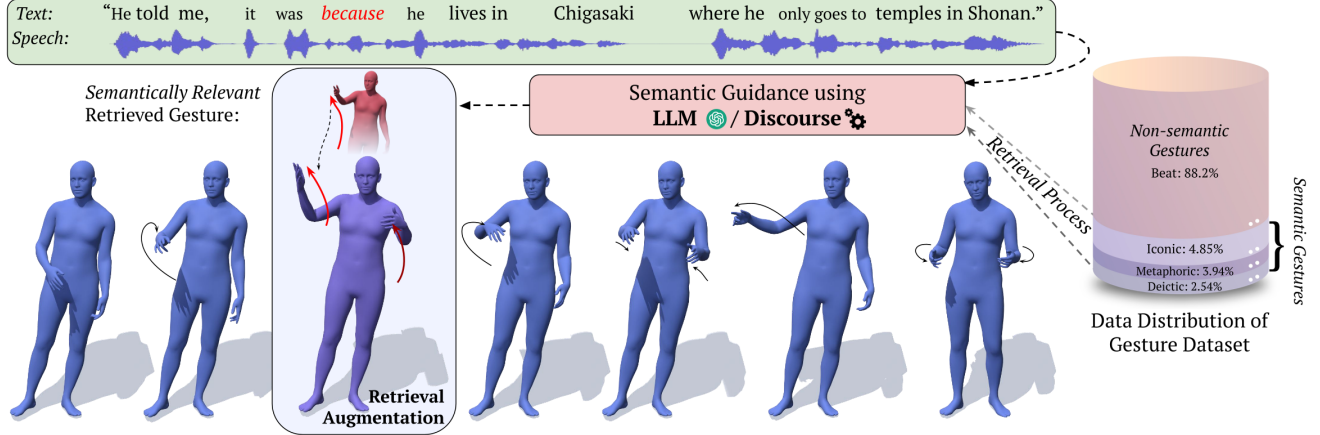


Figure 1. Our RAG-GESTURE approach produces **semantically meaningful gestures** by leveraging explicit knowledge to retrieve exemplar gestures from the sparse semantic data [39] and **guiding the diffusion-based generation process through Retrieval Augmentation**.

Abstract

Non-verbal communication often comprises of semantically rich gestures that help convey the meaning of an utterance. Producing such semantic co-speech gestures has been a major challenge for the existing neural systems that can generate rhythmic beat gestures, but struggle to produce semantically meaningful gestures. Therefore, we present RAG-GESTURE, a diffusion-based gesture generation approach that leverages Retrieval Augmented Generation (RAG) to produce natural-looking and semantically rich gestures. Our neuro-explicit gesture generation approach is designed to produce semantic gestures grounded in interpretable linguistic knowledge. We achieve this by using explicit domain knowledge to retrieve exemplar motions from a database of co-speech gestures. Once retrieved, we then inject these semantic exemplar gestures into our diffusion-based gesture generation pipeline using DDIM inversion and retrieval guidance at the inference time without any need of training. Further, we propose a control paradigm for guidance, that allows the users to modulate the amount of influence each retrieval insertion has over the generated sequence. Our comparative evaluations demonstrate the validity of our approach against recent gesture generation approaches. The reader is urged to explore the results on [our project page](#).

1. Introduction

Human communication is a complex and multifaceted process that involves both verbal and non-verbal elements. Non-verbal communication comprises *co-speech gestures*, which are defined as body and hand movements that are temporally aligned and semantically integrated with the speech [43]. Gestures convey information in tandem with speech and language, and can carry additional meaning that enhances the semantic construct of the message. Generating meaningful full-body gestures is therefore important for the communicative efficacy of virtual humans in telepresence and content-creation domains.

McNeill [43, 44] categorizes gestures into Semantic gestures are context-driven and bear a specific meaning that complements the utterance [6, 43]. Rule-based approaches retrieve these semantic gestures and combine them for generation, while neural methods leverage the data to learn the synthesis process. However, the former tends to result in an unnatural outcome and the latter struggles to generate semantic gestures because these gestures, while being a part of the existing datasets, occur rather sporadically compared to the rhythmic beat gestures [52].

In this work, we address this issue by framing the problem as a Retrieval Augmented Generation (RAG) task. We propose RAG-GESTURE – a diffusion-based gesture syn-

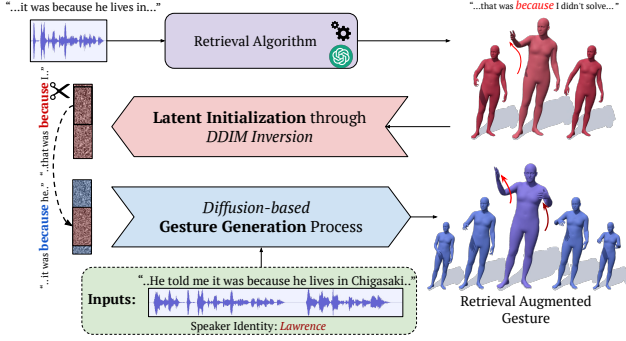


Figure 2. **Overview.** Our approach retrieves example gestures on semantically important words in speech and inserts those examples into the generated gesture by using them to guide the generation.

thesis approach that generates natural and semantically rich gestures by retrieving context-appropriate exemplars from a database and injecting them into the gesture generation process. Motivated by the insights from Neff [49], we decompose the generation problem into two tasks: *specification*, which ascertains “what gesture” to produce for “which word” in the speech, and *animation*, which determines “how to generate” that specified gesture.

To solve the specification problem, we explicitly retrieve the relevant exemplar motions from a gesture database using either of the two retrieval algorithms proposed in Sec. 3.4. More specifically, we show that the chain-of-thought reasoning abilities of Large Language Models [53, 73] can be exploited to extract which parts of the utterance would likely be semantically gestured, and what ‘type’ of semantic gesture (*iconic, deictic, metaphoric, etc*) the phrase would invoke. We also demonstrate that one could, alternatively, use specific linguistic elements to base the retrieval on. In particular, we focus on discourse connectives such as *because, while, and on one hand*, which have been shown to affect the gestural patterns [9, 45].

The animation problem, in our context, requires integrating the retrieved gesture motions (corresponding to a subset of the sentence) into the overall motion generated for the full sentence/utterance (see Fig. 1). To this end, we design an inference-time approach to augment the gesture generation capability of diffusion models with retrieved gestures clips. Our approach allows us to surgically integrate the semantic gestures into the generation process of the base model, while not affecting the synthesis for the remaining parts of the utterance.

Through comprehensive numerical and perceptual evaluations, we demonstrate the effectiveness of our proposed neuro-explicit approach compared with purely neural state-of-the-art approaches like [39, 40, 51, 68] that typically struggle to synthesize semantically meaningful gestures. To summarize, our core contributions are as follows:

- We introduce RAG-GESTURE – a diffusion-based ap-

proach which leverages RAG to generate natural-looking and semantically rich gestures by injecting meaningful motion exemplars into the gesture generation process.

- We introduce multiple retrieval algorithms to extract semantically relevant exemplars from the sparse semantic data distribution by leveraging either the linguistic structure of utterances or LLM-based gesture type prediction.
- Further, we propose a novel inference-time RAG approach for our framework and demonstrate how gesture exemplars can be transferred onto the generated motion locally through the combination of Latent Initialization and Retrieval Guidance technique, which provides a way to control the influence of the retrieval insertion.

2. Related Work

We first provide an overview of gesture synthesis approaches, followed by a discussion of relevant approaches from RAG literature. Lastly, we discuss recurrent patterns of gestures used to express discourse structure, which motivates the design of the retrieval-based method.

2.1. Co-Speech Gesture Synthesis

Co-speech gestures consist of hand and arm movements that occur in sync with speech and help convey meaning and structure discourse. Earlier methods in co-speech gesture generation [10, 61] perform well on semantics and appropriateness of gestures [52]. These systems are intent-driven and utilize heuristics to retrieve the most appropriate gesticulation units for given speech, resulting in high semantic alignment. However, since they insert predefined units into animation, the synthesis quality of these systems tends to be unnatural. To mitigate the lack of naturalness, recent learning-based methods [2, 18, 24, 34] employ deep networks to convert speech input to gesture output and aim to ensure smooth motions. These methods are mostly data-driven and rely on learning speech-to-gesture matching through large scale gesture datasets [17, 20, 39, 65]. Most of these recent methods [1, 34, 64, 66] suffer from reduced communicative efficacy and low semantic alignment, due to bad generalization to semantic co-speech gestures in the dataset. Furthermore, generation and reproduction of semantically appropriate gestures corresponding to speech still remains a challenge for these systems [67].

