

Deception Detection in Group Video Conversations using Dynamic Interaction Networks

Srijan Kumar,¹ Chongyang Bai,² V.S. Subrahmanian,² Jure Leskovec³

¹Georgia Institute of Technology

²Dartmouth College

³Stanford University

srijan@gatech.edu, chongyang.bai.gr@dartmouth.edu, vs@dartmouth.edu, jure@cs.stanford.edu

Abstract

Detecting groups of people who are jointly deceptive in video conversations is crucial in settings such as meetings, sales pitches, and negotiations. Past work on deception in videos focuses on detecting a single deceiver and uses facial or visual features only. In this paper, we propose the concept of Face-to-Face Dynamic Interaction Networks (FFDINs) and Negative Dynamic Interaction Networks (NDINs) to model the interpersonal interactions within a group of people. The use of FFDINs and NDINs enables us to leverage network relations in detecting group deception in video conversations for the first time. We use a dataset of 185 videos from a deception-based game called Resistance. We first characterize the behavior of individual, pairs, and groups of deceptive participants and compare to non-deceptive participants. Our analysis reveals that pairs of deceivers tend to avoid mutual interaction and focus their attention on non-deceivers. In contrast, non-deceivers interact with everyone equally. We propose the notion of Negative Dynamic Interaction Networks (NDINs) to capture the notion of missing interactions. We create an algorithm to detect deceivers from NDINs extracted from videos that are just one minute long. We show that our method outperforms recent state-of-the-art computer vision, graph embedding, and ensemble methods by at least 20.9% AUROC in identifying deception from videos.

Introduction

Web-based face-to-face video conversations have become a pervasive mode of work and communication throughout the world, especially since the COVID-19 pandemic. Important tasks, such as interviews, negotiations, deals, and meetings, are all happening through video call platforms such as Microsoft Teams, Google Meet, Facebook Messenger, Zoom and Skype. Furthermore, video content have become a central theme in social media and video conversations have also become an integral part of social media platforms, including on Facebook, WhatsApp, and SnapChat. Deception and disinformation in all these settings can be disruptive, counterproductive, and dangerous.

The problem of accurately and quickly identifying whether a group of people is being deceptive is crucial in

many settings. Specifically, consider the scenario where a group of deceivers work together to fool a group of unsuspecting users, but the latter do not know who the deceivers are. This occurs in practice, for instance, when defectors are present in security teams, when liars are present in sales teams, and when people from competing firms infiltrate an organization.

While there has been significant research in identifying individual deceivers in real-world face-to-face interactions (Wu et al. 2018; Ding et al. 2019; Gogate, Adeel, and Hussain 2017), little is known about how groups of deceivers work together in the online setting. Current deception research is largely limited to analysing the audio-visual behavior of a single deceiver using voice signatures, gestures, facial expressions, and body language. In contrast, research on social media analytics has extensively studied the impact of individual and team of deceivers (Kumar et al. 2017; Kumar, Spezzano, and Subrahmanian 2015; Kumar, Zhang, and Leskovec 2019; Addawood et al. 2019; Wu et al. 2017; Keller et al. 2017), but those findings do not translate to the case of video-based face-to-face group deception.

The behavioral characteristics when multiple deceivers operate simultaneously is drastically different from the case of a single deceiver. This is primarily because the behavior of one deceiver can influence the behavior of the other deceivers. For instance, when one deceiver lies to a potential target, other partners of the deceiver may show certain facial reactions which might be leveraged to predict that deception is going on. Moreover, multiple simultaneous deceivers, each deceiving a little bit, may be more successful in deceiving victims than a single deceiver alone. There is little work on studying the behavioral patterns of groups of deceivers, which is a gap that we bridge.

We conduct the first network-based study of group deception in face-to-face discussions. We analyse the verbal behavior, non-verbal behavior, and inter-personal interaction behavior when multiple deceivers are present within a group of people. We elicit deceptive behavior in the form of a multi-person face-to-face game called *Resistance*¹. Resistance is a social role-playing card-based party game, where

¹Many variants of Resistance such as Mafia and Werewolf are played around the world by thousands of people.

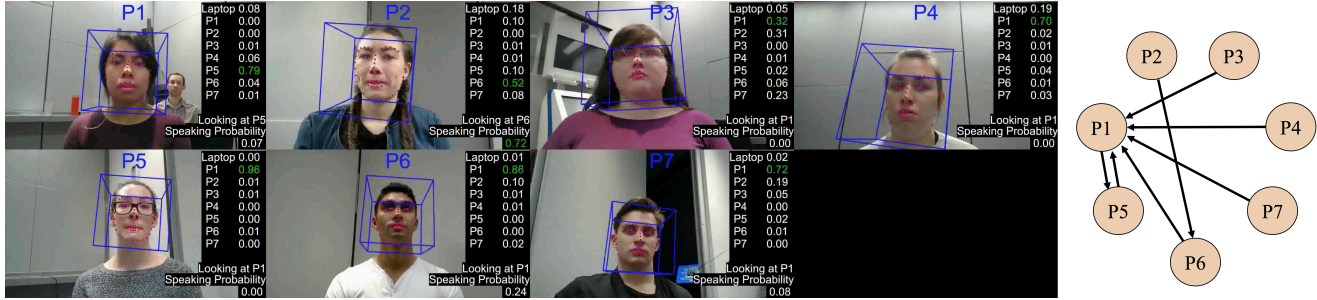


Figure 1: Given a group video conversation (left), we extract face-to-face dynamic interaction networks (right) representing the instantaneous interactions between participants. Participants are nodes and interactions are edges in the network. In this work, dynamic interaction networks are used to characterize and detect deception.

a small group (deceivers) tries to disrupt the larger group (non-deceivers) working together. In this study, we use a dataset of 26 Resistance games, each 38-minutes long on average (Bai et al. 2019).

We propose the concept of a Face-to-Face Dynamic Interaction Network (FFDIN for short) which captures instantaneous interactions between the participants. Participants are nodes and interactions are edges in the network. FFDINs include both verbal (who talks to whom) and non-verbal (who looks at whom) interactions. One FFDIN is extracted per second of the video. An example is shown in Figure 1. As discussions unfold over the course of the game, FFDINs evolve rapidly over time and their dynamics can provide valuable clues for detecting deception. We use the dynamic FFDINs extracted from the videos of the Resistance games in this work. Even though 26 games does not plenty, our dataset consists of 59,762 FFDINs in total.

We conduct a series of analysis on these networks which reveal novel behaviors of deceivers, extending the research done in social sciences (Driskell, Salas, and Driskell 2012; Baccarani and Bonfanti 2015; Vrij 2008). In particular, we find that deceivers who are less engaged (as measured by the number of participants they interact with, how often they speak, and who listens to them) are more likely to lose the game. On the other hand, deceivers successfully deceive others when they are as engaged in the game as non-deceivers, thus adeptly camouflaging themselves. Across all games, we also found that deceivers interact significantly more with non-deceivers than with other deceivers, echoing previous findings by (Driskell, Salas, and Driskell 2012). In contrast, as non-deceivers do not know the identity of other participants, they interact equally with everyone.

