

# My Team Will Go On: Differentiating High and Low Viability Teams through Team Interaction

HANCHENG CAO, Stanford University, United States

VIVIAN YANG, Stanford University, United States

VICTOR CHEN, Stanford University, United States

YU JIN LEE, Stanford University, United States

LYDIA STONE, Stanford University, United States

N'GODJIGUI JUNIOR DIARRASSOUBA, Texas Tech, United States

MARK E. WHITING, University of Pennsylvania, United States

MICHAEL S. BERNSTEIN, Stanford University, United States

Understanding team viability — a team's capacity for sustained and future success — is essential for building effective teams. In this study, we aggregate features drawn from the organizational behavior literature to train a viability classification model over a dataset of 669 10-minute text conversations of online teams. We train classifiers to identify teams at the top decile (most viable teams), 50th percentile (above a median split), and bottom decile (least viable teams), then characterize the attributes of teams at each of these viability levels. We find that a lasso regression model achieves an accuracy of .74–.92 AUC ROC under different thresholds of classifying viability scores. From these models, we identify the use of exclusive language such as 'but' and 'except', and the use of second person pronouns, as the most predictive features for detecting the most viable teams, suggesting that active engagement with others' ideas is a crucial signal of a viable team. Only a small fraction of the 10-minute discussion, as little as 70 seconds, is required for predicting the viability of team interaction. This work suggests opportunities for teams to assess, track, and visualize their own viability in real time as they collaborate.

CCS Concepts: • **Human-centered computing** → **Empirical studies in collaborative and social computing**.

Additional Key Words and Phrases: Team viability; team dynamics; predictability; collaboration; remote work

## ACM Reference Format:

Hancheng Cao, Vivian Yang, Victor Chen, Yu Jin Lee, Lydia Stone, N'godjigui Junior Diarrassouba, Mark E. Whiting, and Michael S. Bernstein. 2020. My Team Will Go On: Differentiating High and Low Viability Teams through Team Interaction. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3, Article 230 (December 2020), 27 pages. <https://doi.org/10.1145/3432929>

Authors' addresses: Hancheng Cao, hanchcao@stanford.edu, Stanford University, United States; Vivian Yang, Stanford University, United States; Victor Chen, Stanford University, United States; Yu Jin Lee, yujin512@stanford.edu, Stanford University, United States; Lydia Stone, Stanford University, United States; N'godjigui Junior Diarrassouba, Texas Tech, United States; Mark E. Whiting, markew@seas.upenn.edu, University of Pennsylvania, United States; Michael S. Bernstein, msb@cs.stanford.edu, Stanford University, United States.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2020 Association for Computing Machinery.

2573-0142/2020/12-ART230 \$15.00

<https://doi.org/10.1145/3432929>

## 1 INTRODUCTION

Even high performing teams can be miserable. Whether or not a team succeeds at their goals, their team environment can be joyless, toxic, or even downright malicious. *Team performance* is a widely studied metric of successful collaboration [11, 48, 62, 68], but it is only one piece of the picture. *Team viability* – the capacity of a team for sustainable growth and future success [10] – is a less studied yet equally critical component of a sustainable, successful team [47]. A team with high viability can adapt to internal and external changes [16]. A team which suffers a loss of viability will see its members resist, resent, or even actively undermine their teams [41, 49, 61, 67]. A loss of viability can be so disastrous that the members fracture and do not want to work together again [69]. Consequently, it is critical for teams to evaluate and sustain their viability in order to promote continued engagement and future success. In particular, remote teams are more susceptible to antisocial behavior and developing conflicts [30]. As remote collaboration becomes increasingly prevalent and necessary, it becomes even more crucial to study team viability, which reflects both the members' satisfaction with their teammates and the members' behavioral intent to remain in the team [28]. Since team viability is not necessarily skill-based or task-specific, not just remote teams but any team can strive to improve their viability.