A few recent approaches directly addressed this problem. Kucherenko *et al.* [35] use speech to predict the gesture properties, which inform a probabilistic model on what type of gesture to generate. In comparison, our framework proposes to use RAG to follow the gesture type information instead of taking gesture properties as conditioning input. Liang *et al.* [38] decouple semantic irrelevant cues like rhythmic beats and volume to generate plausible gestures. Mughal *et al.* [47] generate semantic gestures by providing word-level semantic control and demonstrate how

their diffusion model uses it to sample more semantically meaningful gestures. SemanticGesticulator [71] is a generative retrieval framework which retrieves semantic gestures from a curated motion library and aligns them with rhythmic gestures. Compared to their approach, ours does not rely on a curated library. Rather, it introduces linguistics-based explicit algorithms to retrieve semantics from the existing data [39]. Moreover, we perform retrieval augmentation using a diffusion model at inference time rather than fine-tuning our gesture generator to follow the retrieval.

2.2. Retrieval Paradigms for Generative Methods

Employing database retrieval for improved animation performance is a prevalent paradigm in motion and gesture synthesis. Classical [8, 11] and learning-based motion matching approaches [28, 29] have been applied for better character control in motion synthesis [58]. Diffusion-based text-to-motion synthesis frameworks have also dealt with the problem of semantic control at a global level [68] or at the local keyframe level [23, 30]. Zhang *et al.* [68] introduce RAG to diffusion-based motion synthesis by using retrieval based on global text similarity for text-to-motion tasks. In contrast to the text-to-motion task, gesture generation is affected by language content and prosody at varying levels of semantic granularity. That means one can neither perform global semantic matching between the retrieved and generated motion [68] nor simply copy and paste retrieved motion parts onto the keyframe locations [23] without considering the differences in speech contexts.

Gesture synthesis approaches have also utilized database retrieval both in classical and learning based paradigms. Early approaches [5, 33, 50] utilize statistics to retrieve gestures from a database of gesticulation units. A few learning-based approaches combine rule-based retrieval and deep learning-based synthesis to improve upon the semantic quality of gesture generation. ExpressGesture [19] introduces a database-driven framework to ensure that generated gestures retain the expressive and defined gesture form. Habibie *et al.* [25] employ nearest neighbor search to retrieve most appropriate gestures and provide a style-control mechanism over the generative framework. In comparison, our retrieval approach is not limited to style or select keywords but is grounded in explicit rules for semantic retrieval that are driven from linguistic and gestural structure.

2.3. The effect of Linguistic Dynamics on Gestures

The goal of our retrieval approach is to extract semantic gestures from the database. Our proposed design is therefore motivated by the relation between language dynamics and co-speech gestures. Laparle *et al.* [36] demonstrate how gesture features like hand shape or orientation are important to distinguish between topics in discourse. Gestures can also indicate the relations between different parts of the

message. For example, a contrast relation can be expressed by the lexical marker *on the one hand...on the other hand*, framing the two contrasting arguments as items placed on separate hands [26]. The raised index finger draws the attention of the interlocutor to new and important topics, and is commonly used to express exceptions or concessions (which can be expressed linguistically using *however*) [7, 31, 32].

In sum, gestures can be used to convey the semantic structure of a message, but existing generation works fail to reproduce these context-driven patterns [67], which results in repetitive gestures. We attempt to leverage the linguistic information in our retrieval framework. This framework can, in turn, improve semantic grounding of generated gestures through RAG based insertion.

3. Approach

Our primary goal is to generate semantically meaningful co-speech gesture sequences that are consistent with the content of the corresponding speech. This goal can be partly achieved by our latent-diffusion based gesture generation model, which is trained to synthesize gestures while being conditioned on the corresponding modalities (refer to Sec. 3.1). This data-driven gesture generation approach is capable of generating natural looking and beat-aligned gesture sequences from the input speech.

However, this basic generation framework is not sufficient to exhibit rich semantics in the generated gestures. Therefore, we propose additional steps that ensure the solution to two key problems – *animation* (refer to Sec. 3.2 and Sec. 3.3) and *specification* (refer to Sec. 3.4) of semantically rich gestures [49]. Our approach is illustrated in Fig. 3.

3.1. Gesture Generation with Latent-Diffusion

As the first step, we build a robust generation framework based on latent-diffusion modeling [55], which serves as our base model to generate plausible co-speech gesture sequences. The gesture sequence $\mathbf{x} \in \mathbb{R}^{N \times d_{\mathbf{x}}}$ consists of N frames of human motion with $d_{\mathbf{x}}$ representing a combination of J full body joints represented using the 6D rotation representation of [74], FLAME model parameters for face motion, root translation and foot contact labels [40]. Instead of training our diffusion framework on raw motion \mathbf{x} , we encode motion using a VAE encoder, which provides separate encodings for different body parts [47]. These part-wise VAE encodings are then used to train our conditional latent-diffusion model that generates co-speech gestures.

Decoupled Gesture Encoding. Based on the observations that each body region has different relations with speech [40] and scale differences between them affect the generation quality [47], we decouple the gesture sequence

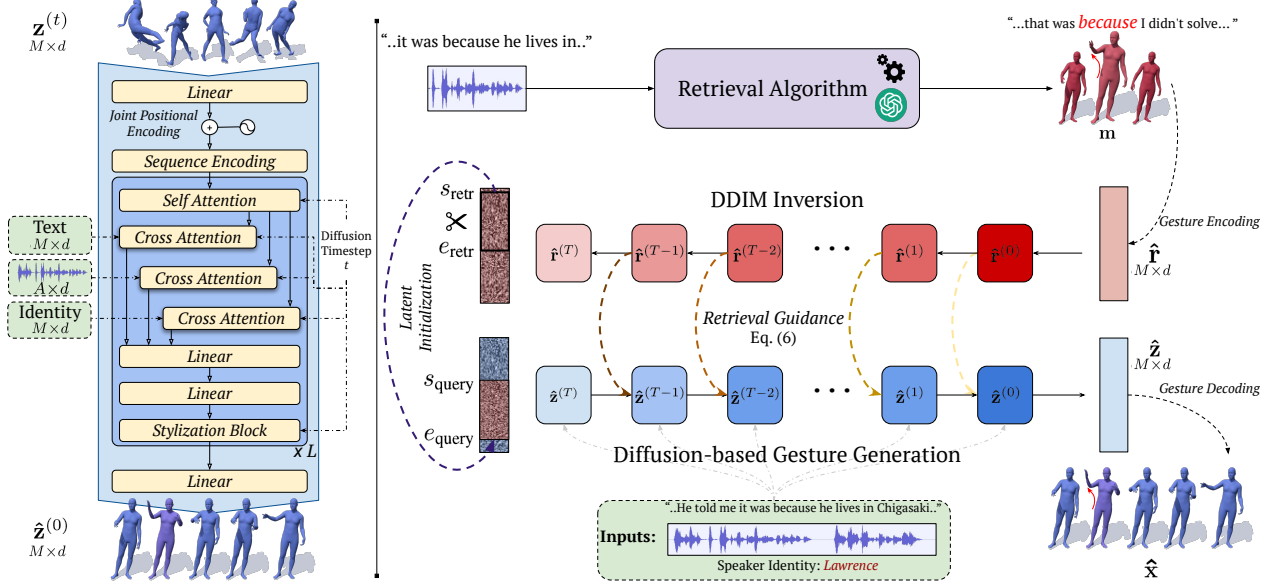


Figure 3. **RAG-GESTURE Framework.** Our approach leverages a **diffusion model** which predicts clean sample $\hat{\mathbf{z}}^{(0)}$ from noisy gesture sample $\mathbf{z}^{(t)}$. We then utilize retrieval algorithms (Sec. 3.4) to modify the gesture sampling at inference time by inserting the retrieved motion through Latent Initialization (Sec. 3.2) and further controlling the sampling process through Retrieval Guidance (Sec. 3.3). This results in a **sampld motion** which follows the **semantic retrieval**.

