

GestureCoach: Rehearsing for Engaging Talks with LLM-Driven Gesture Recommendations

Ashwin Ram
Saarland University, Saarland
Informatics Campus
Saarbrücken, Germany
ashwinram10@gmail.com

Varsha Suresh
Saarland University, Saarland
Informatics Campus
Saarbrücken, Germany
vsuresh@lst.uni-saarland.de

Artin Saberpour Abadian
Saarland University, Saarland
Informatics Campus
Saarbrücken, Germany
saberpour@uni-saarland.de

Vera Demberg
Saarland University, Saarland
Informatics Campus
Saarbrücken, Germany
vera@coli.uni-saarland.de

Jürgen Steimle
Saarland University, Saarland
Informatics Campus
Saarbrücken, Germany
steimle@cs.uni-saarland.de

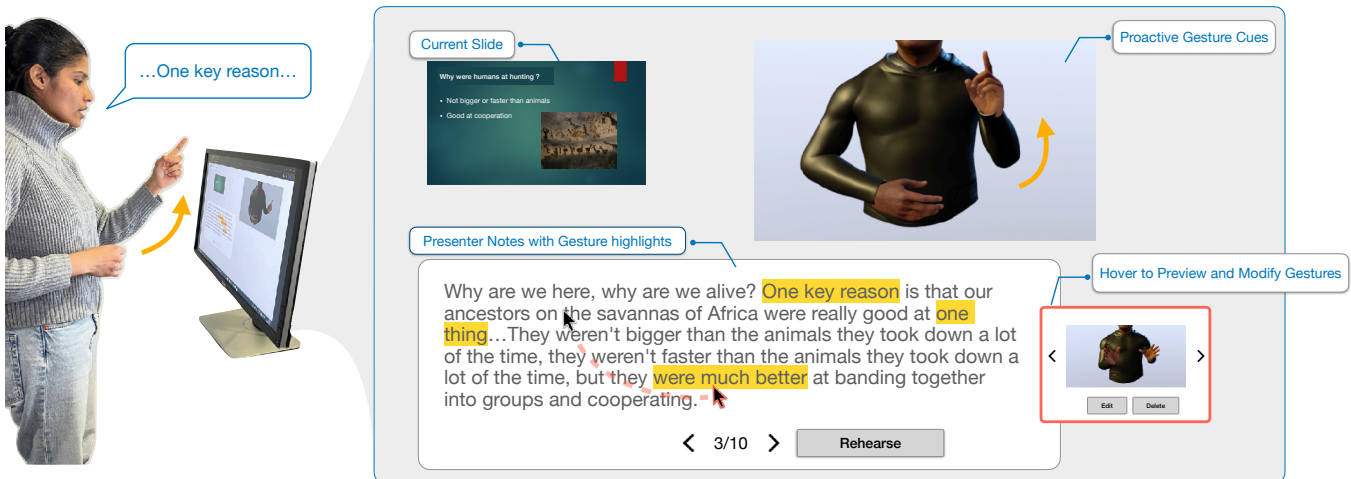


Figure 1: *GestureCoach* guides speakers to perform gestures while rehearsing their talk. A gesture recommendation model predicts text segments in the presenter notes that should be emphasized with gestures and retrieves relevant semantic gestures for each segment. During rehearsal, the system highlights the segments and proactively cues a video clip of the gesture by tracking users' speech, allowing them to integrate the gesture smoothly into their talk. Hovering over a segment allows users to preview and modify the associated gesture with alternate suggestions from the model.

Abstract

This paper introduces *GestureCoach*, a system designed to help speakers deliver more engaging talks by guiding them to gesture effectively during rehearsal. *GestureCoach* combines an LLM-driven gesture recommendation model with a rehearsal interface that proactively cues speakers to gesture appropriately. Trained on experts' gesturing patterns from TED talks, the model consists of two modules: an emphasis proposal module, which predicts *when* to gesture by identifying gesture-worthy text segments in the presenter

notes, and a gesture identification module, which determines *what* gesture to use by retrieving semantically appropriate gestures from a curated gesture database. Results of a model performance evaluation and user study (N=30) show that the emphasis proposal module outperforms off-the-shelf LLMs in identifying suitable gesture regions, and that participants rated the majority of these predicted regions and their corresponding gestures as highly appropriate. A subsequent user study (N=10) showed that rehearsing with *GestureCoach* encouraged speakers to gesture and significantly increased gesture diversity, resulting in more engaging talks. We conclude with design implications for future AI-driven rehearsal systems.



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

UIST '25, Busan, Republic of Korea

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2037-6/25/09

<https://doi.org/10.1145/3746059.3747705>

CCS Concepts

• Human-centered computing → Interactive systems and tools.

Keywords

Gesture, LLM, Generative AI, Presentation Talks

ACM Reference Format:

Ashwin Ram, Varsha Suresh, Artin Saberpour Abadian, Vera Demberg, and Jürgen Steimle. 2025. GestureCoach: Rehearsing for Engaging Talks with LLM-Driven Gesture Recommendations. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*, September 28–October 01, 2025, Busan, Republic of Korea. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3746059.3747705>

1 Introduction

Delivering a “great” talk requires more than just good content; it needs an engaging delivery with the right blend of non-verbal cues, such as gestures [15, 25, 37]. In particular, semantic gestures that complement the spoken content [36], such as tracing the outline of a circle in the air while saying ‘a round table’ or moving hands upward while saying ‘rising prices’, play a critical role in enhancing the dynamism of the speech [6, 23, 45] and facilitating comprehension [17, 23, 35].

However, speakers, especially novices, often struggle with using semantic gestures in their talks [49, 60]. They may either completely neglect to use gestures or rely heavily on beat gestures [36]—simple, repetitive hand movements that do not enhance the meaning of the content. While some existing training systems [3, 10, 47, 54] track users’ motion and provide generic feedback to use more or fewer gestures, they do not offer specific guidance on what semantic gesture to use or when to use it for effective emphasis. Prior research has shown that verbal qualities, such as voice modulation, can be improved with rehearsal and then effectively applied during live talks [56]. Inspired by this work, we aim to develop a rehearsal system specifically designed to improve semantic gesture usage during talks.

Designing a rehearsal system that can give targeted guidance for semantic gestures, however, presents various technical and design challenges. The key difficulty lies in developing a gesture recommendation model that predicts appropriate text segments in the presenter notes that should be emphasized with gestures (henceforth referred to as gesture regions) and retrieves suitable semantic gestures for effective emphasis. Predicting whether a text segment is a gesture region depends on several factors, including its broader context within the text, the topic of the talk, and implicit knowledge expert speakers use to engage an audience. Furthermore, the suggested gestures must extend beyond the limited gesture vocabulary of novices, offering access to the more refined and expressive gestures used by experts. From a design perspective, these gesture recommendations must be delivered proactively during rehearsal, allowing users to practice and seamlessly incorporate them into their talk.

As a step towards tackling these challenges, we propose *GestureCoach*, the first system capable of providing targeted guidance to perform gestures during rehearsal. *GestureCoach* combines a gesture recommendation system with a rehearsal interface that delivers real-time gesture cues, helping speakers naturally integrate gestures while speaking. The gesture recommendation system utilizes recent advances in AI chaining, by combining two specialized

Large Language Model (LLM) modules, each targeting a specific aspect of gesture-based emphasis. First, the *emphasis proposal* module takes presenter notes as input and determines appropriate gesture regions. For this, we leverage the foundational knowledge of LLMs and further fine-tune it with annotated data from 10 TED talk videos of expert speakers. Second, the *gesture identification* module selects the specific semantic gesture to emphasize each predicted gesture region in a data-driven manner, retrieving relevant examples from a curated database of gestures performed by experts using a Retrieval-Augmented Generation (RAG) framework [29].

We then identified key design requirements for a gesture rehearsal system through an iterative design process, which included developing and evaluating an early Wizard-of-Oz prototype in a pilot study with five participants. This study highlighted the need for an effective Human-AI collaborative interface design; one that enables users to proactively engage with and co-create the AI’s gesture recommendations. Based on these findings, we developed a desktop-based rehearsal interface (shown in Figure 1) that tracks users’ speech in real-time and proactively cues users during their practice as and when gestures need to be performed. Additionally, the system supports collaboration by allowing users to modify the AI’s recommendations, enabling the creation of personalized gesture palettes tailored to each speaker’s preferences.

We evaluate our system in two ways: 1) technical evaluation of our gesture recommendation model, and 2) user evaluation of the final *GestureCoach* system. Results show that our emphasis proposal module outperforms off-the-shelf LLMs in identifying appropriate gesture regions in presenter notes. To assess the quality of these predictions, particularly false positives, we conducted a user study (N=30) evaluating both the appropriateness of the gesture regions and the suitability of the semantic gestures from the gesture identification module. Participants rated the majority of predictions as highly appropriate and the gestures as well-suited for conveying emphasis. Finally, in a user study (N=10), participants rehearsed for a talk using *GestureCoach*. Results show that speakers preferred using *GestureCoach* over rehearsing with just presenter notes, and they performed significantly more unique semantic gestures after using the system. Users also found the system easy to use and effective, reporting that it encouraged them to use a more diverse range of gestures and deliver more engaging talks. We also identified current limitations and outline opportunities for future work to improve AI-driven rehearsal systems and advance human-AI collaborative interfaces more broadly.

Our contributions include:

- We develop a gesture recommendation model and evaluate its performance for identifying gesture regions and for suggesting semantic gestures given input presenter notes.
- We implement an interactive rehearsal interface designed to help speakers integrate semantic gestures when giving a talk.
- Insights from a user study with *GestureCoach* on the strengths and limitations of AI-driven gesture rehearsal and design implications for future systems in this space.

