

Multimodal Conversation Structure Understanding

Kent K. Chang Mackenzie Hanh Cramer Anna Ho
Ti Ti Nguyen Yilin Yuan David Bamman

University of California, Berkeley
kentkchang@berkeley.edu

Abstract

While multimodal large language models (LLMs) excel at dialogue, whether they can adequately parse the structure of conversation—conversational roles and threading—remains underexplored. In this work, we introduce a suite of tasks and release TV-MMPC, a new annotated dataset, for multimodal conversation structure understanding. Our evaluation reveals that while all multimodal LLMs outperform our heuristic baseline, even the best-performing model we consider experiences a substantial drop in performance when character identities of the conversation are anonymized. Beyond evaluation, we carry out a sociolinguistic analysis of 350,842 utterances in TVQA. We find that while female characters initiate conversations at rates in proportion to their speaking time, they are 1.2 times more likely than men to be cast as an addressee or side-participant, and the presence of side-participants shifts the conversational register from personal to social.

1 Introduction

Multi-party conversation understanding involves identifying and structuring interactions between multiple speakers and recipients in a face-to-face setting (Akhtiamov et al., 2019; Lerner et al., 2022; Korbar et al., 2024). Such conversations typically involve complex floor-claiming (Allan, 2002; Eisner and Charniak, 2008, 2010; Wang et al., 2020) and turn-taking patterns (Sacks et al., 1974; Hawes et al., 2009; Zhou et al., 2018) marked by various pragmatic cues, which is a challenging problem in NLP (Gu et al., 2022; Tan et al., 2023; Penzo et al., 2024). Modeling these interactions—often described in sociolinguistics as ritualized social encounters (Goffman, 1981)—further enables us to shed light on power dynamics and reveal implicit social hierarchies in conversation.

While multimodal large language models (LLMs) have shown potential in reasoning over

complex videos and dialogues (Song et al., 2023; Teng et al., 2023; Zhang et al., 2023; Bhattacharyya et al., 2023; Sun et al., 2024), whether they can adequately parse the structure of conversation—resolve the reply-to relationship between utterances, or attribute roles like speakers and addressees—remains underexplored (Le Minh et al., 2018; Inoue et al., 2025). To bridge this gap:

- We draw on sociolinguistics (Goffman, 1981, 1983; Goodwin, 1981; Clark and Carlson, 1982; Ng and Bradac, 1993) to devise a framework and introduce a suite of tasks (Fig. 1) for capturing the interactional patterns of interlocutors in the multimodal context.
- To support these tasks, we build on TVQA (Lei et al., 2018, 2020) and release TV-MMPC, a human-annotated dataset of 4,378 annotations for speakers and reply-to relationship between utterances, 5,599 addressees, and 3,412 side-participants, available under a CC BY-NC 4.0 license.¹
- We find that while the best audio-visual model outperforms others, its performance in speaker and addressee recognition drops significantly when participants are anonymized, casting doubt on model reliability.
- In our analysis of 350,842 utterances in TVQA, we find that while female characters initiate conversations at rates proportional to their speaking time, they are more likely to be cast as listeners than speakers, and the presence of side-participants shifts the conversational register from personal to social. This demonstrates the utility of this work in characterizing cultural representation on screen.

¹Code is available at <https://github.com/kentchang/tv-mmpc>; annotations can be downloaded at <https://doi.org/10.7910/DVN/4KUKUL>.

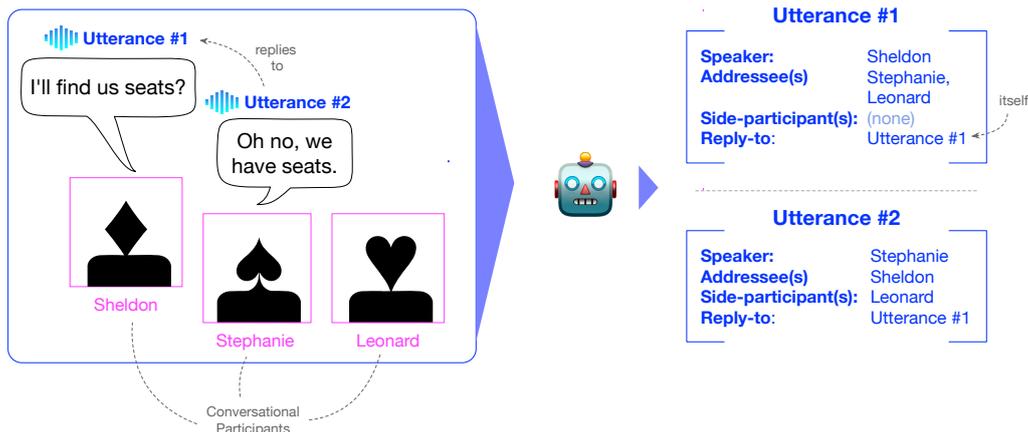


Figure 1: Our proposed structured prediction task for multimodal conversation structure understanding. Grounded in sociolinguistic and conversation analysis (Goffman, 1981, 1983; Ng and Bradac, 1993; Goodwin, 1981; Clark and Carlson, 1982), the task requires predicting, for each utterance in the given clip: the speaker, addressee(s), side-participants, and the utterance it replies to. The example, taken from *The Big Bang Theory* in TVQA (Lei et al., 2018, 2020), illustrates our unified formulation, which treats conversational role attribution and conversation disentanglement as complementary subtasks for modeling the interactional dynamics of dialogue. Further analysis of this example can be found in Appendix §F.3.

	Speakers	Addressees	Side-participants
<i>Per clip</i>			
# Unique	4.18	4.48	2.90
# Avg. Total	21.89	28.00	17.06
<i>Overall</i>			
# Unique	156	168	121
# Total	4,378	5,599	3,412

Table 1: Summary statistics for human-annotated data of 200 randomly sampled clips from TVQA.

2 Data

In constructing our TV-MMPC (Multi-modal Multi-Party Conversation) dataset, we build upon TVQA (Lei et al., 2018, 2020), a large-scale multimodal dataset, widely used for evaluating multimodal models, consisting of 60–90-second video clips from popular TV series.

We devise a data pre-processing pipeline (Fig. 2, detailed in Appendix §A). For human annotation, we randomly sample 200 clips from the following shows in TVQA (50 clips each): *The Big Bang Theory*, *Friends*, *House M. D.*, and *How I Met Your Mother*, which captures diverse conversational dynamics ranging from domestic sitcoms to procedural dramas. To facilitate efficient annotation, we stitch the sampled frames with face captions and use the audio to create a rough approximation of the original clip. Annotations are carried out by co-authors of the paper, who watch the clips on a

	Inter-annotator agreement		
	Pilot	Main	(Δ)
<i>Conversational roles</i>			
Speaker (Acc.)	86.20	97.58	($\uparrow 11.38$)
Addressees (Set F_1)	86.07	92.52	($\uparrow 6.45$)
Side-part. (Set F_1)	82.87	85.43	($\uparrow 2.56$)
<i>Conversational threads</i>			
Linking (F_1)	86.07	97.11	($\uparrow 11.03$)
1–NVI	83.85	92.87	($\uparrow 9.01$)
1–1	77.78	89.25	($\uparrow 11.48$)
EM F_1	35.94	70.56	($\uparrow 34.62$)

Table 2: Annotation quality (pilot vs. main round).

dedicated annotation interface and annotate conversational roles and reply-to for each utterance, following the guidelines (cf. Appendix §F). Statistics for the annotated dataset can be found in Table 1.

We report inter-annotator agreement in Table 2 based on a pilot round (10 clips, 221 lines) and a main round (12 clips, 268 lines), where we observe a clear improvement. Agreement is computed as the average of all pairwise comparisons among the four primary annotators, using the same set of evaluation metrics for experiments (described in Appendix §C).

3 Conversation structure understanding

Our framework for conversation structure understanding is grounded in conversation analysis:² we

²For additional related work, see Appendix §B.

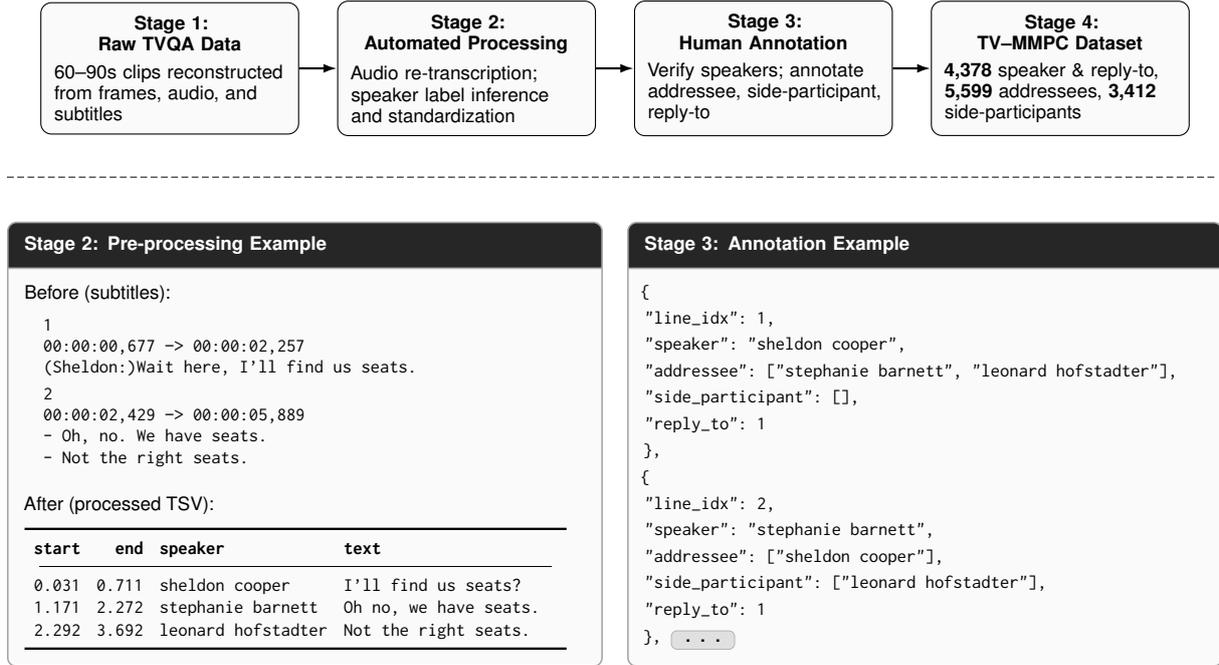


Figure 2: Data creation pipeline. *Top*: four stages from raw TVQA clips to final annotations. *Bottom*: samples from automated preprocessing (Stage 2) and human annotation (Stage 3), drawn from clip ID `s02e09_seg02_clip_04`.

view dialogue not merely as an information exchange, but as a ritualized social encounter (Goffman, 1981) governed by norms of turn-taking and floor-claiming, which are often mediated by non-verbal cues like gaze and posture. To model such conversational dynamics, we introduce the following complementary tasks:

Conversational role attribution. We draw on the standard tasks of speaker and addressee recognition (Jovanovic and op den Akker, 2004; Ouchi and Tsuboi, 2016) and describe a more general task of conversational role attribution. According to Goodwin (1981), conversational roles are observably cast or projected by the *speaker* s_t at time t , regardless of what happens at $t + 1$ (e.g., addressee might not hear or understand the speaker, and that does not impact our annotation). We revise the taxonomy proposed in Clark and Carlson (1982); Clark and Schaefer (1987) by organizing participant roles along three binary dimensions, summarized in Table 3: whether a participant is addressed by the speaker, ratified as a member of the conversational group, and known to be attending or perceiving the utterance. This reformulation clarifies distinctions among roles and highlights their dependency: *addressees* are explicitly addressed, ratified as participants, and known to be listening; *side-participants* are not addressed but still ratified and perceptually engaged. *Bystanders*, by contrast,

Role	addressed	ratified	known
Addressee	+	+	+
Side-participants	–	+	+
Bystanders	–	–	±

Table 3: Participant role matrix.

are neither addressed nor ratified, and may not be known to be present, which, unlike in Clark and Schaefer (1987), includes *overhearers*.