\mathbf{x} into four different body regions: $\mathbf{x}_u \in \mathbb{R}^{N \times 6J_u}$, $\mathbf{x}_h \in \mathbb{R}^{N \times 6J_h}$, $\mathbf{x}_f \in \mathbb{R}^{N \times 100}$ and $\mathbf{x}_l \in \mathbb{R}^{N \times (6J_l + 3 + 4)}$ for upper body, hands, face and a combined representation for lower body and translation respectively. We train separate time-aware VAEs [47] for each body region, which can encode and decode the motion for each region: $\mathbf{z}_i = \xi_i(\mathbf{x}_i)$, $\mathbf{x}'_i = \mathcal{D}_i(\mathbf{z}_i)$. These VAEs are trained using a combination of standard reconstruction/geometric losses for motion and the KL-Divergence loss for the VAE latent space (see supplemental). The resulting gesture representation is a set of decoupled encodings: $\mathbf{z} = \{\mathbf{z}_u, \mathbf{z}_h, \mathbf{z}_f, \mathbf{z}_l\}$. This set can be concatenated to form a gesture representation $\mathbf{z} \in \mathbb{R}^{M \times d_z}$, where $M < N$ is the length of time-compressed encoding corresponding to each \mathbf{x}_i and d_z is embedding length of the representation. Refer to the supplemental for more details.

Conditional Diffusion for Gesture Generation. Once we have obtained \mathbf{z} as our gesture representation, we can frame the gesture synthesis task as that of conditional diffusion [27, 55], where the conditioning set \mathbf{C} comprises of audio, text and speaker embeddings. In this framework, the *forward diffusion* process consists of Markovian chain of successive noising steps, where Gaussian noise ϵ is added to the encoded gesture $\mathbf{z}^{(0)}$ for T timesteps until $\mathbf{z}^{(T)} \sim \mathcal{N}(0, \mathbf{I})$. For gesture generation, the *reverse process* is employed which iteratively denoises $\mathbf{z}^{(T)}$ to generate the gesture representation $\hat{\mathbf{z}}^{(0)}$. The process utilizes a neural network $f_\theta(\mathbf{z}^{(t)}, t, \mathbf{C})$ for denoising, that is trained

to predict clean state $\hat{\mathbf{z}}^{(0)}$ with the following objective:

$$\min_{\theta} \mathbb{E}_{t \sim [1, T], \mathbf{z}^{(0)}, \epsilon} \|\mathbf{z}^{(0)} - f_\theta(\mathbf{z}^{(t)}, t, \mathbf{C})\|_2^2 \quad (1)$$

During inference, we utilize DDIM [57] sampling for efficient generation of the gesture representation i.e. we generate a sample $\hat{\mathbf{z}}^{(t-1)}$ from $\hat{\mathbf{z}}^{(t)}$ using:

$$\hat{\mathbf{z}}^{(t-1)} = \sqrt{\bar{\alpha}_{t-1}} \hat{\mathbf{z}}^{(0)} + \sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \epsilon_\theta(\hat{\mathbf{z}}^{(t)}) + \sigma_t \epsilon^{(t)} \quad (2)$$

Here $\bar{\alpha}_t$ controls rate of diffusion and σ_t controls stochasticity such that $\sigma_t = 0$ results in deterministic sampling.

Our denoising network consists of a transformer decoder network with L decoder layers containing a separate cross-attention head for each modality from the conditioning set \mathbf{C} . The resulting activations from each head are combined using a linear layer in each decoder layer. Diffusion timestep t is passed to the decoder layers through Stylization block [69] after every self-attention, cross-attention and linear layer. We encode speech signal using wav2vec2 [4] and construct frame-aligned representation for transcription using BERT embeddings for each word [13]. Additionally, we create identity-specific speaker embeddings for input.

3.2. Latent Initialization through DDIM Inversion

So far, our base neural generation framework is purely data-driven and lies on one end of neural-to-explicit spectrum. It resembles the majority of state-of-the-art gesture generation frameworks [3, 51, 72], which rely on learned patterns to sample relevant gestures for corresponding speech

and struggle to sample the semantic ones. On the other hand of the spectrum, pure retrieval based approaches can naïvely “paste” raw exemplar motions onto the generated ones which hurts the naturalness of motion. For instance, one could use diffusion inpainting [41] to force our diffusion framework to follow the retrieval (See Sec. 4.3 for analysis), achieving sub-par results. Therefore, we carefully infuse the retrieved gestures into the diffusion sampling process by performing the retrieval transfer in the latent space of the trained diffusion model and provide better sampling path for its generation through retrieval guidance (Sec. 3.3). This preserves the quality of the base model, adapts it to generate semantically rich gestures and provides further control over the influence of retrieval insertion.

In our framework, all retrieval algorithms (discussed later in Sec. 3.4) must output an exemplar motion \mathbf{m} , which can be encoded as the *retrieval* $\mathbf{r} \in \mathbb{R}^{M \times d}$ through our gesture encoder (Sec. 3.1). We also have access to the text and audio corresponding to the retrieved motion. Our goal is to extract the motion chunk for the retrieved semantic gesture around a word of interest (e.g. ‘because’ in Fig. 3) and insert it into the motion around the corresponding text in the input text. Specifically, let the retrieved motion contain the semantically relevant gesture chunk in the frame window $(s_{\text{retr}}, e_{\text{retr}})$ for the marked word. The gesture in this window must be transferred onto the corresponding word location $(s_{\text{query}}, e_{\text{query}})$ in the to-be-generated sequence, also referred to as *query* $\hat{\mathbf{z}}$ (see Fig. 3).

Here, the challenge is *how* to insert the retrieved gesture chunk such that it is seamlessly integrated into the final inference output in terms of temporal and semantic alignment. We perform this insertion through a combination of DDIM [14, 46, 57] inversion based latent initialization and a retrieval guidance objective. DDIM inversion lets us represent the retrieval \mathbf{r} into the latent space of our base diffusion model, which we then use to initialize our generation.

By reversing Eq. (2) and keeping $\sigma_t = 0$ for deterministic sampling path, we can invert $\mathbf{r}^{(0)}$ as follows:

$$\hat{\mathbf{r}}^{(t+1)} = \sqrt{\bar{\alpha}_{t+1}}\hat{\mathbf{r}}^{(0)} + \sqrt{1 - \bar{\alpha}_{t+1}}\epsilon_{\theta}(\hat{\mathbf{r}}^{(t)}) \quad (3)$$

This iterative process inverts the clean $\mathbf{r}^{(0)}$ to full noise $\hat{\mathbf{r}}^{(T)}$, which represents retrieved gesture in the diffusion latent space and provides the initial starting point for DDIM sampling such that it ends up at $\hat{\mathbf{r}}^{(0)} \approx \mathbf{r}^{(0)}$.

To transfer the retrieved gesture onto the query gesture generation, we slice the time axis of the latent $\hat{\mathbf{r}}^{(T)}$ to extract the relevant parts using the retrieval window $[s_{\text{retr}} : e_{\text{retr}}]$ and then, place these parts onto the query window in $\hat{\mathbf{z}}^{(T)}$.

$$\hat{\mathbf{z}}^{(T)}[s_{\text{query}} : e_{\text{query}}] \leftarrow \hat{\mathbf{r}}^{(T)}[s_{\text{retr}} : e_{\text{retr}}] \quad (4)$$

The retrieval augmented $\hat{\mathbf{z}}^{(T)}$ can be represented as $\hat{\mathbf{z}}_{\text{retr}}^{(T)}$.

3.3. Retrieval Guidance

As a result of latent initialization step, $\hat{\mathbf{z}}_{\text{retr}}^{(T)}$ comes out to be the starting point for gesture sampling process, which can be used to generate a sequence that follows the retrieved semantic gesture around the marked word. However, there is still no control over how strongly the retrieval is followed and how the diffusion sampling process goes from $\hat{\mathbf{z}}_{\text{retr}}^{(T)}$ to $\hat{\mathbf{z}}_{\text{retr}}^{(0)}$. Therefore, we present *Retrieval Guidance* mechanism, that provides modulation control over the retrieval-augmented diffusion generation process from $\hat{\mathbf{z}}_{\text{retr}}^{(T)}$ to $\hat{\mathbf{z}}_{\text{retr}}^{(0)}$. We utilize the sequence of inverted diffusion latents $(\hat{\mathbf{r}}^{(0)}, \hat{\mathbf{r}}^{(1)}, \dots, \hat{\mathbf{r}}^{(T-1)}, \hat{\mathbf{r}}^{(T)})$ and form our guidance objective. At a given diffusion timestep t , guidance objective and its corresponding latent update are as follows:

$$G_{\text{retrieval}} = \left\| \hat{\mathbf{z}}_{\text{retr}}^{(t)}[s_{\text{query}} : e_{\text{query}}] - \hat{\mathbf{r}}^{(t)}[s_{\text{retr}} : e_{\text{retr}}] \right\|_2^2 \quad (5)$$

$$\tilde{\mathbf{z}}_{\text{retr}}^{(t)} \leftarrow \hat{\mathbf{z}}_{\text{retr}}^{(t)} - \lambda \nabla_{\hat{\mathbf{z}}_{\text{retr}}^{(t)}} G_{\text{retrieval}} \quad (6)$$

The updated $\tilde{\mathbf{z}}_{\text{retr}}^{(t)}$ can be used further in the diffusion sampling by acting as input to the denoising network $f_{\theta}(\tilde{\mathbf{z}}^{(t)}, t, \mathbf{C})$. This guidance paradigm ensures that latent $\hat{\mathbf{z}}_{\text{retr}}^{(t)}$ follows $\hat{\mathbf{r}}^{(t)}$ at the generation timestep t .