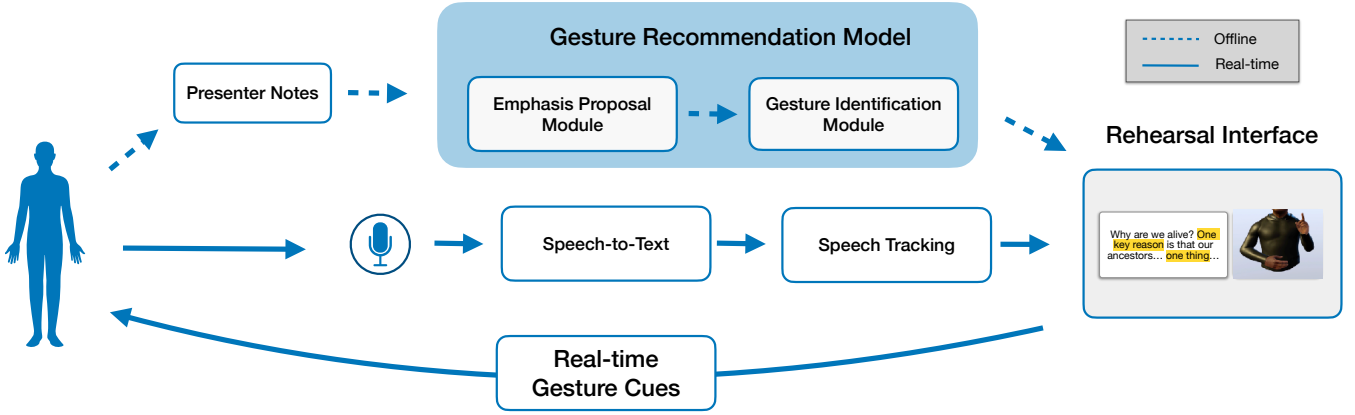


Figure 2: Overview of the *GestureCoach* system. Users’ presenter notes are processed by a gesture recommendation model in the backend. A frontend rehearsal interface displays the recommended gesture regions and gestures. During rehearsal, users’ speech is tracked and gesture cues are proactively delivered to provide real-time gesture guidance.

2 Related Work

2.1 Co-speech Gestures for Communication

Co-speech gestures are hand and arm movements that naturally accompany speech and play a vital role in communication [22, 35, 36, 48]. These gestures are broadly classified into four types: beat, iconic, metaphoric, and deictic gestures [36]. Beat gestures are rhythmic movements that align with speech prosody to emphasize its structure. In contrast, iconic, metaphoric, and deictic gestures are semantic gestures that are directly related to the linguistic content being spoken, enhancing content delivery during communication [17, 23, 35]. Iconic gestures physically resemble the objects or actions they represent, such as mimicking an object’s shape or reenacting a movement. Metaphoric gestures convey abstract concepts through spatial metaphors, like cupped hands symbolizing “holding” an idea. Deictic gestures involve pointing to indicate referents, whether physically, such as gesturing toward an object, or abstractly, such as indicating an imaginative space.

Semantic gestures, which are the focus of this study, have been empirically shown to significantly benefit listener comprehension, especially when they depict motor actions or convey non-redundant information [17]. These gestures also help speakers express their intentions through complementary movements, making the message clearer and more engaging for listeners [16, 24, 36]. For example, moving one’s hands from side to side can indicate contrast when saying “on the other hand” [16]. Moreover, encouraging speakers to use hand gestures has been found to increase narrative length and influence prosodic features, such as fundamental frequency and intensity, thereby improving intonation [9]. These findings show that semantic gestures play an important role in making communication and presentations more effective.

2.2 Systems for Rehearsing Talks

Tutoring systems aimed at improving task-related skills span various domains, including machine task training in manufacturing [19, 33, 50], surgical training [12, 13], and sports training through real-time feedback [14, 31, 34]. In the domain of public speaking,

such as giving talks or presentations, training systems have predominantly focused on vocal delivery, offering feedback on aspects like intonation and speaking pace [10, 54–56]. Despite the well-established role of gestures in enhancing communicative effectiveness, there has been relatively little research dedicated to systems that support the development of gesturing skills while giving talks. Although a few talk rehearsal systems also assist presenters in improving their gestures, they typically provide high-level guidance rather than detailed recommendations. For example, some approaches monitor a speaker’s body posture to detect closed positions or limited gestural variation, and provide suggestions to gesture more or less [3, 10, 47]. While this in-situ feedback encourages speakers to incorporate or reduce gestures, it does not provide specific guidance on *when* to gesture or *what* gestures to use for effective emphasis.

Recent advancements in generative AI, particularly LLMs, have opened up new possibilities for building more intelligent interactive systems. In HCI research, the versatility of LLM-based tools has been demonstrated across diverse domains such as writing assistance [57, 62], image generation [4], programming [2, 21], cognitive support [64], choreography [32], content creation in AR/VR [11, 51, 53], and natural language-driven physical design [44]. More recently, LLMs have been applied to presentation training; for example, Park et al. introduced a system that generates diverse virtual audiences to provide feedback on talks [43].

Building on this foundation, our work introduces an LLM-driven rehearsal system that offers fine-grained, proactive guidance for integrating meaningful semantic gestures. By analyzing presenter notes, the system identifies points where gestures would be most impactful and recommends specific gestures based on patterns drawn from expert speakers. This approach provides more targeted support than existing systems, enhancing speakers’ ability to deliver expressive and engaging presentations.

2.3 Modeling Semantic Gestures

Modeling semantic gestures in co-speech gesture synthesis is challenging due to the complex interaction between language and human motion [42, 59]. Unlike beat gestures, semantic gestures are sparse and exhibit a long-tailed distribution, making it challenging for models to learn their patterns effectively [26, 40, 58]. There has been active research on modeling semantic gestures for co-speech gesture synthesis [8, 26, 30, 39, 40, 61, 63], with most efforts focusing on generating gestures from speech.

To specifically incorporate semantic gestures into the gesture generation process, some works predict gesture properties such as gesture type information (deictic, metaphoric, and iconic) and gesture phase information from speech, in order to condition or retrieve semantic gestures for co-speech gesture generation [26, 40]. However, when it comes to rehearsing a presentation where the user only has presenter notes and not the actual speech, applying these methods becomes challenging. A recent work predicts semantic gestures using labels such as "hands chop" and "throat cut" in location of the text transcript, using a predefined motion library [61]. However, the semantic gestures in this approach are limited to the library's set, which does not cover domain-relevant gestures, such as those performed while giving presentation talks.

In this work, we aim to disentangle the tasks of predicting *when* to gesture and *what* semantic gesture to perform [41]. Linguistic research suggests that semantic gestures often co-occur with specific spoken discourse markers, such as discourse connectives (e.g., "however," "for example") that indicate topic shifts, and stance markers (e.g., "absolutely" "probably") that reflect the speaker's level of certainty [7, 28, 38, 52]. For instance, raised index fingers tend to co-occur with contrastive discourse markers like "but," "however," that signal exceptions [20]. These insights suggest that there are detectable patterns that can help determine when to gesture directly from text, which we address by fine-tuning an LLM to predict these gesture regions. Determining what semantic gesture to perform, on the other hand, is more challenging as there are no clear rules on what gesture best fits a given context. Hence, we guide our modeling using a data-driven approach, dynamically retrieving contextually appropriate gestures from a corpus of expert gestures using RAG.

3 GestureCoach: Overview

We present *GestureCoach*, a rehearsal system designed to help speakers practice their presentations while integrating semantic gestures effectively. Unlike previous training systems that offer generic cues to increase or decrease gesturing during talks [3, 10, 47], *GestureCoach* provides targeted guidance by analyzing expert speakers' behavior to determine both *when* a speaker should gesture during their talk and *what* semantic gesture they should use for effective emphasis.

Figure 2 provides an overview of our system. To start with, the speaker provides their presenter notes as input. *GestureCoach* then processes the presenter notes to identify specific text segments where gestures can enhance clarity and emphasis, i.e., gesture regions, and suggests appropriate semantic gestures for each region. As the speaker rehearses by reading the notes aloud, the interface actively tracks their speech to identify their position in the notes

and proactively plays a video clip of the semantic gesture to be performed in real-time. In the following sections we elaborate on the two main components of *GestureCoach*: 1) a backend gesture recommendation model (Section 4), and 2) a frontend rehearsal interface (Section 5).

4 GestureCoach: Recommendation Model

Developing an effective gesture recommendation model that predicts both *when* a speaker should gesture and *what* semantic gesture to use for emphasis presents a significant challenge. Gesture placement and selection are not merely a function of the words spoken; rather, they require a deeper knowledge of the contextual meaning, discourse structure, and audience engagement strategies that skilled speakers intuitively employ.

Given recent advancements in language modeling, we hypothesized that zero-shot prompting of an off-the-shelf LLM such as GPT-4o could effectively identify gesture regions in presenter notes where speakers should gesture. However, our preliminary testing revealed two key challenges. First, the LLM predicted large segments of text as gesture regions, likely because it lacks knowledge of expert gesturing patterns. This lack of precision can make it difficult for users to determine exactly when to gesture while speaking. Second, it provided generic or abstract descriptions of what gestures to perform rather than specific, actionable suggestions. This lack of clarity, especially for novice speakers, can make it challenging to translate the recommendations into meaningful physical movements.

To address these issues, we developed a gesture recommendation model that accurately predicts specific parts of the presenter notes that should be emphasized with gestures, and provides corresponding semantic gesture in the form of short (1-3 seconds) video clips. Figure 3 shows an overview of our model. It comprises of 1) a gesture database created from TED talk videos, by annotating the text segments in the talk transcript where speakers performed semantic gestures in the video. 2) an *Emphasis Proposal* module, trained on this database to identify gesture regions in the presenter notes that should be highlighted with semantic gestures. 3) a *Gesture Identification* module that retrieves a suitable semantic gesture for each predicted gesture region from a database of recorded gesture videos.

4.1 Gesture Database

To create a gesture database, we first collected TED talk videos of expert speakers who performed semantic gestures. The database consists of annotated gesture regions in the talk transcripts where the speaker performed semantic gestures. Each gesture region is also associated with a video clip of a humanoid avatar performing the semantic gesture.

