

Enhancing Spoken Discourse Modeling in Language Models Using Gestural Cues

Varsha Suresh¹, M. Hamza Mughal², Christian Theobalt^{1,2}, Vera Demberg^{1,2}

¹Saarland University, ²Max Planck Institute for Informatics, Saarland Informatics Campus

Correspondence: {vsuresh, vera}@lst.uni-saarland.de, {mmughal, theobalt}@mpi-inf.mpg.de

Abstract

Research in linguistics shows that non-verbal cues, such as gestures, play a crucial role in spoken discourse. For example, speakers perform hand gestures to indicate topic shifts, helping listeners identify transitions in discourse. In this work, we investigate whether the joint modeling of gestures using human motion sequences and language can improve spoken discourse modeling in language models. To integrate gestures into language models, we first encode 3D human motion sequences into discrete gesture tokens using a VQ-VAE. These gesture token embeddings are then aligned with text embeddings through feature alignment, mapping them into the text embedding space. To evaluate the gesture-aligned language model on spoken discourse, we construct text infilling tasks targeting three key discourse cues grounded in linguistic research: discourse connectives, stance markers, and quantifiers. Results show that incorporating gestures enhances marker prediction accuracy across the three tasks, highlighting the complementary information that gestures can offer in modeling spoken discourse. We view this work as an initial step toward leveraging non-verbal cues to advance spoken language modeling in language models.

1 Introduction

Non-verbal cues, such as gestures, eye contact, body posture, and facial expressions, play a crucial role in communication, functioning as a secondary channel alongside language (Krauss et al., 1996; Fichten et al., 1992; McNeill, 1992). Linguistic research highlights the importance of gestures in spoken discourse (Goldin-Meadow et al., 1993; Khosrobeigi et al., 2022; Kita et al., 2017; Holler et al., 2009; Cassell et al., 2001; Ram et al., 2025), with studies showing that they help structure conversations, often marking topic shifts (Kendon, 1995; Quek et al., 2002), and convey the speaker’s attitudes and certainty about their propositions

(Roseano et al., 2016; Andries et al., 2023; Debras, 2015).

Most language models dealing with spoken language data, however, rely solely on textual content and overlook the rich information conveyed through gestural cues. We believe that incorporating this additional non-verbal modality will enhance the language model’s capabilities of modeling spoken discourse. Several studies have explored combining modalities such as image and speech with language models (Tang et al., 2024; Tsimpoukelli et al., 2021; Team, 2024; Wang et al., 2022). However, incorporating gesture modality into language modeling remains underexplored. A recent work has taken steps in this direction by defining gesture tokens by encoding the spatial position of a speaker’s wrists using coarse grid tokenization on each 2D frame (Xu and Cheng, 2023). However, gestures often contain more fine-grained information involving coordinated movements of the wrist and shoulder joints.

In this work, we aim to integrate gesture motion sequences into language models and assess their impact on spoken discourse modeling in language models. Given 3D human motion sequences and corresponding spoken text, we first identify atomic units (tokens) of gestures. Inspired by recent works on co-speech gesture synthesis, we use a VQ-VAE-based architecture to learn these gesture tokens (Liu et al., 2024a). These gesture representations capture fine-grained motion details as they are trained to reconstruct 3D human motion. Once these tokens are obtained, we perform feature alignment to map the feature spaces of the gesture token embeddings to the input space of the language model. This alignment is achieved by training a projection via a joint masked gesture prediction and masked language modeling objective. Finally, we fine-tune the gesture-aligned language model for discourse-based tasks using low-rank adaptation (Hu et al., 2022).

To evaluate whether the proposed method of incorporating gestures into language models enhances spoken discourse modeling, we propose three text infilling tasks based on key markers in spoken discourse: (i) Discourse Connectives (Kendon, 1995; Quek et al., 2002), (ii) Quantifiers (Lorson et al., 2024), and (iii) Stance Markers (Roseano et al., 2016; Andries et al., 2023; Debras, 2015), which linguistic research has shown to frequently co-occur with gestural cues. The text infilling task (Wu et al., 2019), similar to the Cloze task in linguistics (Taylor, 1953), involves predicting missing words based on its surrounding context. In our approach, we mask markers from the three categories during fine-tuning and train the language model to predict them. To achieve this, we follow the connective generation pipeline from Liu and Strube (2023). Our experiments demonstrate that incorporating gestures alongside text improves the marker prediction performance across all three tasks. To further validate these findings, we conduct ablation studies and error analysis, offering deeper insights into the model’s performance.

The main contributions of this work are as follows:

1. We propose a framework that integrates gestural cues into language models and evaluate their potential to improve the language modeling of spoken discourse.
2. We design three linguistically grounded text infilling tasks to evaluate the performance of spoken discourse modeling.
3. We provide insights on when joint modeling of language and gestures can be helpful based on an in-depth error analysis.

2 Related Work

2.1 Evidence of Co-Speech Gestures from Linguistic Research

Co-speech gestures are hand and arm movements that temporally co-occur with speech (Kendon, 1972; McNeill, 1992; Shattuck-Hufnagel and Prieto, 2019; Ebert, 2024). These gestures can often be broadly classified into two types: rhythmic beat gestures, which are driven by prosody, and semantic gestures, which are context-driven and carry the underlying meaning of the linguistic content (McNeill, 1992). Semantic gestures, in particular, have been shown to frequently co-occur with certain

types of spoken discourse markers like discourse connectives, quantifiers and stance markers.

Discourse connectives, such as *however* or *for example*, are used by speakers to signal relationships between different parts of a message. They have been shown to significantly influence gestural patterns associated with lexical markers and across the broader discourse context (McNeill et al., 2014). Gesture features like hand shape or its orientation, help differentiate topics in discourse (Laparle, 2021; Laparle et al., 2024; Calbris, 2011; Hinnell, 2019). For example, raised index fingers are commonly used to express exceptions or concessions (usually signified in language by the words *however* or *but*) (Bressem and Müller, 2014; Kendon, 2004; Inbar, 2022).

Stance markers are generally used to express a speaker’s feelings, attitudes, perspectives, or positions (Barbara et al., 2024; Alghazo et al., 2021; Schneider and Barron, 2014). Stance can be categorized based on its function in communication. One such category is epistemic stance, which reflects the speaker’s degree of confidence or commitment toward a proposition. Within this category, hedges indicate a reduced level of certainty or commitment, often using words such as *probably* or *may* (Wei et al., 2021; Deng and He, 2023; Liu and Tseng, 2021), whereas boosters strengthen commitment and certainty through expressions like *must*, *really*, and *certainly* (Deng and He, 2023; Liu and Tseng, 2021). Multiple studies have found that speakers frequently use gestures to convey epistemic stance (Andries et al., 2023), such as palm-up gestures to indicate high certainty (Marrese et al., 2021) and shoulder shrugs for low certainty (Roseano et al., 2016; Borràs-Comes et al., 2011; Debras, 2015).

