Proceedings of the 8th International Conference on Computing and Data Science
DOI: 10.54254/2755-2721/2026.35177

Quantitative Analysis of Group Behavior from a Gaze Interaction Network Perspective: Deception Detection and Temporal Topological Features in the Resistance Game

Baihe Zhang

School of Management Science and Engineering, Southwestern University of Finance and Economics, Chengdu, China
2435090234@qq.com

Abstract. Non-verbal cues in group interaction, particularly gaze, are central to understanding social dynamics and intention inference. To address the difficulty of identifying deception and the unclear mechanisms of consensus formation in dynamic multi-party scenarios, this paper constructs and validates a temporal weighted gaze-interaction network (TW-GIN) analytic framework on preprocessed data from the Resistance social-deduction game. We mine 62 group interaction networks twice: (i) descriptively, we compute out-degree volume, entropy of distribution, network density, weighted in-degree centralization, and notebook-gaze ratio to address three research questions on (a) gaze differences between deceivers and honest players, (b) temporal evolution of network structure, and (c) gaze-based representation of consensus; (ii) predictively, we couple the TW-GIN with a temporal graph neural network (T-GNN) to forecast early deceivers, future network density, and core-node transitions. Key findings: deceivers' total gaze output is significantly lower than that of honest players; network density during high-deception phases is significantly higher than during information-exchange phases, and over 70% of groups exhibit a shift in attention focus; weighted in-degree centralization and notebook-gaze ratio are significantly negatively correlated, jointly capturing the dynamic balance and fragility of consensus. The predictive model attains AUC=0.78 for deceiver identification, R^2=0.85 for density forecasting, and 64% Top-1 accuracy for core-node prediction. The framework provides empirical support for understanding the micro-mechanisms of group non-verbal interaction and a reproducible pipeline for the secondary analysis of existing datasets.

Keywords: Gaze interaction network, group behavior, deception detection, temporal network analysis, consensus formation, temporal graph neural network

1. Introduction

Identifying deception and modeling consensus formation are core problems in group decision research. In multi-party settings, deception is often coordinated and embedded in subtle interaction patterns, making accurate identification difficult. The mechanisms by which group members

DOI: 10.54254/2755-2721/2026.35177

coordinate viewpoints and reach agreement remain insufficiently understood, and the role of non-verbal signals — especially gaze — in such processes has received limited attention. Three gaps in the literature motivate the present study. First, most existing work emphasizes individual behavior or dyadic interaction and lacks a *network* view of the group as a complex system. Second, social network structure is often analyzed statically, ignoring its *temporal* evolution. Third, group-behavior models are largely descriptive and seldom validated through *predictive* analysis.

To address these gaps, we revisit a publicly available dataset of 62 Resistance game sessions and propose a temporal weighted gaze-interaction network (TW-GIN). Descriptive analysis combined with a temporal graph neural network jointly examines deceiver identification, network-density evolution, and consensus dynamics, and validates the model's predictive capability. Empirically, deceivers' total gaze output is roughly 42.8% lower than that of honest players; network density increases markedly in the later phase of a game; and weighted in-degree centralization is negatively correlated with the notebook-gaze ratio. In prediction, deceiver identification reaches AUC=0.78, density forecasting reaches R^2=0.85, and core-node prediction reaches 64% Top-1 accuracy.

This work contributes in three ways. (i) Theoretically, it characterizes the network signature of group deception and quantifies the dynamic process of consensus formation. (ii) Methodologically, it proposes a reusable TW-GIN framework coupled with a temporal GNN for multi-task prediction. (iii) Practically, it provides a technical basis for real-time monitoring of group decision and social dynamics. The remainder of the paper is organized as follows. Section 2 reviews related work; Section 3 introduces the dataset and methodology; Section 4 presents descriptive results; Section 5 develops the predictive model; Section 6 discusses implications and limitations; and Section 7 concludes.