Conversation disentanglement. Here, we want to resolve the conversation structure and detect the start of individual sub-conversations or *conversational threads* (Elsner and Charniak, 2010; Gu et al., 2022). We follow the formulation in Chang et al. (2023a): Given an utterance of interest (UOI) u_t , the task is to identify its *parent utterance* u_p , which is the utterance that u_t directly replies to: $u_t \rightarrow u_p$. In other words, we find the most relevant preceding utterance in the conversational thread; if the utterance marks the start of a new thread, we say it replies to *itself*. The thread is the transitive closure of such pairwise links.

Formally, given an input video clip V , we first extract the sequence of utterances $\mathcal{U} = \{u_1, u_2, \dots, u_{|\mathcal{U}|}\}$, where each utterance $u_i \in \mathcal{U}$ corresponds to a segment of the video defined by a start time $t_{s,i}$ and an end time $t_{e,i}$, $t_{e,i} \leq t_{s,i+1}$. We also identify the set of participants

line	text	speaker	addr.	side-part.	reply-to
11	Have fun.	Leonard	Penny	Sheldon, Amy	11 (SELF)
12	Thanks.	Penny	Leonard	Sheldon, Amy	11
13	Are you ready?	Penny	Amy	Sheldon, Leonard	13 (SELF)
14	Uh-huh.	Amy	Penny	Sheldon, Leonard	13

Table 4: Example of conversational roles and threads.

$\mathcal{P} = \{p_1, p_2, \dots, p_{|\mathcal{P}|}\}$ appearing in the video, which are derived from face recognition or other identity tracking methods applied to V . For each UOI, the goal is to predict its speaker, addressees, side-participants, and the utterance to which it directly replies. A multi-party conversation structure solver \mathcal{G} maps the current utterance u_i and relevant context to these structural elements:

$$\mathcal{G} : (u_i, \mathcal{U}, \mathcal{P}) \mapsto \{u_p, S_i, \mathcal{A}_i, \mathcal{E}_i\}, \quad (1)$$

where *reply-to utterance* $u_p \in \{u_1, \dots, u_{i-1}\}$ is the parent utterance that u_i directly replies to, identified by a reply-to function $R : u_i \mapsto u_p$ such that $p < i$. If u_i initiates a new thread, $u_p = u_i$; *speaker* $S_i \in \mathcal{P}$ is the participant who produced utterance u_i ; *addressees* $\mathcal{A}_i \subseteq \mathcal{P}$ is the set of participants primarily addressed by S_i during u_i ; *side-participants* $\mathcal{E}_i \subseteq \mathcal{P}$ refers to the set of participants actively listening or engaged in the conversation involving u_i , but not directly addressed by S_i ; note that $\mathcal{A}_i \cap \mathcal{E}_i = \emptyset$. In this formulation, the *bystanders* are participants who are neither addressees nor side-participants, but nevertheless belong to the same thread and can thus be determined heuristically.

Example. Consider this scene from *The Big Bang Theory* (season 10, episode 13, clip 6). It features four characters in the apartment living room: Leonard, Sheldon, Penny, and Amy. Penny and Amy are preparing to leave for a trip while Leonard and Sheldon are seeing them off. Our example annotations (see Table 4) for this exchange highlight the dynamic nature of conversational interactions.

In terms of *conversational roles*, in Line 11, Leonard is the *speaker*. He directs his utterance to Penny, making her the *addressee*. Sheldon and Amy, who are present and part of the ratified group but not directly spoken to, are classified as *side-participants*. The most critical moment for role attribution occurs in Line 13: after replying to Leonard, Penny becomes the speaker and pivots

her *attention*: through non-verbal cues like gaze and turning her body, she selects Amy as the new addressee. This action simultaneously relegates Leonard and Sheldon to the role of side-participants for this new exchange.

The reply-to column in the table captures two distinct, interleaved *conversational threads* within this short exchange: In Thread 1 (ll. 11–12), Leonard’s “Have fun” is marked SELF because it initiates a new social action (i.e., a parting pleasantry) rather than responding to prior dialogue. Penny’s response in Line 12, “Thanks,” directly replies to Leonard’s utterance, and is therefore linked to its parent, Line 11. This completes the first thread. In Thread 2 (ll. 13–14), Penny’s question, “Are you ready?”, is also marked SELF. Although it follows her reply to Leonard, it does not respond to it. Instead, it starts a new sub-conversation focused on the logistics of departure, shifting the group’s attention. Amy’s “Uh-huh” in Line 14 is a direct reply to this question, linking to its parent, Line 13, and completing the second thread.

Further details and examples can be found in the annotation guidelines in Appendix §F.3.

4 Experiments

In this section, we evaluate six popular audio-visual models and vision–language models on TV-MMPC (§4.2), using the standard set of metrics (described in Appendix C). We analyze the impact of different modalities (§4.3), explore resource-efficient approaches for task adaptation (§4.4), and discuss the effects of anonymization (§4.5).

4.1 Heuristic baseline

To establish a baseline, we use Whisper (Radford et al., 2022) and pyannote (Bredin et al., 2020) to obtain word-level timestamps and speaker labels. Since pyannote only provides generic speaker labels, we determine the speaker by aggregating the frequency of each face appearing at the word level, then assigning the speaker role to the face that appears most frequently within each sentence. For each utterance u_i , we consider faces present in the context window $[i-1, i]$. Among these, the most frequently occurring face *excluding* the speaker is designated as the *addressee*, while all remaining faces are labeled as *side-participants*. For the linking task, the previous utterance is treated as the parent utterance (i.e., u_i always replies to u_{i-1}).

	Conversational roles			Conversational threads			
	Speaker	Addressees	Side-part.	Linking	Clustering		
	Acc.	Set F ₁	Set F ₁	F ₁	1–NVI	1–1	EM F ₁
Heuristic baseline	34.67 [29.13–40.29]	19.49 [16.15–23.07]	36.98 [31.97–42.14]	92.67 [90.61–94.60]	83.34 [79.24–87.41]	76.20 [70.08–82.25]	31.93 [19.52–45.92]
Vision–language models (image and text)							
LLaMA 4 Scout	49.23 [43.09–55.33]	38.47 [32.94–44.15]	42.70 [35.04–50.76]	87.69 [85.07–90.19]	82.60 [78.51–86.63]	76.13 [70.19–81.96]	31.01 [18.80–43.96]
GPT-4.1 mini	55.76 [49.04–62.32]	46.28 [40.38–52.21]	53.72 [46.36–61.32]	81.40 [78.42–84.25]	78.61 [74.68–82.50]	75.82 [70.90–80.64]	21.68 [11.89–32.49]
o4-mini	53.39 [46.72–60.12]	49.37 [43.69–55.09]	56.88 [48.99–64.71]	84.65 [81.50–87.70]	78.15 [74.13–82.17]	74.98 [69.89–80.02]	24.87 [14.79–35.83]
Gemini 2.0 Flash	51.78 [45.45–58.08]	43.40 [37.94–48.93]	53.87 [46.45–61.44]	85.46 [82.67–88.14]	80.49 [76.52–84.43]	77.18 [72.08–82.21]	24.37 [13.65–36.26]
Audio-visual LLMs (video with audio, image, and text)							
Qwen 2.5-Omni 7B	32.24 [27.04–37.58]	23.03 [18.45–27.91]	25.94 [20.22–32.12]	60.36 [56.06–64.57]	66.66 [62.75–70.67]	62.19 [57.51–66.99]	10.03 [4.53–16.57]
Reka-Flash	41.48 [36.60–46.49]	35.71 [30.77–40.69]	15.56 [11.99–19.52]	86.67 [84.38–88.86]	81.87 [77.91–85.73]	74.41 [68.40–80.32]	27.00 [15.50–39.50]
Gemini 2.0 Flash	78.60 [74.21–82.61]	68.11 [63.29–72.86]	57.68 [51.20–64.18]	89.51 [87.63–91.32]	85.21 [81.73–88.58]	80.33 [75.32–85.15]	34.60 [23.73–46.01]

Table 5: LLM zero-shot performance on conversational roles and threads prediction. Metrics are reported with 95% confidence intervals (CIs) from 10,000 bootstrap resamples.

4.2 LLM zero-shot performance

We evaluate the zero-shot performance of two types of models on our tasks: vision–language models (LLaMA 4 Scout, GPT-4.1, o4-mini, and Gemini 2.0 Flash)³ and audio-visual LLMs, including Reka-Flash (Reka Team et al., 2024), Qwen 2.5 Omni (Xu et al., 2025), and Gemini 2.0 Flash with reconstructed clips (cf. §2). We select both open-weight and closed, general-purpose models (i.e., not trained only for captioning or segmentation) to represent distinct approaches to multimodal understanding: vision–language models process sampled frames and text but lack native audio processing, while audio-visual models can process signals from text, audio, and vision.

In our setup, the multimodal LLM \mathcal{F} performs a structured prediction o_i for each $u_i \in \mathcal{U}$:

$$o_i = (o_{s_i}, o_{a_i}, o_{e_i}, o_{r_i}) \sim \mathcal{F}(C_i, \mathcal{P}_{V_i}, Q_{u_i}) \quad (2)$$

where C_i is the context of u_i , $\mathcal{P}_{V_i} = \{p_1, p_2, \dots, p_{|P|}\}$ is the set of identified participants associated with C_i , which corresponds to the cast list of the episode from which the clip is taken. Participants include the text of character names

³<https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E>, <https://openai.com/index/gpt-4-1/>, <https://openai.com/index/introducing-o3-and-o4-mini/>, <https://cloud.google.com/vertex-ai/generative-ai/docs/models/gemini/2-0-flash>

and, where available, the best face crop between $t_{s,i}$ and $t_{e,i}$. Q_{u_i} is the question associated with u_i : it includes the line index i , the start and end timestamps $t_{s,i}$ and $t_{e,i}$, and the question text itself, which prompts the model for conversation structure analysis. The output o_i contains the index of the reply-to utterance $o_{r_i} \leq i$, identities of the predicted speaker s_i , addressees a_i , side-participants e_i ; note $o_{a_i}, o_{e_i} \subseteq \mathcal{P}_{V_i}$. Given the design of models, we need to structure the overall prompt differently for each class of models, particularly what context information C_i to put in the model:

Vision–language models. The standard setup in video understanding often involves sampling visual frames from a video while discarding its audio. Here, we use a modified context $C_i = \mathcal{U}_i^*$, and $\mathcal{U}_i = \{(v_j^*, u_j, f_j)\}_{j=1}^{|\mathcal{U}|}$: v_j^* is the final frame sampled from V_i in the interval $[t_{s,j}, t_{e,j}]$, interleaving with the words of u_j ; this is further augmented with audio-derived speaker features f extracted using Librosa (McFee et al., 2025): speech rate (slower or faster), pitch range (wide or narrow), pitch mean (lower or higher), and spectral centroid (darker or brighter), which collectively aim to capture the acoustic profile of the speaker.