To control the amount of guidance at each step, one can control the number of iterations for latent update at each diffusion timestep t . As an example, the generation from $\hat{\mathbf{z}}_{\text{retr}}^{(0)}$ without any guidance will be the weakest form of retrieval augmentation where gestures are constrained to match only at the starting timestep $t = 0$. On the other hand, if one were to make sure retrieval is followed to its strongest extent, $\tilde{\mathbf{z}}_{\text{retr}}^{(t)}$ can be updated at each diffusion timestep t such that $G_{\text{retrieval}} \rightarrow 0$. Similarly, one can perform more latent updates when t is closer to T and lesser or no updates when $t \rightarrow 0$, in order to only perform guidance at the start of the sampling process. Refer to Sec. 4.3 for analysis.

3.4. Retrieval for Semantic Gesture Extraction

Having discussed *how* to insert a retrieved gesture chunk into the gesture generation process, we now discuss retrieval algorithms that utilize linguistic context to determine *which* words in the query utterance are likely to invoke semantic gestures and *what* are those gestures. These algorithms retrieve semantically relevant gestures using domain knowledge of language and gesture type information. However, our approach is not limited to these algorithms and it can be extended using any retrieval paradigm for gestures.

At the core, each retrieval algorithm ranks examples in a database (the BEAT2 dataset [40] in our case) according to a specific focus. During ranking, some of the steps are common across algorithms and are driven by over-arching themes in gesture understanding. For example, examples from the same speaker typically have less

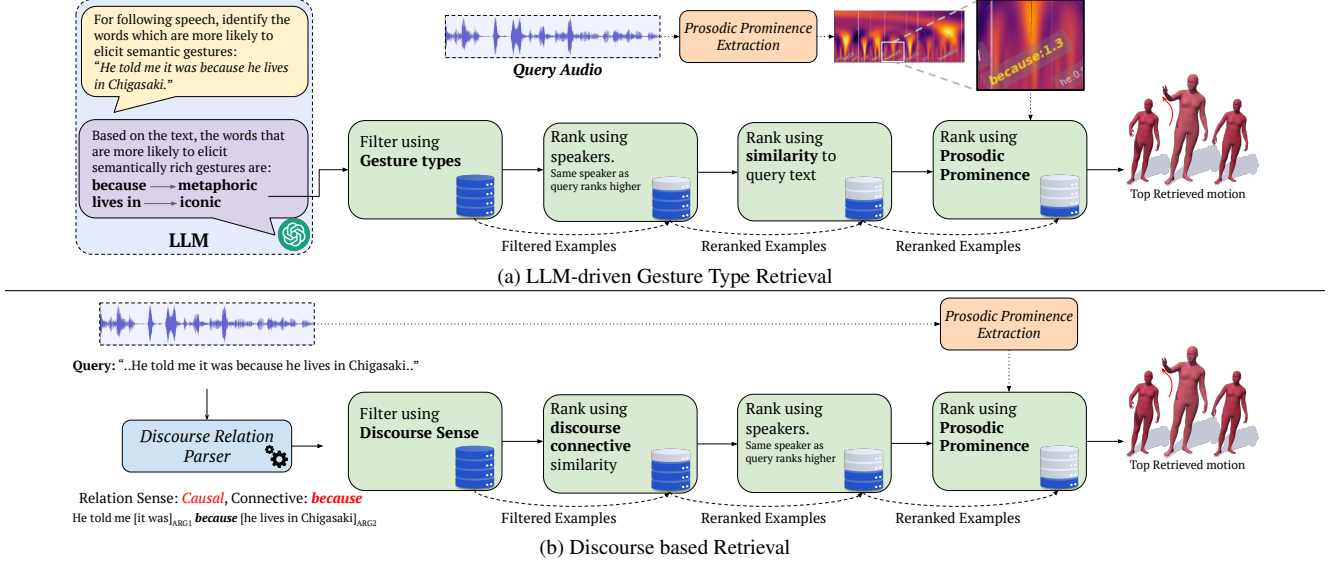


Figure 4. **Retrieval Algorithms.** Each algorithm parses the relevant semantic information (gesture types from LLM or discourse relations) and extracts gestures from a database by filtering examples using that information. Moreover, it also considers textual and prosodic context.

speaker-dependent variation and should be ranked higher. Another common step across algorithms is ranking based on prosodic prominence extraction [16]. Certain elements of speech typically stand out due to variations in pitch, loudness, or duration, referred to as *prosodic prominence* [60, 62]. Prosodic prominence and gestures are closely connected: co-speech gestures often align with prosodically stressed parts of speech [54]. Therefore, we identify utterances in the database that match the prominence values of the marked word in the input speech [59]. An overview of individual steps is provided in Fig. 4.

LLM-based Gesture Type Retrieval. One of the commonly used classification for semantic gestures divides these gestures into *iconic*, *metaphoric* and *deictic* gestures [43]. Kucherenko *et al.* [35] demonstrate the benefit of utilizing these gesture properties in the synthesis framework, as gesture type information can decouple the specification problem from the animation. Therefore, we leverage the reasoning abilities of an LLM [53] and prompt it to identify words which may elicit a semantically meaningful gesture. We also use it to predict the gesture type on each predicted word. This results in a word-to-gesture type mapping, which we can use to retrieve examples from the database that contains type labels. Details on prompts and LLM are given in the supplemental material.

Given the word-to-gesture type mapping, we take an identified word and its corresponding gesture type and filter examples in the database based on gesture type. We leverage the semantic gesture type annotations in the BEAT2 dataset [40] for this task. We then perform ranking based on speaker similarity, followed by ranking according to the text

feature similarity between the query and database examples in order to match the semantic context. Lastly, we re-rank the top semantically similar examples based on prosodic prominence values. Fig. 4a illustrates the steps involved.

Discourse-based Retrieval. A text has a structure that organizes its information into a coherent flow. This structure is achieved through *discourse relations*, semantic-pragmatic links such as CAUSE-CONSEQUENCE and CONTRAST, which hold between clauses and sentences [42, 56]. These relations can be signaled by lexical cues referred to as connectives, such as *on the other hand* or *because*. For example, in an utterance from BEAT2 dataset [40]: “I’ll go shopping *if* I’m not that tired”, the word “*if*” marks the conditional relation between two sentence segments. Prior works have demonstrated that discourse connectives have an effect on the gestural patterns [9, 45].

We leverage connectives to retrieve the relevant semantic gestures that co-occur with specific discourse relations. However, connectives can be ambiguous. For example, the connective *since* can signal a causal relation or a temporal one. We therefore retrieve the gesture co-occurring with a connective with the same meaning, but not necessarily the same word. For example, we can retrieve a gesture co-occurring with *because* to inform the generation of a gesture co-occurring with a causal *since*. Additionally, it allows us to retrieve more samples for less frequent connectives.

In sum, we find examples in the database that carry the same discourse relational sense as the query and rank again based on the word similarity of the connective. We conclude the algorithm with speaker-based and prosodic similarity ranking. The steps are shown in Fig. 4b.

	1 Speaker				All Speakers			
	FID↓	BeatAlign→	L1Div→	Diversity→	FID↓	BeatAlign→	L1Div→	Diversity→
GT		0.703	11.97	127		0.477	7.29	110
CaMN [39]	0.604	0.711	9.97	107	0.512	0.200	5.58	98
EMAGE [40]	0.570	0.793	11.41	124	0.692	0.284	6.06	88
Audio2Photoreal [51]	1.02	0.550	12.47	145	0.849	0.326	6.24	99
ReMoDiffuse [68]	0.702	0.824	12.46	123	1.120	0.218	5.06	116
Ours (w/ Discourse)	0.879	0.730	12.62	129	0.447	0.417	9.03	114
Ours (w/ LLM & Gesture Type)	0.808	0.734	11.97	122	0.487	0.514	9.94	118

Table 1. **Comparison with state-of-the-art methods trained on BEAT.** We demonstrate superior performance, especially when generalizing well across multiple speaker identities.