The annotation was done by two researchers (authors) who viewed a TED talk video and marked text segments in the transcript where the speaker performed semantic gestures with their arms/hands. A gesture was identified as semantic if it deviated from their natural beat gesture and was representative of the major types of semantic gestures identified in literature (Section 2.1), i.e., iconic, metaphoric, and deictic gestures. In particular, our annotation aimed to capture the semantic correspondence between a text segment and

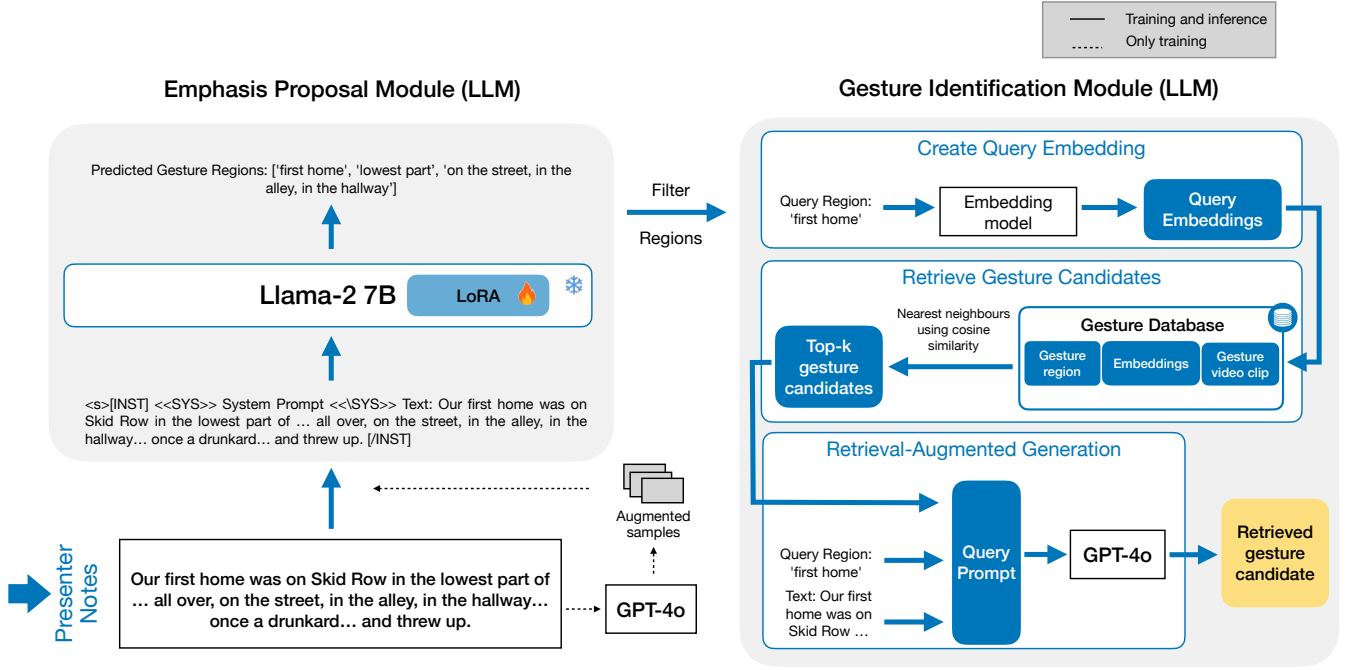


Figure 3: Gesture recommendation model architecture. It consists of two LLM-based modules: 1) The *emphasis proposal* module is fine-tuned on data from expert speakers to predict gesture regions in the presenter notes. The predicted regions are filtered and passed to 2) the *gesture identification* module which selects the most suitable semantic gesture from the gesture database using Retrieval-Augmented Generation.

the associated gesture rather than the text spanning the duration of the gesture. Given the subjective nature of the labeling task, we followed a two-step approach to ensure labeling consistency. Both researchers first calibrated by independently labeling three TED videos and resolving any labeling conflicts through discussions. The final interrater agreement was Cohen’s $\kappa = 0.88$, indicating a strong level of agreement. The remaining videos were annotated by one of the researchers.

We used the original transcripts uploaded by the TED YouTube channel for annotation. For each TED video, the text transcript was divided into chunks of approximately 100 words, with each chunk forming a sample. Within each sample, text segments where the speaker performed semantic gestures were annotated. A total of 10 TED talk videos (listed in Appendix A) were analyzed, resulting in 252 samples and 994 annotated gesture regions. For each video, we randomly split 80% of the data for training and 20% for testing.

The annotated gesture regions have an average length of 3.10 (± 2.48) words with an average of 3.81 (± 2.38) gestures regions in a sample. The gesture video clips in the database were created by enacting the gestures while wearing a Quest Pro headset, with the motion mirrored onto a humanoid avatar using the Movement SDK¹. In total, we recorded 343 clips, each 1–3 seconds long, at 1080p resolution and 30 frames per second.

¹<https://developers.meta.com/horizon/documentation/unity/move-overview/>

4.2 Emphasis Proposal

This module takes presenter notes as input and identifies gesture regions that should be emphasized with semantic gestures. For this, we fine-tune an LLM (Llama-2 7B) using a parameter-efficient fine-tuning technique known as Low-Rank Adaptation (LoRA) [18], in which a small number of task-specific trainable parameters (adapter layers) are inserted into the frozen transformer layers. These adapter-layers are trained to learn our specific task of identifying gesture regions. This approach has two-fold benefits. First, using LoRA helps us achieve this task with a few million parameters in a modularized fashion, allowing us to balance the rich knowledge from the pre-trained LLM with the new information learned from the data. Second, this helps reduce the sensitivity towards the system prompt, which can be a challenge in zero-shot prompting, where performance often depends on prompt engineering.

4.2.1 Fine-tuning with Data Augmentation. Given the limited amount of training data available for fine-tuning, we employed data augmentation using synthetic samples generated by LLMs. Specifically, we prompted GPT-4o [1] to generate novel text transcripts and their corresponding gesture regions, ensuring they were structurally and semantically similar to the original annotated sample but adapted to different TED talk topics. To maintain data quality, all synthetic samples underwent manual verification to ensure that the generated gesture regions were either semantically equivalent or verbatim to the original annotations. We randomly selected 70% of the videos for data augmentation, and for each training sample in these videos,

we generated five additional augmented samples, increasing the total training dataset size to 881 samples. This augmentation process introduced greater variability into the dataset, enabling the model to better learn the underlying patterns of semantic gestures in text.

We use Llama-2 7B² as the base LLM and integrate LoRA adapters with $r = 128$, $\alpha = 256$ for fine-tuning using the PEFT³ library. The base LLM remains frozen, and only the LoRA parameters are trained. Training is conducted for 20 epochs with a learning rate of $2e - 03$ using a single NVIDIA H100 GPU. For fine tuning and subsequent inference, we pass the input text using the standard prompt format⁴ (seen in Figure 3) of Llama 2 models.

4.2.2 Filtering the Predicted Regions. We include a filtering module that ensures predictions from the *emphasis proposal* module align with the original presenter notes before passing them to the *gesture identification* module. This step is important because the generative nature of LLMs can lead to predictions where the gesture regions do not match the original text verbatim. Such differences can range from minor ones (e.g., a missing apostrophe) to complete paraphrasing. The filtering process first attempts a verbatim match between the predicted gesture region and the presenter notes. If no exact match is found, we apply a sliding window approach to search for the most similar gesture region in the notes. Specifically, we compute the cosine similarity between the embedding of the predicted gesture region and the embeddings of each text window in the presenter notes. We set a similarity threshold of 75%, above which a match is accepted. If no sufficiently similar match is found, the prediction is discarded. We used bge-base-en-v1.5-gguf⁵ as the embedding model.

4.3 Gesture Identification

The gesture identification module determines the most appropriate semantic gesture to emphasize the gesture regions predicted by the emphasis proposal model. Selecting the correct semantic gesture for a given gesture region requires understanding the types of gestures available, their contextual meaning, and which gestures are most suitable for emphasis in a particular context.

To systematically model these dependencies, we adopt a Retrieval-Augmented Generation (RAG) framework [29], which enables an LLM to access a knowledge base of semantic text–gesture relationships. First, we generate semantic vector embeddings for the labeled gesture regions in the database using a language embedding model. For each gesture region predicted by the emphasis proposal module, we retrieve the most relevant candidates from the vector database based on embedding similarity. We then provide the LLM with the presenter’s notes and the retrieved top-k gesture candidates, prompting it to select the most contextually appropriate gesture for emphasizing the given gesture region. The video clip associated with the selected candidate is then used as the gesture cue for that gesture region.

We developed a custom RAG framework using the bge-base-en-v1.5-gguf model to get semantic vector embeddings, FAISS⁶ to

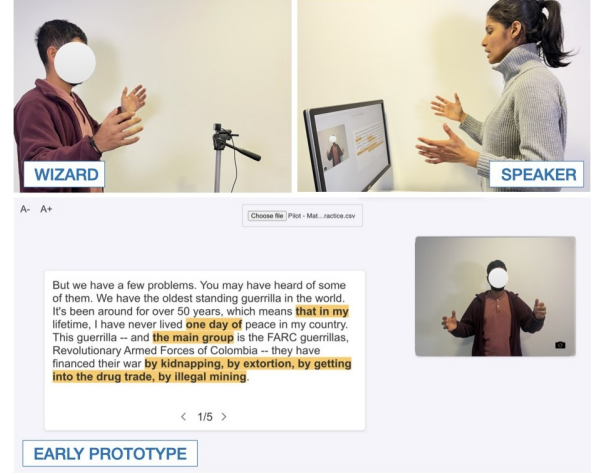


Figure 4: Early Wizard-of-Oz prototype that simulates a real-time gesture cueing interface for rehearsing talks.

create and efficiently search over the vector database, and the GPT-4o model to generate the final gesture candidate. We retrieve the top-3 gesture candidates using the established approximate K nearest neighbor search with cosine score as the similarity measure [27]. The prompts for LLMs in both modules can be found in Appendix B.

5 GestureCoach: Rehearsal Interface

Our objective was to identify the essential components and potential benefits of a gesture rehearsal interface. Following Buchenau’s experience prototyping protocol [5], which emphasizes gaining insights through users’ direct interactions with functional prototypes, we developed an initial Wizard-of-Oz prototype of the interface and conducted a pilot study with users.

5.1 Early Prototype

Our initial prototype was a minimalist web interface developed using ReactJS as shown in Figure 4. It consisted of two windows. The first window displayed the presenter notes that users would rehearse. Text corresponding to the gesture regions were highlighted in the presenter notes. The second window displayed the semantic gesture cue that the user should perform to emphasize the highlighted region. In our prototype, a “wizard” (one of the authors) enacted these cues in real-time via a webcam, displaying them in the gesture cue window. This setup simulated a just-in-time system that monitors speech and provides guidance to gesture just before users reach the highlighted region. Additionally, the windows were adjustable, allowing users to rearrange and resize them for comfort.

5.2 Pilot Study

We recruited 5 participants (3 males, 2 females, aged 25 to 45 years) who volunteered for the study. Four of them self-identified as novice speakers who had given 1-2 formal talks in the past and one as an expert who had given more than 10 talks. After an introduction and walkthrough of the prototype, participants were asked to rehearse

²<https://huggingface.co/meta-llama/Llama-2-7b>

³<https://huggingface.co/docs/peft/en/index>

⁴<https://www.llama.com/docs/model-cards-and-prompt-formats/meta-llama-2/>

⁵<https://huggingface.co/CompendiumLabs/bge-base-en-v1.5-gguf>

⁶<https://github.com/facebookresearch/faiss>

for a talk, emphasizing the highlighted regions in the presenter notes by mimicking gesture cues displayed in the gesture window.