Audio-visual LLMs. In this setup, $C_i = (V_i, \mathcal{U}_i)$, where V_i is the input video clip containing utterance u_i , i.e., its audio-visual context. It is the same

Modality	Roles			Threads	
	Spkr.	Addr.	Side-part.	Linking	Clustering
	Acc.	Set F ₁	Set F ₁	F ₁	EM F ₁
Text	36.40	36.46	40.28	90.67	30.73
Text + Visual	51.78	43.40	53.87	85.46	24.37
Text + Visual + Audio	78.60	68.11	57.68	89.51	34.60

Table 6: Modality ablation of Gemini 2.0 Flash.

clip that annotators watch (described in §2). The transcription \mathcal{U}_i consists of utterances spoken in V_i ($u_i \in \mathcal{U}_i$), each represented as a sequence of words with associated start and end timestamps: $\mathcal{U}_i = \{(u_j, t_{s,j}, t_{e,j})\}_{j=1}^{|\mathcal{U}_i|}$, where each element corresponds to j -th utterance in the clip, and the utterance-level start and end times $t_{s,j}$ and $t_{e,j}$.

The results are reported in Table 5.⁴ All vision-language models we consider perform significantly better than the heuristic baseline across all metrics, other than on the reply-to task. Another exception is LLaMA 4 Scout, which falls short on side-participant attribution, although it achieves better results on utterance linking than other models, with a statistically significant margin over GPT-4.1 mini. Among the audio-visual language models we consider, Gemini 2.0 Flash outperforms others by a considerable margin; Qwen 2.5-Omni 7B and Reka-Flash do not show statistically significant improvements over the heuristic baseline on most metrics. These results suggest the need for further research into multimodal models, particularly in developing methods that more effectively leverage and integrate different input modalities.⁵

4.3 Impact of modalities

Ablation. To quantify the specific contribution of individual modalities, we perform an ablation study on our best-performing model, Gemini 2.0 Flash (Table 6). While adding visual information improves role attribution over text alone, the full audio-visual context is required for high accuracy in speaker and addressee identification.

Error analysis. To better understand the factors driving performance variation in conversational role attribution, we examine the correlations between interpretable, clip-level features (across textual, acoustic, and visual modalities) and F₁ scores. We employ Spearman’s ρ to assess monotonic relationships: a positive ρ indicates that higher feature

values consistently co-occur with higher performance ranks.

Fig. 3 displays the signed explained rank-variance: $R^2 = \text{sign}(\rho) \cdot \rho^2 \times 100$, where solid bars denote statistical significance ($p < 0.05$). Unsurprisingly, the number of participants has the strongest negative impact, reducing F₁ scores across all roles (3–8% variance). Clip length shows negligible influence. For textual features, we consider mean sentence length (measured by tokens separated by whitespace) and type–token ratio; neither shows significant correlation, suggesting lexical diversity alone does not drive performance. Acoustic features (pitch, spectral centroid) positively correlate with side-participant attribution ($\rho = +0.17$ and $+0.15$), which indicates that speech clarity aids in identifying peripheral roles. Visual features reveal role-specific dynamics: max face coverage positively correlates with speaker attribution ($\rho = +0.15$), suggesting that accurate face detection benefits this task, while raw face counts negatively correlate with side-participant detection ($\rho = -0.23$), thereby confirming that visual clutter degrades performance.

Overall, performance relies heavily on acoustic and visual fidelity rather than textual signals. We suggest that future work look beyond text to ground conversation structure in the multimodal context, and for future audio-visual LLMs, cross-modal cues (e.g., face–participant identity alignment, prosodic salience, etc.) have the potential to enable better parsing of conversation structure.

4.4 Resource-efficient supervised fine-tuning

Given the increasing popularity of audio-visual LLMs and the growing number of open-source implementations in this space, we turn to explore task-specific adaptation through supervised fine-tuning (SFT). In this subsection, we describe our approach using Low-Rank Adaptation (LoRA) to fine-tune Qwen 2.5-Omni 7B within the LLaMA Factory framework (Zheng et al., 2024). We employ a leave-one-show-out evaluation strategy: training the model on all *but* one show, and testing it on the held-out show. This process is repeated for each show in the dataset. The results are reported in Table 5, which we compare the performance before and after LoRA SFT on the same TV series. After fine-tuning, despite the constraint in resources, the model has a statistically significant boost on both addressee and side-participant recognition.

⁴Additional experimental details are in Appendix §D.

⁵Error analysis for Gemini can be found in Appendix §E.

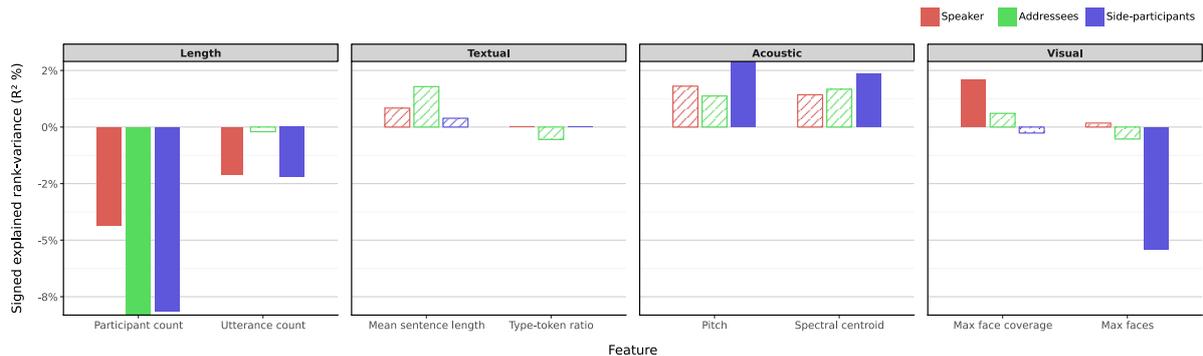


Figure 3: Signed explained rank variance from Spearman’s ρ between clip-level features and F_1 scores for individual conversational roles. Bars indicate direction and magnitude of correlation: solid ones are significant ($p < 0.05$).

	Qwen 2.5-Omni 7B		
	Zero-shot	LoRA SFT	(Δ)
<i>Conversational roles</i>			
Speaker (Acc.)	32.24	41.14	($\uparrow 8.90$)
Addressees (Set F_1)	23.03	34.90	($\uparrow 11.87$)
Side-part. (Set F_1)	25.94	57.69	($\uparrow 31.74$)
<i>Conversational threads</i>			
Linking (F_1)	60.36	60.40	($\uparrow 0.04$)
1-NVI	66.66	72.82	($\uparrow 6.16$)
1-1	62.19	71.20	($\uparrow 9.01$)
EM F_1	10.03	32.00	($\uparrow 21.97$)

Table 7: Performance of LoRA SFT of Qwen 2.5-Omni 7B compared to zero-shot. Boldface indicates statistically significant (non-overlapping 95% CIs) differences.

4.5 Effects of anonymization

Prior work has shown that anonymizing speakers can influence downstream performance in text-based dialogue understanding (Chen et al., 2022; Chang et al., 2024). Motivated by this, we examine whether anonymizing conversational participants has a similar impact in our multimodal setting. To do so, for each clip, we create a mapping based on its individual episode cast list: we assign a random letter to each character within a clip (e.g., “Sheldon Cooper” becomes Character C). We then apply the same frame captioning procedure described in §2, but we replace those actual character names with their anonymized labels.

The difference in performance is presented in Table 8. We observe significant drops in the performance for speaker and addressee recognition, by 64.92 and 52.37 points, respectively. In light of this, future work might assess the degree to which the performance drop can be attributed to the phenomenon of memorization in the audio-visual LLMs, which is relevant in the context of vision-

	Gemini 2.0 Flash (audio-visual)		
	Original	Anonymized	(Δ)
<i>Conversational roles</i>			
Speaker (Acc.)	78.60	13.68	($\downarrow 64.92$)
Addressees (Set F_1)	68.11	15.73	($\downarrow 52.37$)
Side-part. (Set F_1)	57.68	46.06	($\downarrow 11.62$)
<i>Conversational threads</i>			
Linking (F_1)	89.51	88.44	($\downarrow 1.08$)
1-NVI	85.21	84.56	($\downarrow 0.65$)
1-1	80.33	80.14	($\downarrow 0.18$)
EM F_1	34.60	33.49	($\downarrow 1.11$)

Table 8: Impact of anonymization on Gemini (bold indicates non-overlapping 95% CIs).

language models (Jayaraman et al., 2024) and ASR models (Wang et al., 2024). This is more than a straightforward case of test set contamination (where test data overlaps with training data), but a subtler one: models, when solving a task, might leverage their parametric knowledge, among other things, acquired from exposure to cultural artifacts and information about them, potentially across different modalities. Such memorization raises further questions about model reliability and cultural representation, which has been explored in the context of literary texts (Chang et al., 2023b; Walsh et al., 2024) and can be relevant to audio-visual LLMs.

5 Analysis

We posit that a robust structural understanding of conversation serves as the necessary scaffolding for analyzing power dynamics and cultural representation. To demonstrate the affordance of this work, we apply our best-performing model to analyze the clips in TVQA we do not annotate, which comprises 350,842 utterances across four TV series, and present three case studies below.

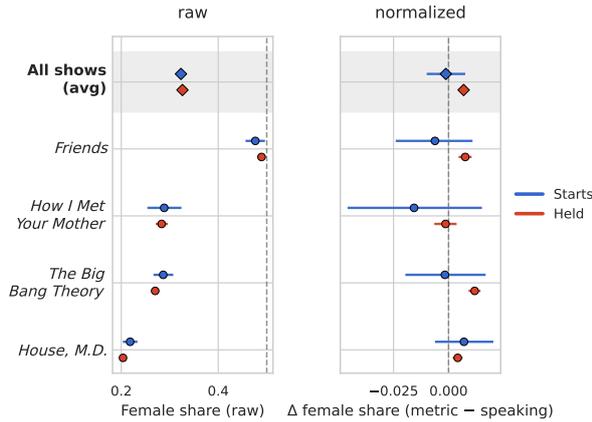


Figure 4: Female share of starting vs. holding conversational threads, aggregated by show and overall. *Left*: raw percentage of threads started or held by women. *Right*: normalized difference (Δ) between female share of metric and speaking time within each clip. Points are bootstrapped means (95% CIs).

5.1 Gender and floor claiming

How is the act of claiming and holding conversational floor distributed between genders? To explore this, we measure two aspects of conversational agency across 9,780 conversation-initiating utterances with known speaker gender: who *starts* new threads mid-clip, and who *holds* the thread by being the recipient of a reply. Given the well-documented disparity in speaking time in media (Lauzen, 2019; O’Meara, 2016), we analyze both raw counts and per-clip metrics normalized by lifting each character’s share of speaking time.

Fig. 4 presents these dynamics. In raw terms, the female share of starting threads is low (31.8%), confirming that men dominate this action in absolute numbers. However, the right panel (the normalized difference) reveals a more nuanced picture. For thread *starters*, the average delta is not statistically different from zero (mean $\Delta = -0.12\%$, $p = 0.79$), indicating that women initiate threads at a rate proportional to their speaking time. In contrast, for thread *holders*, the average delta is small but statistically significant and positive (mean $\Delta = +0.69\%$, $p < 0.001$). This suggests a dynamic where female characters not only exhibit agency in starting conversations but also receive slightly more conversational uptake in the form of replies than their speaking time would predict.