4. Experiments

We first compare our approach with neural state-of-the-art baselines through quantitative (Sec. 4.1) and perceptual evaluation (Sec. 4.2). Further, we demonstrate how our RAG approach improves upon the diffusion-based RAG baseline. Lastly, we validate our method design through extensive ablations and user studies (Sec. 4.3).

Specifically, we compare with recent data-driven gesture generation approaches that are CaMN [39] (LSTM-based), EMAGE [40] (transformer-based) and Audio2Photoreal [51] (diffusion-based). Furthermore, to validate RAG performance for gesture synthesis, we not only compare our method with ReMoDiffuse [68] (re-trained for gesture synthesis), but we also analyze it further by tailoring their approach and training it with our retrieval algorithms. We provide seed motions for baselines [39, 40], but do not use seed motions to generate our results.

Evaluation Dataset. We evaluate our performance on train/val/test split from BEAT2 dataset [40], which contains 25 speakers. Unlike baseline approaches for BEAT2 which report results only for a set of speakers, we perform evaluation on the test set of one speaker (*Scott*) and all speakers in order to evaluate performance on large-scale multi-speaker data. Our test set contains 265 utterances chunked into 10-sec sequences. We retrain methods on BEAT2 if needed.

4.1. Quantitative Evaluation

We evaluate our method on established metrics: Beat-Alignment [37], FID [65], L1 Divergence and Diversity, which measure different aspects of the motion quality. We tabulate the quantitative results in Tab. 1, where the results are divided on the basis of number of speakers used for training and evaluation. We observe that our approach achieves state-of-the-art metric performance when trained with all speakers. This also entails that our RAG-based diffusion approach generalizes and scales well across large-scale data with multiple speaker identities. For the case of single speaker training, we observe best performance in terms of diversity and L1 Divergence and second-best in Beat Alignment score. However, our diffusion-based




Figure 5. **Results of Perceptual Evaluation.** A2P: Audio2Photoreal [51] GT: Ground Truth. * denotes p-value < 0.05

method has lower FID compared to transformer and LSTM based methods [39, 40] that utilize ground-truth seed and can be over-fitted to a small-scale data of single speaker, while not being able to generalize to all speakers.

4.2. Perceptual Evaluation

Evaluating gesture synthesis performance is a highly challenging task due to the stochasticity in gesture production and subjective nature of gesture perception [52, 63]. Therefore, we perform three different user studies for comparison and ablative analysis (Sec. 4.3). First, we evaluate the synthesis quality compared to the baseline methods. We investigate the quality of gestures based on well-established measures of *naturalness* and *appropriateness* in relation to speech [48]. We ask the participants to perform pairwise comparisons between generations from baselines and our method. Results compare RAG-GESTURE with discourse and LLM-based retrievals against state-of-the-art methods (Fig. 5). We observe that our method achieves better perceptual quality for both naturalness and appropriateness. For LLM based approach, we are marginally below the ground-truth preference, underscoring that our method produces not only natural, but also highly appropriate gestures.

4.3. Ablative Analysis

Improvement over RAG Baselines. One of our RAG-driven baselines (ReMoDiffuse [68]) follows global text similarity for retrieval and proposes to train the diffusion network to follow the retrievals. In contrast, we design our retrieval to match local context around the semantically important word for gesticulation and do not




Figure 6. **Perceptual Comparison** between training-based base-lines and our inference-time approach. *: p-value < 0.05

	FID↓	BeatAlign→	L1Div→	MPJPE (mm)↓
GT		0.477	7.29	
ReMoDiffuse [68]	1.120	0.218	5.06	200.5
ReMoDiffuse [68] + Our Retrieval	0.525	0.414	6.909	34.6
RAG-GESTURE	0.447	0.471	9.03	35.0

Table 2. Quantitative Comparison between RAG baselines.




Figure 7. Study Preference comparing **LLM and Discourse** algorithms.

require training/fine-tuning for RAG unlike existing approaches [68, 71]. We evaluate these differences by comparing our RAG-GESTURE (Discourse) with (i) base ReMoDiffuse as well as (ii) a modified version of their diffusion architecture, which uses gesture encodings as input instead of raw motion, along with the proposed discourse-based retrieval algorithm for RAG (see supplementary for details). Note that this architecture is similar to SemanticGesticulator [71] at a higher level, which also performs retrieval merging in encoded gestures during fine-tuning.

We evaluate this quantitatively in Tab. 2 and perform an additional user study (Fig. 6). Moreover, we also report MPJPE between the generated gesture and retrieved motion during the retrieval insertion window to measure the retrieval following performance. Firstly, we observe that our modified version of ReMoDiffuse which uses local semantic retrieval during training performs better than the base model, both in terms of perceptual quality and metrics. Moreover, we also see higher faithfulness to the semantic retrievals through low MPJPE which entails that properly specifying the retrieved motion on the exact word affects the gesture quality. Secondly, our inference-time RAG-GESTURE is better than the training-based model, even after utilizing our retrieval algorithm. This underscores our method’s superiority and the importance of having a good gesture animation performance.

LLM-based Gesture Type vs. Discourse Retrieval. For evaluating perceptual differences between LLM-based gesture type retrieval and discourse-based retrieval, we conduct an additional user study where we ask the participants to evaluate differences between generated gestures from two algorithms in terms of naturalness and appropriateness (Fig. 7). Interestingly, we find that evaluators do not show any strong preference between the two algorithms.

	FID↓	BeatAlign→	L1Div→	MPJPE (mm)↓
GT		0.477	7.29	
No RAG	0.519	0.447	8.644	59.6
RAG by Inpainting [41]	0.446	0.450	8.457	28.1
RAG by LI only	0.589	0.486	9.228	54.6
RAG-GESTURE (LI+RG)	0.447	0.471	9.038	35.0

Table 3. **Extent of Semantic Retrieval Insertion.** LI: Latent Initialization; RG: Retrieval Guidance.




Figure 8. **Ablative Study** which illustrates preference of RAG-GESTURE over varying degrees of insertions. *: p-value < 0.05

Extent of Retrieval Augmentation Since our approach can control the influence of the retrieved gesture onto the generated one, we analyze how much influence is optimal to synthesize semantically accurate yet natural looking gestures. We compare RAG-GESTURE’s gesture prediction against varying degrees of retrieval insertion (Tab. 3), where “Inpainting” means naively pasting the retrieval over the gesture during synthesis and “No RAG” represents no retrieval insertion. We conduct another user study with a third “Retrieval Appropriateness” question which tries to gauge how semantically appropriate the generated gesture was during the insertion window Fig. 8. Interestingly, we observe increasing preference of our retrieval-guided model as we move towards both extremes of the insertion. In terms of gesture perception, evaluators neither prefer generation without RAG (purely data-driven) nor completely strong insertion through Inpainting (purely retrieval-based). Lastly, our model achieves good balance of retrieval following in terms of MPJPE, prosodic alignment and diversity.

5. Conclusion

In this work, we present RAG-GESTURE – a novel approach that leverages interpretable retrieval algorithms to introduce inference-time RAG to a diffusion-based gesture generator. Therefore, it combines advantages of the pure retrieval-based and fully data-driven approaches to synthesize natural yet semantically meaningful gestures. Perceptual evaluation demonstrates that this type of neuro-explicit approach is better than either of the two approaches and existing state-of-the-art methods. Moreover, we demonstrate that LLM-driven gesture type prediction and explicit discourse relations can mitigate the lack of semantic understanding in current gesture generation architectures. Lastly, our framework can be extended to synthesize gestures that exhibit task-specific gestural patterns like referential or

emotion-specific gestures.

Acknowledgements. This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – GRK 2853/1 “Neuroexplicit Models of Language, Vision, and Action” - project number 471607914. The third author of this study (MS) was supported through the NWO-funded project “Searching for meaning: Identifying and interpreting alternative discourse relation signals in multi-modal language comprehension” (VI.Veni.231C.021).