We used 5 samples from the database as presenter notes. Participants first read through and understood the text in the presenter notes, before presenting each note with the cued gestures. For rehearsal, they were asked to rehearse each note three times. After the task, participants completed a 15-minute semi-structured interview to discuss the system’s usability and share feedback. We also observed participants’ interactions throughout the study. Each session lasted approximately 45 minutes.

5.3 Design Insights & Final Interface

All participants recognized the potential of a gesture rehearsal interface as it helped them successfully incorporate the suggested gestures while practicing the talk. Novice speakers, in particular, found it useful for breaking out of a static posture and making their presentations more dynamic: *"It motivates me to do gestures, makes me more relaxed to do gestures"* (P4). Highlighting gesture regions was an effective way for conveying when to gesture. While expert speakers found this highlighting to be sufficient to improve their rehearsal, novice speakers relied on the displayed gesture cues to understand what gestures to use for emphasis. To minimize context switching, they often placed the gesture window next to the notes window. While the interface was useful in its current form, users identified several issues and suggested improvements to enhance the interface. We discuss four key insights, which translate to our final interface.

- **Gesture cue visualization.** Early in our design process, we chose to demonstrate the gestures rather than representing them symbolically. While lightweight symbols such as emojis or gesture icons may be quicker for users to interpret during rehearsal, it can also be challenging to convey the spatial relationships and movement dynamics of gestures with such static representations. Our pilots indicate that users were generally comfortable following gesture demonstrations presented in a third-person, upper-body view. This perspective allowed them to clearly observe arm and hand movements while understanding their spatial relation to the upper body, aiding effective gesturing. However, two participants reported that guidance from a real human felt distracting, inducing a sense of social evaluation, particularly due to eye contact and facial expressions: *"Its a bit embarrassing ...I feel like someone's looking at me... it make me more self-aware and affects my practice"* (P2). Hence, in our final interface we use a animated humanoid avatar (upper-body only) captured from below the face, focusing on arm and hand motions to emphasize the gesture.
- **Onset of gesture cues.** Effective integration of gestures into speech requires timely presentation of gesture cues. In our pilots, we explored the optimal time for delivering these cues by varying the onset time at which the wizard triggered the cue relative to a highlighted region. We found that presenting gesture cues 4–5 words before the user reached a highlighted region was sufficient for seamless integration.
- **Modifying predicted gestures.** Users occasionally found that some of the gesture regions or the suggested gesture did not align with their expectations due to individual communication

styles or cultural differences in gesturing norms. Certain gestures may carry different meanings across cultures, making them less suitable or even inappropriate for some users: *"I don't point like this [pointing index finger at audience]... it seems rude... would be good to have other suggestions [for gestures] here"* (P4). To address this, we enable users to customize the assigned gestures by selecting from two alternative gestures retrieved by the gesture identification model.

- **Supporting preview of gestures.** During their initial rehearsal, 3/5 participants indicated difficulty integrating the suggested gestures naturally while speaking. They also struggled with smooth transitions between consecutive gestures. This challenge was particularly pronounced for gestures that participants had not encountered before: *"If I haven't seen the gesture before, I slow down or stop speaking to see the gesture and do it"* (P1). To address this challenge, we introduced a gesture preview feature. By hovering over the highlighted region, users can review and familiarize themselves with the suggested gestures before starting their rehearsal. It also allows users to anticipate transitions between gestures, making their delivery more fluid and natural.

5.3.1 Final Interface. Our final interface (Figure 1) builds upon our early prototype, incorporating the features identified from our pilots. Users begin by uploading their presenter notes. The gesture recommendation model analyses the notes, predicts gesture regions, and retrieves relevant gestures for each predicted region from the database. These regions are highlighted in the notes and displayed alongside the current presentation slide. Users start their rehearsal by clicking the "Rehearse" button. As they speak, an automatic speech recognition module tracks their progress through the notes. This enables the system to provide real-time gesture cues in the gesture window, offering timely guidance during rehearsal.

The real-time speech tracking module is implemented in Python using the RealtimeSTT library⁷ for live speech-to-text conversion. It aligns spoken words with the presenter’s notes using a lexical similarity matching algorithm. Each time the *rehearse* button is pressed, the module resets the speech flow index to the start of the current slide’s notes. During tracking, each newly recognized word prompts a comparison between the last three spoken words and all three-word sequences within a sliding window of the speech notes. The window spans two words before and ten words after the current speech flow index.

Similarity is measured using the Levenshtein similarity ratio [46], which calculates the minimum number of character edits needed to convert one string to another:

$$\sigma = \left(1 - \frac{\text{Levenshtein Distance}(s_1, s_2)}{|s_1| + |s_2|} \right) \times 100$$

Here, a 100% score indicates identical strings. The speech flow index is updated to the position of the phrase with the highest similarity score, provided it meets a minimum threshold of 50%, which we found to be effective in our pilot study. If no match exceeds this threshold, the index remains unchanged, treating the spoken input as irrelevant to the notes.

⁷<https://github.com/KoljaB/RealtimeSTT>

	Precision		Recall		F1		No. of Predicted gesture regions	Length of Predictions (No. of words)
	DM	SM	DM	SM	DM	SM		
GPT-4o ¹	0.034	0.039	0.043	0.047	0.038	0.043	4.18 \pm 1.67	15.41 \pm 8.52
Llama-2 7B	0.057	0.081	0.062	0.085	0.059	0.083	4.1 \pm 1.95	8.20 \pm 6.52
Our Model	0.159	0.206	0.185	0.233	0.171	0.219	4.45 \pm 1.26	2.92 \pm 1.42

Table 1: Model performance comparison for held-out test samples using direct matching (DM) and semantic matching (SM) schemes.

6 Model Evaluation

We evaluated the performance of our gesture recommendation model in two ways. First, we compared our fine-tuned LLM model with two baseline approaches to evaluate the extent to which our model helps predict gesture regions. We followed this with a more holistic evaluation of our approach through a user study where users subjectively rated the suitability of the predicted gesture regions and retrieved gestures for different samples.

6.1 Model Performance

6.1.1 Method. We evaluate the *emphasis proposal* on the held-out test data from the database (random 80:20 split; 56 samples containing 211 gestures) and on an unseen TED talk video (29 samples containing 256 gestures). We compare our model with two state-of-the-art LLMs: 1) OpenAI GPT-4o model 2) Llama-2 7B. The prompt was kept the same across all models and can be found in the Appendix B.

6.1.2 Metrics. We compared the recall, precision, and F1 scores across the models using two matching schemes.

- **Direct matching:** This is a strict matching where a predicted gesture region is considered a true positive if there is a verbatim overlap with a ground truth.
- **Semantic matching:** This is a more relaxed matching where predicted gesture regions are considered true positives if there is a 75 percent match in similarity between its vector embedding and that of a ground truth. This is an important metric because the LLM is a generative model and the generated predictions could have minor variations from the original text in terms of spelling, missing letters or punctuation. Gesture regions that are not found in the text are considered invalid and not counted as false positives, as the system directly discards them.

Regardless of the matching scheme, it is essential that the predicted gesture region does not substantially exceed the length of its corresponding ground truth. A model that predicts gesture regions as entire sentences or excessively long phrases may achieve high overlap with the ground truth but lacks the granularity needed for precise gestural guidance. We ensure this constraint by verifying that a predicted region extends no more than three words beyond the matched ground truth.

6.1.3 Results. Tables 1 and 2 show that our model outperforms the baseline models across all metrics for both held-out and unseen samples. We see that for held-out samples from our collected dataset (10 videos), our fine-tuned model has an average F1 score of 21.9

¹accessed on 21-05-2024

	Precision		Recall		F1		No. of Predicted gesture regions	Length of Predictions (No. of words)
	DM	SM	DM	SM	DM	SM		
GPT-4o ¹	0.038	0.038	0.002	0.002	0.028	0.028	4.48 \pm 1.13	16.40 \pm 9.44
Llama-2 7B	0.190	0.230	0.101	0.114	0.132	0.152	4.17 \pm 2.21	6.03 \pm 5.02
Our Model	0.348	0.490	0.171	0.233	0.230	0.316	3.86 \pm 1.25	2.93 \pm 1.39

Table 2: Model performance comparison for unseen TED talk video using direct matching (DM) and semantic matching (SM) schemes.

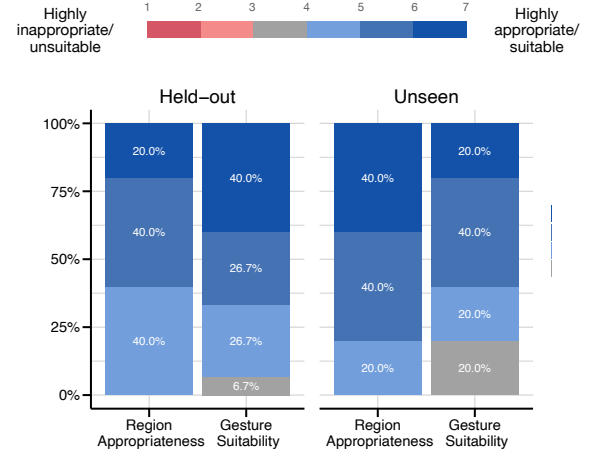


Figure 5: Distribution of users' ratings for the appropriateness of false positives and the suitability of gestures selected by the model. Over half of false positives were rated as *valid* (avg. ratings >5), and 40% as *neutral* (avg. ratings 3.5–5). None were clearly inappropriate (avg. ratings < 3), showing the model's recommendations are sensible in general.

in terms of semantic matching, approximately 3 times higher than base Llama-2 model and 6 times higher than GPT-4o. This is mainly because the baseline models lack knowledge of expert gesturing patterns and hence tends to predict overly long gesture regions (upto 16 words long), resulting in significantly lower performance. In contrast, our model learns from expert data and is able to predict more precise regions (upto 3 words long).

A similar trend is observed on the unseen video, suggesting that our model could generalize well to other TED-style presentations. While the F1 score is higher for this video, a per-video analysis of the mean performance across held-out samples shows that it falls within two standard deviations ($SD = 0.066$), indicating that the result is consistent with overall model behavior.