We introduce the notion of Negative Dynamic Interaction Networks (or NDINs) that captures when two participants *avoid* interacting with one another. We then create an algorithm called DeceptionRank that can detect deceivers even from very short (one minute) video snippets. DeceptionRank derives NDINs from the original interaction networks — two nodes are linked with an edge if their corresponding participants do not interact. Our method initializes the

prior deception scores of each node through a novel process which normalizes and aggregates the verbal and non-verbal behaviors of nodes. It then iteratively runs PageRank on and aggregates scores from the set of negative dynamic interaction networks. This generates a deception score for each node. We show that DeceptionRank outperforms the state-of-the-art computer vision, graph embedding, and ensemble methods by over 20.9% AUROC in detecting deceivers. Moreover, we show that DeceptionRank is consistently the best performing method across different lengths of video segments and regardless of the final outcome of the game.

The dynamic networks dataset along with ground truth of deception are available at: <https://snap.stanford.edu/data/comm-f2f-Resistance.html>.

Dataset Description

Here we create face-to-face dynamic verbal and non-verbal interaction networks by extending the work by Bai et al. (2019). In Bai et al. (2019), face-to-face interactions are extracted from videos of a group of participants playing Resistance game.² The extraction algorithm is a collective classification algorithm that leverages computer vision techniques for eye gaze and head pose extraction. Please refer to the paper for complete details of how the interactions are extracted. Each game has 5–8 participants, out of which a subset are assigned the roles of being deceivers (others are non-deceivers). A participant has the same role throughout the game. The deceivers know who the other deceivers are, but the non-deceivers do not know the roles of any other participant. One participant is part of exactly one game. In total, the dataset has 26 game and 185 participants.

The game has multiple rounds. Each round starts with a free-flowing discussion in which players discuss who the possible deceivers might be. Players cast votes at the end of each round. To win the game, the non-deceivers must collectively identify the deceivers as early as possible — but as they do not know who the deceivers and non-deceivers are, they must identify who is lying. The dominant winning

²[https://en.wikipedia.org/wiki/The_Resistance_\(game\)](https://en.wikipedia.org/wiki/The_Resistance_(game))

Property	Total	Per game
Number of network timeseries	26	1
Temporal length (in seconds)	59,762	2,299
Number of games won by deceivers	14	54%
Number of nodes	185	5–8
Number of look-at edges	689,501	26,519
Number of speak-to edges	26,556	1,021
Number of listen-to edges	25,798	992

Table 1: Statistics of the Resistance dataset.

strategy of the non-deceivers is to be truthful and that of deceivers is to lie and pretend that they are non-deceivers. In our dataset, deceivers win 14 out of 26 games (or 54%). Hence the data is reasonably balanced. We use “DW” to label the games that the deceivers win and “DL” to identify games that the deceivers lose.

Face-to-Face Dynamic Interaction Networks (FFDINs). We create a dataset of FFDINs. Each game is represented as a sequence of interaction networks, with one network snapshot per second. Nodes in a FFDIN represent participants in the corresponding game. Each node has a binary attribute representing its role, i.e., deceiver or non-deceiver (a participant’s role does not change during the game). An edge represents the interaction between a pair of participants during the corresponding second – we will discuss the types of edges considered shortly. The resulting FFDINs have highly dynamic edges because of the free-form discussion and interaction changes over time. All edges are directed and weighted — the weight indicates the strength or probability of the interaction. An example is shown in Figure 1. In total, there are 26 games (network sequences), 185 participants, and 996 minutes of recordings. Table 1 shows the statistics of the Resistance networks. We create three types of networks from the video of each game:

- *Look-At FFDINs* $N_G^t = (V, E_G^t)$ captures non-verbal interactions between participants. The edges at time t represent who-is-looking-at-who during the time duration t . The edge weight $E_G^t(u, v)$ is the probability of participant u looking at participant v at time t .
- *Speak-To FFDIN* $N_S^t = (V, E_S^t)$ captures verbal interactions between participants. The edges represent who speakers are looking at while speaking. At any time point, the edges emanate from speaker nodes. Edge weights represent the probabilities of speakers looking at targets.
- *Listen-To FFDIN* $N_L^t = (V, E_L^t)$ shows who listens to the speaker. The edges are incoming weighted edges directed towards the speaker node at each point in time.

Ethical and IRB Considerations

An extensive set of IRB approvals were obtained by the authors of Bai et al. (2019) to collect the data. IRB review was conducted at the institutions where the data was collected as well as the IRB of the project sponsor. The participants gave permission to the research team of Bai et al. (2019) to record and analyze their videos. After the networks are extracted from Bai et al. (2019), all personally identifiable in-

formation (PII) is stripped and original videos are not used further.

Our data and networks are derived using the interaction information provided as output by Bai et al. (2019), which has no PII. As a result, the FFDINs created in this work do not contain any PII as well. The dynamic networks dataset we release in this work does not have any PII as well.

Characterizing Deceptive Behavior

The goal of this section is to answer two important questions: (i) what are the behavioral characteristics that separate deceivers from non-deceivers? and (ii) what are the factors that distinguish *successful* deceivers from unsuccessful ones? We answer these questions through three research questions.

RQ1: Do deceivers and non-deceivers have distinct gaze patterns?

There is an asymmetry between the knowledge that deceivers and non-deceivers have. In the game, deceivers know who the deceivers and non-deceivers are. In contrast, a non-deceiver only knows her own role and nothing whatsoever about the other participants. Since deceivers know the roles of all participants, a natural question to ask is whether they focus their attention on specific participants? If yes, how does this focus affect their success in deceiving others? This is a key question as prior social science research has shown that frequent/rapid gaze change is linked to low confidence (Rayner 1998) and higher likelihood of both deception (Pak and Zhou 2013) and anxiety (Dinges et al. 2005; Laretzaki et al. 2011), which may be exhibited by users in certain roles. For instance, prior research has found that deceivers are more anxious than non-deceivers (Ströfer et al. 2016).

We analyse the behavior of deceivers in the Look-At networks. A participant u ’s “looking” behavior can be represented as a sequence of consistent gaze periods $[P_{u1}, P_{u2}, \dots, P_{un}]$. A period P_{uk} is a continuous time interval $[P_{uk}^0, P_{uk}^1]$ with a single gaze target T_{uk} , i.e., the recipient of u ’s highest weight outgoing edge at every time step in the interval. Suppose we use D_{uk} to denote the duration of participant u ’s k^{th} period P_{uk} . Thus, the duration sequence for u ’s gaze behavior is represented as $[D_{u1}, D_{u2}, \dots, D_{uk}]$, where $D_{u1} + D_{u2} + \dots + D_{uk} = T$, i.e. $D_{u1}, D_{u2}, \dots, D_{uk}$ partition the time interval T .

Gaze Entropy. We calculate the entropy of u ’s gaze behavior in the game as the entropy of the set $\{\frac{D_{u1}}{T}, \frac{D_{u2}}{T}, \dots, \frac{D_{uk}}{T}\}$. Specifically, $H_u = -\sum_{i=1}^k \frac{D_{ui}}{T} \log(\frac{D_{ui}}{T})$. A high entropy value H_u means that u changes his/her focus of attention very frequently, indicating more engagement in the game. Conversely, a low H_u means longer periods of gaze towards the same participant, indicating lower engagement with the rest of the group.