Team viability is typically measured through self-reported questionnaires [18], limiting its effectiveness to after-the-fact instances of low viability. Assessing team viability at early stages of team interaction can offer actionable insights for improving ongoing collaborative efforts, as opposed to collecting retrospective lessons and regrets in the case of post-hoc evaluation. Therefore we ask: can we differentiate more viable and less viable teams based on a glimpse of their interaction? What attributes of team interaction are indicative of high viability? If this were possible, we could aid teams in diagnosing problematic starts and aid them in righting themselves.

To address these questions, we study a dataset of 669 online, 10-minute long, remote team interactions [69, 70], in which each team completed a collaborative task and a post-hoc team viability instrument [18]. We then use this data to judge the efficacy of particular signals in classifying whether a given interaction has a viability above or below the viability thresholds of top decile, median, or bottom decile. Since we are interested in which features differentiate the behavior of more viable and less viable interactions rather than the particular viability score they achieve, we focus on the task of classification.

We design features drawn from organizational behavior literature and compare multiple machine learning models for classification. For instance, we measure cosine similarity of TF-IDF vectors based on unigrams and bigrams to characterize topic alignment in the conversation, and analyze subjectivity and sentiment of the messages to characterize linguistic output using LIWC (Linguistic Inquiry Word Count) [60, 66]. Furthermore, features that cannot be easily detected by algorithms, such as psychological safety [13, 13, 23, 24], were coded by Amazon Mechanical Turk workers. We employ a lasso (L1 regularized) regression model to identify the most predictive features, achieving the highest accuracy of .92 AUC ROC for the classification of the top decile viability teams, .75 AUC ROC for the task with a median split, and .74 AUC ROC for the classification of the lowest decile viability teams. These results remain robust even when controlling for team performance.

Given these models, we highlight a few behavior patterns that we found as being important signals for viability. First, we find that use of exclusive language, such as 'but', 'either', 'except', and addressing via the second person, such as 'you' and 'your', are significant predictors of higher viability across all three tasks, which suggest teams whose members actively respond to and engage with the ideas of others are more likely to be viable. Sarcastic humor, anxiety terms, and higher average reading level were additionally significant in distinguishing the top decile teams from the rest, suggesting that the most viable teams are engaging in nuanced and complex interactions, in

which context and focus matter. On the other hand, the use of words indicating sadness as defined by LIWC and the negativity text sentiment of the most negative individual were associated with the lowest decile viability teams.

Many of the significant features in our models, such as those reflecting word choice and sentiment, were computationally derived rather than human-coded. This result hints at the possibility of enabling teams to conduct real-time assessment of their viability. Furthermore, we demonstrate that the relationship between viability and team interaction is consistent over time, yet the pattern differs between the first time a team convenes and when it meets again even if they do not recognize their teammates [69], which suggests that viability is a dynamic measure that could be monitored continuously. Finally, we find that models using features computationally derived from the first 70 seconds of an interaction can achieve a comparable level of performance to the models utilizing full 10 minute long interaction, implying that signals of team viability arise even in the first few moments of interaction.

While previous works have focused on predicting and understanding the determinants of team performance [11, 48], team viability is an equally critical goal for promoting effective collaboration. Evaluating viability is especially important in remote teams because they lack a shared context [54], are more likely to have difficulty developing mutual understanding [25], are less cohesive than face-to-face groups [63], and exhibit more competitive behavior [57], which render them more prone to developing conflicts [30]. In this paper, we present models for assessing viability via classification of online team interaction and explore potential signals of team viability in text interaction data. We make our classification algorithm available for others to use.<sup>1</sup>

## 2 BACKGROUND AND RELATED WORK

Teams are complex dynamic systems that develop as members interact and respond to the unfolding of situational demands over time in a given context [37]. On one hand, a positive team experience can help members come together to produce a whole that is greater than the sum of its parts. On the other hand, a negative team experience can sabotage success to the point where members vow to never work together again. Given the critical role that teams and teamwork play in our economy and society, how might we better understand the viability of any given team?