5.2 Gender and conversational roles

Beyond who starts or holds a thread, how are men and women positioned within the broader structure

of conversations? We analyze the distribution of conversational roles across over 929,000 instances in our dataset. Fig. 5 breaks down this relationship. Panel (a), showing $P(\text{gender}|\text{role})$, confirms that men are the numerical majority in every role. However, Panel (b), which shows $P(\text{role}|\text{gender})$, reveals a structural disparity: Men are more likely than women to occupy the *speaker* role (33.0% of their roles vs. 28.7% for women). Conversely, this pattern inverts for listening roles: women are more likely than men to be cast as an *addressee* (39.4% vs. 38.7%) and a *side-participant* (31.9% vs. 28.3%). This skewed distribution positions male characters as primary speakers and female characters more frequently as the audience. Panel (c) formalizes this disparity using a multinomial logistic regression. With speaker as the reference category, the model estimates the log-odds of a person being an addressee or side-participant, controlling for gender and show-level fixed effects:

$$\log \frac{P(\text{role} = j)}{P(\text{role} = \text{speaker})} = \beta_{j,0} + \beta_{j,\text{female}} \cdot \mathbb{I}(\text{female}) + \sum_{s=1}^{S-1} \gamma_{j,s} \cdot \mathbb{I}(\text{show}_s), \quad (3)$$

where $j \in \{\text{addressee}, \text{side-participant}\}$ and $\mathbb{I}(\cdot)$ is an indicator function, and the odds ratio for being female is given by $\exp(\beta_{j,\text{female}})$. These odds ratios are 1.19 for addressee ($p < 0.001$, 95% CI [1.08, 1.39]) and 1.20 for side-participant ($p < 0.001$, 95% CI [1.02, 1.64]). This suggests that, controlling for show-specific effects, women are significantly more likely than men to be cast as a listening role than as the primary speaker.

5.3 Audience design

Speakers dynamically adapt their linguistic style based on their audience; to investigate whether such *audience design* (Bell, 1984; Giles et al., 1991) occurs, we identify distinctive lexical items using a stratified analysis based on the weighted log-odds ratio with an informative Dirichlet prior (Monroe et al., 2017). First, we place an informative Dirichlet prior on each term t , where the prior count $\alpha_t = C^* \cdot p_t$ is weighted by the term’s background frequency p_t . Second, the prior strength C^* is set via Empirical Bayes calibration: we permute group labels within each show and select the C^* from a grid that results in the null z-scores having a standard deviation closest to 1. Finally, with the calibrated prior, we compute a

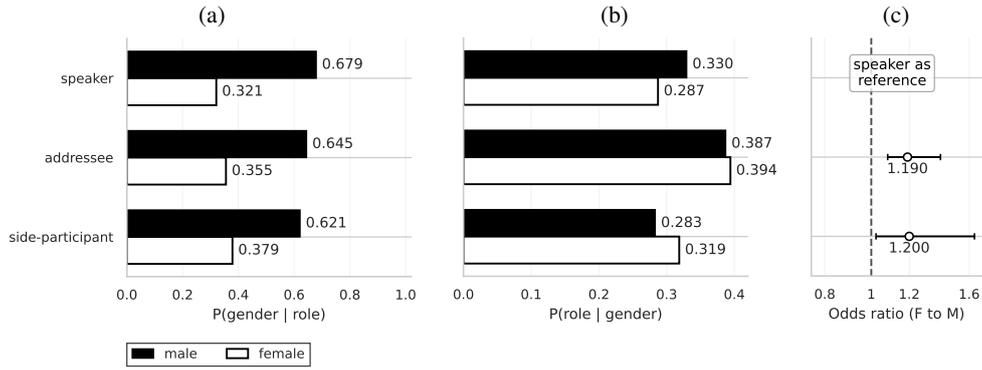


Figure 5: The relationship between gender and conversational roles. (a) $P(\text{gender} | \text{role})$: the probability that a participant in a given role is female or male; (b) $P(\text{role} | \text{gender})$: the probability distribution of roles for each gender; (c) odds ratios from a multinomial regression, with speaker as the reference.

No side-participants vs. side-participants present			
more private		more public	
term	score	term	score
i've	2.54	happened	-2.41
liver	2.18	ball	-2.17
job	2.13	lovely	-2.15
dum	2.02	possible	-2.02
went	2.02	list	-2.00
idiot	1.96	next	-1.99
without	1.87	princeton	-1.92
straight	1.81	fluid	-1.86
hire	1.80	how's	-1.82
first	1.79	ice	-1.74

Table 9: Distinctive terms illustrating a shift from private to public conversational register. The table lists the top 10 terms most associated with whether side-participants are present, ranked by their final z-scored log-odds ratios. Positive scores indicate terms characteristic of private, dyadic speech, while negative ones indicate terms characteristic of public, group speech.

per-show z-score ($\hat{\zeta}_{ts}$) for each term. This score is the log-odds ratio, $\hat{\delta}_{ts}$, scaled by its variance, $\sigma^2(\hat{\delta}_{ts}) \approx (y_{tsa} + \alpha_t)^{-1} + (y_{tsb} + \alpha_t)^{-1}$. These per-show scores are then aggregated across all k shows using equal-weight Stouffer's method to produce a final, overall z-score, $Z_t = 1/\sqrt{k} \sum_{s=1}^k \hat{\zeta}_{ts}$.

Table 9 highlights the results of this analysis by comparing speech in private, dyadic settings (i.e., no side-participants) with speech in more public, group settings (i.e., side-participants present). The linguistic patterns reveal a clear shift in register: language in private contexts trends towards the transactional and personal, featuring first-person experiences (i've), direct tasks (job), and even confrontational terms (idiot). In contrast, language used when side-participants are present is more social and performative, including politeness

markers (lovely) and terms that manage group conversation (happened, how's). Taken together, this lexical analysis represents an instance of audience design, where speakers would alter their vocabulary in response to their interlocutors.

6 Conclusion

Conversations are usually structured by roles—who is speaking, who's being addressed, and who's listening—and unfold in threads that break with changes in speaker floor or topical focus. In this work, we address the limited evaluation of multimodal LLMs on conversation structure by introducing a unifying framework for conversational role attribution and thread disentanglement. To support this work, we present TV-MMPC, a human-annotated dataset built on TVQA, with 4,378 speaker and reply-to annotations, 5,599 addressee labels, and 3,412 side-participant tags.

Our experiments reveal that while Gemini 2.0 Flash outperforms baselines by leveraging audiovisual context, its performance drops significantly under anonymization, opening new avenues for future work. To demonstrate the affordance of this work, we carry out a sociolinguistic analysis and uncover a communicative hierarchy where female characters are 1.2 times more likely to be cast as listeners and observe that side-participants shift dialogue registers from personal to performative.

We hope this work can enable future work to improve multimodal LLMs and provide more evidence of how social roles and power dynamics are constructed—and reinforced—in cultural artifacts. Code is available at <https://github.com/ken-tchang/tv-mmpc>; annotations can be downloaded at <https://doi.org/10.7910/DVN/4KUKUL>.

Limitations

This work introduces a complex set of tasks that required developing detailed annotation guidelines and extensive annotators, and ensuring quality entails a time-consuming and labor-intensive process. We leave large-scale crowdsourcing or other approaches to scale up annotation, which this work enables along with the annotation guidelines, for future work. To maximize representativeness, we sample clips from TVQA entirely at random. This is not exhaustive of the original dataset and may overlook certain conversational patterns.

While the shows in our data feature diverse casts, the conversational dynamics are scripted and directed primarily for American audiences, which inherently privileges Western communicative norms. Cues discussed in this work (e.g., direct eye contact to determine addressees) are also signals that are culturally specific.

We rely extensively on metadata from TMDb and IMDb for actor and character identification; while we believe they are generally reliable and accurate, these sources can include uncredited characters or omit minor roles. The effectiveness of our TVQA processing pipeline is limited by the capabilities of ASR and face recognition models, which is beyond the scope of this work.

Ethical considerations

The dataset presented in this paper is derived from the existing TVQA dataset TVQA (Lei et al., 2018, 2020), which consist of video clips from TV shows and manually created QA pairs. Since the source dataset has already undergone prior curation and is widely used in the research community without known misuse concerns, we believe it poses low risk for dual-use or misuse. No personally identifiable information or sensitive private data is included in either the original or our derived dataset.

Acknowledgments

The research reported in this article was supported by resources provided by the Google Gemma Academic Program.

References

- AI@Meta. 2025. Llama 4 model card.
- Oleg Akhtiamov, Ingo Siegert, Alexey Karpov, and Wolfgang Minker. 2019. Cross-corpus data augmentation for acoustic addressee detection. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 274–283, Stockholm, Sweden. Association for Computational Linguistics.
- James Allan. 2002. [Introduction to Topic Detection and Tracking](#). In James Allan, editor, *Topic Detection and Tracking: Event-based Information Organization*, pages 1–16. Springer US, Boston, MA.
- David Bamman, Rachael Samberg, Richard Jean So, and Naitian Zhou. 2024. Measuring diversity in hollywood through the large-scale computational analysis of film. *Proceedings of the National Academy of Sciences of the United States of America*, 121(46):e2409770121.
- Allan Bell. 1984. Language style as audience design. *Language in society*, 13(2):145–204.
- Aanisha Bhattacharyya, Yaman K Singla, Balaji Krishnamurthy, Rajiv Ratn Shah, and Changyou Chen. 2023. A video is worth 4096 tokens: Verbalize videos to understand them in zero shot. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9822–9839, Singapore. Association for Computational Linguistics.
- Herve Bredin and Guilhem Gelly. 2016. Improving speaker diarization of TV series using talking-face detection and clustering. In *Proceedings of the 24th ACM International Conference on Multimedia*.
- Herve Bredin, Ruiqing Yin, Juan Manuel Coria, Gregory Gelly, Pavel Korshunov, Marvin Lavechin, Diego Fustes, Hadrien Titeux, Wassim Bouaziz, and Marie-Philippe Gill. 2020. Pyannote.audio: Neural building blocks for speaker diarization. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE.
- Kent K Chang, Danica Chen, and David Bamman. 2023a. Dramatic conversation disentanglement. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4020–4046, Toronto, Canada. Association for Computational Linguistics.
- Kent K Chang, Mackenzie Cramer, Sandeep Soni, and David Bamman. 2023b. Speak, memory: An archaeology of books known to ChatGPT/GPT-4. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7312–7327, Singapore. Association for Computational Linguistics.
- Kent K Chang, Anna Ho, and David Bamman. 2024. Subversive characters and stereotyping readers: Characterizing queer relationalities with dialogue-based relation extraction. In *Proceedings of the Computational Humanities Research Conference 2024*, Aarhus, Denmark. CEUR Workshop Proceedings.