All scores are normalized per game by subtracting the mean score of all participants in the game. Thus, after normalization, a positive (negative) gaze entropy of a participant p means that p shifts her gaze more (less) often than average. To compare the overall behavior of deceivers and

non-deceivers, we average the normalized entropy score of all deceivers’ entropy across all games — and likewise do the same with non-deceivers. Furthermore, since the behavior of participants can vary dramatically based on the game’s outcome (i.e., whether deceivers win or lose), we aggregate the scores for DW (Deceivers Win) vs DL (Deceivers Lose) games separately.

Gaze Reciprocity. We define the reciprocity of u ’s gaze in the k^{th} period P_{uk} as the average looking-at probability of u ’s target T_{uk} towards u during the same time period.

Specifically, $R_{uk} = \frac{\sum_{t=P_{uk}^0}^{P_{uk}^1} E_G^t(T_{uk}, u)}{D_{uk}}$. A high reciprocity means that u ’s targets pay attention to u , while a lower reciprocity indicates that u ’s targets ignore u ’s gaze. The average reciprocity of u in the game is the average reciprocity across all its periods, weighted by the duration of the period. We normalize and aggregate reciprocity scores to zero-mean as we did with entropy.

Findings. Figure 2(left) and (right) respectively compare the gaze entropy of participants and their gaze reciprocity. The figures report the mean scores across participants and the 95% confidence interval of the score distribution. Independent two-sample t-test is used to compare distributions throughout the paper. We observe that the behavior of deceivers depends heavily on the outcome of a game.

Finding (F1): In DW games, deceivers and Non-Deceivers look at similar numbers of speakers and are looked at to similar extents by other participants. Deceivers and non-deceivers have similar entropy and reciprocity scores (both $p > 0.05$) in DW games. Thus, deceivers win when they successfully camouflage themselves by imitating the nonverbal behavior of non-deceivers.

Finding (F2): In DL games, deceivers look at fewer participants than non-deceivers and are also looked at less by other participants. Both gaze entropy and gaze reciprocity are significantly lower for deceivers in DL games ($p < 0.001$), showing that they have a steadier gaze and receive less attention compared to non-deceivers. This indicates that deceivers are easily identified and lose when they are less engaging as compared to the rest of the participants.

RQ2: Does the verbal behavior of group of deceivers differ from non-deceivers?

The way in which people speak has previously been shown to indicate deception (Baccarani and Bonfanti 2015; Beslin and Reddin 2004). However, verbal characteristics of a coordinated group of deceivers is less well known. We now compare the speaking patterns of deceivers and non-deceivers. We use the Speak-To FFDIN network for the analysis. As before, we compare the differences partitioned by the final game outcome.

Finding (F3): Deceivers speak less than non-deceivers. We established this by first computing the fraction of time slices in which a participant u speaks. Both speaking excessively and being anomalously quiet have previously been noted to be indicators of deception (Wiseman 2010; Vrij 2008).

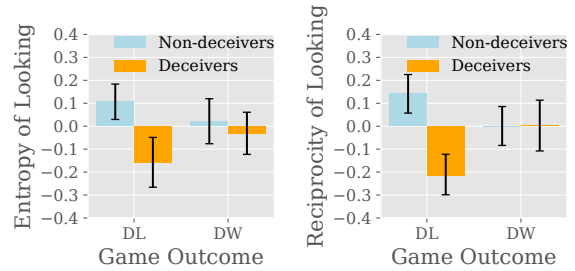


Figure 2: *Non-engaging deceivers are unsuccessful in deceiving others.* (Left) Deceivers have a lower entropy of looking at others compared to non-deceivers in DL (Deceivers Lose) games, while this difference does not exist in DW (Deceivers Win) games. (Right) Deceivers get less attention, measured in terms of reciprocity, in DL games. In both plots, we observe that deceivers and non-deceivers have similar scores in DW games, as deceivers successfully camouflage themselves.




Figure 3: *Deceivers tend to speak less than non-deceivers.* This difference is more pronounced in DL games.

Our experiments show that deceivers speak less than non-deceivers regardless of the game outcome. As shown in Figure 3, the differences are statistically significant both in DW games (0.08 vs -0.13, $p < 0.05$) and in DL games (0.07 vs -0.10, $p < 0.001$). This shows that non-deceivers are more vocal in all games.

Finding (F4): In DL games, deceivers get less attention while speaking than non-deceivers. We can infer the attention a speaker u is getting based on how many other participants are looking at u while u is speaking. We define the average attention that a participant u gets as the average weighted in-degree of u in the Listen-To FFDINs: $\frac{1}{T} \sum_t \sum_{(v,u) \in E_L^t} E_L^t(v, u)$, where T is the number of networks in which u is a speaker.

Figure 4 (left) shows that in DL games, less attention is paid to deceivers when they speak, compared to non-deceivers. However, the story is different in DW games—players in both roles receive a similar amount of attention.

Finding (F5): In DL games, deceiver speakers are reciprocated less than non-deceivers. We also looked at the




Figure 4: *Deceivers get less attention in the games they lose.* (Left) We observe that when deceivers speak in DL games, other participants have lower likelihood of looking at the deceiver as compared to when a non-deceiver is speaking. (Right) Similarly, the target of speakers is less likely to look back at deceivers than at non-deceivers in DL games. These differences are not present in DW games as deceivers are equally engaged and central to discussions.

gaze behavior of the person who is being spoken to. How often do they pay attention to the speaker? Specifically, when person u talks to v , does v look back at u ? This is a sign of trust and respect (Ellsberg 2010; Derber 2000).

We define the reciprocity of u 's target T_{ut} at time t as the edge weight from T_{ut} to u at the same time. We calculated the average reciprocity of participant u in the entire time period $t \in [1, T]$ as $\frac{1}{T} \sum_{t=1}^T E_L^t(T_{ut}, u)$. We compared the average reciprocity of deceivers and non-deceivers in Figure 4 (right). We see that in DL games, deceivers are not as frequently reciprocated as non-deceivers. This suggests that in DL games, other participants pay less attention to and trust deceivers less. This, however, is not the case in DW games where both deceivers and non-deceivers are given equal attention by listeners.

Summarizing, we find that non-deceivers are highly vocal and more active and central compared to deceivers in DL games. This is not the case in DW games, where the engagement and importance of deceivers and non-deceivers is equivalent. This shows that deceivers are successful in deceiving others when they are as engaging as the non-deceivers in the game and camouflage their behavior well.

RQ3: Do deceivers interact differently with other deceivers vs non-deceivers?

Since deceivers know the role of all participants in the game, do they focus their attention on specific individuals? Past survey based social science studies (Driskell, Salas, and Driskell 2012) conclude that deceivers are unlikely to respond to each other. We develop competing hypothesis about this. The first hypothesis is that the deceivers interact more with deceivers in order to cooperate and deceive other participants. The alternate hypothesis states that deceivers interact less with each other in order to avoid being identified by non-deceivers.