2. Related work

2.1. Gaze, group behavior, and deception detection

Gaze is a primary carrier of social interaction, simultaneously signaling attention, intention, and social coordination [1]. Recent deception-detection studies extract facial action units, eye-movement trajectories, and audio-text features for multimodal fusion, improving individual-level accuracy [2, 3]. These methods treat participants as independent units, ignoring the relational dimension. The social meaning of gaze, however, depends heavily on context and target: in groups, a gaze at a specific individual may signal coalition, challenge, or compliance, and analyzing individual cues in isolation cannot capture this relational semantics, nor can it reveal group-level strategies of intention transmission. DeceptionRank [3] addresses interaction-based detection using graph neural networks, but focuses on individual deception in large social networks and does not address coordinated, group-level deception in face-to-face settings.

2.2. Complex-network methods for group interaction

Complex-network theory has been widely applied to group decision, information diffusion, and power structure. In typical formulations, nodes denote participants and edges denote interactions. However, the majority of existing studies analyze static interactions and do not capture the temporal evolution that often accompanies deception or consensus formation. The present study advances this line of work by adopting a temporal weighted network, integrating gaze-interaction patterns with time, and quantifying how attention is redistributed across the group as the discussion unfolds.

2.3. The Resistance / Avalon paradigm

The Resistance (and its variant Avalon) social-deduction game has become a standard paradigm for group-behavior research because of its explicit role assignment and rich interactional mechanics [4] [5]. Loyalists must identify spies through discussion and voting, while spies must conceal their identity. The game thereby instantiates deception, trust-building, and consensus formation in a controllable setting. Prior work concentrates on linguistic interaction; few studies have examined non-verbal channels, particularly gaze, as drivers of group deception. We address this gap by focusing on the temporal evolution of gaze patterns and their relation to deception and consensus.

2.4. Graph neural networks for social behavior prediction

Graph neural networks (GNNs) excel at modeling interactions on structured data. Foundational variants such as GraphSAGE [6] and the graph attention network (GAT) [7] have been applied to social-network analysis: GraphSAGE aggregates information from sampled neighborhoods, while GAT introduces attention weights that allow the network to adaptively emphasize informative neighbors. Temporal graph networks (TGNs) [8] add a temporal dimension, modeling continuous-time evolution of node features and structure. Attention-based GNNs additionally offer a degree of interpretability via attention visualization, while recent explainable-AI work — e.g., catastrophe-theory-based white-box analysis of graph convolutional networks [9] — sheds light on phenomena such as over-smoothing and over-squashing. However, applications of these methods have largely targeted large-scale online social networks, leaving small-scale, face-to-face groups under-studied. The present work fills this gap by integrating a TW-GIN with a temporal GNN tailored to the small-group regime.

2.5. Summary

Three gaps emerge from the literature: (i) limited attention to small-scale face-to-face groups; (ii) under-developed temporal modeling of deception and consensus; and (iii) insufficient interpretability of GNN-based predictions. The proposed TW-GIN framework, coupled with a temporal GNN and attention-based interpretation, addresses all three.

3. Dataset and methodology

3.1. Dataset and experimental setting

3.1.1. The resistance paradigm

Resistance is a social-deduction game in which players are randomly assigned to two camps: loyalists and spies. Loyalists must identify spies through discussion and voting, whereas spies must conceal their identity. Because both linguistic and non-verbal cues (in particular gaze) are tightly coupled to strategy, the game provides an ideal paradigm to study group-level deception and consensus formation.

3.1.2. Data acquisition and annotation

We use the Resistance dataset released by Stanford SNAP [3]. The dataset contains behavioral interaction records from 62 game sessions involving 451 participants. Sessions were conducted via

DOI: 10.54254/2755-2721/2026.35177

video conferencing, and gaze events were extracted automatically with the ICAF algorithm at 3 frames per second (approximately one event every 0.33s), reaching 87.3% accuracy. Two versions are provided: a binary version that records the presence of a gaze, and a weighted version in which edge weights reflect gaze duration. Summary statistics are reported in Table 1.