References

- [1] Chaitanya Ahuja, Dong Won Lee, Yukiko I. Nakano, and Louis-Philippe Morency. Style transfer for co-speech gesture animation: A multi-speaker conditional-mixture approach. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [2] Simon Alexanderson, Gustav Eje Henter, Taras Kucherenko, and Jonas Beskow. Style-controllable speech-driven gesture synthesis using normalising flows. *Computer Graphics Forum*, 2020. 2
- [3] Tenglong Ao, Zeyi Zhang, and Libin Liu. Gesturediffuclip: Gesture diffusion model with clip latents. *ACM TOG*, 42(4): 1–18, 2023. 4
- [4] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460, 2020. 4
- [5] Kirsten Bergmann and Stefan Kopp. Gnetic—using bayesian decision networks for iconic gesture generation. In *Intelligent Virtual Agents: 9th International Conference, IVA 2009 Amsterdam, The Netherlands, September 14-16, 2009 Proceedings 9*, pages 76–89. Springer, 2009. 3
- [6] Jessica A. Bernard, Zachary B. Millman, and Vijay A. Mittal. Beat and metaphoric gestures are differentially associated with regional cerebellar and cortical volumes. *Human Brain Mapping*, 36(10):4016–4030, 2015. 1
- [7] Jana Bressemer and Cornelia Müller. A repertoire of German recurrent gestures with pragmatic functions. In *Body–language– communication: An international handbook on multi-modality in human interaction [Handbooks of Linguistics and Communication Science 38.2]*, pages 1575–1591. Mouton de Gruyter, 2014. 3
- [8] Michael Büttner and Simon Clavet. Motion matching—the road to next gen animation. *Proc. of Nucl. ai*, 1(2015):2, 2015. 3
- [9] G. Calbris. *Elements of Meaning in Gesture*. John Benjamins Publishing Company, 2011. 2, 6
- [10] Justine Cassell, Hannes Högni Vilhjálmsón, and Timothy Bickmore. Beat: The behavior expression animation toolkit. In *SIGGRAPH Conference Proceedings*, 2001. 2
- [11] Simon Clavet et al. Motion matching and the road to next-gen animation. In *Proc. of GDC*, page 4, 2016. 3
- [12] Rishabh Dabral, Muhammad Hamza Mughal, Vladislav Golyanik, and Christian Theobalt. Mofusion: A framework for denoising-diffusion-based motion synthesis. In *CVPR*, 2023. 1
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018. 4, 2
- [14] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. 5
- [15] P Kingma Diederik. Adam: A method for stochastic optimization. In *ICLR*, 2015. 2
- [16] Adso Fernández-Baena, Raúl Montaña, Marc Antonijoan, Arturo Roversi, David Miralles, and Francesc Alías. Gesture synthesis adapted to speech emphasis. *Speech communication*, 57:331–350, 2014. 6
- [17] Ylva Ferstl and Rachel McDonnell. Investigating the use of recurrent motion modelling for speech gesture generation. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*, 2018. 2
- [18] Ylva Ferstl, Michael Neff, and Rachel McDonnell. Adversarial gesture generation with realistic gesture phasing. *Computers & Graphics*, 89:117–130, 2020. 2
- [19] Ylva Ferstl, Michael Neff, and Rachel McDonnell. Express-gesture: Expressive gesture generation from speech through database matching. *Computer Animation and Virtual Worlds*, 32(3-4):e2016, 2021. 3
- [20] Saeed Ghorbani, Ylva Ferstl, Daniel Holden, Nikolaus F. Troje, and Marc-André Carbonneau. Zeroeggs: Zero-shot example-based gesture generation from speech. *Computer Graphics Forum*, 42(1):206–216, 2023. 2
- [21] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. Imos: Intent-driven full-body motion synthesis for human-object interactions. In *Eurographics*, 2023. 2
- [22] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. Remos: 3d motion-conditioned reaction synthesis for two-person interactions. In *European Conference on Computer Vision (ECCV)*, 2024. 2
- [23] Purvi Goel, Kuan-Chieh Wang, C Karen Liu, and Kayvon Fatahalian. Iterative motion editing with natural language. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–9, 2024. 3
- [24] Ikhsanul Habibie, Weipeng Xu, Dushyant Mehta, Lingjie Liu, Hans-Peter Seidel, Gerard Pons-Moll, Mohamed Elgharib, and Christian Theobalt. Learning speech-driven 3d conversational gestures from video. In *Proceedings of the International Conference on Intelligent Virtual Agents*, 2021. 2
- [25] Ikhsanul Habibie, Mohamed Elgharib, Kripasindhu Sarkar, Ahsan Abdullah, Simbarashe Nyatsanga, Michael Neff, and Christian Theobalt. A motion matching-based framework for controllable gesture synthesis from speech. In *ACM SIGGRAPH 2022 conference proceedings*, pages 1–9, 2022. 3
- [26] Jennifer Hinnell. The verbal-kinesic enactment of contrast in north american english. *The American Journal of Semiotics*, 35(1-2):55–92, 2019. 3

- [27] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 4
- [28] Daniel Holden. Character control with neural networks and machine learning. *Proc. of GDC 2018*, 1:2, 2018. 3
- [29] Daniel Holden, Oussama Kanoun, Maksym Perepichka, and Tiberiu Popa. Learned motion matching. *ACM Transactions on Graphics (TOG)*, 39(4):53–1, 2020. 3
- [30] Yiming Huang, Weilin Wan, Yue Yang, Chris Callison-Burch, Mark Yatskar, and Lingjie Liu. Como: Controllable motion generation through language guided pose code editing. *arXiv preprint arXiv:2403.13900*, 2024. 3
- [31] Anna Inbar. The raised index finger gesture in Hebrew multimodal interaction. *Gesture*, 21(2-3):264–295, 2022. 3
- [32] Adam Kendon. *Gesture: Visible action as utterance*. Cambridge University Press, 2004. 3
- [33] Michael Kipp. *Gesture generation by imitation: From human behavior to computer character animation*. Universal Publishers, 2005. 3
- [34] Taras Kucherenko, Dai Hasegawa, Gustav Eje Henter, Naoshi Kaneko, and Hedvig Kjellström. Analyzing input and output representations for speech-driven gesture generation. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, page 97–104. Association for Computing Machinery, 2019. 2
- [35] Taras Kucherenko, Rajmund Nagy, Patrik Jonell, Michael Neff, Hedvig Kjellström, and Gustav Eje Henter. Speech2Properties2Gestures: Gesture-property prediction as a tool for generating representational gestures from speech. In *Proceedings of the 21th ACM International Conference on Intelligent Virtual Agents*, 2021. 2, 6
- [36] Schuyler Laparle. Tracking discourse topics in co-speech gesture. In *International Conference on Human-Computer Interaction (HCI)*, page 233–249, Berlin, Heidelberg, 2021. Springer-Verlag. 3
- [37] Ruilong Li, Sha Yang, David A. Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *ICCV*, 2021. 7, 1
- [38] Yuanzhi Liang, Qianyu Feng, Linchao Zhu, Li Hu, Pan Pan, and Yi Yang. Seeg: Semantic energized co-speech gesture generation. In *CVPR*, 2022. 2
- [39] Haiyang Liu, Zihao Zhu, Naoya Iwamoto, Yichen Peng, Zhengqing Li, You Zhou, Elif Bozkurt, and Bo Zheng. Beat: A large-scale semantic and emotional multi-modal dataset for conversational gestures synthesis. *European Conference on Computer Vision*, 2022. 1, 2, 3, 7
- [40] Haiyang Liu, Zihao Zhu, Giorgio Becherini, Yichen Peng, Mingyang Su, You Zhou, Xuefei Zhe, Naoya Iwamoto, Bo Zheng, and Michael J Black. Emage: Towards unified holistic co-speech gesture generation via expressive masked audio gesture modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1144–1154, 2024. 2, 3, 5, 6, 7, 1
- [41] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *CVPR*, 2022. 5, 8
- [42] William C Mann and Sandra A Thompson. Rhetorical Structure Theory: Toward a functional theory of text organization. *Text-Interdisciplinary Journal for the Study of Discourse*, 8(3):243–281, 1988. 6
- [43] David McNeill. *Hand and Mind: What Gestures Reveal About Thought*. University of Chicago Press, 1992. 1, 6
- [44] David McNeill. *Gesture and Thought*. University of Chicago Press, 2005. 1
- [45] David McNeill, Elena Levy, and Susan Duncan. Gesture in discourse. *The Handbook of Discourse Analysis*, 2014. 2, 6
- [46] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6038–6047, 2023. 5
- [47] Muhammad Hamza Mughal, Rishabh Dabral, Ikhsanul Habibie, Lucia Donatelli, Marc Habermann, and Christian Theobalt. ConvoFusion: Multi-modal conversational diffusion for co-speech gesture synthesis. In *Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 3, 4
- [48] Rajmund Nagy, Hendric Voss, Youngwoo Yoon, Taras Kucherenko, Teodor Nikolov, Thanh Hoang-Minh, Rachel McDonnell, Stefan Kopp, Michael Neff, and Gustav Eje Henter. Towards a genea leaderboard—an extended, living benchmark for evaluating and advancing conversational motion synthesis. *arXiv preprint arXiv:2410.06327*, 2024. 7
- [49] Michael Neff. Hand gesture synthesis for conversational characters. *Handbook of Human Motion*, pages 1–12, 2016. 2, 3
- [50] Michael Neff, Michael Kipp, Irene Albrecht, and Hans-Peter Seidel. Gesture modeling and animation based on a probabilistic re-creation of speaker style. *ACM TOG*, 27(1):1–24, 2008. 3
- [51] Evonne Ng, Javier Romero, Timur Bagautdinov, Shaojie Bai, Trevor Darrell, Angjoo Kanazawa, and Alexander Richard. From audio to photoreal embodiment: Synthesizing humans in conversations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1001–1010, 2024. 2, 4, 7, 1, 3
- [52] Simbarashe Nyatsanga, Taras Kucherenko, Chaitanya Ahuja, Gustav Eje Henter, and Michael Neff. A comprehensive review of data-driven co-speech gesture generation. In *Computer Graphics Forum*, pages 569–596. Wiley Online Library, 2023. 1, 2, 7
- [53] OpenAI, Achiam J, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>, 2024. 2, 6
- [54] Patrick Louis Rohrer, Elisabeth Delais-Roussarie, and Pilar Prieto. Visualizing prosodic structure: Manual gestures as highlighters of prosodic heads and edges in english academic discourses. *Lingua*, 293:103583, 2023. 6
- [55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2021. 3, 4