To test these hypotheses, we compare the pairwise interactions between participants, grouped by their roles: deceivers vs. deceivers, deceivers vs. non-deceivers, non-

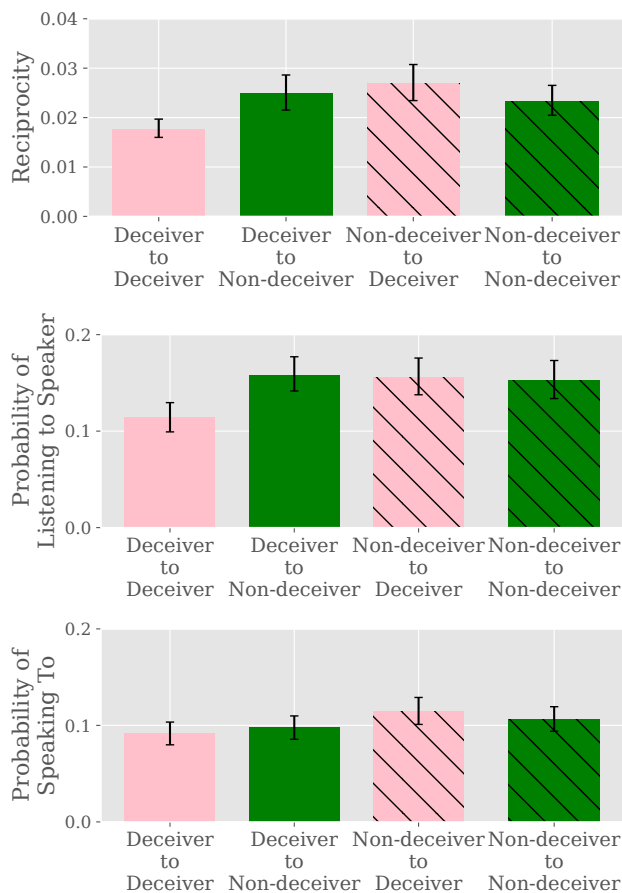


Figure 5: *Deceivers avoid non-verbal interactions with other deceivers.* (Top) Deceivers have lower reciprocity of looking at other deceivers than at non-deceivers. Non-deceivers have similar reciprocity for both. (Middle) Deceivers listen less to other deceivers than to non-deceivers. Non-deceivers have similar values for both. (Bottom) All pairs of participants have similar probability of speaking to one another.

deceivers vs. deceivers, and non-deceivers vs. non-deceivers. Figure 5 compares the average reciprocity of looking, average talk-to probability, and average listen-to probability for all pairs of roles. We aggregate the properties across all games to measure role-specific behavior regardless of the game's outcome. As earlier, we report the mean and the 95% confidence intervals for all properties. We make the following three observations by analysing both verbal and non-verbal behavior of participants.

Finding (F6): Deceivers look less at other deceivers. First, Figure 5 (top) shows that non-deceivers spend similar amount of time looking at deceivers and non-deceivers ($p > 0.05$). However, the looking behavior of deceivers is strikingly distinct—deceivers look less at other deceivers than at non-deceivers ($p < 0.001$). This has several important implications. Since deceivers know the identity of non-deceivers, deceivers spend more time observing non-deceivers (and less time observing their fellow deceivers).

Deceivers may also interact less with other deceivers to avoid ‘guilt-by-association’, i.e., getting caught in case the other deceivers are identified.

In addition, Figure 5 (middle) shows that non-deceivers have a similar probability of listening to both deceptive and non-deceptive speakers ($p = 0.76$). However, this is not the case for deceiver listeners. Deceivers have a lower probability of listening to other deceivers as compared to non-deceivers ($p < 0.05$). This is possibly in order to avoid being suspected of supporting the deceiver.

Finally, Figure 5 (bottom) compares the verbal behavior (as opposed to non-verbal behavior in the previous two paragraphs). Surprisingly, we find that the verbal behavior between all pairs of participants is similar. Since non-deceivers do not know the roles of other participants, they speak equally to non-deceivers and deceivers ($p = 0.48$), as expected. However, it is surprising that deceivers spend equal time talking to both deceivers and non-deceivers ($p = 0.39$). This is in stark contrast to the previous two non-verbal findings. This shows that deceivers consciously adapt their verbal behavior to mimic non-deceivers, but not their non-verbal behavior. Since verbal behavior is noticed by everyone else, deceivers consciously do not exhibit any bias in verbal interaction with other participants to avoid getting caught.

Thus, RQ3 shows that deceivers successfully camouflage their verbal (speaking) behavior, while they are unable to camouflage their non-verbal (looking and listening) behavior. Altogether, deceivers avoid non-verbal interactions with other deceivers.

Network Model for Deception Prediction

In this section, we present DeceptionRank, our PageRank based model that examines FFDINs in order to predict whether a given participant is deceptive or not. Automated detection of deceivers is a challenging task as this is the goal of the non-deceivers. DeceptionRank tries to do this with short duration of videos. In contrast, human participants try to do this throughout the game, which are 38 minutes long, on average, but are still unsuccessful in almost half the games.

DeceptionRank on Negative Dynamic Interaction Networks

DeceptionRank is built on our findings from the previous section that deceivers avoid non-verbal interactions with other deceivers, while non-deceivers do not exhibit this bias. There are four main steps of DeceptionRank: (i) building the network, (ii) initializing node deception priors, (iii) applying network algorithm to obtain node deception scores, and (iv) training a deception classifier.

Building negative dynamic interaction networks. In order to bring this “non-interaction” to the fore, we generate negative interaction networks that capture the pairwise “lack of interactions”. The edges in the negative interaction network connect nodes which avoid interacting with one another. Given a FFDIN $N^t = (V, E^t)$ at time t , where $E^t(u, v) = w_{u,v,t}, \forall u, v \in V$, the associated negative inter-

action network NDIN is given by $N^{t-} = (V, E^{t-})$, where $E^{t-}(u, v) = 1 - w_{u,v,t}$. Note that $E^{t-}(u, v) = 1$ when there is no edge from u to v at time t in the interaction network, i.e., when u does not interact with v .

Initializing node deception priors. First we need to initialize every node’s prior probability of being a deceiver. We introduce a novel technique for initialization based on every node’s verbal and non-verbal features compared to the features of all the nodes in the network. Given a set of feature values $\{x_{1u}, \dots, x_{Fu}\}$ for the F features of a node u , we aim to combine them into an initial deception score $S(u) \in [0, 1]$, which we describe next.

Based on our analysis in the previous section, we build the priors using the following four features that best distinguish between deceivers and non-deceivers: (a) fraction of speaking (FS_u), (b) average entropy of looking (H_u), (c) average in-degree ($E_{G,u}$), and (d) average in-degree while speaking ($E_{L,u}$). Since the feature distributions can vary, we first normalize each feature f ($f \in \{FS_u, H_u, E_{G,u}, E_{L,u}\}$) by linearly scaling it between 0 and 1, corresponding to the minimum and maximum values of the feature. Then we subtract each feature f from 1 because deceivers have lower scores of f than non-deceivers, so $1 - f$ ensures that the deceivers tend to have higher initial scores. Finally, we average each node’s four property scores to get its prior score $S(u)$ for node u :

$$S(u) = 1 - \frac{FS_u + H_u + E_{G,u} + E_{L,u}}{4} \quad (1)$$

This score is used to initialize node priors in the first iteration of our dynamic network algorithm — a higher score would indicate higher prior probability of the node being deceptive.