Although our model significantly outperforms off-the-shelf LLMs, its performance can appear relatively low when compared to traditional classification tasks. However, we argue that these metrics must be interpreted in the context of the task. Even though there are linguistic elements in spoken discourse where, as shown in previous research and also observed in our dataset, expert speakers commonly tend to gesture — for instance, when referring to themselves or the audience, during topic shifts such as discourse connectives, or when listing items — there are also notable speaker

variations. Some speakers, for example, tend to perform iconic gestures such as mimicking scissors for the word "cut." These gestures form the long tail of semantic gestures and account for a very small percentage of the dataset. Due to this imbalanced distribution, the model tends to learn the more frequently occurring gestures. This is a common challenge in machine learning models that deal with long tail distributions. While our model effectively captures broad gesturing patterns shared across speakers, it may fail to predict less frequent or speaker-specific gestures, leading to lower recall.

Similarly, the low precision scores do not necessarily indicate poor performance. Many of the false positives could still be valid gesture regions, suggested based on patterns learned from other speakers. To better understand this, we conduct a user study to subjectively evaluate the model predictions, particularly false positives, in the following section.

6.2 Quality of Predicted Gestures

While our model outperforms off-the-shelf LLMs in predicting gesture regions, the low precision scores raise an important question: Are the detected false positive regions truly inappropriate for gestures, or are some still valid for gesturing? It is possible that a predicted region was not gestured by the original speaker but could still serve as a meaningful emphasis point if accompanied by a suitable semantic gesture. To explore this, we conducted a user study to subjectively evaluate the quality of the gesture regions predicted by our *emphasis proposal* model and the suitability of the gestures selected by our *gesture identification* model.

6.2.1 Method. The study involved viewing video clips of a person who enacted a presentation of the sample text, emphasizing regions predicted by our *emphasis proposal* model using the gestures retrieved by the *gesture identification* model. Participants watched the video clips and evaluated each gesture region on a 7-point Likert scale based on two criteria: (a) the appropriateness of emphasizing the region during the presentation with a gesture, and (b) the suitability of the gesture used to emphasize that region.

We randomly selected five samples (containing 18 gestures) each from the held-out test data in the gesture database and from the unseen video (containing 16 gestures). The video clips were enacted by one of the authors with their face blurred to ensure anonymity and prevent any potential bias in ratings. We recruited 30 (16 female, 13 male, 1 other) volunteers with a mean age of 28.2 ($SD=5.19$) years from the university and randomly assigned them to evaluate either the held-out or the unseen video set. Evaluating a video set took approximately 15 minutes.

6.2.2 Results. Overall, participants found the gesture regions proposed by the *emphasis proposal* model appropriate for emphasis in both held-out ($M = 5.23, SD = 1.59$) and unseen samples ($M = 5.56, SD = 1.43$). Similarly, they rated the semantic gestures retrieved by the *gesture identification* model as well-suited for emphasis in both held-out ($M = 5.44, SD = 1.58$) and unseen samples ($M = 5.08, SD = 1.55$). To assess the impact of false positives (see Figure 5), we define "valid" false positives as samples with an average rating of higher than 5 in both appropriateness and suitability. These can be considered true positives, as users found them to be sensible emphasis points. We found that 53.3% of false positives

in held-out samples and 60% in unseen samples met this criterion. Additionally, 40% of false positives in both cases scored between 3.5 and 5. Such samples can be considered to be "neutral" or slightly positive, implying gesturing at those points would not negatively affect the presentation. Notably, no samples scored below 3, indicating that the model does not suggest completely irrelevant or unsuitable gesture regions. These results highlight the importance of subjective evaluation in generative tasks like gesture prediction, where traditional metrics such as F1 score or segmentation accuracy may not fully capture the quality or effectiveness of the model predictions.

7 User Evaluation

We conducted a user study with 10 participants to evaluate whether rehearsing presentations with *GestureCoach* can help users positively affect their gesturing while giving the talk. We also evaluate the user experience and usability of the system to understand its strengths and limitations.

7.1 Method

We conducted a within-subject study with two counterbalanced conditions: (1) a baseline *Notes* condition, where participants rehearsed using only the presenter notes, and (2) the *GestureCoach* condition, where participants viewed recommended gesture regions and received proactive gesture cues. The rehearsal interface for Notes condition was identical to the *GestureCoach* condition, but without the highlighted regions and gesture cues. To minimize learning effects, we used two separate parts of the transcript from the unseen TED talk video for each condition. We segmented the transcript into 100-word chunks to create the presenter notes for each slide. The slide sets were counterbalanced across the conditions. Participants rehearsed five slides per condition, rehearsing each slide three times. After rehearsal, participants presented the talk as they would in front of an audience, using just the presenter notes without any intervention.

After obtaining informed consent, we briefed participants on the task and asked them to rehearse and deliver two talks as if presenting to an audience. In the *GestureCoach* condition, participants were first introduced to the system and its functionalities. During rehearsal, they could modify or delete the recommended gestures based on their preference. Their presentations after rehearsal were video-recorded for analysis. After each condition, participants reviewed the recorded presentation of their talk. The study concluded with a short survey comparing the two rehearsal methods followed by semi-structured interview to gather qualitative feedback.

7.2 Participants

Our participants included 4 males and 6 females, who were proficient in English with an average age of 25.42 years. Their expertise in giving talks in front of an audience varied; 1 participant identified as beginner (1-2 talks), 2 as intermediate (3-5 talks), 5 as advanced (5-10 talks), and 2 as experts (10+ talks). Each participant spent 1 hour in the study and received 10 Euros for their time.

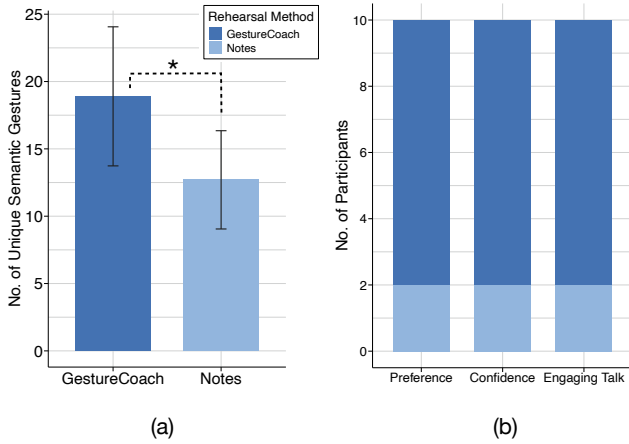


Figure 6: Comparison of gesture usage and user preference after rehearsing with *GestureCoach* vs. Notes. *GestureCoach* led to significantly more unique semantic gesture use and was preferred by 8 of 10 participants, who reported increased confidence after practice and a more engaging final talk.

7.3 Results

7.3.1 Gesture Quantity. We quantified semantic gestures performed by speakers during their talks. An expert in the gesture domain (co-author) blind to the conditions annotated the recorded talks using the same criteria used for creating the gesture database (Section 4.1). Overall, speakers performed slightly higher number of gestures in the *GestureCoach* condition ($M = 30.30, SD = 8.87$) than in the Notes ($M = 25.1, SD = 10.16$) condition, but the difference was not significant. A closer look at speakers' individual data showed that 8 of 10 speakers performed significantly more gestures ($t(7) = 2.99, p = .02$) after rehearsing with *GestureCoach* ($M = 31.75, SD = 8.97$) than with Notes ($M = 21.65, SD = 7.52$). The remaining two users showed an opposite trend, performing more gestures after using Notes ($M = 39, SD = 7.07$) than with *GestureCoach* ($M = 24.5, SD = 7.77$). We discuss potential reasons for this in Section 8.2.2 based on users' qualitative feedback.

7.3.2 Gesture Diversity. We measure gesture diversity as the number of unique semantic gestures performed by speakers during their talk. As seen in Figure 6a, speakers performed considerably higher number of unique gestures in the *GestureCoach* condition ($M = 18.90, SD = 7.21$) than in the Notes ($M = 12.70, SD = 5.10$) condition. A paired samples t-test showed that the difference is significant ($t(9) = 3.27, p = .009$) with a large effect size, Cohen's $d = 1.03$, suggesting that rehearsing with *GestureCoach* helped speakers add diversity to the gestures they used during their talk.

7.3.3 Subjective Evaluation of *GestureCoach*. Overall, speakers found the *GestureCoach* interface intuitive and easy to use with an average SUS score of 90.83 ($SD = 5.86$). Figure 7a shows speakers' evaluation of *GestureCoach* using a 5-point Likert scale. They found that the system recommended appropriate gesture regions for emphasis ($M = 4.20, SD = 0.63$) and that the suggested gestures aligned well with the emphasized text ($M = 3.9, SD = 0.73$). In cases

where the original gesture was unsuitable, speakers found that the alternatives recommended by *GestureCoach* often included a suitable one ($M = 4.6, SD = 0.51$). Using highlights ($M = 4.70, SD = 0.48$) and gesture visualizations ($M = 4.50, SD = 0.70$) to represent the suggestions were found to be easy to understand for users. There was also strong agreement among speakers that rehearsal experiences with these gestures was easy ($M = 4.6, SD = 0.69$).

7.3.4 Subjective Comparison of Rehearsal Methods. Rehearsal with *GestureCoach* was the preferred method for most speakers (8/10), enhancing their confidence (8/10) and enabling them to deliver a more engaging talk (8/10) compared to rehearsing with Notes. As shown in Figure 7b, speakers reported that using *GestureCoach* during rehearsal encouraged them to incorporate gestures more naturally ($M = 4, SD = 1.15$) ($t(9) = 3.87, p < .01$), provided ideas for gesture use ($M = 4.70, SD = 0.48$) ($W = 55.0, z = 2.80, p < .01$), and helped them use meaningful gestures ($M = 4.20, SD = 0.63$) ($W = 55.0, z = 2.80, p < .01$) at appropriate points ($M = 4.1, SD = 0.56$) ($t(9) = 6.00, p < .001$). It also added diversity to the gestures ($M = 4.30, SD = 0.67$) ($t(9) = 7.79, p < .001$) during their talk. Additionally, some participants observed improvements in their intonation ($M = 3.70, SD = 0.82$) ($t(9) = 1.96, p = 0.081$) and pacing ($M = 3.10, SD = 0.87$) ($W = 6.0, z = 1.60, p = 0.174$) after rehearsing with *GestureCoach*. However, they noted that rehearsing with Notes offered more flexibility ($M = 4, SD = 1.33$) for integrating their own gestures than *GestureCoach* ($M = 3.70, SD = 1.06$), although the difference was not significant.

8 Discussion

In the following section, we discuss the main insights from the user evaluation of *GestureCoach* and derived design implications.