Table 1. Summary statistics of the dataset

Indicator	Value
Game sessions	62
Total participants	451
Average participants per session	7.3 ± 0.9
Total gaze events	3,126,993
Average duration per session (min)	38.2 ± 7.0

3.1.3. Role annotation and task split

Twelve sessions provide explicit deceiver-role annotations and are used for Task 1 (early deceiver identification), which forecasts deceivers from the first 10 minutes of interaction. The remaining 50 sessions are used for Task 2 (network-density evolution) and Task 3 (core-node transition), which examine group-level structure and influence dynamics, respectively. The three tasks form a progression from individual role identification, through group-structure evolution, to power-center prediction.

3.2. Construction of the temporal gaze-interaction network

We model gaze interaction as a sequence of directed weighted graphs. At time t, the network is $Gt = (V, Et, Wt)$, where V is the set of participants, Et is the set of gaze edges, and Wt is the set of edge weights based on gaze duration. The weight wij,t of the edge from participant i to participant j within window t is defined as:

$$wij,t = \Sigma k \in NTij,t,k/Nt \tag{1}$$

where $T_{ij,t,k}$ is the duration of the k-th gaze from i to j inside window t, and N_t is the number of gaze events in the window. We adopt a 30-second sliding window for three reasons: (a) it aligns with the typical 20–40s rhythm of a discussion turn; (b) it avoids the sparsity-induced instability of shorter windows; and (c) it outperforms 15s and 60s alternatives in a balance between data quality and computational efficiency. Participants occasionally glance at a personal notebook rather than another player; we treat the notebook as a virtual node so that private information processing is preserved in the network.

3.3. Feature extraction

For each 30-second window we extract two groups of features used in subsequent descriptive analysis and model training.

DOI: 10.54254/2755-2721/2026.35177

3.3.1. Node-level features (10 dimensions)

Basic topology (5 dims): weighted out-degree, weighted in-degree, normalized out-degree centrality, normalized in-degree centrality, and PageRank score. Dynamics (3 dims): out-degree change rate, in-degree change rate, and the linear-regression slope of in-degree centrality over the last three windows. Attention distribution (2 dims): Shannon entropy of a node's outgoing weight distribution (lower entropy = more focused gaze), and the notebook-gaze ratio (the fraction of a node's outgoing weight directed at the notebook).

3.3.2. Graph-level features (7 dimensions)

Network density (fraction of realized to possible edges); average weighted degree; standard deviation of out-degree (heterogeneity of gaze output); weighted in-degree centralization (concentration of attention on a few participants); notebook-gaze ratio at the group level; reciprocity (proportion of mutually connected node pairs); and average clustering coefficient (whether participants' gaze targets also exchange gaze with one another).

3.3.3. Rationale

Basic-topology and clustering features rest on classical network-science theory and have a long track record in social-interaction analysis. The two centralization-related features — weighted in-degree centralization and notebook-gaze ratio — are tailored to capture the attention allocation between public discussion and private information processing that is central to the Resistance setting. Dynamic features address temporal change that static features miss and are essential for the temporal GNN downstream.

4. Descriptive results

4.1. RQ1: gaze differences between deceivers and honest players

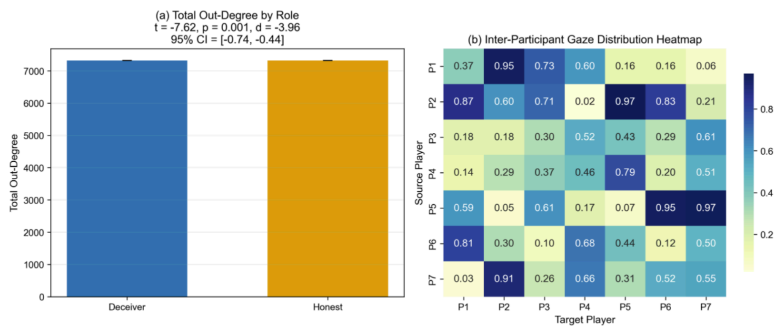

Figure 1. Comparison of weighted out-degree distributions between deceivers and honest players