Obtaining node deception scores from negative networks. To predict whether a participant is a deceiver or not, we extend the PageRank algorithm (Page et al. 1999). By default, PageRank is applicable on static networks, thus, we extend it to apply to dynamic negative interaction network sequences. The method is shown in Algorithm 1. The overall idea is that in each iteration, we aggregate neighborhood scores for each node in all negative networks independently and then aggregate the scores of a node across all the networks. This aggregated score is used in the next iteration.

In detail, we repeat the following three-step procedure until convergence (or until the maximum number of iterations is reached). In the first step, we initialize each node with an initial score in all the networks. Node deception prior scores are used for the first initialization. In the second step, in each negative network E^{t-} , we calculate the score $s^t(v)$ by aggregating node v ’s outgoing neighbors’ scores using the following equation: $s^t(v) = \beta \sum_{(v,u) \in E^{t-}} s^t(u) \cdot w_{v,u,t} + (1 - \beta)s^t(v)$. Here β weighs the importance of a node’s own deception score versus the aggregate of neighbors’ scores. Each neighbor u ’s deception score is weighted by the weight of the outgoing edge from v to u . In the third step, for each node v , we aggregate the scores of v across all the negative networks to get v ’s output score in current iteration, and normalize the scores. Averaging is used as the aggregation function; other aggregation functions, such as recency-weighted averaging, can be used instead, if desired. The normalized

Algorithm 1: DeceptionRank on Negative Dynamic Interaction Networks

Input: Negative dynamic interaction networks $[(V, E^{1-}), (V, E^{2-}), \dots, (V, E^{T-})]$, initial deceptive scores $s(V)$, β , convergence threshold τ , maximum iteration number M

Output: Final deceptive scores $s(V)$

$iter = 0, dif = 1;$

while $dif > \tau$ **and** $iter < M$ **do**

$s^t(V) = s(V), \forall t = 1 \dots T;$

for $t = 1 \dots T$ **do**

foreach $v \in V$ **do**

$r^t(v) = \beta \sum_{(u,v) \in E^{t-}} s^t(u) \cdot w_{u,v,t} + (1 - \beta)s^t(v);$

$r(V) = \sum_t r^t(V)/T;$

$r(V) = r(V) / \|r(V)\|_2;$

$dif = \|(r(V) - s(V))\|_2;$

$iter = iter + 1;$

$s(V) = r(V);$

return $s(V)$

output scores are used as the initial scores in the next iteration. After convergence, the final deceptiveness scores $s(V)$ are returned.

Final classification. Finally, we train a binary classifier for predicting whether a participant is deceptive or not. For a node u , the features we use are u 's final deceptiveness scores and u 's four behavior properties, namely the average fraction of speaking, average entropy of looking, average in-degree, and average in-degree while speaking.

Experiments

In this section, we compare the performance of DeceptionRank with state-of-the-art vision and graph embedding baselines. We show that DeceptionRank outperforms these baselines by at least 20.9% in detecting deceivers.

Experiment Setup

The prediction task is: given a video segment of a game, predict the roles (deceiver or non-deceiver) of all participants in the game.

Since there are only 26 games in the dataset, we augment the dataset by segmenting long videos into several smaller videos. The roles of the players remain the same in the videos after segmenting. We split the games into 1 minute long video segments (we study the effect of segment length on prediction performance later). This results in a dataset with 2781 data points.

To ensure that there is no leakage of ground-truth labels, we split the dataset into training and test set according to games (not by video segments). Every player in our data participated in exactly one game — so we never train on data of a player in one game and use that to predict whether he is deceptive or not in another. We conduct all our experiments with 5-fold cross-validation, where all clips of a game belong to the same fold.³ This ensures that two segments of

³We do 5-fold cross validation rather than 10-fold because of

the same game can not appear in both training and test sets. Further, we split the data by participants as well, so a participant can only be either in training or test across all segments. In each fold, we set 60% participants in the training set and the rest are in the test set (we ensure that at least one deceiver and non-deceiver are in training and test set). Since the task is unbalanced, we report the AUROC, averaged over five folds.

We do not provide any model, either our or the baselines, with the number of deceivers or non-deceivers in the game, so there is no leakage of label distribution. All the experiment setting are the same across all models to ensure fairness.

Baselines

We consider two sets of baselines: vision-based methods and graph embedding methods. All baselines are evaluated in the same setup and dataset as our model.

Computer vision baselines. We compare our method with five computer vision baselines with the same experimental setup as our method (Bai et al. 2019; Baltrusaitis et al. 2018; Demyanov et al. 2015; Bai et al. 2019; Wu et al. 2018). These methods used features extracted from the video, including facial emotion, head and eye movement, facial action units, and time-aggregated features as described below.

(Demyanov et al. 2015) extracts the averaged facial action units (FAUs) features over time. (Baltrusaitis et al. 2018) computes eye movements from the estimated eye ball positions, and uses the movement distributions over time as features. (Wu et al. 2018) extracts the individual dense trajectory features from videos, MFCC features from audio, micro-expression features and text features from transcripts, and uses an ensemble method called late fusion to come up with a joint prediction. Since our dataset doesn't have transcripts and annotated micro-expressions, we remove the text features and replace micro-expressions by FAU features (Demyanov et al. 2015). Lastly, we extract the histograms of emotion features and LiarRank features proposed by (Bai et al. 2019) as other two baselines, where LiarRank captures group information by ranking the feature values in each group as meta-features.

Note that all these methods make predictions for each player individually, without considering interactions between players. Specifically, in these methods, we extract the feature values for each individual player and use them as input to train a binary classifier. All these baseline features are trained with Logistic Regression, Random Forest, Linear SVM and Navie Bayes. We report the best AUROC among these classifiers in Table 2.

Graph embedding baselines. Here we compare our method with dynamic graph embedding based methods. Dynamic graph embedding models have shown incredible success in making predictions for large-scale social networks. In particular, we compare with temporal graph convolution networks (TGCN) (Liu et al. 2019) on the look-at network, speak-to network, and listen-to networks. TGCN model

the small number of games.

Method	Perf.	Improvement Over Baseline
Computer Vision Baselines		
Emotions (Bai et al. 2019)	0.538	39.9%
Movements (Baltrusaitis et al. 2018)	0.549	37.2%
FAUs (Demyanov et al. 2015)	0.569	32.3%
LiarRank (Bai et al. 2019)	0.590	27.6%
Late fusion (Wu et al. 2018)	0.594	26.7%
Graph Embedding Baselines		
TGCN on Look-At (Liu et al. 2019)	0.550	36.9%
TGCN on Speak-To (Liu et al. 2019)	0.538	39.9%
TGCN on Listen-To (Liu et al. 2019)	0.541	39.2%
Ensemble Baseline		
Combining all the above features	0.623	20.9%
Proposed Method		
DeceptionRank	0.753	-

Table 2: Our proposed method DeceptionRank outperforms state-of-the-art vision, graph embedding, and ensemble baselines in the task of predicting deceivers from 1 minute clips of videos. DeceptionRank outperforms all baselines by at least 20.9% AUROC in prediction performance.

combines graph convolution network with LSTM. Given a sequence of networks and the ground-truth training node labels, TGCN trains two-layer GCN models on individual networks. All individual networks share the same GCN parameters. Mean pooling is used to aggregate neighborhood node information in the GCN. The sequence of output scores per node corresponding to the sequence of graphs is fed as input into an LSTM. All nodes share the same LSTM parameters. The final output of the LSTM is used to predict the training nodes’ ground-truth label. The model is trained in an end-to-end manner, where the GCN and LSTM model parameters are trained to accurately predict the node labels. We experimented with other variants of temporal graph models (Zhou et al. 2018; Goyal et al. 2018), which gave similar performance.