Quantifiers are generally used in spoken discourse to indicate an amount or degree (e.g., *many*, *some*, *all*) and can also express exact numerals (Feiman and Snedeker, 2016). Lorson et al. (2024) observed that gestural cues help interpret vague quantities like *several*. Their findings suggest that people associate larger horizontal gestures with higher and smaller gestures with lower quantities, indicating a spatio-numerical relationship.

The findings discussed above highlight the co-occurrence of gestures with respect to three different types of spoken discourse markers. While it is probable that other spoken discourse elements also correlate with gestural cues, our research specifically concentrates on these three types of markers as the primary focus for constructing our text in-

filling tasks and evaluating the potential of joint modeling of language and gestures.

2.2 Gesture-Enhanced Language Modeling

Gestures and language form an integrated system in both production and comprehension, supported by cognitive and psycholinguistic research. McNeill’s (McNeill, 1992, 2005) growth point theory proposes that gestures and speech originate from a shared cognitive source called a “growth point,” which represents the initial unit of thought during language formulation. From this point, imagery (gesture) and linguistic structure (speech) emerge together, shaping meaning in an interdependent way. This suggests that gestures are active contributors to discourse and conceptualization, not merely supportive tools. In comprehension, studies show that gestures and speech function as a unified communicative system. When gestures conflict with speech, listeners experience more processing errors and delays, highlighting the close coupling of gesture with linguistic understanding (Kelly et al., 2015). These insights highlight the importance of integrating gestures and language to better understand communication.

Various works have explored the integration of multiple modalities, such as video, images, and speech, into language models (Tsimpoukelli et al., 2021; Li et al., 2023; Tang et al., 2024; Zhang et al., 2023; Hassid et al., 2024; Merullo et al., 2022; Liu et al., 2024b). However, incorporating human motion sequences, which captures gestures, into language modeling is still underexplored. As one of the first works in this direction, Abzaliev et al. (2022) proposed to align gestures with language through CLIP-style embeddings and utilize binary classification to predict the presence of specific word categories like discourse markers.

More recently, there have been efforts to tokenize gestures in a manner similar to language tokens in order to utilize them with transformer-based language models. Xu and Cheng (2023) proposed a method to tokenize gestures using coarse grid-based locations of the speaker’s hands, employing these tokens for language-modeling tasks. Similarly, Xu et al. (2024) explores learning gesture tokens through VQ-VAE (Van Den Oord et al., 2017) and combining it with LLMs to perform a ‘gesture-translation task’ for co-speech gesture synthesis. These approaches overlook the rich motion information contained in co-speech gestures. The reliance on coarse grid locations in the former lim-

its its ability to capture language-dependent motion cues like quick hand flicks, while the latter treats gesture tokens as language tokens, and omits motion details contained in the codebook embedding.

To incorporate gesture information into language models, we utilize a gesture encoding process developed for gesture synthesis. Synthesis methods either represent gestures through continuous latent representations (Mughal et al., 2024, 2025) or tokenized embeddings (Liu et al., 2024a; Ao et al., 2023; Xu et al., 2024). The tokenized gesture representations are learned through VQ-VAE based training strategy to capture motion information in the resulting embeddings. These gesture tokens are suitable to integrate with transformer-based language models and our approach demonstrates the effectiveness of utilizing this information-rich representation of gestures with the language model to improve spoken discourse modeling.

3 Approach

Figure 1 illustrates our approach, where the first step involves building the Gesture Tokenizer. This is followed by the feature alignment stage, where the codebook embeddings are aligned with the input embedding space of the language model. The aligned embeddings are subsequently used for fine-tuning on spoken discourse marker prediction tasks.

3.1 Gesture Tokenization

Speech-aligned body motion cues provide an information dense signal for downstream language modeling. Therefore, we propose to encode this signal into our tokenized gesture representation. The gesture sequence \mathbf{x} consists of N frames, with each frame represented by the rotations of J upper body joints. We convert \mathbf{x} to a tokenized representation \mathbf{z}^q through the VQ-VAE framework (Van Den Oord et al., 2017). Given \mathbf{x} , we divide the gesture sequence into M chunks and encode each using time-aware transformer encoders (Mughal et al., 2024), resulting in a sequence of latent representations $\mathbf{z} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_M\}$. Then, in the quantization step, we quantize each latent $\mathbf{z}_i \in \mathbb{R}^d$ to \mathbf{z}_i^q using a codebook of size K , which results in the tokenized latent sequence $\mathbf{z}^q = \{\mathbf{z}_1^q, \mathbf{z}_2^q, \dots, \mathbf{z}_M^q\}$. Finally, we utilize a single transformer decoder to reconstruct the gesture sequence $\hat{\mathbf{x}}$ from the tokenized latent sequence \mathbf{z}^q . This pipeline is trained on motion reconstruction task where we use stan-

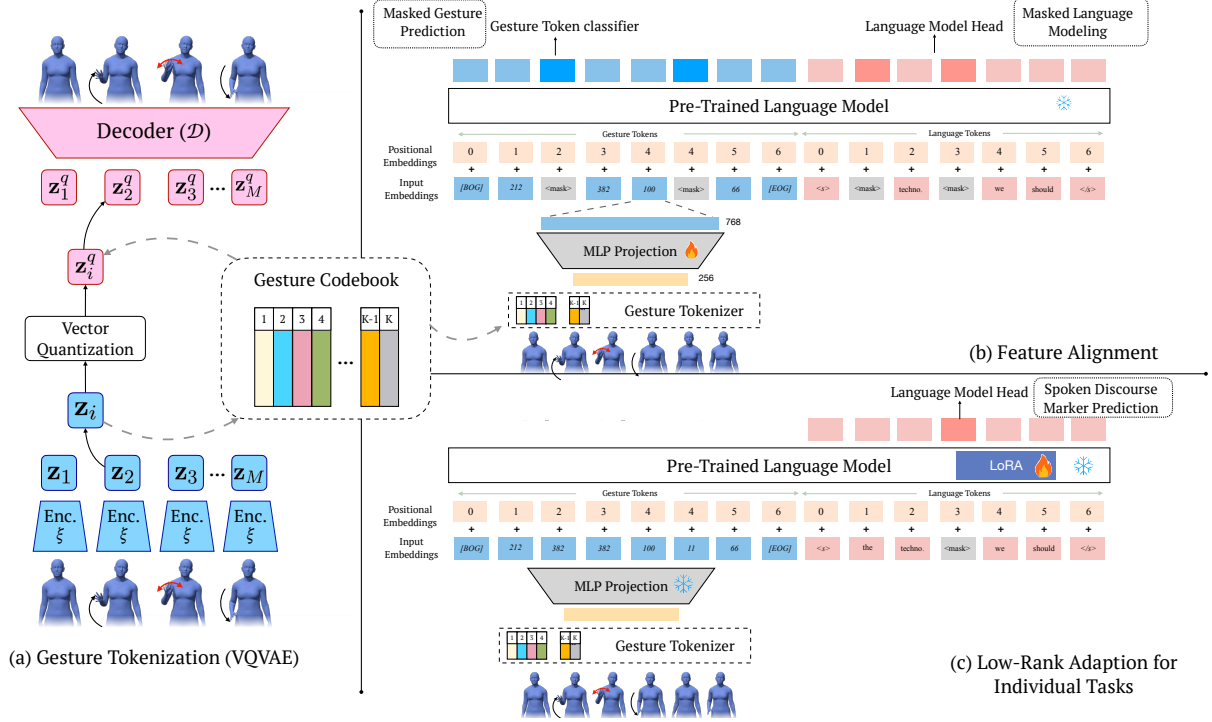


Figure 1: Overall framework of our approach for integrating gestures into language models.