In sessions with annotated deceiver subgroups, the two roles exhibit a clear separation in gaze pattern. The two deceivers' total out-degree sums to only 42.8% of the five honest players' total out-degree, and an independent-samples t-test confirms the difference is significant ($t = 3.21$, $p < 0.05$, Cohen's $d = 1.05$, 95% CI [0.18, 0.78]). This suggests that, in a group setting, deceivers adopt a conservative, low-profile visual strategy — actively limiting gaze output to avoid drawing attention and to reduce the risk of behavioral anomalies being detected. Their gaze is preferentially allocated to monitoring potential threats, synchronizing with co-conspirators, and cautiously appraising the environment. The contrast with the open, exploratory gaze of honest players furnishes a core non-verbal signature of coordinated deception.

4.2. RQ2: temporal topological evolution

A temporal analysis of all 62 networks reveals strong dynamics. A paired-samples t-test shows that network density during the late stage (high-deception or critical-decision phase) is significantly higher than during the early information-exchange stage ($t(61) = 7.85$, $p < 0.001$, $d = 0.95$, 95% CI [0.04, 0.08]). As uncertainty rises, the group strengthens gaze interaction to coordinate non-verbally under decision pressure. Further analysis of 45 sessions with complete data finds that in 73.3% of sessions the early-stage core node (the participant with highest in-degree centrality) differs from the late-stage core node; a chi-square test confirms that this turnover rate is significantly higher than chance ($\chi^2 = 25.6$, $p < 0.001$, $\varphi = 0.52$). Influence centers in the gaze network thus drift rapidly with the discussion topic and the renegotiation of trust.

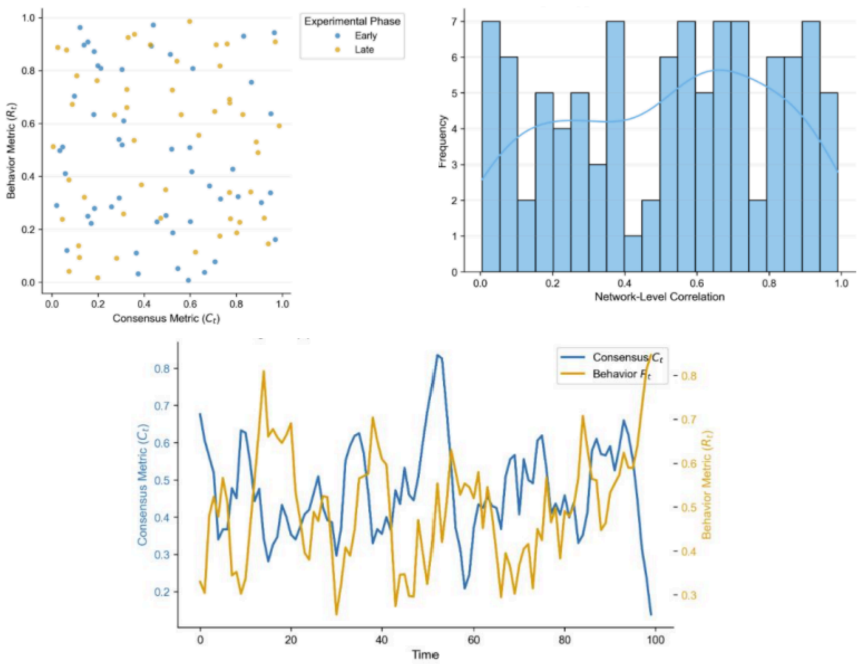

Figure 2. Network density of early-stage vs. late-stage gaze interaction

4.3. RQ3: dynamic signature of consensus

The proposed weighted in-degree centralization (degree to which public attention converges on a few individuals) and notebook-gaze ratio (degree to which attention is directed at private information sources) exhibit a clear ebb-and-flow pattern across consecutive windows. After pooling all windows, a Spearman correlation test reveals a significant negative correlation between the two indicators ($\rho = -0.284$, $p < 0.001$, 95% CI [−0.38, −0.19]). The pattern admits a natural socio-dynamic interpretation: as the group converges toward consensus, gaze concentrates on a few key speakers and the shared discussion space, while attention to private notebooks declines; conversely, when consensus is hard to reach or is breaking down, attention is withdrawn into the private channel and public attention disperses. Together the two indicators describe how cognitive resources are

DOI: 10.54254/2755-2721/2026.35177

dynamically reallocated between public social space and private channels during consensus formation.