Ensemble baseline. We create an ensemble baseline model to combine the strengths of all the baselines. For each node, we concatenate its baseline scores from all the vision and graph embedding classifiers described above. This generates one feature vector per node, which is used as the node’s input to a Logistic Regression classifier to make the prediction.

Predictive Performance of DeceptionRank

Here we compare the performance of DeceptionRank with the baselines. Table 2 shows the cross-validation performance (in terms of AUROC) of all methods on 1 minute segments. We report the performance as average AUROC scores and their 95% bootstrapped confidence interval.

First, we observe that DeceptionRank significantly outperforms all other methods, by at least 20.9%. DeceptionRank has an AUROC of 0.753 with its 95% confidence interval ranging from 0.721 to 0.789. Second, among baselines, the ensemble performs the best. We note that when used individually, vision-based baselines outperform baselines that




Figure 6: DeceptionRank performs better than all features across all segment length. DeceptionRank has stable performance for all segment lengths too.

use graph embeddings. We attribute this difference in performance to the small size of the dataset, leading to lower performance of deep-learning based GCN methods. Finally, LiarRank and Late Fusion outperform the other baselines. This is likely due to the fact that LiarRank was designed to identify deceivers in networks, while late fusion combines audio and transcripts with visual features. However, we remind readers that DeceptionRank generates the best performance compared to all baselines, suggesting that FFDINs and Negative Interaction Networks, together with the DeceptionRank algorithm generate excellent value in terms of performance.

Performance vs. Segment Length

The preceding experiment shows the performance of both DeceptionRank and the baselines using data from segments that are 1 minute long. We now study the impact of the input segment length on the performance of predictors. We vary segment length from 1 minute to 14 minutes. For each segment length, we randomly sample 100 segments from each game. As before, we follow a 5-fold cross-validation setting. Finally, we compare the average AUROC and report the 95% confidence interval of performance.

As the ensemble baseline performed the best among all baselines as shown in the previous experiment, we compare DeceptionRank with this ensemble baseline model. Other individual vision and graph embedding baselines have lower performances.

Figure 6 shows the results of varying the segment length. We show that DeceptionRank outperforms the best baseline for all segment lengths considered. The margin is large when the segment lengths considered are small — once the segment lengths are over 10 minutes long, the performance of DeceptionRank and the baselines is similar. It is important to note that DeceptionRank’s performance is stable across segment lengths, while the baselines have diminished performance when the segments are short. These findings illustrate the robustness of our model with respect to the input durations.




Figure 7: *DeceptionRank performs better than all features across DL and DW games. Model performances are better in DL games compared to DW games.*

Performance in DL vs. DW Games

Here we compare the performance of the models according to the game outcome. Specifically, we compare model performance for DL games vs. DW games. Recall that 14 out of 26 games (= 54%) were won by deceivers. Since our analysis in the previous sections show that the deceivers differ significantly, we evaluate if that affects the models’ performance across these games. As the ensemble baseline performed the best among all baselines, we compare DeceptionRank against the ensemble model.

We follow the default setting as outlined earlier for the experiments. We evaluate the models on 1 minute segments. We report the performance as average AUROC scores and their 95% bootstrapped confidence interval. We randomly sample 100 segments from each game. As before, we follow a 5-fold cross-validation setting.

When evaluating the model performance for DL games, we consider only DL games, i.e., we train and test on DL games only. Similarly, only DW games are considered when evaluating model performance for DW games.

Figure 7 shows the results of according to the game outcome. First, we see that regardless of the game outcome, DeceptionRank performs better than the ensemble model. Second, we note that both models perform better in DL games compared to their performance in the DW games. The explanation for this observation is that the behavior of deceivers in DL games is significantly different from that of non-deceivers. This makes it easier for the machine learning algorithms to distinguish between them. On the other hand, deceivers and non-deceivers have similar behavior in DW games, which makes it comparatively harder for the algorithms to identify them. These findings show the robustness of our model across game outcomes.

In summary, the results in this section show that DeceptionRank outperforms the baseline methods by at least 20.9% in identifying deceivers in groups. Deceivers can be identified effectively even with short segments and better in DL games.

Related Work

We now summarize related work that was not already discussed earlier.

Face-to-face deception. There has been much research on predicting whether an individual is deceptive from facial and body cues (Ding et al. 2019; Randhavane et al. 2019; Wang et al. 2020) with extensions that also include audio and linguistic cues (Gogate, Adeel, and Hussain 2017; Wu et al. 2018). Ding et al. (2019) proposed a deep CNN model to fuse face and body cues together which can be trained with limited data via meta-learning. Randhavane et al. (2019) focused on deception prediction by feeding dynamic 3D gestures to an LSTM. They also identified deceptive behaviors (e.g. looking around, hands in pockets). Wang et al. (2020) proposed the attention mechanism with 3D CNN to identify individual deceptive facial patterns. Wu et al. (2018) combined visual cues (e.g. micro-expression and video trajectory features) with voice and transcript data using a highly effective late fusion mechanism. However, there is limited work on predicting deception in multi-person face-to-face interaction settings. Chittaranjan and Hung (2010) pioneered face to face deception detection from conversation cues (e.g. speaking turns and number of interruptions). Pak and Zhou (2013) used eye gaze to detect deception, Sapru and Boulard (2015) used facial action units as features, and Bai et al. (2019) proposed the notion of Liar-Rank within groups to better capture group-based deception.

However, most of these techniques do not work on group deception and do not consider inter-personal interactions in order to predict deception. We are the first to do so and to show that the Negative Dynamic Interaction Networks we propose are highly effective at detecting deception. This paper does not focus on computer vision — rather it builds on the technique proposed by Bai et al. (2019) to extract the visual focus of attention of people — from this, we build NDINs for deception detection.

Deception detection on social media using social and interaction networks. Extensive research has been done to detect deception using interaction and social networks. Many of these works have focused on web and social media domains. These include methods to detect fake news (Liu and Wu 2018; Horne and Adali 2017), rumors (Zeng, Starbird, and Spiro 2016; Li et al. 2016), fake reviews (Mukherjee et al. 2013; Li et al. 2015), spammers (Wu et al. 2017), and coordinated activity (Kumar et al. 2017; Subrahmanian et al. 2016). Liu and Wu (2018) focused on the early detection of fake news on social media by modeling information spread on the network. Kumar et al. (2018); Hooi et al. (2016); Rayana and Akoglu (2015) leveraged the reviewer-product network to identify fake reviews and fraudulent reviewers in e-commerce platforms. Wu et al. (2017) used sparse learning to detect spammers as a social community from both online social interactions and TF-IDF content features. Kumar et al. (2017) analyzed the behaviors of sock-puppets (users with multiple accounts to manipulate public opinions) in social media through multiple perspectives including their social networks, posting patterns, posting contents etc. Subrahmanian et al. (2016) developed a mix of language, network, and neighborhood features in order to

identify both influence bots and botnets as part of a DARPA challenge.