standard VQ-VAE losses during training (see Appendix A.2 for details). The resulting Gesture Tokenizer provides an encoding pipeline for the motion cues which can enhance spoken discourse modeling in language models.

3.2 Feature Alignment

Existing multimodal approaches use various modality fusion architectures, such as cross-attention layers for modality interaction (Li et al., 2023; Alayrac et al., 2022), or unified transformers that tokenize non-textual data (e.g., visual, auditory) to integrate with transformer-based language models (Liu et al., 2024b; Tsimpoukelli et al., 2021; Tang et al., 2024; Wang et al., 2024). Our work builds on the latter, with the feature alignment stage aiming to map gesture embeddings into the input embedding space of language models.

For feature alignment, we first obtain paired text and gesture data. For each spoken sentence, we extract the corresponding gesture token embeddings by processing the temporally co-occurring 3D human motion sequence via the gesture tokenizer. We then project the gesture token embeddings using an MLP projector to map them to the input space of the language model. The projected gesture embeddings are concatenated with the input text embeddings and fed into the pre-trained language model, as shown in Figure 1b. Importantly, only the MLP

projector’s parameters are updated during training, while the language model and all other components remain frozen.

For training, we randomly mask 30% of both gesture and text tokens and use a joint objective that combines Masked Gesture Prediction (\mathcal{L}_{MGP}) and Masked Language Modeling (\mathcal{L}_{MLM}). For MGP, we train a K -class gesture token classifier via Cross-Entropy loss to predict the correct gesture codebook index for masked gesture tokens, similar to MLM. The final loss for feature alignment (\mathcal{L}_{FA}) is given by, $\mathcal{L}_{FA} = \mathcal{L}_{MGP} + \mathcal{L}_{MLM}$ where

$$\mathcal{L}_{MGP} = -\frac{1}{|M_g|} \sum_{j \in M_g} \log \frac{\exp(x_{y_j})}{\sum_{k \in K} \exp(x_k)}$$

$$\mathcal{L}_{MLM} = -\frac{1}{|M_t|} \sum_{i \in M_t} \log \frac{\exp(x_{y_i})}{\sum_{v \in V} \exp(x_v)}$$

Here M_g and M_t denotes the sets of masked gesture and text tokens, respectively. The vocabulary size of the language model is represented by V and K denotes the gesture codebook size. The true text and gesture tokens are given by y_i and y_j respectively and x is the logits.

To ensure temporal alignment between gestures and spoken words, we assign positional embeddings to gesture tokens based on their corresponding text tokens, so that each gesture token’s po-

sition matches that of the text token it co-occurs with (as illustrated with the case of token *we* in Figure 1). More details regarding computing the gesture positional embedding are provided in the Appendix A.3.

3.3 Fine-tuning for Spoken Discourse Marker Prediction

We fine-tune the gesture-aligned pre-trained language model on text infilling tasks targeting three spoken discourse markers: discourse connectives, quantifiers, and stance markers. Details on the label construction for each task are provided in the Evaluation Section 4.2. For each marker prediction task we obtain a discrete set of n markers: $L_{task} = \{l_1, l_2, \dots, l_n\}$ where L_{task} represents the subset of the vocabulary relevant to a specific category of markers. We follow the discourse connective generation pipeline outlined in (Liu and Strube, 2023). For this task, the input consists of a sequence of text tokens in the following format: $T = \langle s \rangle \mathbf{t}_1 \langle \text{mask} \rangle \mathbf{t}_2 \langle /s \rangle$, where the target marker in the sentence is replaced by $\langle \text{mask} \rangle$ and \mathbf{t}_1 and \mathbf{t}_2 denote the tokens before and after it. The embedding $h_{\langle \text{mask} \rangle}^L$ obtained after L layers of the transformer model is passed through the Language Modeling Head, similar to MLM. We use the probabilities obtained across the subset vocabulary L_{task} pertaining to the task to perform marker prediction. For fine-tuning we use cross-entropy loss and apply low-rank adaptation (Hu et al., 2022). As shown in Figure 1c, we freeze all model components except the adapter layers in the language model during training.

4 Evaluation

4.1 Dataset

We train and evaluate our method on the BEAT2 dataset (Liu et al., 2024a), which contains 60 hours of monologue gesture recordings spanning 25 speakers. We use the train/val/test split from the BEAT2 dataset for training the gesture tokenization stage and language modeling stages. We use Whisper (Radford et al., 2023) to obtain the text transcriptions from speech utterances in BEAT2.

4.2 Tasks

To understand whether incorporating gestures helps in spoken discourse modeling in language models, we construct three text infilling tasks focusing on three lexical markers: discourse connectives, quan-

tifiers, and stance markers. To identify the subset of discourse connectives, we process the BEAT2 dataset transcripts using discopy (Knaebel, 2021), which provides ground truth labels for explicit discourse connectives¹. In contrast, quantifiers and stance markers are less ambiguous, so we rely on labels sourced from previous studies for both categories (quantifiers: Madusanka et al., 2023; Feiman and Snedeker, 2016; Lorson et al., 2024; Barwise and Cooper, 1981; and stance markers: Barbara et al., 2024; Alghazo et al., 2021; Schneider and Barron, 2014). Additionally, only markers occurring more than 30 times in the BEAT2 dataset are considered for this study, following the approach of (Liu and Strube, 2023). The labels used in this study are as follows:

1. Discourse Connectives: *after, also, although, and, as, because, but, for example, however, if, if then, or, since, so, then, though, when, while.*
2. Quantifier: *all, each, enough, entire, few, little, less, many, more, most, much, no, one, some, three, two, whole.*
3. Stance markers: *actually, almost, amazing, especially, extremely, happy, important, may, maybe, might, must, probably, really, very.*

The training distribution of these labels are shown below in Figure 2.