While team viability and team performance are correlated with a strong and positive relationship for teams completing routine activities [10], teams that perform well can still have disastrously low viability, as exemplified by the eventual fracture between Kahneman and Tversky [44]. Moreover, viability and performance also tend to have a weaker relationship when the team encounters dynamic situations that require different strategies and skills [36]. Thus, team viability should be considered and studied as a distinct feature from team performance.

### 2.1 Team Viability

When team viability is absent or lost to such an extent that the team chooses not to work together in the future, a phenomenon called team fracture occurs [69]. Literature shows that interventions for improving the psychological state of a team and its members can be effective [40]. If we can recognize that a team has low viability or is dropping in viability, it would be possible to try to help them before they fail [32]. Thus, there exists an important need for being able to recognize that a team needs help.

Currently, there are two ways that this need may be addressed, and both of them have significant drawbacks. The first mechanism for recognizing that a team needs help is through self-reporting from individual team members or the team as a collective. For example, an individual may raise

---

<sup>1</sup><https://viability.stanford.edu>

complaints about their team to the team manager or other parties. This mechanism encounters three problems. First, self-reporting depends on the *ability* of an individual to recognize low or falling team viability. People are often remarkably bad at recognizing their own state [48], and may not realize the decline of team viability in the moment. Second, the *responsibility* of speaking up about problems with team viability entirely falls on the reporting individual. Social loafing and diffusion of responsibility are obstacles to self-reporting even if an individual recognizes problems, particularly when a task is unattractive. Third, *potential consequences* present significant disincentives for the reporting party. Self-reporting may potentially lead to interventions and improvement [40], but it could also lead to total disintegration of collaboration or the ostracization of the reporting individual from the team and even the entire organization [20]. Indeed, even in the best case scenario where the overall viability of a team increases as a result of an individual's report, the social status of that individual is still likely to suffer [55]. Because of all of these reasons, individuals often may not act to prevent the impending failure of a team.

The second existing mechanism for recognizing that a team needs help is through external involvement, often taking the form of third-party actors, such as managers, coaches, and counselors, being called to step in via processes such as quarterly team assessments [46]. External involvement allows for the external evaluation of a team, and third-party actors often have incentive to report and act upon problems with a team's viability [12]. These factors account for some of the difficulties of relying on self-reporting as a solution. However, external involvement has its own drawbacks. First, introducing an evaluating party who is not directly in the team introduces potential inaccuracies in their evaluation [47]. When this party is related to the team, e.g. in the case of a team manager [5], they may bring biases and misaligned incentives of their own about their team in competition with other teams, or the characteristics of individual members of their team [12]. When this evaluating party is entirely separate from the team, e.g. in the case of a counsellor or a team-building professional [4], they may not be able to understand the complete context in which the team lives, and their evaluations may be distorted as well. Second, even if we disregard whether an accurate evaluation of the team is possible, an external human evaluation of the team costs resources, which presents a barrier to entry for smaller teams or organizations. Finally, by the time a third-party actor is summoned to deal with team issues, it may already be too late to matter [38].

Current ways of recognizing that a team needs help with viability are not ideal — a gap exists in the literature, with regards to a systematic investigation of the features of team viability and its possible predictability. Team interactions have been the subject of countless rounds of observation and analysis. One point of consensus in prior research is that predicting team behaviour is difficult, often because whether the team dynamics are effective is not obvious even to the team members themselves [48]. In this work, we set out to answer a related research question:

**RESEARCH QUESTION 1 (RQ1).** *Can we differentiate high and low viability teams by analyzing team interactions?*

## 2.2 Antecedents of team viability

Existing literature explores predicting team fracture and team performance, but not team viability, using field studies and crowdsourcing. Certain features emerge as strong predictors of team performance in field studies, such as team minimum agreeableness and team mean conscientiousness, openness to experience, collectivism, and preference for teamwork [11]. Interesting insights derived from prior research are that linguistic cues and conversational patterns extracted from the first 20 seconds of a team discussion are predictive of whether it will be a productive discussion [52], and that the type of task matters for the consistency of team success or fracture [69]. Notably, existing literature has focused on a number of broad feature categories, where there is an overlap between

the categories. In this paper, we do not aim to identify patterns on the level of the categories, but we do draw upon trends in these categories to select our features.