Though these papers use the concept of networks to study deception, they have not been applied to the video-based discussion settings. Compared to social networks, FFDINs extracted from videos are very different due to two important reasons. First, the edges in FFDINs represent instantaneous verbal and non-verbal interactions, and thus, the edges are highly dynamic. On the other hand, edges in social networks are relatively long-term and stable. Second, FFDINs have very few number of nodes and edges, while social networks have millions of nodes and edges. Both these major differences call for new methods that can work on small-scale but highly dynamic networks. Our DeceptionRank method bridges this gap.

Deception detection in games. The work most related to us are deception analysis from online chat-based mafia games (Pak and Zhou 2015; Yu et al. 2015). Pak and Zhou (2015) builds a reply network over time, and hypothesizes several deceptive patterns such as centrality and nodes similarity. They conduct statistical analysis of the relationship to deception, but don't predict deceivers. Yu et al. (2015) builds a rule-based attitude network from the chat logs, and clusters nodes into subgroups of deceivers and non-deceivers. The clustering quality is measured by purity and entropy. Their results highly depend on the quality of chat logs and specific rules, which limit the application scope. Moreover, neither of these directly predicts the role for each player as we do. Niculae et al. (2015); Azaria, Richardson, and Kraus (2015) studied the conversational properties patterns in a games to detect deception. However, none of the above works have studied deception in face-to-face video communication, which is the gap we bridge.

Discussion and Conclusion

To the best of our knowledge, this is the first paper to use network analysis methods in order to predict who is being deceptive in a video-based group interaction setting.

Using a dataset based on the well-known Resistance game, we propose the concepts of a Face-to-Face Dynamic Interaction Network (FFDIN) and a Negative Dynamic Interaction Network (NDIN). We propose the DeceptionRank algorithm, and show that DeceptionRank beats several baselines including ones based on computer vision and graph embedding in detecting deceivers.

Relevance to the web and social media community. Our research sheds light on group deception in video-based conversations. While our work focuses on conversations in a social game, such conversations are commonplace in everyday communications via video call apps such as Microsoft Teams, Google Meet, Facebook Messenger, Zoom and Skype, and form an integral part of social media platforms like Facebook, SnapChat, and WhatsApp. The inputs are look-at, speak-to and listen-to interactions between people, which can be extracted from these videos from web and social media platforms. The proposed methods of Negative Dynamic Interaction Network and DeceptionRank can then be applied to identify deceivers in those videos. Since there are no such web and social media datasets with ground-truth

of deception, we leave the experimental evaluation of our methods on web and social media video-based deception for future work. Techniques in our work have the potential to improve the safety and integrity of social media and web-based communication platforms.

Use of one dataset. Currently, there are no other datasets of face-to-face video conversations with ground truth of group deception on which we can test our model. This is because creating such a dataset is an extremely difficult and time-consuming effort. This is highlighted in the paper (Bai et al. 2019) from which we have derived our dataset — it took the authors about 18 months to collect and process the 26 videos. Testing generalizing capability of our method beyond the current dataset will require creation of new datasets, a huge task which can be conducted in the future.

The only other comparable video dataset is given by (Pérez-Rosas et al. 2015) which contains 56 people and spans 57 minutes, but each video only has one person. There is no group interaction or deception. So, we can not use it for our task. By comparison, the dataset we use contains 185 participants and spans 1000 minutes. Our dataset is significantly larger.

Generalizing beyond a game. Although the setting we study in this work is in a social game setting, the discussions are free-form and the participants can deceive others as they want. No instructions or training was provided to them about how to deceive. Thus, the findings in this paper should represent the general properties of how deceivers operate in groups. Importantly, the Negative Dynamic Interaction Network and DeceptionRank method are general and can be applied to any setting involving interactions between groups of people.

Future work. Future work can expand this study to multiple games and settings such as sales meetings, business negotiations, job interviews, and more. The methods of Negative Dynamic Interaction Networks and DeceptionRank can be tested for deception detection in other datasets, including social network datasets. Finally, our methods can be used to study other social affects, such as leadership, trust, liking, and dominance.

Acknowledgements

We gratefully acknowledge the support of NSF under Nos. OAC-1835598 (CINES), OAC-1934578 (HDR), CCF-1918940 (Expeditions), IIS-2030477 (RAPID), IIS-2027689 (RAPID); DARPA under No. N660011924033 (MCS); ARO under Nos. W911NF-16-1-0342 (MURI), W911NF-16-1-0171 (DURIP); Stanford Data Science Initiative, Wu Tsai Neurosciences Institute, Chan Zuckerberg Biohub, Amazon, JPMorgan Chase, Docomo, Hitachi, JD.com, KDDI, NVIDIA, Dell, Toshiba, United-Health Group, Adobe, Facebook, Microsoft, and IDEaS Institute. J. L. is a Chan Zuckerberg Biohub investigator.

References

Addawood, A.; Badawy, A.; Lerman, K.; and Ferrara, E. 2019. Linguistic cues to deception: Identifying political