4.3 Implementation Details

Our approach consists of three main steps: gesture tokenization, feature alignment, and low-rank adaptation. To represent the body motion during tokenization, we utilize 6D rotation representation (Zhou et al., 2019) for each joint in the upper body, which contains $J = 13$ joints. For the encoding-decoding pipeline, we set the number of motion chunks to $M = 8$ per gesture sequence, with each sequence spanning $N = 32$ at 15 fps. The codebook size $K = 512$ with embedding length $d = 256$ for the gesture tokenizer. We use the pre-trained RoBERTa-base model (Liu, 2019) as our language model, given its state-of-the-art performance on discourse-based classification tasks (Costa and Kosseim, 2024; Zhao et al., 2023; Zhou et al., 2022). To construct the paired text-gesture data, we introduce special tokens similar to those

¹Reported F1 score of 95.6 for explicit connective identification in Section 23 of PDTB 2.0




Figure 2: Frequency distribution of markers across three tasks: Discourse Connectives, Quantifiers, and Stance Markers in the BEAT2 training data.

used in language models, such as $\langle s \rangle$ and $\langle /s \rangle$, to mark gesture sequence boundaries. Specifically, we reserve unused codebook indices for [BOG] (beginning of gesture) and [EOG] (end of gesture), as illustrated in Figure 1. For feature alignment, we train an MLP projector with two dense layers and GeLU activation. During fine-tuning, we employ LoRA adapters ($r = 128$, $\alpha = 256$) for each task while keeping all other model components frozen, as shown in Figure 1c. All reported results in Tables 1, 2, and 3 are averaged over five random seed runs, with standard deviations included. Further details on hyperparameter choices and training resources are provided in the Appendix A.1.

5 Results

5.1 Comparison with Existing Approaches

This section addresses two key questions. *First*, how do models that incorporate gesture information compare with Text-only baseline model across the three tasks. *Second*, we evaluate how our approach to integrate gestures with language models compares to prior works that also model language and gestures jointly. We refer to our modeling pipeline as GestureLM in the tables.

We compare our approach with Xu and Cheng (2023), who introduce a grid-based tokenization scheme and a Mixed Modal architecture to integrate gestures with language models. To adapt this approach to our dataset, we apply a 3D-to-2D projection to obtain grid-based tokens and evaluate their modeling pipeline on our tasks. To enable a more direct comparison and better understand the impact of different gesture tokenization strategies, we also compare with their tokenization scheme integrated into our GestureLM pipeline. Additionally, we compare our approach to GesTran (Xu et al., 2024) in how they encode gestures. They augment the language model’s vocabulary with VQ-VAE codebook indices for co-speech gesture generation, however, they do not leverage the learned motion

information in the codebook embeddings. In contrast, our approach utilizes gesture embeddings, which inherently capture motion information, and maps them into the language space. More details regarding implementation are provided in the Appendix A.1.

Overall, the results indicate that incorporating gestures alongside text enhances the model’s performance in spoken discourse modeling. This improvement is particularly pronounced in terms of the F1 score, with an average gain of 4.8% across all three tasks. The results are statistically significant for discourse tasks ($p = 0.01$) and quantifiers ($p = 0.03$). The stance markers task, however, showed a high standard deviation, leading to a p-value of 0.16. After running additional experiments with three more random seeds, the p-value for stance markers decreased to 0.08, suggesting a possible trend.

We also find that gestures play an important role in the prediction of markers that are less frequently used and located in the long tail of the distribution as seen in Figure 3. This is especially relevant and important to capture given that the rare words often convey more specific meaning than frequent ones like *and*. A detailed per-class performance analysis is presented in the Error Analysis section 5.4 to provide further insights. We also find that using learned gesture embeddings through VQ-VAE-based tokenization yields better performance compared to existing gesture encoding approaches.

5.2 Validating the Impact of Gestures via Adversarial Evaluation

Inspired by multimodal translation studies analyzing the necessity of visual input for multimodal translation (Wu et al., 2021), we conduct a similar analysis by replacing pre-trained gesture embeddings with adversarial inputs such as random vectors from a normal distribution simulating uninformative gesture representations. Additionally, we

	Discourse		Quantifier		Stance	
	Acc (%)	F1	Acc (%)	F1	Acc (%)	F1
Text-only baseline	60.4 \pm 2.0	47.5 \pm 2.0	69.4 \pm 3.7	65.2 \pm 2.9	50.6 \pm 4.8	46.5 \pm 6.9
Mixed Modal (Xu and Cheng, 2023)	34.8 \pm 1.3	17.4 \pm 0.5	31.7 \pm 0.3	28.2 \pm 0.6	33.4 \pm 2.4	24.0 \pm 3.2
GestureLM with grid-based tokens (Xu and Cheng, 2023)*	55.3 \pm 13.6	41.3 \pm 21.7	70.5 \pm 3.0	65.4 \pm 4.0	47.9 \pm 2.5	44.5 \pm 3.6
GestureLM with codebook indices (Xu et al., 2024)*	54.4 \pm 13.0	39.2 \pm 20.3	68.4 \pm 2.4	63.9 \pm 2.8	46.5 \pm 6.5	41.7 \pm 5.1
GestureLM (Ours)	61.2 \pm 1.5	51.1 \pm 1.7	74.8 \pm 2.6	70.4 \pm 3.1	52.8 \pm 1.7	52.2 \pm 4.3

Table 1: Comparison with Existing Approaches. The values in the table represent the averaged over five random seed runs, with standard deviations. *ablations to compare different tokenization schemes with our modeling pipeline.

	Discourse		Quantifier		Stance	
	Acc (%)	F1	Acc (%)	F1	Acc (%)	F1
Text-only baseline	60.4 \pm 2.0	47.5 \pm 2.0	69.4 \pm 3.7	65.2 \pm 2.9	50.6 \pm 4.8	46.5 \pm 6.9
GestureLM with random	58.6 \pm 4.4	48.1 \pm 3.6	73.2 \pm 4.5	68.4 \pm 2.8	47.8 \pm 3.1	48.6 \pm 6.3
GestureLM with only positional	48.9 \pm 16.1	29.3 \pm 24.3	50.2 \pm 33.2	41.9 \pm 37.1	34.9 \pm 15.0	26.7 \pm 22.8
GestureLM (Ours)	61.2 \pm 1.5	51.1 \pm 1.7	74.8 \pm 2.6	70.4 \pm 3.1	52.8 \pm 1.7	52.2 \pm 4.3

Table 2: Validating the impact of gesture embeddings through adversarial evaluation.

experiment with using only positional embeddings while setting gesture embeddings to zero. This helps determine whether the performance gains stem from gesture duration, which may correlate with word length rather than semantic content.

From Table 2, we observe that for all three tasks, using pre-trained gesture embeddings consistently improves F1 scores compared to random vectors or positional embeddings alone. These findings suggest that pre-trained gesture embeddings provide meaningful information rather than acting as a form of regularization. We believe that such adversarial evaluation is a crucial step for future research to ensure that the added modality genuinely contributes to language modeling.