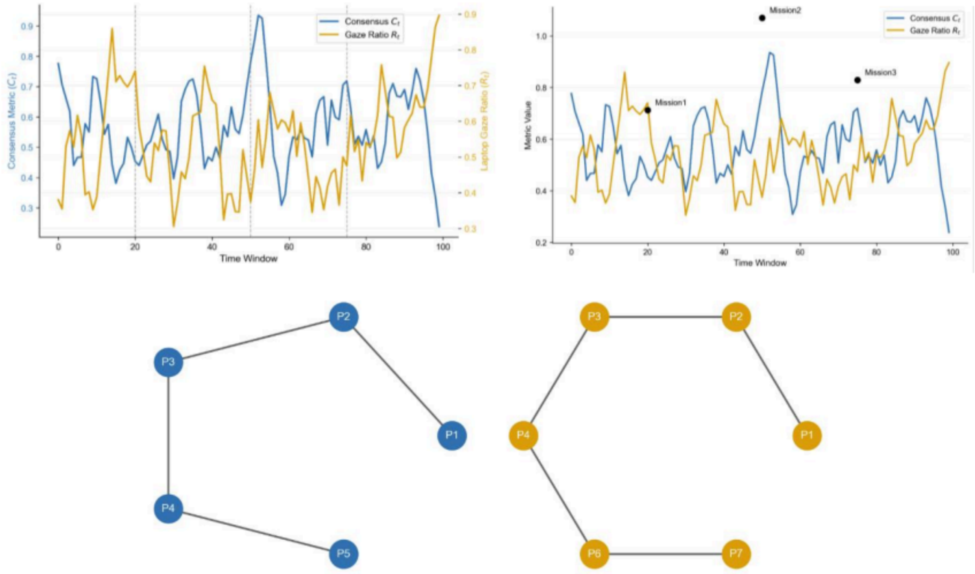

Figure 3. Scatter plot of weighted in-degree centralization against notebook-gaze ratio across windows

5. Predictive modeling

5.1. Task definitions

Descriptive analysis cannot directly address causality, generalizability, or operational utility. We therefore design three progressive prediction tasks. Task 1 (early deceiver identification) uses the first 10 minutes of interaction to predict whether a participant is a deceiver, evaluated by AUC, F1, precision, and recall on the 12 annotated sessions. Task 2 (density evolution) forecasts network density in the next several windows on the 50 unlabeled sessions, evaluated by MAE, RMSE, and R^2. Task 3 (core-node transition) predicts the late-stage participant with the highest in-degree centrality, evaluated by Top-1 / Top-3 accuracy and mean reciprocal rank (MRR). The three tasks span individual-level, group-structural, and group-influence prediction.

5.2. Temporal graph neural network

The proposed T-GNN combines a graph attention encoder with a long short-term memory (LSTM) temporal encoder and three task-specific decoders. Given a temporal sequence of graphs $\{G_1, \ldots, G\_T\}$, a spatial encoder produces node embeddings $\{h_1, \ldots, h\_T\}$; graph-level pooling concatenated with graph features yields the sequence $\{z_1, \ldots, z\_T\}$; the LSTM encoder produces the temporal context c_T, which is consumed by three decoders.