- [56] Ted J M Sanders, Wilbert P M S Spooren, and Leo G M Noordman. Toward a taxonomy of coherence relations. *Discourse Processes*, 15(1):1–35, 1992. [6](#)
- [57] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. [4](#), [5](#)
- [58] Sebastian Starke, Paul Starke, Nicky He, Taku Komura, and Yuting Ye. Categorical codebook matching for embodied character controllers. *ACM Transactions on Graphics (TOG)*, 43(4):1–14, 2024. [3](#)
- [59] Antti Suni, Juraj Šimko, Daniel Aalto, and Martti Vainio. Hierarchical representation and estimation of prosody using continuous wavelet transform. *Computer Speech & Language*, 45:123–136, 2017. [6](#)
- [60] Jacques Terken and Dik Hermes. The perception of prosodic prominence. In *Prosody: Theory and experiment: Studies presented to Gösta Bruce*, pages 89–127. Springer, 2000. [6](#)
- [61] Marcus Thiebaux, Stacy Marsella, Andrew N Marshall, and Marcelo Kallmann. Smartbody: Behavior realization for embodied conversational agents. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*, pages 151–158, 2008. [2](#)
- [62] Michael Wagner and Duane G Watson. Experimental and theoretical advances in prosody: A review. *Language and cognitive processes*, 25(7-9):905–945, 2010. [6](#)
- [63] Pieter Wolfert, Nicole Robinson, and Tony Belpaeme. A review of evaluation practices of gesture generation in embodied conversational agents. *IEEE Transactions on Human-Machine Systems*, 52(3):379–389, 2022. [7](#)
- [64] Hongwei Yi, Hualin Liang, Yifei Liu, Qiong Cao, Yandong Wen, Timo Bolkart, Dacheng Tao, and Michael J Black. Generating holistic 3D human motion from speech. In *CVPR*, 2023. [2](#)
- [65] Youngwoo Yoon, Woo-Ri Ko, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. Robots learn social skills: End-to-end learning of co-speech gesture generation for humanoid robots. In *2019 International Conference on Robotics and Automation (ICRA)*, 2019. [2](#), [7](#)
- [66] Youngwoo Yoon, Bok Cha, Joo-Haeng Lee, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. Speech gesture generation from the trimodal context of text, audio, and speaker identity. *ACM TOG*, page 1–16, 2020. [2](#), [1](#)
- [67] Youngwoo Yoon, Pieter Wolfert, Taras Kucherenko, Carla Viegas, Teodor Nikolov, Mihail Tsakov, and Gustav Eje Henter. The genea challenge 2022: A large evaluation of data-driven co-speech gesture generation. In *Proceedings of the 2022 International Conference on Multimodal Interaction*, pages 736–747, 2022. [2](#), [3](#)
- [68] Mingyuan Zhang, Xinying Guo, Liang Pan, Zhongang Cai, Fangzhou Hong, Huirong Li, Lei Yang, and Ziwei Liu. Remodiffuse: Retrieval-augmented motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 364–373, 2023. [2](#), [3](#), [7](#), [8](#), [1](#)
- [69] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. [4](#)
- [70] Wanyue Zhang, Rishabh Dabral, Thomas Leimkühler, Vladislav Golyanik, Marc Habermann, and Christian Theobalt. Roam: Robust and object-aware motion generation using neural pose descriptors. *International Conference on 3D Vision (3DV)*, 2024. [2](#)
- [71] Zeyi Zhang, Tenglong Ao, Yuyao Zhang, Qingzhe Gao, Chuan Lin, Baoquan Chen, and Libin Liu. Semantic gesticulator: Semantics-aware co-speech gesture synthesis. *ACM Trans. Graph.*, 2024. [3](#), [8](#)
- [72] Weiyu Zhao, Liangxiao Hu, and Shengping Zhang. Diffugesture: Generating human gesture from two-person dialogue with diffusion models. In *International Conference on Multimodal Interaction*, 2023. [4](#)
- [73] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023. [2](#)
- [74] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5745–5753, 2019. [3](#)

Retrieving Semantics from the Deep: an RAG Solution for Gesture Synthesis

Supplementary Material

We first provide additional evaluation of the method in terms of multi-modality and per-speaker analysis. Then we discuss limitations and provide user study details. Lastly, we discuss evaluation metrics and implementation details.

6. Per-speaker FID/Multimodality comparison with SOTA

To perform a robust evaluation on speaker generalizability of our framework, we provide per-speaker FID and Multi-modality metrics in Tab. 4. We observe that our framework achieves best FID for large number of the speakers, which shows that our method generalizes well to the speaker specific patterns and idiosyncrasies despite taking no seed gestures at input. CaMN [39], achieves lower FID score overall because it uses seed input from the ground truth data which results in lower scores. Due to the same reason, EMAGE [40] also gets lower FID score.

However, CaMN and EMAGE do not perform well on Multi-modality [12]. This metric computes Euclidean distance between different generated gestures for a single speech input. Since, both of these methods always generate same gestures for a given speech input, they perform worse in terms of Multi-modality, which makes them less ideal for diverse gesture generation. In contrast, our approach gets the best results showing diverse gesture generation capabilities of our model (Tab. 5).

7. Details on User Study

For the perceptual evaluation of gesture generation capability, 56 participants were shown a randomly sampled set of 16 forced-choice questions. Our study consists of four sections corresponding to four different user studies. First section focuses on the comparison with the state-of-the-art approaches. Each question comprises of a side-by-side animation of our method along with one of EMAGE [40], Audio2Photoreal [51], RemoDiffuse [68] or the ground-truth. Second section is a short section where we compare LLM-driven gesture generation with discourse-based synthesis. Third, we perform pair-wise comparisons between results from one of the RAG baselines (derived from ReMoDiffuse [68]) and our method. Lastly, we ablate different approaches to perform RAG and control its extent.

In first three sections, we try to evaluate naturalness and appropriateness. Specifically, we ask two questions (a) “Which of the two gestures look natural?” and (b) “Which of them looks appropriately aligned to what the person is saying?”. In the last section, we add an additional question which focuses on gauging the semantic appropriateness

of gestures in the retrieval window, with a goal of evaluating the RAG capability. We highlight the identified words from the retrieval algorithms and add an additional question: “Which of the two have better gestures in the highlighted section, especially at the capitalized word in the prompt?”

8. Limitations

Our method relies on the sparse semantic data which is extracted from the BEAT2 dataset [40]. This creates a scarcity of good exemplars which can be used during the database matching steps. As a result, the LLM-based Gesture Type sometimes struggles to find good contextual matches for each gesture type and word identified by the LLM.

Secondly, our method combines explicit rule-based algorithms with a neural generation framework. In rare cases, these algorithms can fail in edge cases and result in an incorrect retrieved example. However, our learned framework on top of the retrieval algorithms mitigates this and ignores out-of-distribution motion exemplars since it has been trained to produce only those gestures which match the speech. This also is affected by the extent of retrieval augmentation.

9. Evaluation Metrics

FID. We employ the Frechet Inception Distance (FID) metric inspired by Yoon *et al.* [66], which is also known as FGD. We use the autoencoder network provided by BEAT2 [40] to get the gesture encodings for FID evaluation and do not retrain our own network.