5.3 Ablations

5.3.1 Masking % during feature alignment

We select the masking percentage in the feature alignment stage based on the lowest validation loss from the masked language modeling objective. A lower loss reflects better-aligned gesture embeddings, leading to improved language modeling. Table 4 shows the relationship between masking percentage and validation loss. As expected, the loss increases with higher masking percentages, but also below 30% because as the masking percentages become lower, the model does not effectively learn as the task becomes easier.

5.3.2 Model Components

We also ablate the design choices in our modeling pipeline to understand their effect on the task performance. In Table 3, we first compare the use

of absolute positional embeddings with relative ones and observe an improvement in model performance across all three tasks. We also test the model without the feature alignment stage, that is, without mapping the representation spaces, and find a significant drop in performance, highlighting the importance of this stage in effectively fusing the two modalities.

5.4 Error Analysis

Overall, we find that gesture incorporation enhances the performance of language models, especially for underrepresented markers in the training data. However, our error analysis also highlights certain challenges, especially in disambiguating confusable markers. To further investigate this, we analyze the class-wise task performance of the models using a relative confusion matrix (Pommé et al., 2022), shown in Figure 3. This matrix captures the difference in predictions between the two comparison models which in our case is GestureLM vs. Text-only baseline. In terms of interpretation, a stronger red intensity indicates more predictions by the Text-only model for a given cell, while a stronger blue intensity signifies more predictions by the GestureLM. Note that the rows are arranged in descending order of marker frequency in test samples, from top to bottom. Rows and columns with minimal difference are omitted for space considerations, and the complete confusion matrices are available in the Appendix A.5.

In the discourse connective prediction task, we observe improved performance particularly for temporal discourse relations such as *after* and *while*.

	Discourse		Quantifier		Stance	
	Acc (%)	F1	Acc (%)	F1	Acc (%)	F1
GestureLM	61.2 ± 1.5	51.1 ± 1.7	74.8 ± 2.6	70.4 ± 3.1	52.8 ± 1.7	52.2 ± 4.3
GestureLM w/o relative positional encoding	59.6 ± 2.3	47.6 ± 2.9	73.8 ± 6.7	69.1 ± 3.2	50.6 ± 3.9	50.8 ± 3.9
GestureLM w/o feature alignment	60.2 ± 4.1	46.7 ± 5.4	68.8 ± 4.6	65.1 ± 5.5	43.0 ± 34.7	34.7 ± 13.8

Table 3: Ablation w.r.t Model Components

Masking percentage	10%	30%	50%	80%
loss (\downarrow)	1.4	0.3	0.4	4.9

Table 4: Masking percentage vs validation loss from the feature alignment stage

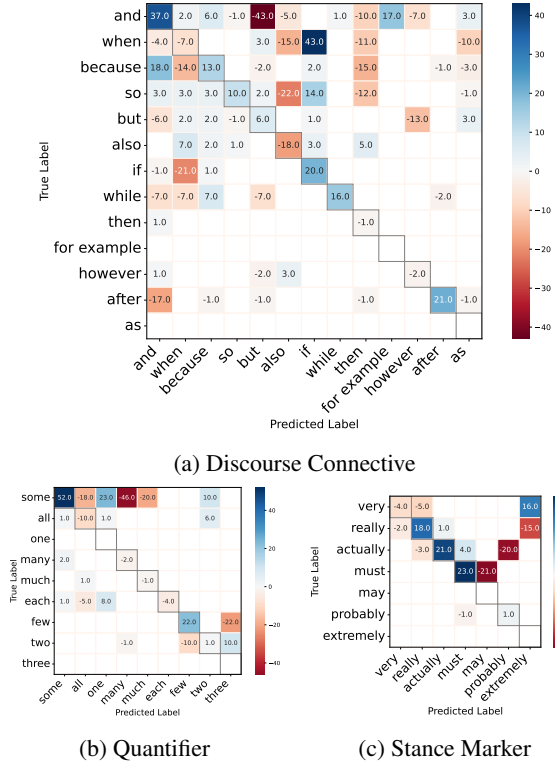


Figure 3: Relative Confusion Matrices comparing GestureLM and Text-only baseline. The matrix highlights differences in class-wise predictions, with red indicating more Text-only predictions and blue signifying more by GestureLM.

This improvement can be attributed to the presence of gestures associated with these connectives near their utterance as seen in the samples in the Figure 4a. These gestures likely help to disambiguate these connectives from more commonly used ones, such as *and* (refer to the training distribution in Figure 2), which the Text-only model often confuses temporal connectives with, as seen in Figure 3a (red cells in the *and* column). Linguistic research further supports that speakers use spatial gestures to express temporal concepts such as *after* and *be-*




Figure 4: Some samples where GestureLM performs better than Text-only model. We see that the semantic gestures co-occur with the spoken discourse markers, potentially leading to improved prediction performance. See the text for detailed description and more examples are shown in the Appendix A.4.

fore (Casasanto and Jasmin, 2012). Similarly, for contrastive connectives like *but*, gestures such as a raised hand during the utterance of *but* (shown in the Figure 4b) potentially helps GestureLM distinguish it from *and*.

Similarly, in the quantifier task, we see in Figure 3b that incorporating gestures helps with under-represented classes such as *few* and *some*. Particularly, in the case of *some*, while the Text-only models confuse it with higher represented classes such as *all* and *much*, the gesture models tend to confuse it with numbers such as *two* or *one*. One possible explanation is that similar wrist movements can be used to represent both types of quantifiers, as the specific finger joint movements that convey numerical information are not captured. Future work could incorporate data from finger joints to better distinguish between these classes.

In stance marker prediction, incorporating gestures helps disambiguate *must* with *may*, which Text-only models often confuse. *Must*, which conveys assertiveness belongs to epistemic stance which is found to co-occur with gestural cues (Andries et al., 2023; Marrese et al., 2021). For instance, in Figure 4c, we see the speaker performing a palm-down gesture which could help the gesture-enhanced models in better disambiguation. Additionally, we see in Figure 3c that both models tend to confuse the most commonly occurring class, *really*, with *very* and *extremely*. Unlike previous tasks, these mistakes are acceptable as the words are often interchangeable. More examples are in Appendix A.4

6 Conclusion

In this work, we presented a framework that integrates gestural cues into language models to enhance the modeling of spoken discourse. To evaluate our approach, we constructed linguistically grounded text infilling tasks focused on three types of spoken discourse markers: discourse connectives, quantifiers and stance markers. Our results demonstrate that incorporating gestural information improves the accuracy of spoken discourse marker prediction in language models. We view this work as a starting step toward multimodal spoken discourse modeling in language models. Future research should investigate whether gestures can also contribute in cases where meaning is not explicitly conveyed through language but is instead expressed solely through gesture (e.g., implicit coherence relations), which would require annotated data for evaluation.