**2.2.1 Work Pattern Features.** From proposing ideas to implementing decisions, equal participation is a key aspect of constructive teams [53]. Well-balanced participation is a challenge in virtual, chat-based discussions because they lack in-person cues. For example, in offline interactions, individuals can demonstrate that they are listening to a person speak by nodding their head in agreement, or maintaining eye contact. However, in online interactions, individuals can only convey their listening by sending written responses [39]. This limitation opens the door to many problematic scenarios as it can be exploited by internet trolls who take pleasure in upsetting others, as well as socially-insensitive individuals who simply neglect to respond [34]. When a person does not see evidence that they are being heard, they feel dissatisfied and frustrated. Such frustration makes people less willing to participate and contribute, sometimes disengaging completely.

A team's performance is also an antecedent to its viability [6, 45], but the relationship is complicated by the stresses of achieving performance in some situations [12]. Producing high quantities of ideas can be equated to high performance since teams that are diverse, insightful, and creative are better able to cultivate a rich flow of ideas. Another way to evaluate a team's performance is its ability to reach conclusions efficiently, as cohesive teams tend to collectively process ideas and converge to a final decision more quickly than non-cohesive teams [21].

**2.2.2 Semantic Features.** Constructive communication has an ameliorative effect on groups, reducing conflict and increasing cooperation [7]. For example, positive interactions induce a pleasant overall mood, which is associated with greater cognitive effort and more constructive methods of dealing with disagreement. The foundations of a healthy team atmosphere are also commonly established during a team's first interactions, since initial communication and interaction behaviors set the standard for team culture [65]. In particular, friendly introductions and greetings between group members create social bonds, which lay the foundation for benevolence and trust [27]. To sustain this trust, team members often engage in supportive communication in accomplishing a task (e.g. celebrating achievements, expressing appreciation of each other's contributions). Additionally, since online discussions do not convey the same richness of emotion and reaction that face-to-face communication enables, they are susceptible to misunderstandings when utterances contain sarcasm or irony [50]. Due to the use of sarcasm, what might look like a negative utterance is, in fact, a positive one (and vice versa). Consequently, sarcastic or ironic utterances online can affect team members' perceptions of positivity and interaction behaviors.

Rudeness and aggression can quickly sabotage a healthy conversation (e.g., ad hominem attacks, name-calling, knee-jerk contradiction) [22]. Instead, healthy team dynamics center on agreeableness, respect, and support. From softening the perceived force of a message to allowing all parties to save face, politeness can shape the course of online interactions for the better [72]. Inter-member civility also encourages social connection and rapport, fostering amicable conversation between members and helping groups maintain cohesion throughout a discussion.

**2.2.3 Topic Features.** Some topical contexts naturally lend themselves to antisocial behavior and discord. Specifically, topics like politics, religion, and ideology are breeding grounds for conflict [72]. Religion and philosophy account for over half of the conflict on Wikipedia [35], while 14% of Americans say they have been harassed online due to their political views [23]. Moreover, people tend to feel more strongly about ideological issues and take matters personally [3]. Such emotional investment makes individuals more susceptible to friction and hostility, with decision-making becoming more emotional than logical.