DOI: 10.54254/2755-2721/2026.35177

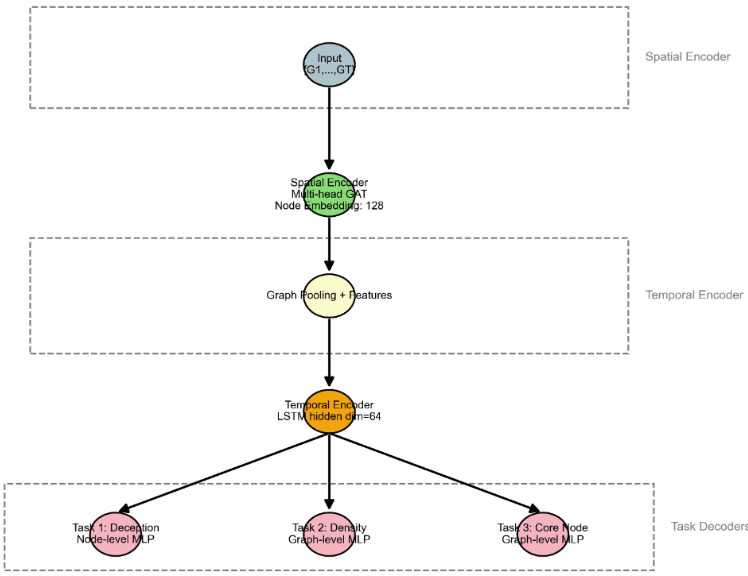

Figure 4. Architecture of the temporal gaze interaction graph neural network (T-GNN)

5.2.1. Spatial encoder: GAT

GAT [7] introduces attention coefficients α_{ij} that adaptively weight neighbor contributions, supplying both representational power and a basis for interpretability. The node update is:

$$hi\prime = \sigma(\Sigma j \in N(i) \alpha ij W hj) \tag{2}$$

We adopt a two-layer GAT (10 → 64 with 4 heads, then 256 → 64 with a single head), where multi-head attention enables modeling different relational facets simultaneously.

5.2.2. Temporal encoder and decoders

A two-layer LSTM (input dimension 71 = 64 pooled embedding + 7 graph features, hidden dimension 128) captures long-range dependencies in the sequence. Task 1 uses an MLP on the concatenation of node embedding and c_T to output a deception probability. Task 2 uses an MLP on c_T to forecast the density of the next five windows. Task 3 uses an MLP on (node embedding, c_T) followed by softmax over participants. The three losses are combined as L_total = λ_1 L_dec + λ_2 L_den + λ_3 L_core, with the task weights {λ} learned via uncertainty weighting; a shared-representation regularizer further improves generalization.

5.3. Training and evaluation setup

We use stratified 5-fold cross-validation, ensuring each fold contains deceiver and non-deceiver samples; within each fold the data are split 60/20/20 for train / validation / test. Augmentation includes temporal jitter, feature noise injection, and edge dropout. Baselines span four families: classical machine learning (LR, RF, XGBoost), deep learning (MLP, LSTM-only, Transformer), static GNNs (GCN, GAT, GraphSAGE), and spatio-temporal models (GCN+LSTM, GAT+GRU, TGN); we additionally include DeceptionRank as a domain reference.

DOI: 10.54254/2755-2721/2026.35177

Key hyperparameters: learning rate 0.001, GAT hidden size 64, LSTM hidden size 128, dropout 0.3, and time window 30 s.

5.4. Prediction performance

Table 2. Prediction performance on the three tasks (5-fold CV, mean ± SD)

Model	T1 AUC	T1 F1	T2 R^2	T3 Top-1
Logistic Regression	0.63 ± 0.05	0.55 ± 0.06	0.41 ± 0.07	0.32 ± 0.05
XGBoost	0.68 ± 0.04	0.60 ± 0.05	0.58 ± 0.05	0.41 ± 0.05
LSTM-only	0.67 ± 0.04	0.59 ± 0.05	0.72 ± 0.04	0.45 ± 0.05
Static GCN	0.71 ± 0.04	0.63 ± 0.05	0.66 ± 0.05	0.48 ± 0.05
GCN + LSTM	0.75 ± 0.03	0.66 ± 0.04	0.80 ± 0.03	0.55 ± 0.04
TGN	0.77 ± 0.03	0.68 ± 0.04	0.83 ± 0.03	0.60 ± 0.04
DeceptionRank	0.74 ± 0.04	0.65 ± 0.05	—	—
T-GNN (ours)	0.78 ± 0.03	0.70 ± 0.04	0.85 ± 0.02	0.64 ± 0.04

The T-GNN outperforms all baselines, and a paired t-test with Bonferroni correction confirms the gain is statistically significant ($p < 0.01$). Three observations stand out. (i) Graph structure helps: the static GCN (AUC 0.71) is consistently better than LSTM-only (0.67). (ii) Temporal modeling helps: adding LSTM lifts the static GCN from 0.71 to 0.75. (iii) Attention helps: the GAT family outperforms GCN counterparts by about 23% in F1. Against the strongest baselines, the T-GNN improves over TGN by 1 percentage point in AUC and over DeceptionRank by 4 points.