8.1 Summary of Findings

8.1.1 Rehearsing with *GestureCoach* Helps Deliver More Engaging Talks. Overall, participants reported that rehearsing with *GestureCoach* enhanced their ability to deliver a more dynamic and expressive presentation. Speakers perceived an increase in gesture usage and the suitability of gestures used after rehearsing with *GestureCoach*. This perception is also supported by our analysis of speakers' gesture usage, which showed a 48.8% increase in the diversity of gestures used when speakers rehearsed with *GestureCoach* compared to using Notes.

Feedback from speakers indicates that the gesture recommendations provided by *GestureCoach* during rehearsal played a key role in encouraging them to do more gestures during their talk. As P1 described, focusing on gestures when rehearsing with just notes was challenging: "[Without *GestureCoach*] I know I need to move [my hands], but it's difficult to think about what would be a good gesture, so I end up just doing random gestures". In contrast, P1 explained that *GestureCoach* "helped me reflect on where to emphasize and what kind of gestures I could use... it's like someone giving feedback". The recommendations provided practical examples and demonstrated how to use gestures more intentionally.

While the primary objective of *GestureCoach* was to enhance users' gesturing during presentations, some speakers also reported a synergistic effect on their speech delivery. For instance, P7 felt that rehearsing with *GestureCoach* improved their intonation: "When it's

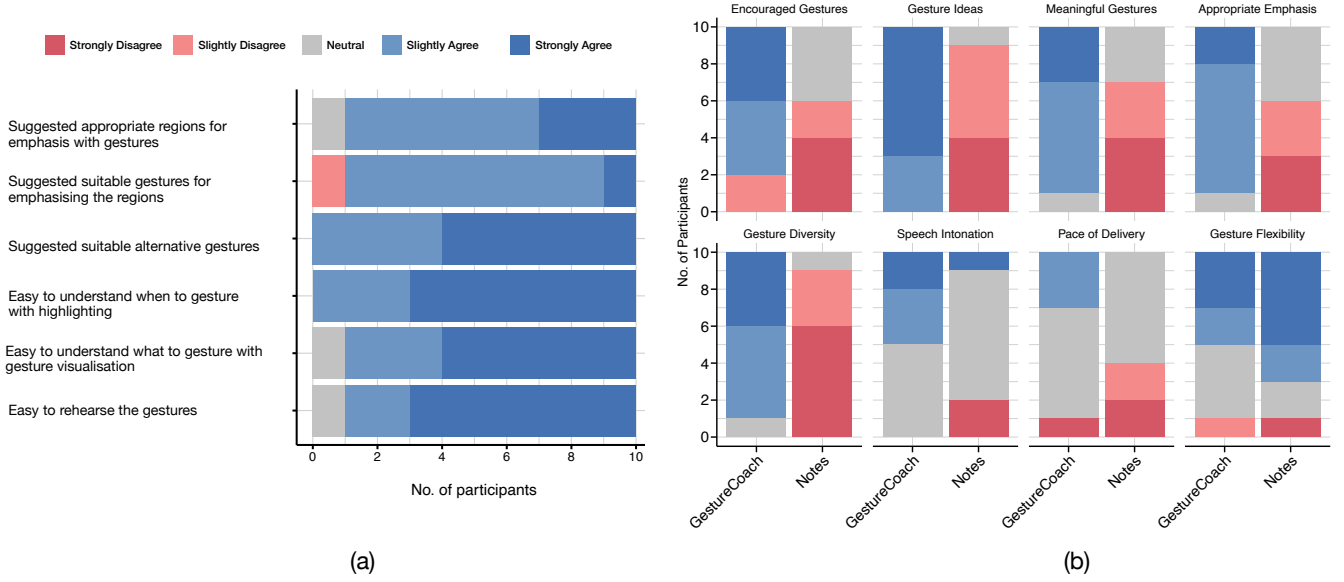


Figure 7: (a) Subjective evaluation of *GestureCoach* features indicate its usefulness and ease of use for rehearsal, (b) Results of subjective comparison between *GestureCoach* and *Notes*.

highlighted I also stress that point in my speech such that it matches the gesture". This positive impact on intonation occurs because effectively performing a gesture with the intended emphasis requires speakers to slow down and emphasize the associated words, allowing the gesture's intended meaning to be conveyed more effectively. Thus, by encouraging intentional emphasis in both gestures and speech, *GestureCoach* helped speakers deliver more expressive and engaging talks.

8.1.2 Enhancing Speakers' Gesture Vocabulary. A key goal of *GestureCoach* was to enhance speakers' gesture repertoire by providing access to the rich gesture vocabulary used by experts. Participant feedback strongly indicates that *GestureCoach* successfully achieved this goal by encouraging speakers to go beyond their usual, simplistic gestures and explore a wider range of expressive movements. The ability to quickly preview gestures and explore alternative ones suggested by *GestureCoach* was particularly valuable in this regard. We observed that speakers frequently browsed through these gesture recommendations, even when the initially recommended gesture was already meaningful: "I'm curious on what other gestures I could use... It gives me options" (P1). This exploration evolved into an active reflection process, where speakers considered the suitability of each gesture and selected one that aligned with their expectations: "It made me think more on what gesture I should use.. I consider the options and select the best" (P10). Through this process, speakers were exposed to new gestures, adding diversity to their repertoire: "I normally do a lot of simplistic gestures, but it [*GestureCoach*] gave me really iconic or unique gestures that fit the context and add value to the talk" (P5).

8.1.3 Learnability of Gestures. Speakers found it easy to incorporate the gestures within their talk during rehearsal and later recall many of them without any highlighted regions or gesture cues. P2

attributed this to the strong semantic relevance between the gestures and the emphasized text: "The suggested gestures contextually make sense, so it was easy to learn and do it later". This meaningful connection allowed speakers to naturally associate the gestures with the corresponding text in their presenter notes, making the gestures easier to remember and perform. These findings suggest that targeted semantic gesture recommendations could have significant potential to improve gesture training compared to generic feedback about the speaker's amount of gesturing [3, 10, 47].

8.1.4 Structured Rehearsal Improved Confidence. Traditional rehearsal often left speakers uncertain about what gestures to incorporate, relying primarily on gestures they could recall from their previous talks. As P4 explained, "The gestures I use are too random when I rehearse on my own. It's different in each iteration". This lack of structure was challenging for moderately experienced speakers and even more so for those who are naturally less expressive. In contrast, *GestureCoach* enabled speakers to critically reflect on the gestures they intended to use during their presentations. This reflection process allowed them to create an initial mental "sketch" of their talk, which structured their rehearsal process: "gave me a consistent set of gestures that I can start with" (P4). This structured rehearsal made it easier to rehearse with *GestureCoach* and ultimately enhanced participants' confidence in delivering their talks post-rehearsal.

8.2 Design Implications

8.2.1 Personalization via Human-AI Collaborative Authoring. Participants' ratings of the gesture recommendations provided by *GestureCoach* indicate that they generally found the gesture regions to be appropriate and effectively paired with meaningful gestures.

However, participants did not merely accept the originally suggested gestures passively. Instead, the process of identifying the best gesture to use for a recommended region was iterative and involved active reflection and modification. Through this iterative process, participants engaged in a collaborative authoring experience with *GestureCoach*, combining inputs from the system with their own experiences to create a gesture palette that felt unique and personalized to their presentation style. These findings align with insights from evaluations of recent AI-driven interfaces [44, 51], which emphasize the importance of providing customization controls and ensuring that the human remains actively involved in the creative process. Future research could further enhance this collaboration by enabling users to refine the system’s recommendations through prompts or demonstrations, and support gesture palette creation by actively learning from user selections.

8.2.2 Hybrid Rehearsal to Integrate with Natural Gestures. While *GestureCoach* helped speakers learn and incorporate diverse gestures into their presentations, expert speakers reported an unintended side effect: training with the system appeared to suppress their natural “beat” gestures, particularly during rehearsal. For these speakers, the recommendations and structured guidance provided by *GestureCoach* proved to be a double-edged sword. On the one hand, the system accelerated the rehearsal process and offered valuable inspiration, as noted by P4: “*It speeds up the process my rehearsal and gives me these ideas... But, it also constrains me because I focus too much on the suggested regions and getting them right that I don’t focus much on other gesturing at parts of my speech*”. To address this, these speakers expressed that they would use *GestureCoach* in a hybrid manner during rehearsal: beginning with the system to explore gesture suggestions, then continuing to rehearse using only the highlighted regions. This approach would allow them to integrate *GestureCoach*’s suggestions while adapting them to their own natural gesturing style. Future rehearsal systems could better support this process by observing speakers’ gesturing and providing cues only at targeted points in a unobtrusive manner. In addition, participants expressed a preference for being able to highlight new text segments they felt should be emphasized and to receive corresponding gesture recommendations from the system.

8.2.3 Modeling Considerations for Gesture Recommendation. We approached the gesture recommendation problem as two distinct tasks: deciding *when* to gesture and determining *what* gesture to use. The *when* component was driven by observable and well-established patterns such as gestures frequently occurring at topic changes, during list enumerations, or at points of emphasis. These patterns make it easier to model it as a finetuning task. In contrast, the *what* component was more complex due to high variability. Semantic gestures follow a long-tail distribution, where a single gesture can suit many different spoken contexts. For such problems, a RAG-based approach may be a better choice as it enables the model to retrieve grounded examples during inference rather than relying entirely on mappings learned through direct fine-tuning.

8.3 Limitations & Future work

8.3.1 Need for verbatim presenter notes. Our gesture recommendation model requires users to provide a complete transcript of

their talk in order predict gesture regions and potential gestures. During our pilots with the early prototype, novice speakers indicated that they indeed prefer to create such detailed presenter notes and closely adhere to them, particularly for timed presentations. In contrast, expert speakers typically rely on brief bullet points as a reference rather a full script. Future research should explore ways to relax the requirement for presenter notes while still capturing enough information to make accurate gesture predictions.

8.3.2 Guidance for subtleties in gestures. Our system currently provides only a general representation of a gesture’s overall motion but does not capture finer details such as the speed or intensity with which the gesture should be performed, which may depend on the context of the talk. Additionally, the system does not account for other aspects of body language, such as head movements or body posture, which are often used to enhance talks. Future gesture recommendation systems should aim to model these nuances and incorporate them into gesture cues to provide more comprehensive guidance during rehearsal.