- Mingda Chen, Zewei Chu, Sam Wiseman, and Kevin Gimpel. 2022. SummScreen: A Dataset for Abstractive Screenplay Summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8602–8615, Dublin, Ireland. Association for Computational Linguistics.
- Herbert H Clark and Thomas B Carlson. 1982. Hearers and Speech Acts. *Language*, 58(2):332–373.
- Herbert H Clark and Edward F Schaefer. 1987. [Concealing one’s meaning from overhearers](#). *Journal of memory and language*, 26(2):209–225.
- Tri Dao. 2024. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations (ICLR)*.
- Micha Elsner and Eugene Charniak. 2008. [You Talking to Me? A Corpus and Algorithm for Conversation Disentanglement](#). In *Proceedings of ACL-08: HLT*, pages 834–842, Columbus, Ohio. Association for Computational Linguistics.
- Micha Elsner and Eugene Charniak. 2010. [Disentangling Chat](#). *Computational Linguistics*, 36(3):389–409.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, and 1331 others. 2023. Gemini: A family of highly capable multimodal models. *arXiv [cs.CL]*.
- Howard Giles, Nikolas Coupland, and Justine Coupland. 1991. Accommodation theory: Communication, context, and consequence. In Howard Giles, Justine Coupland, and Nikolas Coupland, editors, *Contexts of Accommodation*, pages 1–68. Cambridge University Press, Cambridge.
- Erving Goffman. 1981. *Forms of Talk*. University of Pennsylvania Press.
- Erving Goffman. 1983. The Interaction Order: American Sociological Association, 1982 Presidential Address. *American sociological review*, 48(1):1–17.
- Charles Goodwin. 1981. *Conversational organization: Interaction between speakers and hearers*. Academic Press.
- Jia-Chen Gu, Zhenhua Ling, Quan Liu, Cong Liu, and Guoping Hu. 2023. GIFT: Graph-induced fine-tuning for multi-party conversation understanding. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11645–11658, Toronto, Canada. Association for Computational Linguistics.
- Jia-Chen Gu, Chongyang Tao, and Zhen-Hua Ling. 2022. Who says what to whom: A survey of multi-party conversations. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, California. International Joint Conferences on Artificial Intelligence Organization.
- Jia-Chen Gu, Chongyang Tao, Zhenhua Ling, Can Xu, Xiubo Geng, and Daxin Jiang. 2021. MPC-BERT: A Pre-Trained Language Model for Multi-Party Conversation Understanding. 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 3682–3692, Online. Association for Computational Linguistics.
- Timothy Hawes, Jimmy Lin, and Ps Resnik. 2009. [Elements of a Computational Model for Multi-Party Discourse: The Turn-Taking Behavior of Supreme Court Justices](#). *JASIST*, 60:1607–1615.
- Qingqiu Huang, Yu Xiong, Anyi Rao, Jiaze Wang, and Dahua Lin. 2020. MovieNet: A holistic dataset for movie understanding. In *Computer Vision – ECCV 2020*, pages 709–727. Springer International Publishing.
- Koji Inoue, Divesh Lala, Mikey Elmers, Keiko Ochi, and Tatsuya Kawahara. 2025. An LLM benchmark for addressee recognition in multi-modal multi-party dialogue. *arXiv [cs.CL]*.
- Md Mohaiminul Islam and Gedas Bertasius. 2022. Long movie clip classification with state-space video models. *arXiv [cs.CV]*.
- Bargav Jayaraman, Chuan Guo, and Kamalika Chaudhuri. 2024. Déjà vu memorization in vision–language models. In *Advances in Neural Information Processing Systems*, volume 37, pages 50722–50749. Curran Associates, Inc.
- Natasa Jovanovic and Rieks op den Akker. 2004. Towards Automatic Addressee Identification in Multi-party Dialogues. In *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue at HLT-NAACL 2004*, pages 89–92, Cambridge, Massachusetts, USA. Association for Computational Linguistics.
- Jiho Kim, Woosog Chay, Hyeonji Hwang, Daeun Kyung, Hyunseung Chung, Eunbyeol Cho, Yohan Jo, and Edward Choi. 2024. DialSim: A real-time simulator for evaluating long-term dialogue understanding of conversational agents. *arXiv [cs.CL]*.
- Bruno Korbar, Jaesung Huh, and Andrew Zisserman. 2024. Look, listen and recognise: Character-aware audio-visual subtitling. *arXiv [cs.CV]*.
- Jonathan K Kummerfeld, Sai R Gouravajhala, Joseph J Peper, Vignesh Athreya, Chulaka Gunasekara, Jatin Ganhotra, Siva Sankalp Patel, Lazaros C Polymenakos, and Walter Lasecki. 2019. [A Large-Scale Corpus for Conversation Disentanglement](#).

- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with PagedAttention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Martha M Lauzen. 2019. [It’s a man’s \(celluloid\) world: Portrayals of female characters in the top grossing films of 2018](#). *Center for the Study of Women in Television and Film, San Diego State University*, pages 411–428.
- Thao Le Minh, N Shimizu, Takashi Miyazaki, and K Shinoda. 2018. Deep learning based multi-modal addressee recognition in visual scenes with utterances. *International Joint Conference on Artificial Intelligence*, abs/1809.04288.
- Jina Lee and Stacy Marsella. 2011. Modeling side participants and bystanders: The importance of being a laugh track. In *Intelligent Virtual Agents*, Lecture notes in computer science, pages 240–247. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara Berg. 2018. TVQA: Localized, compositional video question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1369–1379, Brussels, Belgium. Association for Computational Linguistics.
- Jie Lei, Licheng Yu, Tamara Berg, and Mohit Bansal. 2020. TVQA+: Spatio-temporal grounding for video question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8211–8225, Online. Association for Computational Linguistics.
- Paul Lerner, Juliette Bergoënd, Camille Guinaudeau, H Bredin, Benjamin Maurice, Sharleyne Lefevre, Martin Bouteiller, Aman Berhe, Léo Galmant, Ruiqing Yin, and C Barras. 2022. Bazinga! a dataset for multi-party dialogues structuring. *International Conference on Language Resources and Evaluation*.
- Bo-Ru Lu, Yushi Hu, Hao Cheng, Noah A Smith, and Mari Ostendorf. 2022. Unsupervised learning of hierarchical conversation structure. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5657–5670, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Kartikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. 2023. EgoSchema: A diagnostic benchmark for very long-form video language understanding. *arXiv [cs.CV]*.
- Brian McFee, Matt McVicar, Daniel Faronbi, Iran Roman, Matan Gover, Stefan Balke, Scott Seyfarth, Ayoub Malek, Colin Raffel, Vincent Lostanlen, Benjamin van Niekirk, Dana Lee, Frank Cwitkowitz, Frank Zalkow, Oriol Nieto, Dan Ellis, Jack Mason, Kyunghyun Lee, Bea Steers, and 77 others. 2025. *librosa/librosa*: 0.11.0.
- Robert McKee. 2016. *Dialogue: The art of verbal action for page, stage, and screen*. Hachette UK.
- Burt L Monroe, Michael P Colaresi, and Kevin M Quinn. 2017. Fightin’ words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political analysis: an annual publication of the Methodology Section of the American Political Science Association*, 16(4):372–403.
- Sik H Ng and James J Bradac. 1993. *Power in Language: Verbal Communication and Social Influence*. SAGE Publications.
- Jennifer O’Meara. 2016. [What “The Bechdel Test” doesn’t tell us: examining women’s verbal and vocal \(dis\)empowerment in cinema](#). *Feminist Media Studies*, 16(6):1120–1123.
- OpenAI, Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, A J Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, and 400 others. 2024. GPT-4o system card. *arXiv [cs.CL]*.
- Hiroki Ouchi and Yuta Tsuboi. 2016. Addressee and Response Selection for Multi-Party Conversation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2133–2143, Austin, Texas. Association for Computational Linguistics.
- Nicolò Penzo, Maryam Sajedinia, Bruno Lepri, Sara Tonelli, and Marco Guerini. 2024. Do LLMs suffer from multi-party hangover? a diagnostic approach to addressee recognition and response selection in conversations. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11210–11233, Miami, Florida, USA. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. Robust speech recognition via large-scale weak supervision. *arXiv [eess.AS]*.
- Reka Team, Aitor Ormazabal, Che Zheng, Cyprien de Masson d’Autume, Dani Yogatama, Deyu Fu, Donovan Ong, Eric Chen, Eugenie Lamprecht, Hai Pham, Isaac Ong, Kaloyan Aleksiev, Lei Li, Matthew Henderson, Max Bain, Mikel Artetxe, Nishant Relan, Piotr Padlewski, Qi Liu, and 7 others. 2024. Reka core, flash, and edge: A series of powerful multi-modal language models. *arXiv [cs.CL]*.
- Joseph Roth, Sourish Chaudhuri, Ondrej Klejch, Radhika Marvin, Andrew Gallagher, Liat Kaver, Sharadh Ramaswamy, Arkadiusz Stopczynski, Cordelia Schmid, Zhonghua Xi, and Caroline Pantofaru. 2020. Ava active speaker: An audio-visual dataset for active speaker detection. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4492–4496. IEEE.

- Harvey Sacks, Emanuel A Schegloff, and Gail Jefferson. 1974. [A Simplest Systematics for the Organization of Turn-Taking for Conversation](#). *Language*, 50(4):696–735.
- Rudrabha Sharma and Shrikanth Narayanan. 2023. Audio-visual activity guided cross-modal identity association for active speaker detection. *IEEE Open Journal of Signal Processing*, 4:225–232.
- Enxin Song, Wenhao Chai, Guan hong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe Chi, Xun Guo, Tian Ye, Yanting Zhang, Yan Lu, Jenq-Neng Hwang, and Gaoang Wang. 2023. MovieChat: From dense token to sparse memory for long video understanding. *arXiv.org*.
- Guangzhi Sun, Wenyi Yu, Changli Tang, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, Yuxuan Wang, and Chao Zhang. 2024. video-SALMONN: speech-enhanced audio-visual large language models. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*, Vienna, Austria. JMLR.org.
- Chao-Hong Tan, Jia-Chen Gu, and Zhen-Hua Ling. 2023. Is ChatGPT a good multi-party conversation solver? In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4905–4915, Singapore. Association for Computational Linguistics.
- Ruibo Tao, Zhipeng Pan, Rohan Kumar Das, Xubo Qian, Mike Zheng Shou, and Haizhou Li. 2021. Is someone speaking? exploring long-term temporal features for audio-visual active speaker detection. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3927–3935.
- Yiyang Teng, Chenguang Song, and Bin Wu. 2023. Recognizing social relationships in long videos via multi-modal character interaction. *IEEE Signal Processing Letters*, 30:573–577.
- Paul Vicol, Makarand Tapaswi, Lluís Castrejon, and Sanja Fidler. 2018. MovieGraphs: Towards understanding human-centric situations from videos. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE.
- Melanie Walsh, Anna Preus, and Maria Antoniak. 2024. Sonnet or not, bot? poetry evaluation for large models and datasets. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15568–15603, Miami, Florida, USA. Association for Computational Linguistics.
- Lun Wang, Om Thakkar, and Rajiv Mathews. 2024. Unintended memorization in large ASR models, and how to mitigate it. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE.
- Weishi Wang, Steven C H Hoi, and Shafiq Joty. 2020. [Response Selection for Multi-Party Conversations with Dynamic Topic Tracking](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6581–6591, Online. Association for Computational Linguistics.
- Chao-Yuan Wu and Philipp Krahenbuhl. 2021. Towards long-form video understanding. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. 2025. Qwen2.5-omni technical report. *arXiv [cs.CL]*.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hang Zhang, Xin Li, and Lidong Bing. 2023. Video-LLaMA: An instruction-tuned audio-visual language model for video understanding. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 543–553, Singapore. Association for Computational Linguistics.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. LlamaFactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand. Association for Computational Linguistics.
- Xiangyang Zhou, Lu Li, Daxiang Dong, Yi Liu, Ying Chen, Wayne Xin Zhao, Dianhai Yu, and Hua Wu. 2018. [Multi-Turn Response Selection for Chatbots with Deep Attention Matching Network](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1118–1127, Melbourne, Australia. Association for Computational Linguistics.
- Rongxin Zhu, Jey Han Lau, and Jianzhong Qi. 2021. [Findings on Conversation Disentanglement](#). In *Proceedings of the The 19th Annual Workshop of the Australasian Language Technology Association*, pages 1–11, Online. Australasian Language Technology Association.

A Data pre-processing for annotation

To complement Fig. 2, we include further descriptions of our data pre-processing pipeline below:

Re-transcribe the audio on the sentence-level. The original subtitles lack clear sentence boundaries, which makes them unsuitable for annotation aimed at understanding conversational structure. To address this, we transcribe the raw audio using Whisper (Radford et al., 2022) and align the

transcription with the original subtitle data from TVQA.