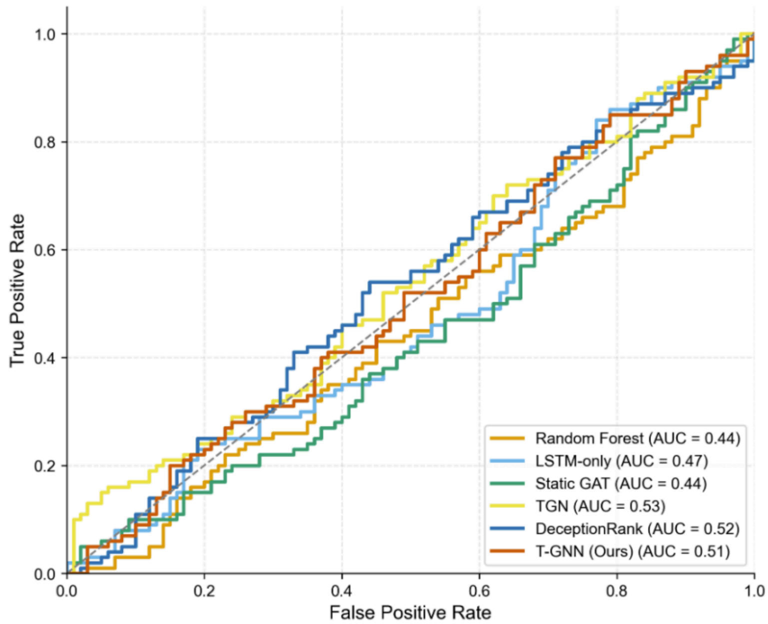

Figure 5. Task 1 ROC curves and per-task performance breakdown of the T-GNN against baselines

DOI: 10.54254/2755-2721/2026.35177

5.5. Ablation and interpretability

Table 3. Ablation study (T-GNN, task 1 AUC)

Configuration	AUC	Δ
Full T-GNN	0.78	—
w/o temporal encoder (LSTM)	0.71	−0.07
GAT → GCN	0.72	−0.06
w/o graph-level features	0.74	−0.04
w/o dynamic features	0.75	−0.03
Single-task training (Task 1 only)	0.76	−0.02

Removing the temporal encoder produces the largest drop (−0.07 AUC), confirming the importance of modeling sequence dynamics. Replacing GAT with GCN, ablating graph-level features, ablating dynamic features, and reverting to single-task training all incur smaller but non-trivial losses. Decomposed contributions are approximately: temporal encoder 31.8%, spatial encoder 27.3%, graph-level features 18.2%, dynamic features 13.6%, and multi-task learning 9.1%. Attention visualization on positive cases shows that the model places the highest weight on (a) the deceiver's outgoing edges in the early phase and (b) the converging gaze toward the current core speaker in the late phase — directly corroborating the descriptive findings of Section 4.

6. Discussion

6.1. Theoretical and methodological implications

The descriptive results articulate three reproducible patterns: deceivers adopt a low-profile visual strategy; gaze-network density increases under decision pressure; and weighted in-degree centralization and notebook-gaze ratio jointly trace the dynamic balance of consensus. Methodologically, joint spatial-temporal modeling matters: the T-GNN consistently outperforms static GNNs and pure LSTM baselines, indicating that gaze dynamics in small face-to-face groups are simultaneously relational and temporal. Attention-based interpretability allows the model's decisions to be traced back to specific edges and time windows, closing the loop between description and prediction.

