

Modeling Multimodal Social Interactions: New Challenges and Baselines with Densely Aligned Representations

Sangmin Lee¹ Bolin Lai² Fiona Ryan² Bikram Boote¹ James M. Rehg¹
¹University of Illinois Urbana-Champaign ²Georgia Institute of Technology
 {sangminl,boote,jrehg}@illinois.edu {bolin.lai, fkryan}@gatech.edu

Abstract

Understanding social interactions involving both verbal and non-verbal cues is essential for effectively interpreting social situations. However, most prior works on multimodal social cues focus predominantly on single-person behaviors or rely on holistic visual representations that are not aligned to utterances in multi-party environments. Consequently, they are limited in modeling the intricate dynamics of multi-party interactions. In this paper, we introduce three new challenging tasks to model the fine-grained dynamics between multiple people: speaking target identification, pronoun coreference resolution, and mentioned player prediction. We contribute extensive data annotations to curate these new challenges in social deduction game settings. Furthermore, we propose a novel multimodal baseline that leverages densely aligned language-visual representations by synchronizing visual features with their corresponding utterances. This facilitates concurrently capturing verbal and non-verbal cues pertinent to social reasoning. Experiments demonstrate the effectiveness of the proposed approach with densely aligned multimodal representations in modeling fine-grained social interactions. Project website: <https://sangmin-git.github.io/projects/MMSI>.

1. Introduction

Real-world social interactions involve intricate behaviors between multiple people. People communicate not only through verbal cues (e.g., language) but also through non-verbal cues (e.g., gesture, gaze). While spoken language conveys explicit meaning, inferring the full social context from language alone can sometimes be ambiguous. Non-verbal cues can often play a crucial role in clarifying these subtle social nuances and providing additional context. Consequently, comprehensively understanding social interactions involving multimodal social cues is essential to interpret social situations appropriately.

There have been attempts to investigate social behaviors

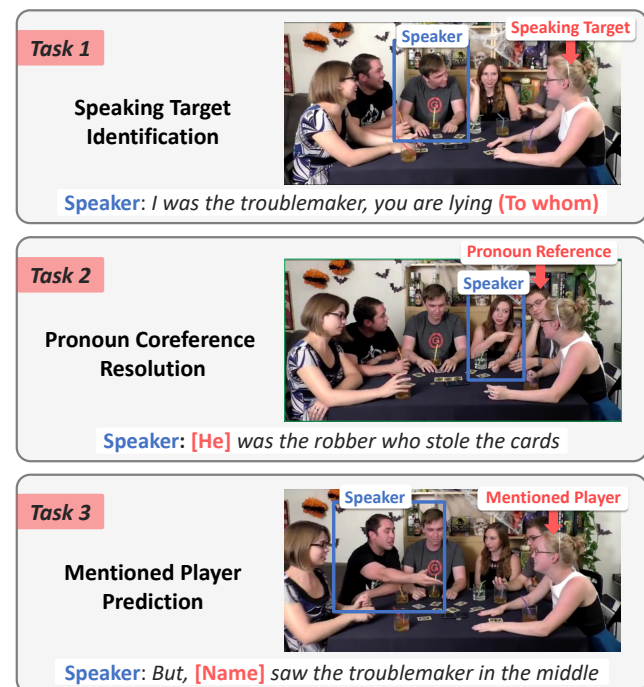


Figure 1. Concepts of the proposed three social tasks in multi-party environments: speaking target identification, pronoun coreference resolution, and mentioned player prediction.

in multimodal aspects by jointly considering language and visual cues. Some works tried to learn the relationships between spoken language and visual gestures for gesture generation [1, 2, 32] and gesture-language grounding [29]. Other multimodal approaches utilized the interconnection between spoken language and visual expressions to recognize human sentiment and emotion [9, 21, 22, 39, 42, 53]. However, these works mainly focus on single-person behaviors or rely on holistic visual representations in multi-party settings, rather than modeling the fine-grained dynamics of social interactions among multiple people. Modeling such intricate multi-party dynamics requires understanding the relationships between language and visual cues at an individual level, rather than using global representations.

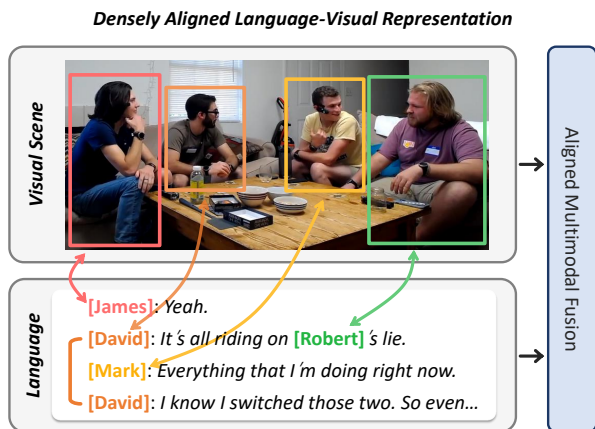


Figure 2. Concept of densely aligned language-visual representations. People are matched in the language and visual domains.

Recently, a multimodal work [28] addressed social behaviors in social deduction games, which provide an effective testbed for studying multimodal social interactions. In these games, players take on roles and attempt to deduce the roles of their opponents by engaging in communication, deception, inference, and collaboration. These games encompass rich social interactions including verbal and non-verbal cues in multi-party settings. Lai *et al.* [28] leveraged language and visual cues to predict persuasion strategies at the utterance level. However, their work has limitations in modeling multi-party interactions in terms of its task and methodology. Although persuasion strategies emerge in communication, the task primarily focuses on understanding the social behaviors of a single person rather than the dynamics among people. Moreover, their approach is limited in distinguishing and recognizing fine-grained interactions because it utilizes holistic visual representations for the entire scene, despite the presence of multiple people.

To address these issues, we introduce three new multimodal tasks that encompass the dynamics of multi-party interactions, along with a novel baseline model. We extend the social deduction game datasets [28] with extensive data annotations to curate new social tasks focused on identifying referents in multi-party conversations. Appropriately identifying referents is important because it significantly affects interpreting the social intent and context of conversations. Figure 1 shows the overall concepts of our curated social tasks. The three tasks are as follows.