7 Limitations

7.1 Models and Representation

First, our model captures only the upper-body joints up to the wrist, excluding finger joint movements. As a result, gestures relying on finger joints, such as distinguishing between *two* and *three* cannot be disambiguated. Future work will aim to incorporate these finer-grained features to improve the model’s ability to capture subtle gesture variations. Second, our approach is currently designed for encoder-only MLM-based models, as they are generally used for token classification. In future work, we aim to explore how gesture-enhanced models can be helpful in decoder-based models that are trained for next-token prediction. Lastly, to

represent body joints in motion, we use 6D rotation representation in this work. However, the choice of other representations, such as joint positions, acceleration etc., could vary the performance of spoken language modeling.

7.2 Data

The dataset used in this study consists of monologues from a single speaker in a fixed context. While it includes speakers from different demographic backgrounds, it lacks annotations regarding cultural or demographic attributes and does not encompass a range of social interaction settings. This limitation affects the generalizability of our findings, as gesture use can vary significantly across cultures, languages, and communicative contexts. Investigating these dimensions would require a large-scale dataset that includes demographic annotations and covers diverse social environments. Such an undertaking is beyond the scope of the present work but represents an important direction for future research. Additionally, many real-world scenarios lack 3D motion capture data and instead rely on 2D video, which may reduce gesture encoding accuracy and limit the effectiveness of integrating gestures into language models.

8 Ethical Statement

This work uses publicly available dataset and uses only the raw 3D joint data for gesture modeling, without utilizing personally identifiable information or video recordings of individuals. However, we acknowledge potential ethical concerns, as gesture data may unintentionally reveal personal attributes, such as gender or physical characteristics, which could pose privacy risks if misused for unintended inferences.

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft, Funder Id: <http://dx.doi.org/10.13039/501100001659>, SFB 1102: “Information Density and Linguistic Encoding”, project number 232722074. M. Hamza Mughal was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - GRK 2853/1 “Neuroexplicit Models of Language, Vision, and Action” - project number 471607914.

References

- Artem Abzaliev, Andrew Owens, and Rada Mihalcea. 2022. Towards understanding the relation between gestures and language. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5507–5520.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736.
- Sharif Alghazo, Mohd Nour Al Salem, Imran Alrashdan, and Ghaleb Rabab’ah. 2021. Grammatical devices of stance in written academic english. *Heliyon*, 7(11).
- Fien Andries, Katharina Meissl, Clarissa de Vries, Kurt Feyaerts, Bert Oben, Paul Sambre, Myriam Vermeerbergen, and Geert Brône. 2023. Multimodal stance-taking in interaction—a systematic literature review. *Frontiers in Communication*, 8:1187977.
- Tenglong Ao, Zeyi Zhang, and Libin Liu. 2023. Gesturediffuclip: Gesture diffusion model with clip latents. *ACM Transactions on Graphics*, 42(4):1–18.
- Siu Wing Yee Barbara, Muhammad Afzaal, and Hesah Saleh Aldayel. 2024. A corpus-based comparison of linguistic markers of stance and genre in the academic writing of novice and advanced engineering learners. *Humanities and Social Sciences Communications*, 11(1):1–10.
- Jon Barwise and Robin Cooper. 1981. Generalized quantifiers and natural language. In *Philosophy, language, and artificial intelligence: Resources for processing natural language*, pages 241–301. Springer.
- Joan Borràs-Comes, Paolo Roseano, Maria del Mar Vanrell, Aoju Chen, and Pilar Prieto. 2011. Perceiving uncertainty: facial gestures, intonation, and lexical choice. *Proceedings of GESPIN*.
- Jana Bressem and Cornelia Müller. 2014. A repertoire of German recurrent gestures with pragmatic functions. In Cornelia Müller, Alan Cienki, Ellen Fricke, Silva H. Ladewig, David McNeill, and Jana Bressem, editors, *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.
- G. Calbris. 2011. *Elements of Meaning in Gesture*. Gesture studies. John Benjamins Publishing Company.
- Daniel Casasanto and Kyle Jasmin. 2012. The hands of time: Temporal gestures in english speakers. *Cognitive Linguistics*, 23(4):643–674.
- Justine Cassell, Yukiko I Nakano, Timothy W Bickmore, Candace L Sidner, and Charles Rich. 2001. Non-verbal cues for discourse structure. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*, pages 114–123.
- Nelson Filipe Costa and Leila Kosseim. 2024. Exploring soft-label training for implicit discourse relation recognition. In *Proceedings of the 5th Workshop on Computational Approaches to Discourse (CODI 2024)*, pages 120–126.
- Camille Debras. 2015. Stance-taking functions of multimodal constructed dialogue during spoken interaction. In *Actes Du Colloque Gesture and Speech in Interaction 4*.
- Liming Deng and Ping He. 2023. “we may conclude that:” a corpus-based study of stance-taking in conclusion sections of ras across cultures and disciplines. *Frontiers in Psychology*, 14:1175144.
- Cornelia Ebert. 2024. Semantics of gesture. *Annual Review of Linguistics*, 10(1):169–189.
- Roman Feiman and Jesse Snedeker. 2016. The logic in language: How all quantifiers are alike, but each quantifier is different. *Cognitive psychology*, 87:29–52.
- Catherine S Fichten, Vicki Tagalakakis, Darlene Judd, John Wright, and Rhonda Amsel. 1992. Verbal and nonverbal communication cues in daily conversations and dating. *The Journal of Social Psychology*, 132(6):751–769.
- Susan Goldin-Meadow, Martha Wagner Alibali, and R Breckinridge Church. 1993. Transitions in concept acquisition: using the hand to read the mind. *Psychological review*, 100(2):279.
- Michael Hassid, Tal Remez, Tu Anh Nguyen, Itai Gat, Alexis Conneau, Felix Kreuk, Jade Copet, Alexandre Defossez, Gabriel Synnaeve, Emmanuel Dupoux, et al. 2024. Textually pretrained speech language models. *Advances in Neural Information Processing Systems*, 36.
- Jennifer Hinnell. 2019. The verbal-kinesic enactment of contrast in north american english. *The American Journal of Semiotics*, 35(1-2):55–92.
- Judith Holler, Heather Shovelton, and Geoffrey Beattie. 2009. Do iconic hand gestures really contribute to the communication of semantic information in a face-to-face context? *Journal of Nonverbal Behavior*, 33:73–88.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Anna Inbar. 2022. The raised index finger gesture in Hebrew multimodal interaction. *Gesture*, 21(2-3):264–295.
- Spencer Kelly, Meghan Healey, Asli Özyürek, and Judith Holler. 2015. The processing of speech, gesture, and action during language comprehension. *Psychonomic bulletin & review*, 22:517–523.