Infer and standardize speaker labels. We are interested in *conversational participants*, which, in the context of TV series, are all characters that appear in the clip. For this reason, we focus on post-processing. Some of the subtitles from TVQA contain speaker labels; those labels are usually the first names of the characters in the clip. To facilitate downstream face recognition and annotation, we standardize those labels and map them to actor names by querying TMDb⁶ for the cast list of each episode. This gives us the mapping between the canonical names of the characters and the actors who play them. With this information, we project TVQA speaker labels where possible (i.e., when the label is present and not unknown for the given sentence) and the algorithm described in (Huang et al., 2020, p. 24) to align TVQA subtitles and Whisper transcriptions. For a second pass, we infer speaker labels from the sampled frames and audio: we use the most frequently occurring face during the speech in question, and we recognize the face following the pipeline described in Bamman et al. (2024), which maps the faces to actor names on IMDb,⁷ and the character who dominates the duration of the sentence is treated as the speaker of the sentence. Those speaker labels are verified during the annotation stage; if no face or actor is detected from the automatic pipeline described above, we manually annotate them too.

B Additional related work

Multi-party conversation understanding. Multi-party conversation understanding (i.e., dialogues with more than two participants) can be summarized as solving “who says what to whom” (Gu et al., 2022). This work tends to consider a limited set of modalities, including text- (Lu et al., 2022; Tan et al., 2023; Penzo et al., 2024) or audio-only (Akhtiamov et al., 2019; Lerner et al., 2022) approaches.

Representative tasks include addressee recognition and conversation disentanglement; both require understanding the relationship between speakers—or more generally, conversation participants—and their utterances. They are designed and built upon text-based data, such as chatroom logs (Kummerfeld et al., 2019; Gu

et al., 2021). Works involving audio-visual data focus on speaker activity, identity, and visual co-occurrence (Bredin and Gelly, 2016; Roth et al., 2020; Sharma and Narayanan, 2023) and their long-range dependencies (Tao et al., 2021; Kim et al., 2024), mostly without resolving the semantics of conversation structure, until more recently (Korbar et al., 2024).

Text–vision models leverage image captions or scene descriptions to infer conversational roles in the broader video understanding framework, which suggests a focus on reasoning over the visual aspect of people and objects in the frame that ignores speech (Zellers et al., 2019; Lei et al., 2020). In this work, we believe that multiple modalities are useful in capturing different aspects of the conversation structure, and a robust understanding of conversational dynamics should mirror human understanding of social interactions.

Video understanding. Video understanding involves reasoning over an extended temporal context that includes a sequence of images and/or audio signals. It is fundamental to tasks ranging from movie genre classification to scene graph creation (Huang et al., 2020; Vicol et al., 2018; Wu and Krahenbuhl, 2021; Islam and Bertasius, 2022; Mangalam et al., 2023; Zhang et al., 2023; Sun et al., 2024). Such work focuses on multimodal representations of video and parsing interactions between visible objects, exemplified by visual commonsense reasoning (VCR, Zellers et al., 2019) and TVQA (Lei et al., 2018, 2020), require models to infer plausible interactions between characters in an image (or a sequence of frames sampled from videos) and reason over their actions. However, they typically do not explicitly resolve the underlying conversation structure and often discards speech data (audio or transcription).

Those works converge in recent advances of foundational models: they are natively multimodal (Gemini Team et al., 2023; OpenAI et al., 2024; Reka Team et al., 2024; AI@Meta, 2025), general-purpose models that can also reason over videos. In this work, we assess the capabilities of such models to resolve the conversation structure of such long-form videos. Related work that bears most resemblance to this study does consider multiple modalities for solving conversation structure (addressee recognition) but involves a controlled setting (Le Minh et al., 2018; Inoue et al., 2025). This work represents interactions and settings rep-

⁶<https://www.themoviedb.org/>

⁷<https://www.imdb.com/>

resented on screen, and we discuss the impact of this in §E.

C Evaluation metrics

Following existing work, we use the following metrics to evaluate the performance of models:

C.1 Conversational roles

Accuracy. Speaker identification is evaluated by comparing the predicted speaker label with the true label: $1/N \sum_{i=1}^N \mathbb{1}(y_i = \hat{y}_i)$, where y_i is the true label, \hat{y}_i is the predicted label. This is equivalent to precision @1 used in previous work (Gu et al., 2023).

Set-Based F₁. Existing work uses F₁ scores for addressee recognition, treating it as a multi-class classification task. Our setup allows multiple addressees and side-participants, so we adapt multi-label metrics that compare gold (annotated) and predicted sets of participants: $1/N \sum_{i=1}^N F1_i$, where $\text{Precision}_i = |\hat{Y}_i \cap Y_i| / |\hat{Y}_i|$, $\text{Recall}_i = |\hat{Y}_i \cap Y_i| / |Y_i|$, $F1_i = 2 \cdot \text{Precision}_i \cdot \text{Recall}_i / (\text{Precision}_i + \text{Recall}_i)$, and Y_i is the gold set and \hat{Y}_i is the predicted set for utterance i .

C.2 Conversational thread

Reply-to involves first identifying the pairwise links between an utterance of interest and its parent utterance (child utterance *replies to* the parent), and secondly turning those links into conversational threads. We can see the former as equivalent to binary classification (i.e., whether two utterances exhibit the reply-to relationship) and the second as a clustering problem. The former is evaluated using the F₁ score; the latter on the cluster-based metrics below:

Normalized Variation of Information (NVI). This metric measures the agreement between two sets of clusters. Based on information theory, it asks, “How much information does the predicted clustering give me about the actual (gold) clustering?” A higher score indicates that the predicted and gold threads are more structurally similar. Formally, this metric measures the dissimilarity between clustering assignments C and C' , based on entropy and mutual information:

$$\text{VI}(C, C') = H(C) + H(C') - 2I(C, C') \quad (4)$$

$$\text{NVI}(C, C') = 1 - \frac{\text{VI}(C, C')}{\log_2 N}, \quad (5)$$

where $H(C)$ is the entropy of clustering C , $I(C, C')$ is the mutual information, and N is the total number of utterances. We report $100 \times (1 - \text{NVI})$, so larger is better.

One-to-One Overlap (1-1, Elsner and Charniak, 2008). This is a more forgiving metric that rewards identifying the *essence* of each conversational thread. It finds the best possible one-to-one pairing between the predicted and gold clusters to maximize their shared utterances. It doesn’t require a perfect match, so it gives credit even if a predicted thread is mostly correct but has a few extra or missing utterances. Formally, this metric finds a one-to-one mapping between gold and predicted clusters that maximizes their overlap. It treats the contingency matrix M , where $M_{i,j}$ is the number of shared utterances between gold cluster i and predicted cluster j , as a weighted bipartite graph. The optimal assignment is solved using a min-cost flow algorithm, which identifies the pairing of clusters that yields the greatest total intersection:

$$\text{One-to-One Overlap} = \frac{\max_{\pi} \sum_i M_{i,\pi(i)}}{\sum_{i,j} M_{i,j}}, \quad (6)$$

where π is a one-to-one mapping from gold to predicted clusters.

Exact Match (EM, Kummerfeld et al., 2019). This is our strictest metric; it measures the perfect recovery of threads. A predicted cluster only gets credit if it is an identical match to a gold-standard cluster (containing the exact same set of utterances). This score reflects the model’s ability to precisely identify the complete boundaries of a conversation. Formally, this metric computes precision, recall, and F₁-score based on exact identity between individual gold-standard clusters $c_i \in C_{\text{gold}}$ and predicted clusters $c'_j \in C_{\text{pred}}$. A match is identified for a gold cluster c_i if there exists any predicted cluster c'_j such that $c_i = c'_j$ (i.e., they contain the identical set of elements). Letting N_{match} be the total number of unique gold clusters c_i for which such a match exists, Precision is defined as $P = N_{\text{match}} / |C_{\text{pred}}|$, Recall as $R = N_{\text{match}} / |C_{\text{gold}}|$, and F₁ is their harmonic mean, measuring the perfect recovery of clusters.

D Additional experimental details

We do not change the default temperature of the LLMs that we evaluate. We use the Pydantic schema for structured output (see Fig. 6). Our

```

class ConversationalRoles(BaseModel):
    line_index: int
    reply_to: int
    speaker: str
    addressees: list[str]
    side_participants: list[str]

class ClipRoles(BaseModel):
    clip_roles: list[ConversationalRoles]

```

Figure 6: Pydantic schema.

system instructions for Gemini 2.0 Flash, similar to those used for other models we consider, are in Fig. 9. This represents a succinct summary of our annotation guidelines. For our LLaMA 4 Scout experiments, we run the model in full precision on four H100 GPUs served by vLLM (Kwon et al., 2023), which takes around 3 hours to complete. For supervised fine-tuning with Qwen 2.5-Omni 7B, to minimize computational overhead, LoRA is applied across all layers of the model with a lightweight rank of 2, and a total training steps of 3,125 on one L40S GPU. Each training sample is limited to 4,096 tokens to manage memory and processing time, and the total number of random samples is capped at 1,000. Qwen 2.5-Omni 7B inference is performed with full precision with Flash Attention 2 enabled (Dao, 2024) on four L40S GPUs, which takes around 7 hours. For Qwen 2.5-Omni 7B LoRA SFT, we follow the instructions and configurations provided on <https://github.com/hyouga/LLaMA-Factory/pull/7537>.

E Gemini error analysis

To further illustrate these quantitative findings, we discuss an edge-case clip that models perform badly on: *House M.D.*, season 4, episode 2, segment 2, clip 11:

- *Visual clutter*: The scene features a large number of participants, where a group of four or more doctors discusses a case in a busy hallway. This directly relates to our quantitative finding that participant count is the strongest negative predictor of performance across all conversational roles.
- *Visual-conversational misalignment*: The editing in this scene breaks a simple alignment between sight and sound: for instance, the dialogue of the medical team is heard while the camera follows House, a different, lone character walking through another part of the

hospital. Models that rely heavily on the visual presence of a speaker’s face to assign roles more frequently fail here, which supports our quantitative result that lower face coverage correlates with weaker speaker attribution performance.

- *Speech bleed-over*: The audio from one conversation begins before the speakers are visually introduced, playing over shots of other scenes or characters. This editing technique, where audio precedes the corresponding video, creates significant ambiguity that makes it difficult for a model to correctly segment and attribute the initial utterances. This aligns with our finding that lower acoustic clarity can negatively impact role attribution.

F Annotation guidelines

Note to the Reader. The section outlines the conceptual framework we use to annotate conversational roles and threads in multi-party conversation.⁸ We draw on such traditions as sociolinguistics, linguistic anthropology, and television studies, especially frameworks developed by Clark and Carlson (1982); Clark and Schaefer (1987), Sacks et al. (1974), and Goodwin (1981), who, taken together, lay the foundation of conversation analysis and provide us with the basis of taxonomy on which our role attribution task is based; Goffman (1981), on the face-to-face, social nature of dialogic interactions; and finally, McKee (2016), who analyzes TV dialogues in the familiar set of vocabulary (Austinian speech acts, for instance).

Our annotation scheme is the result of several iterations, beginning with a pilot round and subsequent discussions with annotators trained in rhetoric, cognitive science, film studies, and related disciplines. Based on observations from the pilot, we refine the definitions of utterance boundaries, reply-to structure, and role attribution to better reflect how participants—or characters on screen—engage in face-to-face dialogues in TVQA data.⁹ We hope the annotation guidelines presented below will bring out the nuances of the annotation tasks at hand and their intellectual stakes, as well

⁸Written primarily by AH and KC, with input and guidance from other co-authors.