**2.2.4 Word Choice Features.** Psychological safety refers to a shared belief held by teammates that a team is safe for interpersonal risk taking [24]. Task conflicts that occur in a psychologically safe environment have been found to improve creativity and decision making without damaging interpersonal relations or interactions between teammates [13]. In this way, a highly productive team may be highly critical but also highly psychologically safe. In the opposite direction, a highly critical team that is not psychologically safe can transform how teammates respond to challenges. A multitude of factors can result in teams not being psychologically safe, for example name-calling, physical threats, dismissive remarks, and social punishments associated with rejecting the group norms [13, 23]. In 2017, 41% of Americans said they have suffered online harassment, 27% said they have been called offensive names online, and 22% say that someone intentionally embarrassed them online [23]. Victims of such behaviors often experience mental or emotional stress, and are unwilling to work with their perpetrators again. Such stress is particularly common in online mediums, where anonymity emboldens perpetrators to use words that they likely do not use in in-person interactions. These situations motivate another research question:

**RESEARCH QUESTION 2 (RQ2).** *Which behaviors are most strongly associated with high vs. low viability teams?*

### 2.3 The durability of team viability antecedents

Team interactions evolve over time [29] — what indicates a positive relationship early in a team interaction, might indicate a problematic one if the team is expected to be very familiar due to months of close collaboration. Though rare, studies of longitudinal interactions in teams show compounding effects of low viability [42]. And thus the question of how repeated episodes of work might expose unexpected qualities of team viability. Intriguingly, the dataset analyzed in this work has the unique property of involving repeated interactions between the same team without them knowing it, as well as interactions when they did know they were working with the same prior collaborators. While the interactions in the data are only 10 minutes long each, their repeated nature may shed light on another research question:

**RESEARCH QUESTION 3 (RQ3).** *How does the historic collaboration context influence team viability?*

### 2.4 Early warnings

When studying events with final outcomes, it can be helpful to understand whether that final result can be identified by some patterns or behaviors early in the process rather than after the outcome has already occurred. “Thin slicing” studies show how relatively small and relatively early interactions can be powerful predictors of long-term outcomes — famously by accurately predicting which marriages will fail years later [15, 26], or the grades at the end of a semester [33] — by making judgements based on only a few minutes of interactions. Understanding viability in the future, as opposed to after a team has already come to have low viability, would turn predicting viability into something strikingly useful for teams in practice. This goal motivates a final research question to study:

**RESEARCH QUESTION 4 (RQ4).** *How little duration must be observed to make accurate predictions of team viability?*

In this section we have mapped out the challenges in classifying team viability and motivated a range of theoretically and practically interesting questions to try to answer. We next move in to explaining how we achieve the classification task.



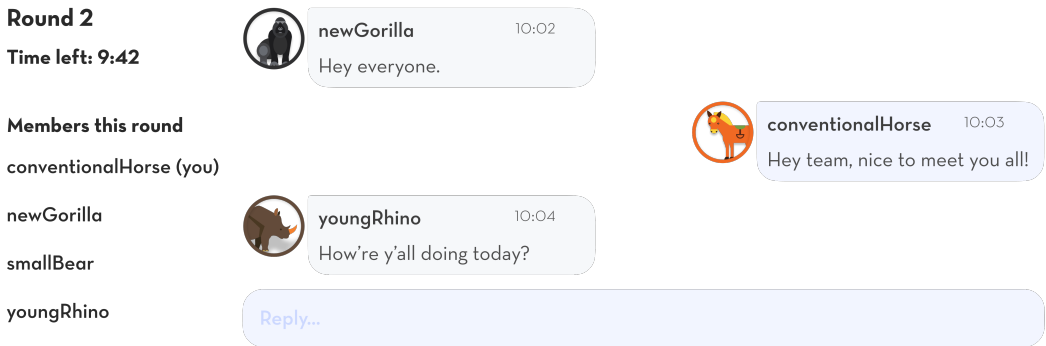


Fig. 1. Participants collaborated in a chat room with anonymous usernames. Our dataset of interaction histories consists of all chat messages, who they were sent by, and a timestamp.

### 3 METHOD

In this paper, we focus on differentiating high and low viability teams through chat transcripts of online team interaction. We operationize our goal as a binary classification task. Specifically, we derive both algorithmic and human-coded features from organizational behavior and behavioral science and apply to them to the interaction data from the teams' chat transcripts. The algorithmic features are drawn from a toolkit of modern natural language processing and computational social science tools; human-coded features are labeled by multiple independent workers on Amazon Mechanical Turk. Using these features, we then seek to classify high or low team viability using machine learning models.