1. **Speaking target identification:** Identifying who a speaker is talking to in a conversation.
2. **Pronoun coreference resolution:** Determining who a pronoun refers to in a conversation.
3. **Mentioned player prediction:** Predicting who is mentioned by name in a conversation.

These tasks are challenging as they require understanding the fine-grained dynamics of player interactions. For in-

stance, when an utterance is accompanied by visual cues like pointing gestures, it is necessary to comprehensively interpret the utterance content and the speaker’s gestures for holistic reasoning. To this end, we need to figure out who the speaker is and who is being pointed at within the visual scene based on the utterance context. This involves matching utterances with the visually identified individuals. Thus, it is required to align language references with player visuals and to derive densely aligned multimodal representations from such aligned features. Figure 2 shows the concept of densely aligned language-visual representations.

To this end, we propose a novel baseline model leveraging densely aligned language-visual representations. We detect and visually track each player in the video to distinguish individual players. By initially aligning player visual positions with their language references, we continuously track player visuals in sync with the dialogue. This alignment allows the model to visually identify the speaker and the other players (*i.e.*, listeners) for given utterances. Encoding speaker gestures and the relative positions of the listeners enables deciphering visual relationships for understanding non-verbal dynamics. We then combine this with the linguistic context of the conversation to obtain densely aligned representations. Consequently, we can effectively predict referents by leveraging densely aligned representations containing both verbal and non-verbal dynamics.

The major contributions of this paper are as follows.

- We introduce new social tasks in multi-party settings via extensive data annotations: *speaking target identification*, *pronoun coreference resolution*, and *mentioned player prediction*. These tasks are challenging as they require understanding the fine-grained dynamics of interactions.
- We propose a novel multimodal baseline model leveraging language and visual cues for understanding multi-party social interactions. To the best of our knowledge, this is the first work to address the multimodal dense alignment between language and visual social cues.

2. Related Work

2.1. Social Behavior Analysis

Analyzing social behaviors has been widely investigated in the fields of computer vision and natural language processing. Various works have focused primarily on analyzing social behaviors from a single-modal perspective. In terms of visual cues, some works proposed gaze target estimation techniques [11, 18, 27, 47, 48] to analyze where a person is looking within a scene. There have also been studies that recognize social gaze patterns between multiple people such as identifying shared attention [16, 20, 36, 45]. Gesture recognition approaches [3, 30, 31, 54, 56] have been researched to identify specific types of human gestures such as shaking hands and thumbs-up. Regarding language cues,

Annotation Type	Utterance Example	YouTube DB		Ego4D DB	
		Count	Krippendorff's α	Count	Krippendorff's α
Speaking Target Identification	<i>Why are you helping the Werewolves out? (To [Name])</i>	3,255	0.922	832	0.907
Pronoun Coreference Resolution	<i>I'm a Villager which makes me think he (\rightarrow [Name]) was the Werewolf</i>	2,679	0.962	503	0.846
Mentioned Player Prediction	<i>I'm the troublemaker and I switched [Name] with somebody</i>	3,360	N/A	472	N/A

Table 1. Summary of annotation details for three new social tasks. We achieve sufficiently high *Krippendorff's alpha* values ($\alpha > 0.8$) for both speaking target identification and pronoun coreference resolution tasks, which indicates the high reliability of our data annotations.

dialogue act recognition methods [8, 34, 40, 41, 46, 50] have been introduced to understand the communicative intent behind utterances in social dialogues. Furthermore, there have been works on sentiment analysis and emotion recognition based on dialogue language [4, 23, 44, 55, 57].

Recently, joint modeling of visual and language modalities has been studied for social behavior analysis. Some works focused on learning the relationships between spoken language and gestures for gesture generation [1, 2, 32] and gesture-language grounding [29]. Liu *et al.* [32] proposed a multimodal model that integrates visual, language, and speech cues via hierarchical manners to synthesize naturalistic gestures. Additionally, the intersection of spoken utterances and visual expressions has been explored for sentiment analysis and emotion recognition [9, 21, 22, 39, 42, 53]. Hu *et al.* [22] proposed a unified feature space to capture the knowledge of sentiment and emotion comprehensively from multimodal cues. There have also been multimodal works for question & answering in social contexts [37, 51, 52].

However, these works mainly focus on the behaviors of a single person or rely on holistic visual features that are not densely aligned to language in multi-party environments. They are unable to model the complex dynamics of interactions, which requires understanding the spatial relationships of multiple people in addition to their utterances. We propose a novel approach leveraging densely aligned language-visual representations to capture the fine-grained dynamics.

2.2. Social Deduction Game Modeling

There have been works investigating computational models for social deduction games where players actively communicate and strategize with one another. Some prior studies have focused on developing game-playing agents and analyzing optimal strategies using game theory [5, 6, 12, 35, 43]. These works aim to model the state of the game computationally but do not address understanding the dialogue and behaviors of players. Chittaranjan *et al.* [10] modeled game outcomes from communication patterns such as player speaking and interrupting behaviors. Bakhtin *et al.*

[14] built an agent that can play diplomacy games by utilizing language models with strategic reasoning. These approaches do not capture verbal and non-verbal multimodal aspects of modeling social behaviors. Recently, Lai *et al.* [28] addressed social behaviors in social deduction games using multimodal representations. They leveraged language and visual cues to predict persuasion strategies at the utterance level such as identity declaration and interrogation.

However, this multimodal work is limited in addressing multi-party dynamics due to the lack of person-level feature recognition. To address this gap, we introduce new tasks in social deduction games that explicitly demand recognizing person-level features. We also propose the corresponding baseline model that captures the multi-party dynamics across both language and visual representations.

3. Proposed Benchmark

3.1. Base Datasets

We extend two social deduction game datasets [28]: YouTube and Ego4D with additional extensive data annotations for curating new social tasks.

YouTube dataset. This dataset was collected from the YouTube video platform by searching keywords of Werewolf social deduction game. It contains 151 games of One Night Ultimate Werewolf, which corresponds to 151 separate videos with 14.8 hours. It consists of videos, transcripts, player roles, voting outcomes, and persuasion strategy labels. The transcripts comprise 20,832 utterances.