8.3.3 Design Space of Gesture Visualizations. In our system, we visualized gesture cues using an animated avatar that demonstrated the gesture motion. However, it is also possible to consider other lightweight alternatives, such as emojis or icons, that are embedded within the presentation script. Such cues maybe easier to reference while speaking and could be particularly helpful in real-time contexts. Future work should explore the design space of gesture visualizations and empirically compare them to understand the trade-offs between gesture comprehension and cognitive load.

8.3.4 Adapting to cultural contexts. The gesture recommendation model selects gestures from a database of semantic gestures collected from expert speakers of diverse cultural backgrounds. However, our studies revealed that some gestures may not be appropriate for users from specific backgrounds. To address this, we provide customization options that allow users to choose alternative gestures for a given region. An improved system could further enhance this adaptability by actively learning user preferences through feedback, enabling more personalized gesture recommendations over time.

8.3.5 Generalisability beyond TED-style presentations. Our *emphasis proposal* module is fine-tuned on data from TED talk videos and shows promising predictive performance on both held-out and unseen presenter notes of TED-style talks. Although we are optimistic that it could generalize to other settings, such as public speeches, differences in gesturing patterns across contexts may not be fully captured. Future work should focus on developing more generalizable models by training on larger, more diverse datasets or creating models that adapt gesture predictions depending on the speaking context.

9 Conclusion

In this paper, we introduced *GestureCoach*, a system designed to proactively guide speakers during rehearsal in using effective, well-timed semantic gestures in their talk. To realise this, we developed an LLM-driven gesture recommendation model that predicts appropriate gesture regions from presenter notes and suggests suitable semantic gestures based on expert speakers’ gesturing patterns. Our

model evaluation shows promising accuracy in predicting gesture regions and suitable gestures. Building on insights from an exploratory study using a Wizard-of-Oz prototype, we then designed a rehearsal interface that proactively cues speakers to integrate the gesture recommendations from the model during rehearsal. An evaluation of our implemented system with users highlighted its efficacy in encouraging speakers to perform more diverse gestures, leading to a more engaging talk. These findings open up new possibilities for AI-driven rehearsal tools to support speaker training and performance.

Acknowledgments

This work has been partly supported by the Deutsche Forschungsgemeinschaft DFG under Grant No. 521601028 within the Priority Program SPP 2199 Scalable Interaction Paradigms for Pervasive Computing Environments, and the Deutsche Forschungsgemeinschaft, Funder Id: <http://dx.doi.org/10.13039/501100001659>, SFB 1102: “Information Density and Linguistic Encoding”, project number 232722074. We also thank the reviewers for their insightful comments that helped to improve the paper.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Tyler Angert, Miroslav Suzara, Jenny Han, Christopher Pondoc, and Hariharan Subramonyam. 2023. Spellburst: A Node-based Interface for Exploratory Creative Coding with Natural Language Prompts. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 100, 22 pages. doi:10.1145/3586183.3606719
- [3] Roghaye Barmaki. 2016. Improving Social Communication Skills Using Kinesics Feedback. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (San Jose, California, USA) (CHI EA '16). Association for Computing Machinery, New York, NY, USA, 86–91. doi:10.1145/2851581.2890378
- [4] Stephen Brade, Bryan Wang, Mauricio Sousa, Sageev Oore, and Tovi Grossman. 2023. Promptify: Text-to-Image Generation through Interactive Prompt Exploration with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 96, 14 pages. doi:10.1145/3586183.3606725
- [5] Marion Buchenau and Jane Fulton Suri. 2000. Experience prototyping. In *Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques*. 424–433.
- [6] JUDEE K. BURGOON, THOMAS BIRK, and MICHAEL PFAU. 1990. Nonverbal Behaviors, Persuasion, and Credibility. *Human Communication Research* 17, 1 (1990), 140–169. doi:10.1111/j.1468-2958.1990.tb00229.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-2958.1990.tb00229.x
- [7] G. Calbris. 2011. *Elements of Meaning in Gesture*. John Benjamins Publishing Company. <https://books.google.de/books?id=hat4PMRx7IUC>
- [8] Qingrong Cheng, Xu Li, and Xinghui Fu. 2024. SIGGesture: Generalized Co-Speech Gesture Synthesis via Semantic Injection with Large-Scale Pre-Training Diffusion Models. In *SIGGRAPH Asia 2024 Conference Papers*. 1–11.
- [9] Alice Cravotta, M Grazia Busà, and Pilar Prieto. 2019. Effects of encouraging the use of gestures on speech. *Journal of Speech, Language, and Hearing Research* 62, 9 (2019), 3204–3219.
- [10] Ionut Damian, Chiew Seng (Sean) Tan, Tobias Baur, Johannes Schöning, Kris Luyten, and Elisabeth André. 2015. Augmenting Social Interactions: Realtime Behavioural Feedback using Social Signal Processing Techniques. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 565–574. doi:10.1145/2702123.2702314
- [11] Fernanda De La Torre, Cathy Mengying Fang, Han Huang, Andrzej Banburski-Fahey, Judith Amores Fernandez, and Jaron Lanier. 2024. LLMR: Real-time Prompting of Interactive Worlds using Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 600, 22 pages. doi:10.1145/3613904.3642579
- [12] Mehrad Faridan, Bheesha Kumari, and Ryo Suzuki. 2023. Chameleoncontrol: Teleoperating real human surrogates through mixed reality gestural guidance for remote hands-on classrooms. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–13.
- [13] Éléonore Ferrier-Barbut, Ignacio Avellino, Geoffroy Canlorbe, Marie-Aude Vitrani, and Vanda Luengo. 2023. Learning With Pedagogical Models: Videos As Adjuncts to Apprenticeship for Surgical Training. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–40.
- [14] Mai Geisen and Stefanie Klatt. 2022. Real-time feedback using extended reality: A current overview and further integration into sports. *International Journal of Sports Science & Coaching* 17, 5 (2022), 1178–1194.
- [15] Susan Goldin-Meadow and Martha Wagner Alibali. 2013. Gesture's role in speaking, learning, and creating language. *Annual review of psychology* 64, 1 (2013), 257–283.
- [16] Jennifer Hinnell. 2019. The verbal-kinesic enactment of contrast in North American English. *The American Journal of Semiotics* 35, 1-2 (2019), 55–92.
- [17] Autumn B Hostetter. 2011. When do gestures communicate? A meta-analysis. *Psychological bulletin* 137, 2 (2011), 297.
- [18] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR* 1, 2 (2022), 3.
- [19] Gaoping Huang, Xun Qian, Tianyi Wang, Fagun Patel, Maitreya Sreeram, Yuanzhi Cao, Karthik Ramani, and Alexander J Quinn. 2021. Adaptutur: An adaptive tutoring system for machine tasks in augmented reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [20] Anna Inbar. 2022. The raised index finger gesture in Hebrew multimodal interaction. *Gesture* 21, 2-3 (2022), 264–295.
- [21] Majeed Kazemitabaar, Justin Chow, Carl Ka To Ma, Barbara J. Ericson, David Weintrop, and Tovi Grossman. 2023. Studying the effect of AI Code Generators on Supporting Novice Learners in Introductory Programming. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 455, 23 pages. doi:10.1145/3544548.3580919
- [22] Adam Kendon. 1972. Some relationships between body motion and speech. *Studies in dyadic communication* 7, 177 (1972), 90.
- [23] Adam Kendon. 1994. Do gestures communicate? A review. *Research on language and social interaction* 27, 3 (1994), 175–200.
- [24] Adam Kendon. 2004. *Gesture: Visible action as utterance*. Cambridge University Press.
- [25] 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*. Vol. 28. Elsevier, 389–450.
- [26] Taras Kucherenko, Patrik Jonell, Sanne Van Waveren, Gustav Eje Henter, Simon Alexandersson, Iolanda Leite, and Hedvig Kjellström. 2020. Gesticulator: A framework for semantically-aware speech-driven gesture generation. In *Proceedings of the 2020 international conference on multimodal interaction*. 242–250.
- [27] Eyal Kushilevitz, Rafail Ostrovsky, and Yuval Rabani. 1998. Efficient search for approximate nearest neighbor in high dimensional spaces. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*. 614–623.
- [28] Schuyler Laparle. 2021. Tracking Discourse Topics in Co-speech Gesture. In *International Conference on Human-Computer Interaction (HCI)*. Springer-Verlag, Berlin, Heidelberg, 233–249. doi:10.1007/978-3-030-77817-0_18
- [29] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems* 33 (2020), 9459–9474.
- [30] Yuanzhi Liang, Qianyu Feng, Linchao Zhu, Li Hu, Pan Pan, and Yi Yang. 2022. Seeg: Semantic energized co-speech gesture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10473–10482.
- [31] Jingyuan Liu, Nazmus Saquib, Zhutian Chen, Rubaiat Habib Kazi, Li-Yi Wei, Hongbo Fu, and Chiew-Lan Tai. 2022. Posecoach: A customizable analysis and visualization system for video-based running coaching. *IEEE Transactions on Visualization and Computer Graphics* (2022).
- [32] Yimeng Liu and Misha Sra. 2024. DanceGen: Supporting Choreography Ideation and Prototyping with Generative AI. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) (DIS '24). Association for Computing Machinery, New York, NY, USA, 920–938. doi:10.1145/3643834.3661594
- [33] Ziyi Liu, Zhengzhe Zhu, Enze Jiang, Feichi Huang, Ana M Villanueva, Xun Qian, Tianyi Wang, and Karthik Ramani. 2023. Instrumentar: Auto-generation of augmented reality tutorials for operating digital instruments through recording embodied demonstration. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [34] Dizhi Ma, Xiyun Hu, Jingyu Shi, Mayank Patel, Rahul Jain, Ziyi Liu, Zhengzhe Zhu, and Karthik Ramani. 2024. avaTTAR: Table Tennis Stroke Training with Embodied and Detached Visualization in Augmented Reality. In *Proceedings of the*