In this section we will describe how the team interaction dataset was processed, how the features were encoded on the interaction data, and how the machine learning algorithms were set up to classify team viability.

#### 3.1 Data processing

We work with data of teams' interactions gathered for a previous study that investigated team fracture and viability [69, 70]. In this dataset, participants recruited from Amazon Mechanical Turk were grouped into teams of 3–8 and remotely collaborated through a synchronous messaging chat room to complete various tasks over several rounds of 10 minutes. The teams are all made up of only US workers who have completed at least 500 tasks on Mechanical Turk (48% female, and 38 years old on average). The identity of members shown to others was manipulated such that any one group of members would work together 4–6 times without knowing that they had collaborated previously. These data include all of the messages exchanged during the discourse associated with unique pseudonyms and the timestamp since the beginning of each interaction. In these interactions, three different types of tasks from McGrath's task circumplex were used [64]: creative, intellectual, and cognitive conflict, although around 80% of the data utilize the creative task. After completion of each task, team members were asked to complete several surveys, including a 14-item viability instrument on a 5 point Likert scale [18].

As just mentioned, teams worked together multiple times in succession without knowing they were in fact interacting with the same collaborators [69, 70]. We achieved this effect through a two directional pseudonym masking technique designed to reset team dynamics between rounds — to participants, it appeared that they were starting a fresh collaboration with new team members. This was validated through an in-sample manipulation check to ensure that collaborators were truly not

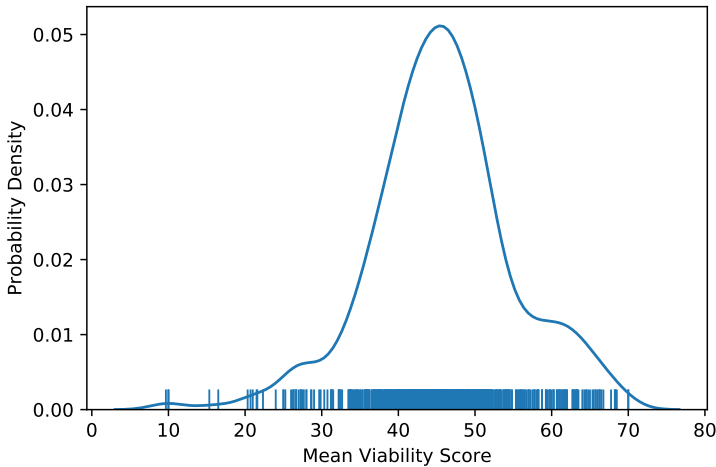


Fig. 2. Distribution of average team viability scores across all teams in our dataset.

aware that they were working together repeatedly. Additionally, among the repeated interactions for each team, there was also one pair of interactions where the team members were knowingly *reconvened* with the same group by using consistent pseudonyms. In the prior experiments, this treatment was used to understand path dependence of viability in teams. In this study, we will leave aside data from these repeated teams for every part of our analysis which does not explicitly require reconvened teams, in order to avoid comparing teams with different lengths of interaction. We do take advantage of this unique aspect of the data to specifically make claims about the properties of viability signals in reconvening teams relating to RQ3. As far as we know, this is the only chat dataset affording comparative analysis across independent rounds and reconvening teams, which is why we selected it for this paper.

In total, we analyzed 669 independent chat interactions of teams, aggregating to a total of 40,024 messages. For each team, we summed each collaborator’s 14 viability scale items into a single individual score with a range of 14–70 and calculated the mean of all of the individual scores to determine a team viability score. Figure 2 illustrates the overall distribution of team viability: the median viability score of all teams is 45.5, the 10th percentile is 34.95, and the 90th percentile is 58.25.