Ego4D dataset. This Ego4D dataset is a subset of Ego4D Social dataset [19]. It has 40 games of One Night Ultimate Werewolf and 8 games of The Resistance: Avalon. It contains 101 separate videos with 7.3 hours. Among them, we leverage 83 videos where we can visually identify individuals for new data annotations. To guarantee the visibility of all players within the frame, this dataset adopts third-person view videos instead of first-person view videos. It also consists of videos, transcripts, player roles, voting outcomes, and persuasion strategy labels. The transcripts contain 5,815 utterances during the games.

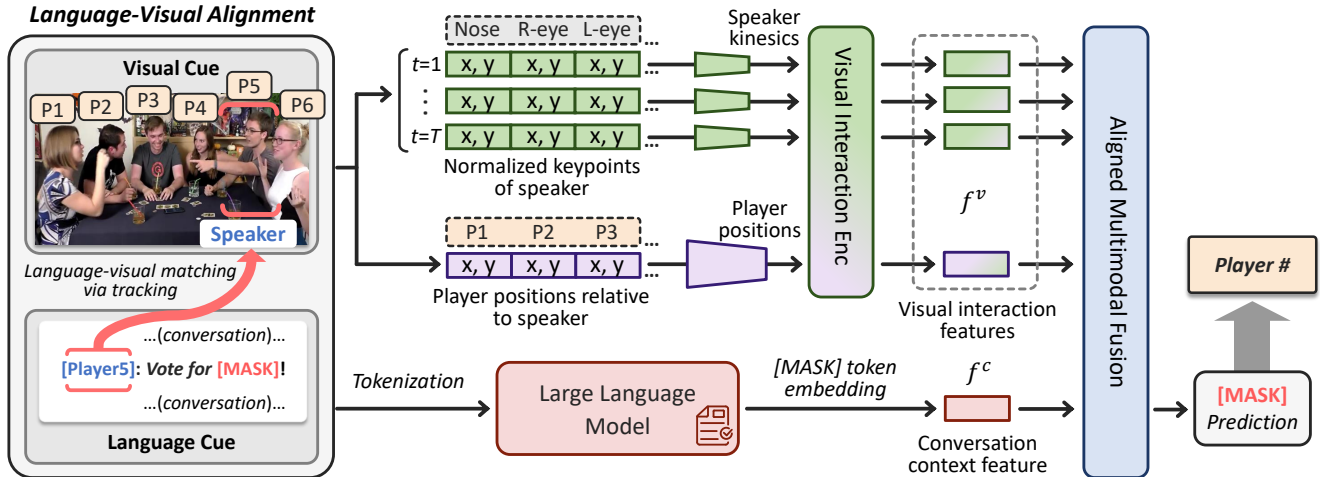


Figure 3. Proposed baseline model for understanding multimodal social interactions to tackle our new social tasks via densely aligned language-visual representations. The model consists of four main parts: language-visual alignment (grey), visual interaction modeling (green & purple), conversation context modeling (red), and aligned multimodal fusion for prediction (blue).

3.2. Data Annotation

To address the fine-grained dynamics of interactions, we design three new tasks in social deduction games: speaking target identification, pronoun coreference resolution, and mentioned player prediction. Annotators reference both transcripts and videos comprehensively to make their annotations in the transcripts. To achieve reliable annotation quality, we initially request three annotators to label subsets of data and measure their annotation agreement using *Krippendorff’s alpha* [26]. After we train the annotators sufficiently with α larger than 0.8, we request the three annotators to label independently for the rest of the data. Note that *Krippendorff’s alpha* > 0.8 is generally considered to indicate good reliability with a high level of agreement [7].

Table 1 shows the summary of our data annotation results. We achieve sufficiently high α for both speaking target identification and pronoun coreference resolution. Note that we could obtain the annotations for the mentioned player prediction automatically. In the training and testing process, we anonymize all names in transcripts (e.g., [“David”, “Alice”, “Thomas”] \rightarrow [“Player1”, “Player2”, “Player3”]). For each task, a test set is constructed using about 20% of the annotations. We split the training and test sets at the video level rather than at the individual utterance level to ensure no overlap in terms of contextual utterances. Annotation details for each task are as follows.

Task1: Speaking target identification. This task aims to predict who a speaker is talking to in a dialogue. To this end, we annotate the labels of who a speaker is speaking to at the utterance level. Since utterances are often directed to all players, we perform the annotation only on the utterances that include [“you”, “your”]. We give our annotators three labeling options: (To Player#), (To Everyone),

and (To Unknown). Based on the annotations, we add “(To [MASK])” at the end of a target utterance.

Task2: Pronoun coreference resolution. This task aims to predict who a pronoun refers to in a dialogue. We conduct the annotation on the third-person pronouns that are used in our dataset, which are [“he”, “she”, “him”, “her”, “his”] in utterances. We give two options to annotators: [Player#] and [Unknown]. We only target the pronouns that indicate a specific player in the game. In terms of modeling the task, we change a pronoun to [MASK] in a target utterance.

Task3: Mentioned player prediction. This task aims to predict who is referred to by their name in a dialogue. Since we know the ground truth names in utterances, it is possible to annotate these mentioned player labels automatically. We modify a mentioned player name with [MASK] in utterances and predict who is referred to in that part.

4. Proposed Approach

The introduced social tasks can be formulated as follows. Let $u = \{u_i\}_{i=k-n}^{k+n}$ denote utterance texts that include the k -th target utterance containing a [MASK] token representing an unknown player, while $v = \{v_t\}_{t=1}^T$ indicates the corresponding T video frames aligned with the utterance timeline. Given u and v , our objective is to optimize the multimodal predictive function $\mathcal{F}(v, u)$ to effectively classify the identity of the player associated with the [MASK].

To this end, we introduce a novel multimodal baseline that leverages densely aligned representations between language and visual cues to capture the fine-grained dynamics of interactions. Figure 3 shows the overall framework of the proposed multimodal baseline. The proposed model consists of four main parts: (i) language-visual alignment, (ii) visual interaction modeling, (iii) conversation context modeling, and (iv) aligned multimodal fusion for prediction.

4.1. Language-Visual Alignment

To achieve densely aligned representations containing both verbal and non-verbal dynamics properly, it is necessary to first establish fine-grained alignments between visual and language cues. To this end, we distinguish individual visuals and match language references with them.