- Adam Kendon. 1972. Some relationships between body motion and speech. *Studies in dyadic communication*, 7(177):90.
- Adam Kendon. 1995. Gestures as illocutionary and discourse structure markers in southern italian conversation. *Journal of pragmatics*, 23(3):247–279.
- Adam Kendon. 2004. *Gesture: Visible action as utterance*. Cambridge University Press.
- Zohreh Khosrobeigi, Maria Koutsombogera, and Carl Vogel. 2022. Gesture and part-of-speech alignment in dialogues. In *Proceedings of the 26th Workshop on the Semantics and Pragmatics of Dialogue*, pages 172–182.
- Sotaro Kita, Martha W Alibali, and Mingyuan Chu. 2017. How do gestures influence thinking and speaking? the gesture-for-conceptualization hypothesis. *Psychological review*, 124(3):245.
- René Knaebel. 2021. *discopy: A neural system for shallow discourse parsing*. In *Proceedings of the 2nd Workshop on Computational Approaches to Discourse*, pages 128–133, Punta Cana, Dominican Republic and Online. Association for Computational Linguistics.
- Robert M Krauss, Yihsiu Chen, and Purnima Chawla. 1996. Nonverbal behavior and nonverbal communication: What do conversational hand gestures tell us? In *Advances in experimental social psychology*, volume 28, pages 389–450. Elsevier.
- Schuyler Laparle. 2021. *Tracking discourse topics in co-speech gesture*. In *International Conference on Human-Computer Interaction (HCII)*, page 233–249, Berlin, Heidelberg. Springer-Verlag.
- Schuyler Laparle, Gaëlle Ferré, and Merel CJ Scholman. 2024. More than one gesture but less than two? inter-stroke dependencies in form and meaning. In *International Conference on Human-Computer Interaction*, pages 245–264. Springer.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Chunhong Liu and Ming-Yu Tseng. 2021. Paradigmatic variation in hedging and boosting: A comparative study of discussions in narrative inquiry and grounded theory research. *English for Specific Purposes*, 61:1–16.
- Haiyang Liu, Zihao Zhu, Giorgio Becherini, Yichen Peng, Mingyang Su, You Zhou, Xuefei Zhe, Naoya Iwamoto, Bo Zheng, and Michael J Black. 2024a. Eimage: 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.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024b. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Wei Liu and Michael Strube. 2023. Annotation-inspired implicit discourse relation classification with auxiliary discourse connective generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15696–15712.
- Yinhan Liu. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 364.
- Alexandra Lorson, Vinicius Macuch-Silva, Christopher Hart, and Bodo Winter. 2024. Gesture size affects numerical estimates in quantifier comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Tharindu Madusanka, Iqra Zahid, Hao Li, Ian Pratt-Hartmann, and Riza Theresa Batista-Navarro. 2023. Not all quantifiers are equal: Probing transformer-based language models’ understanding of generalised quantifiers. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8680–8692.
- Olivia H Marrese, Chase Wesley Raymond, Barbara A Fox, Cecilia E Ford, and Megan Pielke. 2021. The grammar of obviousness: The palm-up gesture in argument sequences. *Frontiers in Communication*, 6:663067.
- David McNeill. 1992. Hand and mind: What gestures reveal about thought.
- David McNeill. 2005. Gesture and thought.
- David McNeill, Elena Levy, and Susan Duncan. 2014. *Gesture in discourse*. *The Handbook of Discourse Analysis*.
- Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. 2022. Linearly mapping from image to text space. In *The Eleventh International Conference on Learning Representations*.
- M. Hamza Mughal, Rishabh Dabral, Merel C. J. Scholman, Vera Demberg, and Christian Theobalt. 2025. Retrieving semantics from the deep: an rag solution for gesture synthesis. In *Computer Vision and Pattern Recognition (CVPR)*.
- Muhammad Hamza Mughal, Rishabh Dabral, Ikhsanul Habibie, Lucia Donatelli, Marc Habermann, and Christian Theobalt. 2024. ConvoFusion: Multi-modal conversational diffusion for co-speech gesture synthesis. In *Computer Vision and Pattern Recognition (CVPR)*.
- Luc-Etienne Pommé, Romain Bourqui, Romain Giot, and David Auber. 2022. Relative confusion matrix: efficient comparison of decision models. In *2022 26th International Conference Information Visualisation (IV)*, pages 98–103. IEEE.

- Francis Quek, David McNeill, Robert Bryll, Susan Duncan, Xin-Feng Ma, Cemil Kirbas, Karl E McCulloch, and Rashid Ansari. 2002. Multimodal human discourse: gesture and speech. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 9(3):171–193.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR.
- Ashwin Ram, Varsha Suresh, Artin Saberpour Abadian, Vera Demberg, and Jürgen Steimle. 2025. Gesturecoach: Rehearsing for engaging talks with llm-driven gesture recommendations. *arXiv preprint arXiv:2504.10706*.
- Paolo Roseano, Montserrat González, Joan Borràs-Comes, and Pilar Prieto. 2016. Communicating epistemic stance: how speech and gesture patterns reflect epistemicity and evidentiality. *Discourse Processes*, 53(3):135–174.
- Klaus P Schneider and Anne Barron. 2014. *Pragmatics of discourse*. De Gruyter Mouton.
- Stefanie Shattuck-Hufnagel and Pilar Prieto. 2019. Dimensionalizing co-speech gestures. In *Proceedings of the international Congress of phonetic sciences*, volume 5, pages 1490–1494.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, MA Zejun, and Chao Zhang. 2024. Salmonn: Towards generic hearing abilities for large language models. In *The Twelfth International Conference on Learning Representations*.
- Wilson L Taylor. 1953. “cloze procedure”: A new tool for measuring readability. *Journalism quarterly*, 30(4):415–433.
- Chameleon Team. 2024. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*.
- Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. 2021. Multimodal few-shot learning with frozen language models. *Advances in Neural Information Processing Systems*, 34:200–212.
- Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Haoran Wang, Mohit Mendiratta, Christian Theobalt, and Adam Kortylewski. 2024. Facegpt: Self-supervised learning to chat about 3d human faces. *arXiv preprint arXiv:2406.07163*.
- Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. 2022. Simvlm: Simple visual language model pretraining with weak supervision. In *International Conference on Learning Representations*.
- Yipu Wei, Jacqueline Evers-Vermeul, Ted M Sanders, and Willem M Mak. 2021. The role of connectives and stance markers in the processing of subjective causal relations. *Discourse Processes*, 58(8):766–786.
- Xing Wu, Tao Zhang, Liangjun Zang, Jizhong Han, and Songlin Hu. 2019. [Mask and infill: Applying masked language model for sentiment transfer](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5271–5277. International Joint Conferences on Artificial Intelligence Organization.
- Zhiyong Wu, Lingpeng Kong, Wei Bi, Xiang Li, and Ben Kao. 2021. Good for misconceived reasons: An empirical revisiting on the need for visual context in multimodal machine translation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6153–6166.
- Chenghao Xu, Guangtao Lyu, Jiexi Yan, Muli Yang, and Cheng Deng. 2024. Llm knows body language, too: Translating speech voices into human gestures. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5004–5013.
- Yang Xu and Yang Cheng. 2023. Spontaneous gestures encoded by hand positions improve language models: An information-theoretic motivated study. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9409–9424.
- Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. 2023. Speechgpt: Empowering large language models with intrinsic cross-modal conversational abilities. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15757–15773.
- Haodong Zhao, Ruifang He, Mengnan Xiao, and Jing Xu. 2023. Infusing hierarchical guidance into prompt tuning: A parameter-efficient framework for multi-level implicit discourse relation recognition. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6477–6492.
- Hao Zhou, Man Lan, Yuanbin Wu, Yuefeng Chen, and Meirong Ma. 2022. Prompt-based connective prediction method for fine-grained implicit discourse relation recognition. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3848–3858.
- Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. 2019. 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.