- 37th Annual ACM Symposium on User Interface Software and Technology. 1–16.
- [35] Fridanna Maricchiolo, Augusto Gnisci, Marino Bonaiuto, and Gianluca Ficca. 2020. Effects of different types of hand gestures in persuasive speech on receivers' evaluations. In *Speech Accompanying-Gesture*. Psychology Press, 239–266.
- [36] David McNeill. 1992. Hand and mind: What gestures reveal about thought.
- [37] David McNeill. 2008. Gesture and thought. In *Gesture and thought*. University of Chicago press.
- [38] David McNeill, Elena Levy, and Susan Duncan. 2014. Gesture in Discourse. *The Handbook of Discourse Analysis* (04 2014). doi:10.1002/9781118584194.ch12
- [39] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 1388–1398.
- [40] M Hamza Mughal, Rishabh Dabral, Merel CJ Scholman, Vera Demberg, and Christian Theobalt. 2024. Retrieving Semantics from the Deep: an RAG Solution for Gesture Synthesis. *arXiv preprint arXiv:2412.06786* (2024).
- [41] Michael Neff. 2016. Hand gesture synthesis for conversational characters. *Handbook of Human Motion* (2016), 1–12.
- [42] Simbarashe Nyatsanga, Taras Kucherenko, Chaitanya Ahuja, Gustav Eje Henter, and Michael Neff. 2023. A comprehensive review of data-driven co-speech gesture generation. In *Computer Graphics Forum*, Vol. 42. Wiley Online Library, 569–596.
- [43] Jeongeon Park and DaEun Choi. 2023. AudiLens: Configurable LLM-Generated Audiences for Public Speech Practice. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23 Adjunct). Association for Computing Machinery, New York, NY, USA, Article 122, 3 pages. doi:10.1145/3586182.3625114
- [44] Wanli Qian, Chenfeng Gao, Anup Sathya, Ryo Suzuki, and Ken Nakagaki. 2024. SHAPE-IT: Exploring Text-to-Shape-Display for Generative Shape-Changing Behaviors with LLMs. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) (UIST '24). Association for Computing Machinery, New York, NY, USA, Article 118, 29 pages. doi:10.1145/3654777.3676348
- [45] Emma Rodero. 2022. Effectiveness, attractiveness, and emotional response to voice pitch and hand gestures in public speaking. *Frontiers in communication* 7 (2022), 869084.
- [46] Sandip Sarkar, Dipankar Das, Partha Pakray, and Alexander Gelbukh. 2016. JUNITMZ at SemEval-2016 task 1: Identifying semantic similarity using Levenshtein ratio. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. 702–705.
- [47] Jan Schneider, Dirk Börner, Peter van Rosmalen, and Marcus Specht. 2015. Presentation Trainer, your Public Speaking Multimodal Coach (ICMI '15). Association for Computing Machinery, New York, NY, USA, 539–546. doi:10.1145/2818346.2830603
- [48] Stefanie Shattuck-Hufnagel and Pilar Prieto. 2019. Dimensionalizing co-speech gestures. In *Proceedings of the international Congress of phonetic sciences*, Vol. 5. 1490–1494.
- [49] Wing Chee So, Sotaro Kita, and Susan Goldin-Meadow. 2013. When do speakers use gestures to specify who does what to whom? The role of language proficiency and type of gestures in narratives. *Journal of psycholinguistic research* 42 (2013), 581–594.
- [50] Ana Stanescu, Peter Mohr, Dieter Schmalstieg, and Denis Kalkofen. 2022. Model-free authoring by demonstration of assembly instructions in augmented reality. *IEEE Transactions on Visualization and Computer Graphics* 28, 11 (2022), 3821–3831.
- [51] Xia Su, Jon E. Froehlich, Eunye Koh, and Chang Xiao. 2024. SonifyAR: Context-Aware Sound Generation in Augmented Reality. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) (UIST '24). Association for Computing Machinery, New York, NY, USA, Article 128, 13 pages. doi:10.1145/3654777.3676406
- [52] Varsha Suresh, M Hamza Mughal, Christian Theobalt, and Vera Demberg. 2025. Enhancing Spoken Discourse Modeling in Language Models Using Gestural Cues. *arXiv preprint arXiv:2503.03474* (2025).
- [53] Felicia Fang-Yi Tan, Peisen Xu, Ashwin Ram, Wei Zhen Suen, Shengdong Zhao, Yun Huang, and Christophe Hurter. 2024. AudioXtend: Assisted Reality Visual Accompaniments for Audiobook Storytelling During Everyday Routine Tasks. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 83, 22 pages. doi:10.1145/3613904.3642514
- [54] M. Iftekhar Tanveer, Emy Lin, and Mohammed (Ehsan) Hoque. 2015. Rhema: A Real-Time In-Situ Intelligent Interface to Help People with Public Speaking. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (Atlanta, Georgia, USA) (IUI '15). Association for Computing Machinery, New York, NY, USA, 286–295. doi:10.1145/2678025.2701386
- [55] H. Trinh, R. Asadi, D. Edge, and T. Bickmore. 2017. RoboCOP: A Robotic Coach for Oral Presentations. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 2, Article 27 (June 2017), 24 pages. doi:10.1145/3090092
- [56] Xingbo Wang, Haipeng Zeng, Yong Wang, Aoyu Wu, Zhida Sun, Xiaojuan Ma, and Huamin Qu. 2020. VoiceCoach: Interactive Evidence-based Training for Voice Modulation Skills in Public Speaking. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3313831.3376726
- [57] Zheer Xu, Shanjing Cai, Mukund Varma T, Subhashini Venugopalan, and Shumin Zhai. 2024. SkipWriter: LLM-Powered Abbreviated Writing on Tablets. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) (UIST '24). Association for Computing Machinery, New York, NY, USA, Article 22, 13 pages. doi:10.1145/3654777.3676423
- [58] Youngwoo Yoon, Woo-Ri Ko, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. 2019. 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)*. IEEE, 4303–4309.
- [59] Youngwoo Yoon, Pieter Wolfert, Taras Kucherenko, Carla Viegas, Teodor Nikolov, Mihail Tsakov, and Gustav Eje Henter. 2022. The GENE Challenge 2022: A large evaluation of data-driven co-speech gesture generation. In *Proceedings of the 2022 International Conference on Multimodal Interaction*. 736–747.
- [60] Haipeng Zeng, Xingbo Wang, Yong Wang, Aoyu Wu, Ting-Chuen Pong, and Huamin Qu. 2022. GesturLens: Visual analysis of gestures in presentation videos. *IEEE Transactions on Visualization and Computer Graphics* 29, 8 (2022), 3685–3697.
- [61] Zeyi Zhang, Tenglong Ao, Yuyao Zhang, Qingzhe Gao, Chuan Lin, Baoquan Chen, and Libin Liu. 2024. Semantic Gesticulator: Semantics-Aware Co-Speech Gesture Synthesis. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–17.
- [62] Zheng Zhang, Jie Gao, Ranjodh Singh Dhaliwal, and Toby Jia-Jun Li. 2023. VISAR: A Human-AI Argumentative Writing Assistant with Visual Programming and Rapid Draft Prototyping. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 5, 30 pages. doi:10.1145/3586183.3606800
- [63] Yihao Zhi, Xiaodong Cun, Xuelin Chen, Xi Shen, Wen Guo, Shaoli Huang, and Shenghua Gao. 2023. Livelyspeaker: Towards semantic-aware co-speech gesture generation. In *Proceedings of the IEEE/CVF international conference on computer vision*. 20807–20817.
- [64] Wazeer Deen Zulfikar, Samantha Chan, and Pattie Maes. 2024. Memoro: Using Large Language Models to Realize a Concise Interface for Real-Time Memory Augmentation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 450, 18 pages. doi:10.1145/3613904.3642450

A List of TED videos

TED talk Title	URL
Jose Miguel Sokoloff: How we used Christmas lights to fight a war TED	youtube.com/watch?v=0Fi83BHqSMA
How to stay calm when you know you'll be stressed Daniel Levitin TED	youtube.com/watch?v=8JPQjjsBBtC
Why good leaders make you feel safe Simon Sinek TED	youtube.com/watch?v=lmyZMtPVoDo
This could be why you're depressed or anxious Johann Hari TED	youtube.com/watch?v=MB51X-np5FE
What you need to know about CRISPR Ellen Jorgensen	youtube.com/watch?v=1BXYSGepx7Q
Why I love a country that once betrayed me George Takei	youtube.com/watch?v=Le8KBFaPwlc
Wendy Suzuki: The brain-changing benefits of exercise	youtube.com/watch?v=BHY0Fzxc0ZE
What You Can Do to Prevent Alzheimer's Lisa Genova	youtube.com/watch?v=tw6Amr6Jov0
How to build your confidence - and spark it in others Brittany Packnett Cunningham	youtube.com/watch?v=b5ZESp0aIU
Why you should define your fears instead of your goals Tim Ferriss	youtube.com/watch?v=5J6jAC6XXAI
How craving attention makes you less creative Joseph Gordon-Levitt ¹	youtube.com/watch?v=3VTsIju1dLI

B Model details

B.1 System Prompt for Emphasis Proposal module

You are an expert in public speaking. The provided Text will be used for a TED talk. Identify the Emphasis Areas in the Text that should be emphasized when giving the TED talk. The extracted text in the Emphasis Areas should exactly match phrases in the given Text.

¹Unseen video used only for testing

Example Prompt: <s>[INST] «SYS» You are an expert in public speaking. The provided Text will be used for a TED talk. Identify the Emphasis Areas in the Text that should be emphasized when giving the TED talk. The extracted text in the Emphasis Areas should exactly match phrases in the given Text. «/SYS» Text: Cortisol is toxic, and it causes cloudy thinking. So part of the practice of the pre-mortem is to recognize that under stress you're not going to be at your best, and you should put systems in place. And there's perhaps no more stressful a situation than when you're confronted with a medical decision to make. And at some point, all of us are going to be in that position, where we have to make a very important decision about the future of our medical care or that of a loved one, to help them with a decision. And so I want to talk about that. [/INST]

B.2 System Prompt for *Gesture Identification* module

You are an expert in public speaking. You are given a TEXT and a QUERY phrase within the text that needs to be emphasized. Additionally, you are provided with CANDIDATE emphasis areas" Your task is to identify the CANDIDATE emphasis phrase that is most similar to the QUERY phrase in the TEXT in terms of emphasis context. Always choose only from the provided list of CANDIDATES.

Example Prompt: You are an expert in public speaking. You are given a TEXT and a QUERY phrase within the text that needs to be emphasized. Additionally, you are provided with CANDIDATE emphasis areas" Your task is to identify the CANDIDATE emphasis phrase that is most similar to the QUERY phrase in the TEXT in terms of emphasis context. Always choose only from the provided list of CANDIDATES. TEXT: First of all, thank you for your attention. There's nothing quite like being in a room full of people like this, where all of you are giving your attention to me. It's a powerful feeling, to get attention. I'm an actor, so I'm a bit of an expert on, well, nothing, really. But I do know what it feels like to get attention – I've been lucky in my life to get a lot more than my fair share of attention.; QUERY: nothing; CANDIDATES: 1) there's nothing 2) no matter, 3) the only thing. Give only the candidate as output.'