We detect and track players visually in video frames over time using AlphaPose framework [17]. Once we initially match player visuals with the player references in the utterances (*i.e.*, assigning each tracking ID to Player#), we can continuously identify players in both visual and language domains. It enables the model to identify the speaker in the video based on the corresponding utterance and to determine the spatial arrangement of the listeners. Building upon this alignment, we can achieve densely aligned language-visual representations from verbal and non-verbal cues in a comprehensive manner. It enables us to tackle our social tasks effectively, allowing for a more nuanced and holistic understanding of multi-party interactions.

4.2. Visual Interaction Modeling

To distinguish individual players in video frames $v = \{v_t\}_{t=1}^T$, we use the human pose keypoints from AlphaPose. Specifically, we extract 17 body keypoints (x, y) for each player. Figure 3 shows the procedure of encoding visual interactions. The upper path (*green*) of Figure 3 indicates encoding a kinesics feature of a speaker, while the middle path (*purple*) represents encoding spatial positions of all players.

First, we use the keypoints of a speaker in the upper path of Figure 3. Among the 17 part keypoints, we leverage [*nose, l-eye, r-eye, l-shoulder, r-shoulder, l-elbow, r-elbow, l-wrist, r-wrist*] closely related to gaze and gesture characteristics. Let $(x_{part}, y_{part})_t^S \in \mathbb{R}^2$ denote the image coordinates of a part at time t . For example, $(x_{nose}, y_{nose})_t^S$ indicates the nose point. To represent human motion in a unified coordinate, we normalize speaker keypoints by subtracting the speaker nose point from each part point. Each point vector $(x_{part}, y_{part})_t^S$ is encoded by an MLP point encoder E_{point} into a part point feature $f_t^{S,part} \in \mathbb{R}^{d_{point}}$ (d_{point} is channel dim). These part point features are concatenated and processed by an MLP kinesics encoder E_{kin} to obtain a speaker kinesics feature f_t^S as follows.

$$f_t^S = E_{kin}([f_t^{S,nose}, f_t^{S,l-eye}, \dots; f_t^{S,r-wrist}]). \quad (1)$$

Since we have multiple time steps, we can obtain $f^S = \{f_t^S\}_{t=1}^T \in \mathbb{R}^{T \times d}$.

In the meantime, the middle path (*purple*) of Figure 3 receives the position of each player. We consider the nose point of each player as their representative position. We normalize their nose points by subtracting the speaker's nose point from them to get their relative positions from

the speaker. We utilize their representative positions at a single time step corresponding to the start of the utterance. Let $(x, y)^{P\#}$ denote the representative position of Player#. Each point vector $(x, y)^{P\#}$ is independently fed to an MLP point encoder E_{point} to get a player point feature $f^{P\#} \in \mathbb{R}^{d_{point}}$. We concatenate the player point features and feed them to an MLP position encoder E_{pos} to get $\tilde{f}^P \in \mathbb{R}^d$. We then make \tilde{f}^P aware of speaker knowledge. To this end, we make a speaker-label feature $f_{label}^S \in \mathbb{R}^d$ by passing a speaker-label one-hot vector through an FC layer. We combine f_{label}^S with \tilde{f}^P to obtain a player position feature f^P . These procedures are formulated as follows.

$$\tilde{f}^P = E_{pos}([f^{P1}, f^{P2}, \dots; f^{PN}]), \quad (2)$$

$$f^P = \text{FC}(\tilde{f}^P + f_{label}^S), \quad (3)$$

where N indicates the maximum player number in the datasets ($N=6$). If the number of players is less than N for the current input data, we apply zero padding to the excess. If a player is temporarily undetected (*e.g.*, offscreen for a short time), we proceed with position encoding by substituting the corresponding player position stored in a buffer to correct the player position.

Based on the speaker kinesics features $f^S = \{f_t^S\}_{t=1}^T$ and player position feature f^P , we encode the visual interaction by capturing speaker kinesics motion with the context of player visual positions. f^S and f^P are passed through a visual interaction encoder E_v sequentially, which has the form of the transformer [49]. E_v allows modeling dependencies between the speaker kinesics and player positions across time via self-attention. Finally, we can obtain visual interaction features $f^v = \{f_t^v\}_{t=1}^{T+1} \in \mathbb{R}^{(T+1) \times d}$ that represent dynamics between the speaker and players based on the speaker kinesics and listener positions.

4.3. Conversation Context Modeling

The lower path (*red*) of Figure 3 shows encoding spoken utterances from players. To incorporate conversation context, we use surrounding utterances including the target utterance. The input to the language path is formulated as.

$$u = [u_{k-n}; \dots; u_{k-1}; u_k; u_{k+1}; \dots; u_{k+n}], \quad (4)$$

where u_k denotes the target k -th utterance, and the others indicate the preceding and following utterances. Note that the target utterance is the one that contains [MASK]. A [CLS] token is inserted in front of u while a [SEP] token is inserted at the end of each utterance in u for language tokenization processing. Note that all player names in utterances are anonymized as “[Player#]”. We leverage pre-trained language models based on masked-language modeling such as BERT [24]. The tokenized sequence of the utterances is fed into the language model. The output feature corresponding to the index of the [MASK] token is then

Method	Densely Aligned?	Speaking Target Identification (%)	
		YouTube	Ego4D
BERT [24]	-	65.8	56.8
BERT + DINOv2 [38]	✗	66.4	58.0
BERT + MViT [15] (Lai <i>et al.</i> [28])	✗	66.9	57.4
BERT-based Our Baseline	✓	72.7	61.9
RoBERTa [33]	-	72.4	63.6
RoBERTa + DINOv2 [38]	✗	72.7	62.5
RoBERTa + MViT [15] (Lai <i>et al.</i> [28])	✗	73.1	64.2
RoBERTa-based Our Baseline	✓	74.5	66.5
ELECTRA [13]	-	65.8	60.8
ELECTRA + DINOv2 [38]	✗	65.3	60.2
ELECTRA + MViT [15] (Lai <i>et al.</i> [28])	✗	64.6	60.8
ELECTRA-based Our Baseline	✓	69.6	64.8

Table 2. Performance comparison results for the speaking target identification task on YouTube and Ego4D datasets.

retrieved. After passing it through an FC layer to match the channel dimension of the visual interaction features, we get a conversation context feature $f^c \in \mathbb{R}^d$ which contains the context around the [MASK].