⁹This data is obtained from the original authors under TVQA Dataset Download Agreement; see https://docs.google.com/forms/d/e/1FAIpQLSftzvAtUSnsmuwpKk5PU9bS71k4F0u_v9Y2vq2nGjRXKoh6tw/viewform

as provide the theoretical scaffolding that informs our design choices.

* * *

In any dialogic interaction, conversational participants manage and direct attention towards each other, during which they also inhabit various roles (speakers, addressees, etc.). For Goffman (1981) in particular, conversations are more than merely linguistic exchanges but *ritualized* social encounters, regulated by norms that govern such phenomena as turn-taking and floor-claiming. These norms extend beyond linguistic structures alone and are often mediated through non-verbal and non-linguistic cues (gaze, posture, etc.).

If we follow Goffman and attempt to uncover the social architecture of interaction underlying conversational interactions, we can potentially shed light on the tacit rules and alignments that structure how people participate in talk; in the context of media studies, this enables us to probe further into issues pertaining to representation: How do characters include or *exclude* others? How do they calibrate intimacy, formality, or confrontation? Who gets to speak without contest? Who must listen without reply?

We address three core questions for capturing the dynamics of conversational interactions:

Who initiates the conversation and controls conversational flow (holds the floor)? This draws on McKee’s notion of speech acts, as characters who initiate dialogue to *perform* an action, and Goffman’s idea that conversation involves the *distribution of attention* among its participants. Characters who get to start a conversation and maintain the floor can be more powerful or otherwise interesting.

Who is explicitly addressed by the speaker? This reflects Goodwin’s and Clark’s emphasis on speakers and hearers, or Goodwin’s more general inquiry of conversation organization, which helps to shed light on interpersonal relationships and interactive dynamics.

Who participates implicitly as side-participants or bystanders? Building on Goffman’s notion of adjacency pairs (Sacks et al., 1974) and Clark’s ratified vs. non-ratified participation, participants who are present but not addressed have the potential of illuminating social hierarchies, inclusion/exclusion patterns, and attentional structures within multi-party conversations.

The annotation guidelines are organized as follows: §F.1: Definitions; §F.2: Annotation procedure; §F.3: Examples; §F.5: Limitations.

F.1 Definitions

This is grounded primarily in Goffman (1981); McKee (2016):

Utterance. An utterance is a single unit of spoken dialogue, roughly equivalent to a *dialogue line*. It is a bounded communicative act performed by a speaker within an interaction driven by the need for a plot, of the speaker, or other characters in the context of a TV show. Unlike a *statement*, which is defined by its propositional content and logical truth conditions, an utterance is defined by its timing, delivery, and interactional role. Utterances need not be syntactically complete sentences, but as *moves* that accomplish social action. A particular type of utterance is the *utterance of interest* (UOI), which is the current utterance being annotated. For ease of annotation, we present individual utterances as one sentence transcribed by Whisper (Radford et al., 2022). They don’t have to be complete sentences, and they don’t have to be a completed turn in and of itself. This allows us to annotate conversational roles at a reasonable granularity: different lines in the same turn can have different role attributions.

Reply-to. The reply-to relation connects an utterance of interest to its most relevant preceding utterance, also known as a *parent* utterance, representing the line to which it responds most directly. This is more prevalent in NLP literature (Zhu et al., 2021), where the task is formalized as a link from a *child* utterance to its *parent*, and the relationship is modeled as directed graphs. The default parent utterance is the immediately preceding line; if there is no logical antecedent utterance, the UOI is the beginning of a new thread (more below). Since we present each utterance at the sentence level, a special form of reply-to is that of a *continuation*: if the speaker is still in the midst of their turn, and the UOI merely continues the previous line, then the previous line is the natural and logical parent utterance. In this view, we might see TV dialogues as linear sequences, where one line necessarily triggers and informs the next, until the current thread runs its course.

Conversational thread. A coherent sequence of utterances sharing a sustained focus (either on a

character or topic), adapted from Goffman’s definition of *focused interaction*. Dramatic scenes often involve multiple, overlapping threads, each driven by the goals and intentions of characters (aligning with McKee’s concept of dramatic intent). Its annotation involves the following steps:

- Given each UOI (u_t), identify the single most relevant preceding utterance as its *parent utterance* (u_p).
- If no suitable preceding utterance exists, annotate the UOI as replying to itself (indicating a new conversational thread).

This annotation captures how dramatic dialogues unfold through pairs of conversational turns, marking shifts in attention, character intention, or conversational control. Those conversational threads (also known as sub-conversations) represent the latent structure of conversations, and the central aim is to segment a broader dialogue into coherent units that represent a stable distribution of attention from conversational participants.

Crucially, our notion of a *thread* is informed by Goffman’s interactional order, wherein conversation is framed as a ritualized social encounter: participants are not merely exchanging information, but are engaging in a tacit choreography governed by socially shared norms of attention, responsiveness, and turn-taking. These practices include obligations to respond when addressed (more below), to respect shifts in floor control, and monitor others’ contributions for relevance, along with a variety of verbal and non-verbal cues. As such, we emphasize *both* topical *and* floor change; indeed, as Goffman notes and demonstrates in his analysis, conversation is not just organized by the semantic content of the utterances. In other words, conversational threading is fundamentally different from, e.g., running topic modeling to cluster utterances: we care about the distribution of *attention* that we can observe, not that of *words*. Disentangling a conversation by correctly resolving reply-to relations serves two complementary ends: it, on the one hand, clarifies the information structure of a multi-party exchange; on the other, the observable structure of who replies to whom and when reflects the rules of engagement that make conversation intelligible and socially meaningful.

Conversational roles. While disentanglement focuses on the structural linkage between utterances, conversational role attribution concerns the

Role	addressed	ratified	known
Addressee	+	+	+
Side-participants	–	+	+
Bystanders	–	–	±

Table 10: Participant role matrix.

dynamic social positions that participants occupy throughout the course of interaction. Each utterance is produced within a shifting *ecology* of roles, an understanding of who is speaking, to whom, and in front of whom is key to interpreting both literal and, again, social meaning of an interaction.

The first role is the *speaker*, the animator, the source of the utterance at time t . Following Goodwin’s conversational analysis (Goodwin, 1981) and Clark’s role taxonomy (Clark and Carlson, 1982; Clark and Schaefer, 1987), we treat conversational roles as those observably *projected at t* : a speaker’s alignment towards others (signaled through gaze, body orientation, lexical choice) casts other participants into roles. These roles are independent of whether the projected utterance ultimately hears, understands, or replies.

We annotate each character present in the scene according to three dimensions, which is summarized in Table 10:

1. **Addressed:** Is the character explicitly spoken to by the speaker?
2. **Ratified:** Is the character recognized by participants as legitimately part of the conversational group?
3. **Perceptually engaged:** Is the character known (or shown visually, via multi-modal cues) to be attending to or perceiving the utterance?

This yields the following roles:

- **Speaker:** the speaker of the line
- **Addressee:** the *intended* recipient of the line; they are directly addressed, ratified as a co-participant, and likely visibly attending to the speaker, who, in turn, likely constructs the utterance with the addressee in mind, and the addressee is expected to respond or acknowledge the utterance in some form
- **Side-participant:** a participant who is ratified, perceptually engaged, but not explicitly

addressed. This term is not really used in Goffman, but we prefer it because it has been used in previous work related to ours (e.g., (Lee and Marsella, 2011)).

- **Bystander**: a participant who is neither explicitly addressed nor ratified; may or may not be perceptually engaged (includes overhearers). If someone is not in any of the roles above, they are automatically bystanders, so we are not annotating them specifically.

F.2 Annotation procedure

Annotation will take place on a dedicated interface (Fig. 7), and annotators will follow these concrete steps when examining each utterance in a video clip: based on the list of participants provided to you (derived from the cast list of the episode):

1. Watch the reconstructed clip
2. Identify the speaker and explicitly addressed participants, which usually involves verifying or correcting existing speaker labels
3. Identify the UOI and determine its parent utterance (or mark as new thread, using SELF)
4. Identify participants who are ratified and perceptually engaged but not explicitly addressed (side-participants). Participants are represented as draggable badges, which can be moved from the original list into labeled fields

We include the following tags to handle unidentifiable participants; they are: *unknown* (not referred to by anyone else, and not identified by our face recognition pipeline, or whose identity is otherwise never revealed to the audience); *crowd* (a group of unidentifiable people, such as public speaking); *OS* (off screen and unknown, which presents a special edge case that we might not consider).

In annotating the clips, you are encouraged to re-watch the clip and spend as long as you require, but we recommend a time limit (of, say, 5 minutes) for each clip if this is your first pass so you will not get stuck. Your judgment should be entirely based on the clip and the clip alone.

Non-dialogic checkboxes. Given the narrative form of the TV series, we introduce the following shorthands to indicate utterances that might not appear in typical face-to-face interactions:

- **Extra-diegetic** (checkbox): to indicate narrators, “previously on the show”, or other kinds of dialogue lines that are not part of any actual conversation
- **Monologue** (checkbox): to better distinguish situations where the character speaks to themselves or other inanimate objects
- **_OS tag**: to indicate when the identity of the character is generally known to the audience but not within the scope of the clip.

Those will give us some basis to filter out utterances that are edge cases as we assess the performance of each model on this task. We also have a couple of other life quality features; see Fig. 8.

F.3 Examples

For a typical example, let’s consider this scene taken from *The Big Bang Theory* (season 2, episode 9, “The White Asparagus Triangulation”, segment 02, clip 04).¹⁰ In it, we observe three characters seated in a row at a movie theater, who are eligible conversational participants: Stephanie Barnett (Sara Rue), on the left, is turned slightly toward the others; Leonard Hofstadter (Johnny Galecki) sits upright, angled more towards Sheldon Cooper (Jim Parsons) on his right:

SHELDON	I’ll find us seats?	(#1)
STEPHANIE	Oh no, we have seats.	(#2)
LEONARD	Not the right seats.	(#3)

Utterance #1 opens the exchange and appears to be phrased tentatively, which suggests a proposal. Its orientation to both Stephanie and Leonard suggests a plural *us* as the intended group of recipients. We don’t see clearly from the visuals (no direct gaze) or vocative, so we assume both addressees are equally implicated. Being the first utterance in the clip, this is also a thread initiator.

Stephanie provides a reply that rejects Sheldon’s offer (utterance #2); it is triggered and necessitated by utterance #1, which is clearly its parent or reply-to utterance. The “we” still invokes group reference, but the syntactic subject now reclaims authority over the seating decision. Here, we see Stephanie looking directly at, and speaking directly to, Sheldon, which casts Leonard as a side-participant.

¹⁰Segment and clip number are given in TVQA.