Beat Alignment Score. Originally introduced to measure alignment of music beats to dance motion, Beat Alignment Score [37] has been adapted for the gesture synthesis task where it aims to measure the correlation between gesture beats and audio beats.

L1 Divergence. This metric (also called L1 variance) measures the distance of all frames in a single gesture generation from their mean. It is helpful in identifying synthesized gestures that are static and unexpressive.

Diversity. It computes the average pairwise Euclidean distance of the gesture generations in the test set.

Multi-modality. This metric requires sampling different gesture motions for a single speech input from the generative model [12]. Then, it computes Euclidean distance be-

	FID↓						Multimodality↑					
	CaMN	EMAGE	Audio2Phoreal	ReMoDiffuse	Ours (w/ Discourse)	Ours (w/ LLM & Gesture Type)	CaMN	EMAGE	Audio2Phoreal	ReMoDiffuse	Ours (w/ Discourse)	Ours (w/ LLM & Gesture Type)
wayne	1.23	2.06	2.32	3.58	1.49	1.59	n/a	n/a	1.1	3.4	3.1	3.7
scott	0.83	1.17	1.02	1.76	0.78	0.83	n/a	n/a	1.5	7.7	5.3	5.4
solomon	1.22	1.42	1.93	2.45	0.92	0.86	n/a	n/a	0.5	3.5	4.3	4.5
lawrence	0.98	1.39	1.13	3.16	0.69	0.66	n/a	n/a	1.9	6.9	5.6	6.3
stewart	0.65	1.26	1.62	1.76	1.49	1.49	n/a	n/a	0.3	0.7	2.4	3.0
carla	0.81	1.40	1.33	2.95	1.63	1.49	n/a	n/a	0.4	1.3	1.4	1.5
sophie	0.92	1.67	1.85	2.76	1.76	1.74	n/a	n/a	0.69	3.5	2.7	3.2
miranda	0.58	0.87	1.10	1.86	1.09	1.34	n/a	n/a	0.4	0.9	1.7	1.9
kieks	1.30	2.62	1.90	7.65	1.74	1.63	n/a	n/a	1.0	2.1	3.6	4.1
nidal	0.40	0.65	0.72	1.74	0.67	0.64	n/a	n/a	0.7	1.9	2.6	3.1
zhao	1.66	2.70	1.96	3.37	1.37	1.32	n/a	n/a	1.3	3.2	3.3	3.4
lu	1.40	2.73	1.92	2.23	1.27	1.16	n/a	n/a	0.7	1.7	2.4	2.6
carlos	0.78	1.47	1.71	2.47	1.95	2.56	n/a	n/a	0.2	2.5	2.9	3.1
jorge	1.49	2.57	1.97	3.55	1.89	1.93	n/a	n/a	0.3	1.7	1.9	2.1
itoi	0.93	1.61	1.34	2.28	1.07	1.32	n/a	n/a	0.8	1.8	3.0	3.1
daiki	0.78	1.78	1.66	3.04	0.91	1.19	n/a	n/a	0.3	2.3	2.3	2.7
li	1.10	1.74	1.17	2.06	0.71	0.79	n/a	n/a	0.6	3.9	2.7	3.0
ayana	1.19	2.03	1.96	4.35	2.09	2.13	n/a	n/a	0.4	1.8	1.9	2.1
luqi	1.25	2.22	1.67	5.21	1.38	1.86	n/a	n/a	0.7	1.7	2.2	2.5
hailing	0.53	1.20	7.53	5.72	2.35	2.79	n/a	n/a	0.3	1.0	2.5	2.6
kexin	1.07	1.70	1.19	1.87	0.92	0.95	n/a	n/a	0.3	1.8	1.6	1.9
goto	0.84	1.32	2.01	2.51	1.45	2.09	n/a	n/a	0.3	2.3	2.0	2.4
yingqing	1.67	2.50	2.00	4.34	1.82	1.74	n/a	n/a	1.1	3.5	3.0	3.2
tiffnay	0.81	1.35	1.09	2.67	0.92	1.11	n/a	n/a	0.3	1.3	1.2	1.6
katya	1.10	2.09	1.57	2.65	1.15	1.23	n/a	n/a	0.6	2.9	2.3	2.6

Table 4. Per Speaker FID/Multimodality

Multi-modality ↑					
CaMN	EMAGE	Audio2Phoreal	ReMoDiffuse	Ours (w/ Discourse)	Ours (w/ LLM)
n/a	n/a	16.9	66.5	69.1	76.7

Table 5. Overall Multi-modality

tween those generated gestures. It probes the diverse sampling capabilities of a generative model.

10. Implementation Details

10.1. Input Representations

Representing speech and its transcription is highly important aspect of diffusion-based gesture modelling process. In our experiments, we found that changing the structure of text embeddings affects gesture understanding during the learning process, which consequently is reflected during the synthesis phase as well. To construct our text representation, we build a per-frame embedding with corresponding word embeddings residing on each frame. We compute the word embeddings by combining sub-word token activations from last 4 layers of BERT model [13].

10.2. Decoupled gesture encoding

We utilize time-aware VAE architecture by Mughal *et al.* [47]. This architecture utilizes separate encoders for frame window chunks of original motion to encode each chunk into an encoding. Then, it jointly decodes all of the chunks together to reconstruct the original motion. We use $N = 150$ representing 10 seconds of motion at 15 frames per second. Moreover, the motion chunk length of our time-aware VAE is 15, making each chunk encoding correspond to 1 second of motion. This results in an chunked gesture encoding of length 10 for each body part. Finally, we concatenate all 4 body part encodings along the time axis with separators in between them, resulting in $M = 40 + 3 = 43$.

We train the VAE on the reconstruction task by utilizing a set of losses to optimize the model. We apply Geodesic Loss on rotation matrices and standard MSE losses on 6D, axis-angle and joint position representation of the motion. Moreover, we also apply additional MSE losses to optimize velocity/acceleration of motion [21]. Lastly, we apply loss on foot contact predictions during VAE training to reduce foot sliding [22, 70].

10.3. RAG-driven Gesture Diffusion model

To optimize our diffusion model, we utilize Adam [15] with a learning rate of $1e - 4$. We utilize “scaled linear” as our β_t schedule and use 1000 steps while training. For inference, we use spaced 50 steps with DDIM scheduler. The transformer network contains 16 attention heads and 8 decoder layers. To better disambiguate body parts in our gesture encoding, we also add a separate sinusoidal positional encoding for body parts.

Retrieved Motion Insertion. In order to insert the retrieved gestures into the query latents, we only consider latents for upper body and hands. The encodings of these body parts are transferred from retrieval to query sample because speech has the most amount of semantic significance on these two body regions in terms of co-verbal gestures.

10.4. Details on LLM Prompting

We utilize OpenAI’s gpt-4o-mini model for semantic gesture type prediction. We provide a system prompt containing a brief explanation of gesture types and a user prompt which contains text from the test dataset and the question.

System Prompt. “You are an expert in human gestures. You need to identify words that may elicit semantically

meaningful gestures(deictic, iconic, metaphoric) and their types: (a) Metaphoric Gesture: Represents abstract ideas or concepts physically, creating a vivid mental image. (b) Iconic Gesture: Mimics the shape or action of the object or concept being described. (c) Deictic Gesture: Points to or indicates a person, object, or location. Format your response as a python list of python tuples of (word, type). For example: [(‘hello’, ‘beat’), (‘world’, ‘iconic’)]”

User Prompt. Identify at most 2 important words which are more likely to elicit semantically meaningful gestures and what are types of those gestures in following text: “TEXT”.

10.5. Baseline Retraining Details

To be consistent with single speaker evaluation on BEAT2 dataset, we utilize released model weights by EMAGE [40]. For other approaches (including ours), we retrain the method on single speaker data belong to speaker “*scott*”. Since there are no available models for the chosen baselines which have been trained on all speakers in BEAT2 dataset, we train all the methods on complete dataset through their provided codebases. Methods which do not contain speaker specific generalizations like Audio2Photoreal [51], are modified to include a speaker embedding along with text and speech embeddings. Moreover, Audio2Photoreal is adapted to support the skeletal format of BEAT2.

Lastly, ReMoDiffuse, originally released for text-to-motion task, is modified for gesture synthesis and their retrieval process is implemented using text feature similarity method. To ablate the training-time RAG, we further modify the ReMoDiffuse architecture to use our gesture encoding and our retrieval algorithms, which is referred to as “ReMoDiffuse+Our Retrieval”. Analysis on this is given in the Ablative Analysis section.