4.4. Aligned Multimodal Fusion

To fuse the aligned visual interaction features $f^v = \{f_t^v\}_{t=1}^{T+1}$ and conversation context feature f^c , we first concatenate them in the sequence dimension along with an [AGG] token for feature aggregation. It can be formulated as follows.

$$f^{v+c} = [\text{[AGG]}; f_1^v; \dots; f_{T+1}^v; f^c]. \quad (5)$$

Note that positional encoding [49] for transformers is applied to the f^v parts. Then, $f^{v+c} \in \mathbb{R}^{(T+2) \times d}$ is processed with a multimodal transformer to encode their joint relationships. We leverage an output multimodal feature $f^m \in \mathbb{R}^d$ from the [AGG] token. Finally, a densely aligned multimodal feature f^m is passed through a classification head consisting of an FC layer and softmax to predict the anonymized player identity \hat{y} (e.g., Player#) for the target [MASK]. We optimize the model using cross-entropy loss between the predicted player \hat{y} and the ground-truth label y .

At training time, we apply permutations to anonymized identities to prevent the model from relying on consistent identities. Specifically, we randomly shuffle the mapping from player names to the anonymized player identities in utterances for every iteration. For example, [“David”, “Alice”, “Thomas”] \rightarrow [“Player1”, “Player2”, “Player3”] \rightarrow [“Player3”, “Player1”, “Player2”]. This mapping permutation from the text domain is also applied to the visual position encoding and ground truth label y to ensure that language and visual cues are consistently aligned. This player permutation learning forces the model to learn more

Method	Densely Aligned?	Pronoun Coreference Resolution (%)	
		YouTube	Ego4D
BERT [24]	-	60.3	47.3
BERT + DINOv2 [38]	✗	58.2	46.4
BERT + MViT [15] (Lai <i>et al.</i> [28])	✗	59.8	46.4
BERT-based Our Baseline	✓	65.9	49.1
RoBERTa [33]	-	69.0	48.2
RoBERTa + DINOv2 [38]	✗	68.6	46.4
RoBERTa + MViT [15] (Lai <i>et al.</i> [28])	✗	69.5	49.1
RoBERTa-based Our Baseline	✓	73.0	52.7
ELECTRA [13]	-	62.5	44.6
ELECTRA + DINOv2 [38]	✗	61.1	42.9
ELECTRA + MViT [15] (Lai <i>et al.</i> [28])	✗	62.1	43.8
ELECTRA-based Our Baseline	✓	67.6	46.4

Table 3. Performance comparison results for the pronoun coreference resolution task on YouTube and Ego4D datasets.

generalizable representations of player interactions that do not depend on specific identifiers during the training time.

5. Experiments

5.1. Implementation

We adopt the language model as pre-trained BERT [24], RoBERTa [33], and ELECTRA [13] which are based on masked-language modeling. The proposed model is trained by Adam optimizer [25] with a learning rate of 5e-6 for the language model and 5e-5 for the other parts. We use a batch size of 16. We leverage about 3 seconds of video frames (frame interval 0.4s) that correspond to the timeline of the utterance. We use the preceding and following 5 utterances for encoding conversation context. The detailed network structures are described in the supplementary material.

5.2. Performance Comparison

We measure the identity classification accuracies for our curated tasks: speaking target identification, pronoun coreference resolution, and mentioned player prediction.

Table 2 shows the experimental results for speaking target identification on YouTube and Ego4D datasets with different language models. We compare our proposed baselines with the recent multimodal model [28] (*i.e.*, Language Model + MViT [15]) for social deduction games. In addition, we further adopt DINOv2 [38] which is a powerful versatile visual feature generally used for various downstream tasks. Note that both comparison methods cannot leverage densely aligned language-visual representations. As shown in the table, these methods are not effective in improving upon the language models alone. This reflects that they are not able to figure out who the speaker is and who their gestures are directed at, in correspondence with

Method	Densely Aligned?	Mentioned Player Prediction (%)	
		YouTube	Ego4D
		BERT [24]	-
BERT + DINOv2 [38]	✗	54.4	47.4
BERT + MViT [15] (Lai <i>et al.</i> [28])	✗	53.1	46.2
BERT-based Our Baseline	✓	58.8	50.0
RoBERTa [33]	-	59.9	50.0
RoBERTa + DINOv2 [38]	✗	60.7	50.0
RoBERTa + MViT [15] (Lai <i>et al.</i> [28])	✗	60.6	51.3
RoBERTa-based Our Baseline	✓	62.5	55.1
ELECTRA [13]	-	55.7	42.3
ELECTRA + DINOv2 [38]	✗	56.1	43.6
ELECTRA + MViT [15] (Lai <i>et al.</i> [28])	✗	55.3	41.0
ELECTRA-based Our Baseline	✓	61.0	46.2

Table 4. Performance comparison results for the mentioned player prediction task on YouTube and Ego4D datasets.

the language domain. In contrast, our baselines leveraging densely aligned language-visual representations consistently enhance the language models for this task.

Table 3 and 4 show the performance comparison results for pronoun coreference resolution and mentioned player prediction, respectively. We follow a similar experimental setup, evaluating our baselines against the language models and the multimodal methods (*i.e.*, Language Models + DINOv2/MViT) across three different language models and two datasets. The results show that the competing multimodal methods fail to achieve substantial improvements over the language baselines. In contrast, our proposed multimodal baseline consistently outperforms both the language models and the other multimodal methods. Our multimodal approach demonstrates the effectiveness of aligned multimodal cues in addressing these social tasks.

5.3. Effects of Visual Features

We conduct ablation studies on visual feature types to analyze the contribution of each component in our baseline model. Table 5 shows the performance results according to the types of encoded non-verbal cues (*i.e.*, gesture and gaze features) for our social tasks. Our final baseline model encodes speaker kinesics using keypoints related to gaze and gesture, specifically [*nose, l-eye, r-eye, l-shoulder, r-shoulder, l-elbow, r-elbow, l-wrist, r-wrist*]. “w/o gesture feature” indicates the model utilizing only the head-related keypoints of [*nose, l-eye, r-eye*] while “w/o gaze feature” employs only the gesture-related keypoints of [*l-shoulder, r-shoulder, l-elbow, r-elbow, l-wrist, r-wrist*]. We adopt the BERT-based baseline for evaluation on YouTube dataset. As shown in the table, the results show that the gesture features are more dominant compared to the gaze features in our setting. The proposed baselines using both gesture and gaze features generally achieve good performances.