- trolls on social media. In *International AAAI Conference on Web and Social Media*, volume 13, 15–25.
- Azaria, A.; Richardson, A.; and Kraus, S. 2015. An agent for deception detection in discussion based environments. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 218–227.
- Baccarani, C.; and Bonfanti, A. 2015. Effective public speaking: a conceptual framework in the corporate-communication field. *Corporate Communications: An International Journal*.
- Bai, C.; Bolonkin, M.; Burgoon, J.; Chen, C.; Dunbar, N.; Singh, B.; Subrahmanian, V. S.; and Wu, Z. 2019. Automatic long-term deception detection in group interaction videos. In *IEEE International Conference on Multimedia and Expo*.
- Bai, C.; Kumar, S.; Leskovec, J.; Metzger, M.; Nunamaker, J. F.; and Subrahmanian, V. 2019. Predicting the visual focus of attention in multi-person discussion videos. In *International Joint Conference on Artificial Intelligence*.
- Baltrusaitis, T.; Zadeh, A. B.; Lim, Y. C.; and Morency, L.-P. 2018. OpenFace 2.0: Facial behavior analysis toolkit. In *IEEE International Conference on Automatic Face and Gesture Recognition*.
- Beslin, R.; and Reddin, C. 2004. How leaders can communicate to build trust. *Ivey business journal* 69(2): 1–6.
- Chittaranjan, G.; and Hung, H. 2010. Are you A werewolf? Detecting deceptive roles and outcomes in a conversational role-playing game. In *international Conference on Acoustics, Speech, & Signal Processing*.
- Demyanov, S.; Bailey, J.; Ramamohanarao, K.; and Leckie, C. 2015. Detection of deception in the mafia party game. In *ACM International Conference on Multimodal Interaction*.
- Derber, C. 2000. *The pursuit of attention: Power and ego in everyday life*. Oxford University Press.
- Ding, M.; Zhao, A.; Lu, Z.; Xiang, T.; and Wen, J.-R. 2019. Face-focused cross-stream network for deception detection in videos. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Dinges, D. F.; Rider, R. L.; Dorrian, J.; McGlinchey, E. L.; Rogers, N. L.; Cizman, Z.; Goldenstein, S. K.; Vogler, C.; Venkataraman, S.; and Metaxas, D. N. 2005. Optical computer recognition of facial expressions associated with stress induced by performance demands. *Aviation, space, and environmental medicine* 76(6): B172–B182.
- Driskell, J. E.; Salas, E.; and Driskell, T. 2012. Social indicators of deception. *Human Factors* 54(4): 577–588.
- Ellsberg, M. 2010. *The power of eye contact: Your secret for success in business, love, and life*. Harper Paperbacks.
- Gogate, M.; Adeel, A.; and Hussain, A. 2017. Deep learning driven multimodal fusion for automated deception detection. In *IEEE Symposium Series on Computational Intelligence*.
- Goyal, P.; Kamra, N.; He, X.; and Liu, Y. 2018. DynGEM: Deep embedding method for dynamic graphs. *arXiv preprint arXiv:1805.11273*.
- Hooi, B.; Shah, N.; Beutel, A.; Günnemann, S.; Akoglu, L.; Kumar, M.; Makhija, D.; and Faloutsos, C. 2016. Birdnest: Bayesian inference for ratings-fraud detection. In *Proceedings of the 2016 SIAM International Conference on Data Mining*, 495–503. SIAM.
- Horne, B. D.; and Adali, S. 2017. This just in: fake news packs a lot in title, uses simpler, repetitive content in text body, more similar to satire than real news. In *International AAAI Conference on Web and Social Media*.
- Keller, F. B.; Schoch, D.; Stier, S.; and Yang, J. 2017. How to manipulate social media: Analyzing political astroturfing using ground truth data from South Korea. In *International AAAI Conference on Web and Social Media*.
- Kumar, S.; Cheng, J.; Leskovec, J.; and Subrahmanian, V. 2017. An army of me: Sockpuppets in online discussion communities. In *Proceedings of the International World Wide Web Conference*.
- Kumar, S.; Hooi, B.; Makhija, D.; Kumar, M.; Faloutsos, C.; and Subrahmanian, V. 2018. Rev2: Fraudulent user prediction in rating platforms. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, 333–341.
- Kumar, S.; Spezzano, F.; and Subrahmanian, V. 2015. Vews: A wikipedia vandal early warning system. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 607–616.
- Kumar, S.; Zhang, X.; and Leskovec, J. 2019. Predicting dynamic embedding trajectory in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1269–1278.
- Laretzaki, G.; Plainis, S.; Vrettos, I.; Chrisoulakis, A.; Palikaris, I.; and Bitsios, P. 2011. Threat and trait anxiety affect stability of gaze fixation. *Biological psychology* 86(3): 330–336.
- Li, H.; Chen, Z.; Mukherjee, A.; Liu, B.; and Shao, J. 2015. Analyzing and detecting opinion spam on a large-scale dataset via temporal and spatial patterns. In *International AAAI Conference on Web and Social Media*.
- Li, Q.; Liu, X.; Fang, R.; Nourbakhsh, A.; and Shah, S. 2016. User behaviors in newsworthy rumors: A case study of twitter. In *International AAAI Conference on Web and Social Media*.
- Liu, Y.; Shi, X.; Pierce, L.; and Ren, X. 2019. Characterizing and forecasting user engagement with in-app action graph: A case study of Snapchat. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Liu, Y.; and Wu, Y.-F. B. 2018. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *AAAI Conference on Artificial Intelligence*.
- Mukherjee, A.; Venkataraman, V.; Liu, B.; and Glance, N. 2013. What yelp fake review filter might be doing? In *International AAAI Conference on Weblogs and Social Media*.

- Niculae, V.; Kumar, S.; Boyd-Graber, J.; and Danescu-Niculescu-Mizil, C. 2015. Linguistic Harbingers of Betrayal: A Case Study on an Online Strategy Game. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*.
- Page, L.; Brin, S.; Motwani, R.; and Winograd, T. 1999. The PageRank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.
- Pak, J.; and Zhou, L. 2013. Eye gazing behaviors in online deception. In *19th Americas Conference on Information Systems*.
- Pak, J.; and Zhou, L. 2015. Temporal Patterns of Structural Deception Behavior in a Massively Multiplayer Online Game. In *2015 48th Hawaii International Conference on System Sciences*, 131–140. IEEE.
- Pérez-Rosas, V.; Abouelenien, M.; Mihalcea, R.; and Burzo, M. 2015. Deception detection using real-life trial data. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, 59–66.
- Randhavane, T.; Bhattacharya, U.; Kapsaskis, K.; Gray, K.; Bera, A.; and Manocha, D. 2019. The Liar’s Walk: Detecting Deception with Gait and Gesture. *arXiv preprint arXiv:1912.06874*.
- Rayana, S.; and Akoglu, L. 2015. Collective opinion spam detection: Bridging review networks and metadata. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 985–994.
- Rayner, K. 1998. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin* 124(3): 372.
- Sapru, A.; and Bourlard, H. 2015. Automatic recognition of emergent social roles in small group interactions. *IEEE Transactions on Multimedia* 17(5): 746–760.
- Ströfer, S.; Ufkes, E. G.; Noordzij, M. L.; and Giebels, E. 2016. Catching a deceiver in the act: Processes underlying deception in an interactive interview setting. *Applied psychophysiology and biofeedback* 41(3): 349–362.
- Subrahmanian, V.; Azaria, A.; Durst, S.; Kagan, V.; Galstyan, A.; Lerman, K.; Zhu, L.; Ferrara, E.; Flammini, A.; and Menczer, F. 2016. The DARPA Twitter bot challenge. *Computer* 49(6): 38–46.
- Vrij, A. 2008. *Detecting lies and deceit: Pitfalls and opportunities*. John Wiley & Sons.
- Wang, L.; Bai, C.; Bolonkin, M.; Burgoon, J.; Dunbar, N.; Subrahmanian, V.; and Metaxas, D. N. 2020. Attention-based Facial Behavior Analytics in Social Communication. In *30th British Machine Vision Conference, BMVC 2019*.
- Wiseman, R. 2010. *59 Seconds: Change Your Life in Under a Minute*. Vintage.
- Wu, L.; Hu, X.; Morstatter, F.; and Liu, H. 2017. Adaptive spammer detection with sparse group modeling. In *International AAAI Conference on Web and Social Media*.
- Wu, Z.; Singh, B.; Davis, L. S.; and Subrahmanian, V. 2018. Deception detection in videos. In *AAAI Conference on Artificial Intelligence*.
- Yu, D.; Tyshchuk, Y.; Ji, H.; and Wallace, W. 2015. Detecting deceptive groups using conversations and network analysis. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*.
- Zeng, L.; Starbird, K.; and Spiro, E. S. 2016. # unconfirmed: Classifying rumor stance in crisis-related social media messages. In *International AAAI Conference on Web and Social Media*.
- Zhou, L.; Yang, Y.; Ren, X.; Wu, F.; and Zhuang, Y. 2018. Dynamic network embedding by modeling triadic closure process. In *AAAI Conference on Artificial Intelligence*.