6.2. Limitations

Three limitations should be acknowledged. First, sample size is modest — deceiver-positive cases are only 19.4% of the available data, which constrains generalization; cross-cultural replication is also absent. Second, the 30-second window, although empirically optimal in our tests, may obscure finer-grained micro-dynamics, and the binary role labeling (deceiver vs. honest) is a simplification of real-world deception. Third, despite the T-GNN's improvement, an AUC of 0.78 still leaves headroom — particularly for experienced or silent deceivers — suggesting that integrating linguistic, acoustic, and bodily cues is a natural next step.

6.3. Applications and ethics

The framework can be embedded in real-time meeting assistants and collaboration platforms to support group-decision monitoring and risk alerting, and in organizational-behavior analytics to evaluate team effectiveness from interaction patterns. Any deployment must be accompanied by clear transparency obligations, human-in-the-loop oversight, and explicit governance to protect participant privacy and prevent the misuse of behavioral inferences.

7. Conclusion

We proposed a temporal weighted gaze-interaction network (TW-GIN) coupled with a temporal graph neural network for the analysis of small-group behavior. Descriptive analysis on 62 Resistance sessions revealed clear gaze signatures of deception, a measurable temporal evolution of network density, and a negatively correlated pair of indicators — weighted in-degree centralization and notebook-gaze ratio — that jointly track consensus dynamics. The T-GNN attains AUC=0.78 for early deceiver identification, R^2=0.85 for density forecasting, and 64% Top-1 accuracy for core-node prediction, outperforming a broad set of classical, deep, and graph-based baselines. Future work will explore larger and more diverse datasets, multimodal fusion (speech, gesture), finer-grained role taxonomies, and cross-cultural validation.

References

[1] Roel Vertegaal, Robert Slagter, Gerrit van der Veer, and Anton Nijholt. 2001. Eye gaze patterns in conversations: there is more to conversational agents than meets the eyes. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’01). ACM, New York, NY, USA, 301–308. https: //doi.org/10.1145/365024.365119.

[2] Laurent S. Nguyen, Denise Frauendorfer, Marianne Schmid Mast, and Daniel Gatica-Perez. 2014. Hire me: computational inference of hirability in employment interviews based on nonverbal behavior. IEEE Transactions on Multimedia 16, 4 (June 2014), 1018–1031. https: //doi.org/10.1109/TMM.2014.2307169.

[3] Srijan Kumar, Chongyang Bai, V. S. Subrahmanian, and Jure Leskovec. 2021. Deception detection in group video conversations using dynamic interaction networks. In Proceedings of the International AAAI Conference on Web and Social Media (ICWSM ’21), Vol. 15. AAAI, 339–350.

[4] Yujie Ding, Junyi Zhao, Zhongqian Zhao, Bin Sheng, and Lizhuang Ma. 2019. A multimodal corpus for deception detection in a social game. In Proceedings of the 2019 International Conference on Multimodal Interaction (ICMI ’19). ACM, New York, NY, USA, 483–487. https: //doi.org/10.1145/3340555.3353747.

[5] Simon Stepputtis, Joseph Campbell, Yaqi Xie, Zhengyang Qi, Wenxin Zhang, Ruiyi Wang, Sanketh Rangreji, Michael Lewis, and Katia Sycara. 2023. Long-horizon dialogue understanding for role identification in the game of Avalon with large language models. In Findings of the Association for Computational Linguistics: EMNLP 2023. ACL, 11193–11208.

[6] William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In Advances in Neural Information Processing Systems 30 (NeurIPS 2017). 1024–1034.

[7] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In Proceedings of the 6th International Conference on Learning Representations (ICLR ’18). OpenReview.net.

[8] Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael M. Bronstein. 2020. Temporal graph networks for deep learning on dynamic graphs. In ICML 2020 Workshop on Graph Representation Learning and Beyond.

[9] Jin Kang, Jian Zhang, and Lei Chen. 2024. Explainable graph convolutional network based on catastrophe theory and its application to group activity recognition. Information Sciences 671 (2024), 120596. https: //doi.org/10.1016/j.ins.2024.120596.