Target Task	Method	Accuracy (%)
Speaking Target Identification	w/o Visual Features	65.8
	w/o Gesture Feature	69.6
	w/o Gaze Feature	70.2
	Our Baseline	72.7
Pronoun Coreference Resolution	w/o Visual Features	60.3
	w/o Gesture Feature	64.9
	w/o Gaze Feature	66.5
	Our Baseline	65.9
Mentioned Player Prediction	w/o Visual Features	54.6
	w/o Gesture Feature	55.8
	w/o Gaze Feature	56.2
	Our Baseline	58.8

Table 5. Effects of non-verbal visual feature types on the performances for three social tasks. We adopt the BERT-based model and conduct evaluations on YouTube dataset.

Target Task	Conversation Context		Accuracy (%)
	Preceding	Following	
Speaking Target Identification	✗	✗	40.2
	✓	✗	59.1
	✓	✓	72.7
Pronoun Coreference Resolution	✗	✗	51.1
	✓	✗	63.4
	✓	✓	65.9
Mentioned Player Prediction	✗	✗	36.2
	✓	✗	47.3
	✓	✓	58.8

Table 6. Effects of the language conversation contexts on the performances for three social tasks. We adopt the BERT-based model and conduct evaluations on YouTube dataset.

5.4. Effects of Conversation Context

Conversational context plays an important role in understanding the meaning and intent behind individual utterances. To analyze the effects of conversation context on our social tasks, we conduct ablation experiments according to the preceding and following contexts of the target utterance. Table 6 presents the results of different context settings using the BERT-based baseline model on YouTube dataset. We compare three variations of the model: one that uses only the target utterance without any additional context, another that incorporates only the preceding context, and our proposed baseline which utilizes both preceding and following contexts. As shown in the table, leveraging both contexts shows the best results for all social tasks. It is noteworthy that the advantage of using the following context is relatively small in the pronoun coreference resolution task compared to the other tasks. This is because the nature of this task is to resolve the reference of pronouns against people that usually appear in the preceding context.

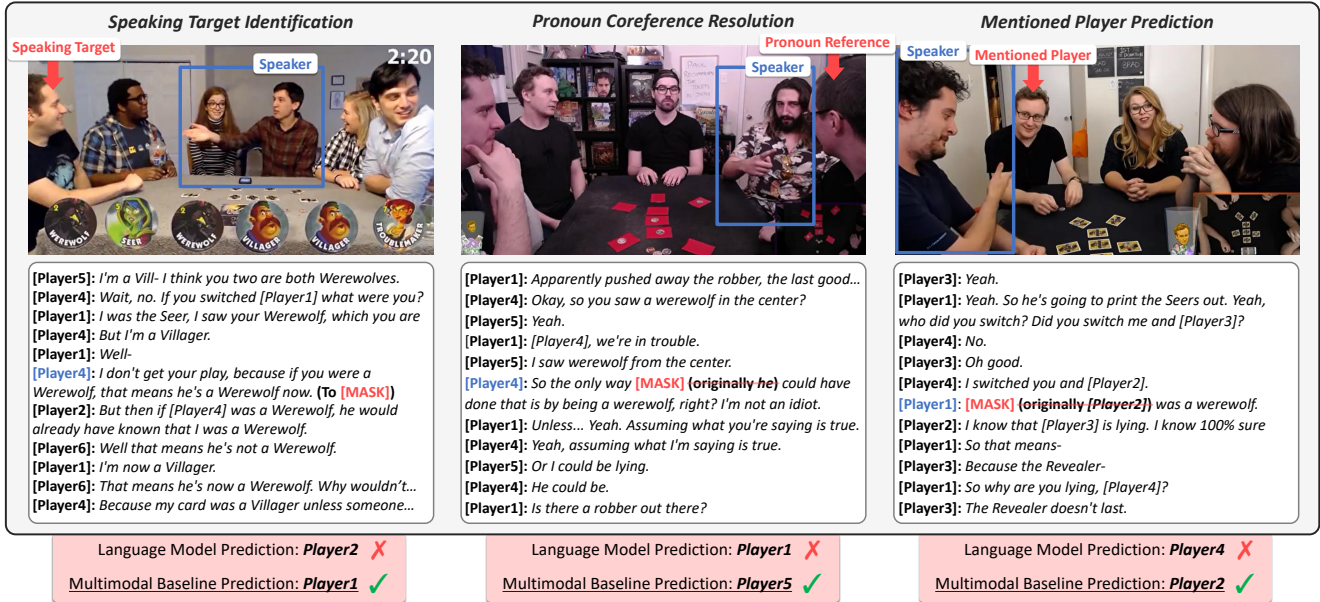


Figure 4. Qualitative results demonstrating the benefit of visual cues for three social tasks. The examples show cases where the language model alone fails, but the proposed multimodal baseline leveraging both language and visual cues correctly predicts the right person. Note that Player# are assigned in ascending order from left to right in the visual scenes of this figure.

Target Task	Player Permutation Learning	Accuracy (%)
Speaking Target Identification	X	70.8
	✓	72.7
Pronoun Coreference Resolution	X	51.5
	✓	65.9
Mentioned Player Prediction	X	52.7
	✓	58.8

Table 7. Effects of player permutation learning on the performances for three social tasks. We adopt the BERT-based model and conduct evaluations on YouTube dataset.

5.5. Effects of Permutation Learning

To validate the effectiveness of our player permutation learning which shuffles anonymized player identities, we conduct ablation experiments by training models with and without permutation. Table 7 shows the experiment results for three tasks with our BERT-based baseline on YouTube dataset. As shown in the table, the permutation learning approach consistently improves the performances for all tasks, implying it helps the model learn more generalizable representations of player interactions. Note that we apply this player permutation learning to all comparison methods in Tables 2, 3, and 4 for fair performance comparisons.