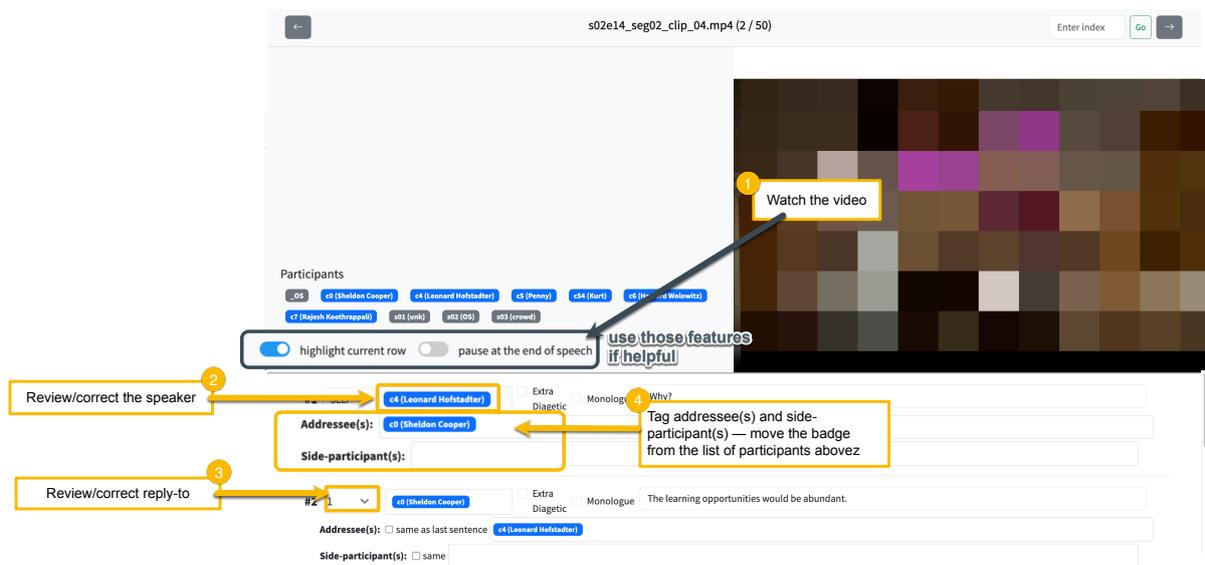


Figure 7: Annotation interface for multimodal conversation structure understanding. Annotators proceed through four steps: a.) watch the reconstructed video clip, b.) verify or correct the speaker label, c.) review and update the reply-to link, and d.) assign addressee(s) and side-participant(s) by dragging the relevant participant badges from the list, derived from the cast list, next to the video.

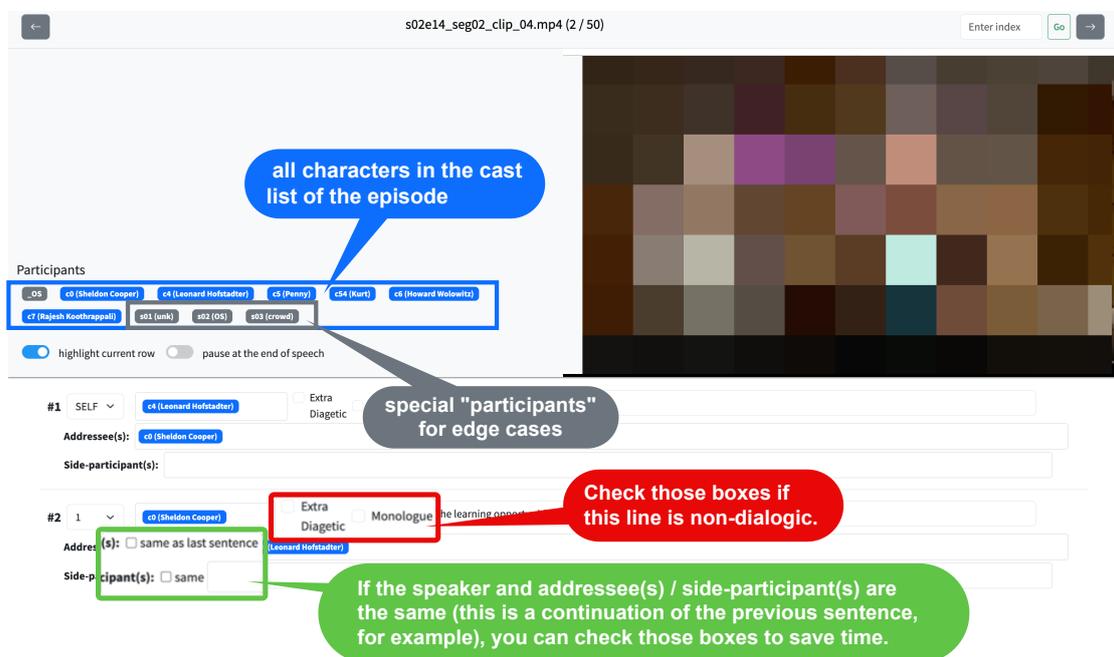


Figure 8: Annotation enhancements for edge cases and workflow efficiency. Annotators can select characters *and* special participant tokens for non-standard utterances. Checkboxes allow them to mark lines as continuations (same role as previous) or non-dialogic (extra-diegetic or monologue), with the goal to minimize redundancy and improve overall annotation consistency.

Leonard's utterance (#3) is a reply to #2, seeing as it builds on Stephanie's claim of having found seats, but disputes its adequacy. Semantically, we can also see this as a response to Sheldon's original proposal because it supports the need to continue looking for the seats, but structurally, we will annotate the reply-to as #2 for two important reasons: a.) there's no reason for Leonard to say this *if it wasn't for* Stephanie's utterance #2, which makes the choice of #1 less justifiable as it also overrides the default reply-to; b.) since threads are transitive closures of pairwise relations, and all the utterances respond to Sheldon's original need to find some seats, those three utterances will form a thread, which signals their connections, and it's important to distinguish between *threads* and a single reply-to.

F.4 Edge cases

Incapacitated side-participants. In this scene from *House M. D.* (season 4, episode 2, segment 2, clip 11), a team of doctors speak about and to a patient who has locked herself into a room and is experiencing hallucinations. Although she is fully conscious, a named speaking character, and within hearing range of the conversations about her, she is not annotated as a side-participant in those conversations about her symptoms as she is preoccupied by her symptoms, not visually depicted to be listening to the conversation and is not capable of interrupting or joining the ongoing conversation. Later in the scene, however, she is labeled as an addressee after responding to her name.

Off-screen side-participants. In this scene from *House M. D.* (season 4, episode 16, segment 2, clip 10), a roomful of doctors speak amongst themselves while one doctor in particular, James Wilson, sits on the side and remains largely unspeaking and unshown. Even though the camera is rarely on him and the other doctors address each other rather than him, he qualifies as a side-participant due to being in the same space and scene as the other interlocutors.

Non-human/inanimate objects and monologues. In this *Friends* scene (season 10, episode 06, segment 2, clip 17), the speaker addresses a duck, which appears on screen. After discussion, we concluded that addressees must be portrayed by a human in order to be counted and left the addressee blank. The label "monologue" distinguishes between this particular *absence* of addressee from

others which arise in different context (i.e., speaker talking to themselves or having no clear audience).

Extra-diegetic moments. In this scene from *How I Met Your Mother* (season 6, episode 9, segment 2, clip 3), speakers cluster around a television to watch a game show, which is then shown on screen with recognizable characters and dialogue. Annotators marked all game show dialogue as extra-diegetic and did not include the TV-watchers as addressees or side-participants.

F.5 Limitations

The practical limitations come from our TVQA post-processing pipeline: for instance, annotators observe that cases where conversations are fast-paced (many different speakers, or many different exchanges between two interlocutors, for instance), Whisper would encounter segmentation issues: an utterance would be misattributed to a previous speaker's line, rather than identified accordingly as an independent utterance (as observed in *House M. D.*, season 5, episode 18, segment 2, clip 20). This error, though infrequent, unsurprisingly suggests that clips featuring sustained, uninterrupted speech from a single speaker, whether in a series of in-scene utterances or as the result of monologic narration overlaid on a series of images, are more likely to be correctly attributed and segmented than those with many voices, even independent of visual face counts. Future studies may focus on the role of interruption, whether verbal, visual or contextual, on multimodal conversational understandings, as well as what constitutes the boundaries of an utterance as it relates to transcription and segmentation.

You are a video analysis assistant. Your task
↳ is to analyze the conversations in a
↳ video clip and its associated subtitles.
↳ For each dialogue line, you will:

- * determine what previous line it is replying
↳ to
- * determine the speaker, addressees, and
↳ side-participants

Here's how to determine the reply-to
↳ relationship between utterances to
↳ resolve conversational threads:

- * The reply-to structure gives us information
↳ about floor-claiming and topical change
↳ within the clip.
- * The character is saying this line because
↳ they want to respond to that previous
↳ line. What previous line is this current
↳ line replying to?
- * If the speaker of the last line is the
↳ same, you can treat it as continuation
↳ and put the index of last line as the
↳ reply-to.
- * If no previous line triggers this line,
↳ then write the current line index,
↳ indicating the current line replies to
↳ itself, which marks the start of a new
↳ conversational thread.

Here's how to determine each role:

- * ****Speaker:**** The character who is speaking
↳ the line. Infer this from lip
↳ movements, body language, and the
↳ context of the dialogue. If a character
↳ finishes one line and immediately starts
↳ another (very short pause), assume it's
↳ the same speaker, UNLESS there's a clear
↳ visual indication of a scene or speaker
↳ change (e.g., a camera cut to a
↳ different person starting to speak).
- * ****Addressee(s):**** The character(s) the
↳ speaker is *directly* addressing. Use
↳ these cues:
 - * ****Eye Contact:**** The most important
↳ cue. Who is the speaker looking at?
 - * ****Body Orientation:**** Is the speaker's
↳ body turned towards a particular person
↳ or group?
 - * ****Dialogue Context:**** Does the line
↳ contain a name, pronoun ("you"), or
↳ clearly refer to a specific individual
↳ or group? ("Hey, John..." or "You all
↳ need to...")
 - * ****Reactions:**** If a character reacts
↳ immediately and strongly to a line
↳ (e.g., nods, responds verbally, shows
↳ surprise), they are likely an addressee.
- * If the speaker seems to be talking to
↳ everyone present, list all characters
↳ who appear to be paying attention.
- * If the speaker is talking to a crowd of
↳ unidentifiable characters, write "crowd".
- * If the speaker is talking to
↳ themselves, or no one in particular,
↳ write "none".
- * ****Side-Participant(s):**** Any character(s)
↳ visible in the scene *during* the line's

- ↳ timeframe* who are *not* the speaker or
↳ addressees. They are present, and their
↳ presence is known to other participants.
↳ They can potentially join the
↳ conversation at any time.
- * If it is not possible to confidently
↳ determine if someone is a
↳ side-participant, write "unknown".
- * If there are no side-participants, write
↳ "none".

****Input:****

You will receive a list of subtitle entries.
↳ Each entry will be a dictionary with the
↳ following keys:

- * **``line_index``:** (int) The index of the
↳ current entry (subtitle line).
- * **``start_time``:** (float) The start time of
↳ the subtitle line in seconds.
- * **``end_time``:** (float) The end time of the
↳ subtitle line in seconds.
- * **``text``:** (string) The text of the dialogue
↳ line.

You will also receive a list of potential
↳ participants for you to assign roles
↳ from. You must pick from this list.

With all this information, analyze the video
↳ segment corresponding to the
↳ `start_time` and `end_time` of each
↳ subtitle entry.

****Output:****

Provide your output in JSON format, mirroring
↳ the structure of the input. For *each*
↳ subtitle entry, add the following keys:

- * **``line_index``:** (int) The line being
↳ analyzed.
- * **``reply_to``:** (int) The line index that
↳ this current line replies to, could be
↳ the same as the current line index or
↳ any previous line index.
- * **``speaker``:** (string) The name of the
↳ speaker. If you cannot determine the
↳ speaker, use "unknown".
- * **``addressees``:** (list of strings) A list of
↳ the names of the addressee(s). This can
↳ be an empty list (`[]`) if there are no
↳ direct addressees, or `["none"]` if the
↳ speaker is speaking generally but to no
↳ one in particular.
- * **``side_participants``:** (list of strings) A
↳ list of the names of the
↳ side-participant(s). This can be an
↳ empty list (`[]`), `["none"]`, or
↳ `["unknown"]`

Figure 9: System instruction to Gemini 2.0 Flash.