A Appendix

A.1 Implementation Details (Cont.)

The three main steps in our approach are gesture tokenization, followed by feature alignment, and then low-rank adaptation. In the gesture tokenizer, the transformer encoder and decoder consist of 8 layers, with each layer containing 4 attention heads. The whole pipeline is trained for 57 epochs with AdamW optimizer with a learning rate of $3e-5$. For the feature alignment step, we have a total of 1.8M trainable parameters (total = 126.6M). We use the AdamW optimizer with a learning rate of $1e-3$, a cosine scheduler with a warmup ratio of 0.03, and train for 20 epochs. Early stopping is based on validation loss, with a batch size of 32. For the fine-tuning stage, we have a total of 4.7M trainable parameters (LoRA), for both Text-only model and GestureLM. For all tasks and settings, we train with a learning rate of $1e-3$ with a cosine scheduler with a warmup ratio of 0.03, and 10 epochs. The best model is chosen based on validation F1, with a batch size of 16. We use the AdamW optimizer with a weight decay of $1e-3$ for this step. The hyperparameters were obtained via manual tuning.

For the comparison approaches, we adapt the Mixed Modal method from (Xu and Cheng, 2023) for the BEAT2 dataset. Specifically, we convert the tokens from 3D to 2D, and to ensure a fairer comparison, we use pre-trained RoBERTa text embeddings instead of training from scratch. We adopt the concatenation fusion approach from their paper, as it was shown to yield the lowest validation loss. Since the main contribution of their paper is the tokenization scheme, we compare their tokenization method with ours (GestureLM with grid-based tokens, as shown in Table 1). For this comparison, we create an Embedding Layer with grid token indices and train these embeddings during our feature alignment stage, as they are randomly initialized. A similar approach is used when incorporating the codebook indices similar to (Xu et al., 2024). For training, we use the same parameters mentioned in the paragraph above.

We run all our experiments on a single NVIDIA Quadro RTX 8000 (48 GB GDDR6) GPU. For implementation we use PyTorch², Roberta weights from HuggingFace transformers³ and AdapterHub for LoRa implementation⁴

²<https://pytorch.org/>

³<https://huggingface.co/>

⁴<https://docs.adapterhub.ml/index.html>

A.2 Losses from the gesture tokenization

Our tokenization pipeline is trained using VQ-VAE for motion reconstruction. Therefore, we apply standard MSE loss on the decoder output $\hat{\mathbf{x}}$ in 6D rotational representation and also apply the reconstruction losses by converting $\hat{\mathbf{x}}$ to axis-angle representation and joint positions. We also apply Geodesic Loss on rotation matrices. Along with that, we apply MSE losses on velocity and acceleration of the motion. The codebook loss includes the VQ objective and the commitment term (Van Den Oord et al., 2017).

$$\mathcal{L}_{rec} = \mathcal{L}_{6D} + \mathcal{L}_{axis-angle} + \mathcal{L}_{joint-pos} + \mathcal{L}_{rot} \quad (1)$$

$$\mathcal{L}_{codebook} = ||\text{sg}[\mathbf{z}] - \mathbf{z}^q||^2 + \beta ||\mathbf{z} - \text{sg}[\mathbf{z}^q]||^2 \quad (2)$$

$$\mathcal{L}_{VQ-VAE} = \mathcal{L}_{rec}(\mathbf{x}, \hat{\mathbf{x}}) + \mathcal{L}_{codebook} + \mathcal{L}_{vel}(\mathbf{x}', \hat{\mathbf{x}}') + \mathcal{L}_{acc}(\mathbf{x}'', \hat{\mathbf{x}}'') \quad (3)$$

Here, sg is the stop gradient operation and the weight of commitment term $\beta = 0.25$.

A.3 Obtaining Gesture Positional Embedding

To obtain gesture positions, we first determine the timestamp of each text token using Whisper. Next, we extract the gesture tokens that fall within this time window. For example, if the word *really* is spoken over 0.5 seconds and each gesture token corresponds to 4 frames at a frame rate of 15 fps, we get approximately 2 (1.875) gesture tokens for this word. We then assign the text token’s position to all the extracted gesture tokens associated with it. Since Whisper only provides time durations for whole words but not individual tokens, we distribute the total duration of multi-token words proportionally based on the number of characters in each token.

A.4 More examples

Figure 5 presents additional examples of gesture sequences accompanied by a detailed explanation in the caption.

A.5 Confusion Matrix

Figure 6a, 6b and 6c illustrates the full confusion matrices for discourse connectives, quantifiers and stance markers respectively.



..meet together — **after** — not seeing each other..

(a) Another example of spatial gesture occurring near the temporal discourse connective *after*. The Text-only model predicts *but*, while GestureLM **correctly** predicts *after*.




..sometimes I feel — I had **two** homes — that made me..

(b) In this example, the speaker lifts the left hand to depict *two* using fingers. The Text-only model predicts *many*, while GestureLM **correctly** predicts *two*.



..government — took **some** — actions to control the networks..

(c) Similar to 5b, the speaker flicks the left hand to depict *some*. The Text-only model predicts *many*, while GestureLM **incorrectly** predicts *two*. Incorporating finger joints can help GestureLM for better disambiguation in both examples.



..about me or backstabbing me -will **actually**, - make me angry..

(d) Gesture occurring near the epistemic stance marker *actually*. The Text-only model predicts *probably*, while GestureLM **incorrectly** predicts *really*. However, in the case of GestureLM this mistake is acceptable as opposed to the Text-model, this further signifies the importance of incorporating non-verbal cues.



..we were both — **very** surprised — to see one another..

(e) Another example of gesture occurring near the epistemic stance marker *very*. Both model **incorrectly** predicts *really* in this case. This gesture is similar in motion to 5d



.. and the users — **must** — be registered..

(f) Another example of palm down gesture happening at stance marker *must* similar to one displayed in the main text.

Figure 5: Examples