5.6. Qualitative Results

Figure 4 shows examples of three social tasks and their qualitative results according to the use of visual cues. We

utilize BERT as the language model for this experiment. As shown in the figure, our multimodal baseline leveraging both language and visual cues in a dense alignment manner can correct the inference when the language model alone fails. The qualitative results show that visual features aligned to utterances provide complementary information to disambiguate referents in social interactions.

6. Conclusion

We introduce three new challenging tasks in social deduction games: speaking target identification, pronoun coreference resolution, and mentioned player prediction - all of which require understanding the fine-grained verbal and non-verbal dynamics between multiple people. We curate extensive dataset annotations for our new social tasks and further propose a novel multimodal baseline that establishes dense language-visual alignments between spoken utterances and player visual features. This approach enables modeling multi-party social interactions through verbal and non-verbal communication channels simultaneously. Experiments show consistent and considerable performance improvements of our multimodal baselines over other approaches without both modalities and without multimodal dense alignment. Furthermore, extensive ablation studies are conducted to validate the effectiveness of our baseline components. We release the benchmarks and source code to facilitate further research in this direction.

Acknowledgement. Portions of this project were supported in part by a gift from Meta.

References

- [1] Chaitanya Ahuja, Dong Won Lee, Ryo Ishii, and Louis-Philippe Morency. No gestures left behind: Learning relationships between spoken language and freeform gestures. In *Findings of Conference on Empirical Methods in Natural Language Processing (Findings of EMNLP)*, pages 1884–1895, 2020. 1, 3
- [2] Chaitanya Ahuja, Pratik Joshi, Ryo Ishii, and Louis-Philippe Morency. Continual learning for personalized co-speech gesture generation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20893–20903, 2023. 1, 3
- [3] Shubhra Aich, Jesus Ruiz-Santaquiteria, Zhenyu Lu, Prachi Garg, KJ Joseph, Alvaro Fernandez Garcia, Vineeth N Balasubramanian, Kenrick Kin, Chengde Wan, Necati Cihan Camgoz, et al. Data-free class-incremental hand gesture recognition. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20958–20967, 2023. 2
- [4] Sara Asai, Koichiro Yoshino, Seitaro Shinagawa, Sakriani Sakti, and Satoshi Nakamura. Emotional speech corpus for persuasive dialogue system. In *Language Resources and Evaluation Conference (LREC)*, pages 491–497, 2020. 3
- [5] Xiaoheng Bi and Tetsuro Tanaka. Human-side strategies in the werewolf game against the stealth werewolf strategy. In *International Conference on Computers and Games (ICCG)*, pages 93–102, 2016. 3
- [6] Mark Braverman, Omid Etesami, and Elchanan Mossel. Mafia: A theoretical study of players and coalitions in a partial information environment. *The Annals of Applied Probability*, 18(3):825–846, 2008. 3
- [7] Jean Carletta. Assessing agreement on classification tasks: the kappa statistic. *Computational Linguistics*, 22(2):249–254, 1996. 4
- [8] Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May, and Jonathan Gratch. Casino: A corpus of campsite negotiation dialogues for automatic negotiation systems. In *Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT)*, pages 3167–3185, 2021. 3
- [9] Feiyu Chen, Jie Shao, Shuyuan Zhu, and Heng Tao Shen. Multivariate, multi-frequency and multimodal: Rethinking graph neural networks for emotion recognition in conversation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10761–10770, 2023. 1, 3
- [10] Gokul Chittaranjan and Hayley Hung. Are you awerewolf? detecting deceptive roles and outcomes in a conversational role-playing game. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5334–5337, 2010. 3
- [11] Eunji Chong, Yongxin Wang, Nataniel Ruiz, and James M Rehg. Detecting attended visual targets in video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5396–5406, 2020. 2
- [12] Robert Chuchro. Training an assassin ai for the resistance: Avalon. *arXiv preprint arXiv:2209.09331*, 2022. 3
- [13] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations (ICLR)*, 2020. 6, 7
- [14] Meta Fundamental AI Research Diplomacy Team (FAIR)†, Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074, 2022. 3
- [15] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6824–6835, 2021. 6, 7
- [16] Lifeng Fan, Yixin Chen, Ping Wei, Wenguan Wang, and Song-Chun Zhu. Inferring shared attention in social scene videos. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6460–6468, 2018. 2
- [17] Hao-Shu Fang, Jiefeng Li, Hongyang Tang, Chao Xu, Haoyi Zhu, Yuliang Xiu, Yong-Lu Li, and Cewu Lu. Alpha-pose: Whole-body regional multi-person pose estimation and tracking in real-time. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 5
- [18] Yi Fang, Jiapeng Tang, Wang Shen, Wei Shen, Xiao Gu, Li Song, and Guangtao Zhai. Dual attention guided gaze target detection in the wild. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11390–11399, 2021. 2
- [19] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18995–19012, 2022. 3
- [20] Matthew W Hoffman, David B Grimes, Aaron P Shon, and Rajesh PN Rao. A probabilistic model of gaze imitation and shared attention. *Neural Networks*, 19(3):299–310, 2006. 2
- [21] Dou Hu, Xiaolong Hou, Lingwei Wei, Lianxin Jiang, and Yang Mo. Mm-dfn: Multimodal dynamic fusion network for emotion recognition in conversations. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7037–7041, 2022. 1, 3
- [22] Guimin Hu, Ting-En Lin, Yi Zhao, Guangming Lu, Yuchuan Wu, and Yongbin Li. Unimse: Towards unified multimodal sentiment analysis and emotion recognition. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7837–7851, 2022. 1, 3
- [23] Wenxiang Jiao, Haiqin Yang, Irwin King, and Michael R Lyu. Higr: Hierarchical gated recurrent units for utterance-level emotion recognition. In *Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT)*, pages 397–406, 2019. 3
- [24] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT)*, pages 4171–4186, 2019. 5, 6, 7

- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 6
- [26] Klaus Krippendorff. *Content analysis: An introduction to its methodology*. Sage publications, 2018. 4
- [27] Bolin Lai, Miao Liu, Fiona Ryan, and James M Rehg. In the eye of transformer: Global-local correlation for egocentric gaze estimation. In *The British Machine Vision Conference (BMVC)*, 2022. 2
- [28] Bolin Lai, Hongxin Zhang, Miao Liu, Aryan Pariani, Fiona Ryan, Wenqi Jia, Shirley Anugrah Hayati, James Rehg, and Diyi Yang. Werewolf among us: Multimodal resources for modeling persuasion behaviors in social deduction games. In *Findings of the Association for Computational Linguistics (Findings of ACL)*, pages 6570–6588, 2023. 2, 3, 6, 7
- [29] Dong Won Lee, Chaitanya Ahuja, and Louis-Philippe Morency. Crossmodal clustered contrastive learning: Grounding of spoken language to gesture. In *Companion Publication of International Conference on Multimodal Interaction (ICMI)*, pages 202–210, 2021. 1, 3
- [30] Yunan Li, Huizhou Chen, Guanwen Feng, and Qiguang Miao. Learning robust representations with information bottleneck and memory network for rgb-d-based gesture recognition. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20968–20978, 2023. 2
- [31] Jianbo Liu, Yongcheng Liu, Ying Wang, Veronique Prinert, Shiming Xiang, and Chunhong Pan. Decoupled representation learning for skeleton-based gesture recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5751–5760, 2020. 2
- [32] Xian Liu, Qianyi Wu, Hang Zhou, Yinghao Xu, Rui Qian, Xinyi Lin, Xiaowei Zhou, Wayne Wu, Bo Dai, and Bolei Zhou. Learning hierarchical cross-modal association for co-speech gesture generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10462–10472, 2022. 1, 3
- [33] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 6, 7
- [34] Ganeshan Malhotra, Abdul Waheed, Aseem Srivastava, Md Shad Akhtar, and Tanmoy Chakraborty. Speaker and time-aware joint contextual learning for dialogue-act classification in counselling conversations. In *ACM International Conference on Web Search and Data Mining (WSDM)*, pages 735–745, 2022. 3
- [35] Noritsugu Nakamura, Michimasa Inaba, Kenichi Takahashi, Fujio Toriumi, Hirotaka Osawa, Daisuke Katagami, and Kousuke Shinoda. Constructing a human-like agent for the werewolf game using a psychological model based multiple perspectives. In *IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–8, 2016. 3
- [36] Chihiro Nakatani, Hiroaki Kawashima, and Norimichi Ukita. Interaction-aware joint attention estimation using people attributes. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10224–10233, 2023. 2
- [37] Sanika Natu, Shounak Sural, and Sulagna Sarkar. External commonsense knowledge as a modality for social intelligence question-answering. In *IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3044–3050, 2023. 3
- [38] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2023. 6, 7
- [39] Georgios Paraskevopoulos, Efthymios Georgiou, and Alexandras Potamianos. Mmlatch: Bottom-up top-down fusion for multimodal sentiment analysis. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4573–4577, 2022. 1, 3
- [40] Libo Qin, Wanxiang Che, Yangming Li, Mingheng Ni, and Ting Liu. Dcr-net: A deep co-interactive relation network for joint dialog act recognition and sentiment classification. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 8665–8672, 2020. 3
- [41] Libo Qin, Zhouyang Li, Wanxiang Che, Minheng Ni, and Ting Liu. Co-gat: A co-interactive graph attention network for joint dialog act recognition and sentiment classification. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 13709–13717, 2021. 3
- [42] Tulika Saha, Aditya Patra, Sriparna Saha, and Pushpak Bhat-tacharyya. Towards emotion-aided multi-modal dialogue act classification. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4361–4372, 2020. 1, 3
- [43] Jack Serrino, Max Kleiman-Weiner, David C Parkes, and Josh Tenenbaum. Finding friend and foe in multi-agent games. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019. 3
- [44] Weizhou Shen, Siyue Wu, Yunyi Yang, and Xiaojun Quan. Directed acyclic graph network for conversational emotion recognition. In *Annual Meeting of the Association for Computational Linguistics and International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, pages 1551–1560, 2021. 3
- [45] Omer Sumer, Peter Gerjets, Ulrich Trautwein, and Enkelejda Kasneci. Attention flow: End-to-end joint attention estimation. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3327–3336, 2020. 2
- [46] Vipul Raheja Joel Tetreault. Dialogue act classification with context-aware self-attention. In *Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT)*, pages 3727–3733, 2019. 3
- [47] Francesco Tonini, Nicola Dall’Asen, Cigdem Beyan, and Elisa Ricci. Object-aware gaze target detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 21860–21869, 2023. 2
- [48] Danyang Tu, Xiongkuo Min, Huiyu Duan, Guodong Guo, Guangtao Zhai, and Wei Shen. End-to-end human-gaze-target detection with transformers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2192–2200, 2022. 2

- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017. [5](#), [6](#)
- [50] Xuewei Wang, Weiyan Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. Persuasion for good: Towards a personalized persuasive dialogue system for social good. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 5635–5649, 2019. [3](#)
- [51] Baijun Xie and Chung Hyuk Park. Multi-modal correlated network with emotional reasoning knowledge for social intelligence question-answering. In *IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3075–3081, 2023. [3](#)
- [52] Amir Zadeh, Michael Chan, Paul Pu Liang, Edmund Tong, and Louis-Philippe Morency. Social-iq: A question answering benchmark for artificial social intelligence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8807–8817, 2019. [3](#)
- [53] AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2236–2246, 2018. [1](#), [3](#)
- [54] Liang Zhang, Guangming Zhu, Lin Mei, Peiyi Shen, Syed Afaq Ali Shah, and Mohammed Bennamoun. Attention in convolutional lstm for gesture recognition. *Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018. [2](#)
- [55] Weixiang Zhao, Yanyan Zhao, Zhuojun Li, and Bing Qin. Knowledge-bridged causal interaction network for causal emotion entailment. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 14020–14028, 2023. [3](#)
- [56] Benjia Zhou, Pichao Wang, Jun Wan, Yanyan Liang, Fan Wang, Du Zhang, Zhen Lei, Hao Li, and Rong Jin. Decoupling and recoupling spatiotemporal representation for rgb-d-based motion recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20154–20163, 2022. [2](#)
- [57] Lixing Zhu, Gabriele Pergola, Lin Gui, Deyu Zhou, and Yulan He. Topic-driven and knowledge-aware transformer for dialogue emotion detection. In *Annual Meeting of the Association for Computational Linguistics and International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, pages 1571–1582, 2021. [3](#)