## 3.2 Feature Collection

**3.2.1 Computationally-derived Features.** To study characteristic differences between high and low viability teams, we first derive features through algorithms via the series of chat messages in a multi-person team interaction. These features can be divided into four categories: work pattern, semantics, topic, and word choice, as motivated by prior literature discussed in Section 2.2. Because there are anywhere from 3 to 8 people in a group, we first calculate the features on the level of each individual team member from the entire chat interaction and consequently produce final team-level statistics of that feature (mean, standard deviation, minimum, maximum) that are used to model the chat discourse.

With work pattern features, we seek to capture higher level behavioral information from how team members are interacting. In task-oriented discussions, well-balanced participation among group members is a key factor present in productive teams [52]. We derive two features in this



Table 1. Computationally-derived features associated with work pattern, semantics, topic, and word choice. Some features are multiplied by four, because the minimum, maximum, mean, and standard deviation were calculated for each.

Feature Group	Number	Specific Features
Work Pattern	2×4	Number of messages per person, number of words per person
Semantics	2×4	Subjectivity, polarity
Topic	1	TF-IDF cosine similarity
Word Choice	25	LIWC (anger, anxiety, sadness, articles, conjunctions, prepositions, inclusive, exclusive, quantifiers, certainty, discrepancies, negation, tentativeness, first person singular, first person plural, second person, indefinite pronouns, adverbs, social), argue, reference, readability

Table 2. Computationally-derived features drawn from previous literature.

Feature	Measurement Method
Average message count, Talkative member message count, Laconic member message count, Variation in message counts, Average word count, Talkative member word count, Laconic member word count, Variation in word count [52]	Statistical measures on team members' chat messages
Average polarity score, Highest polarity individual score, Least polarity individual score, Variation in polarity scores, Average subjectivity score, Most subjective individual score, Least subjective individual score, Variation in subjectivity scores [51]	TextBlob sentiment analysis algorithms evaluated on team members' messages
Average reading level, Highest individual reading level, Lowest individual reading level, Variation in reading levels [58]	Dale-Chall readability score algorithm applied to team members' messages
Anger words, Anxiety words, Sadness words, Articles, Conjunctions, Prepositions, Inclusive language, Exclusive language, Certain language, Disagreement language, Negation language, Tentative language, First person singular, First person plural, Second person, Indefinite pronouns, Adverbs, Relationship-oriented language [19, 66]	LIWC lexical category counts
Collaborator language similarity [60]	Average cosine similarity of TF-IDF vectors based on unigrams and bigrams
Argument discourse markers [1, 2]	Counts of words from <i>ARGUE</i> corpus
Pseudonym use [59]	Counts of referring to others by pseudonym in the chat messages

area: number of messages sent and number of words used by each team member. For each person, we count the number of separate messages he/she sent as well as the total number of words he/she used across all messages for the entire chat discourse.

For semantic features, we applied text-based sentiment analysis as it has been widely used in predicting human behavior as in social media [51]. Specifically, we derive the polarity (whether language is positive or negative) and subjectivity (whether language is objective or subjective) of group members' messages using the natural language processing library TextBlob.<sup>2</sup>

For the topic features, we expanded upon previous work, which show that unigrams and bigrams are particularly effective for disagreement and argument detection [60]. Specifically, we create TF-IDF (term frequency-inverse document frequency) vectors based on unigrams and bigrams as a representation of topics in a person's chat messages. Preprocessing of the messages only involved removing accents to normalize to the Unicode standard, and converting all letters to lowercase; neither stop word removal or other token-level analysis like stemming or lemmatization were used. The average pairwise cosine similarity of these vectors is used as the feature to describe the alignment of topics in different team members' language.

For the word choice features, we applied dictionary-based text analysis, which has often been used to identify signals in language, for example, to study mental health using social media [19]. Specifically, we use LIWC (Linguistic Inquiry Word Count), an extensively human validated lexicon, to count the words that individuals used during the chat discourse that fall into different lexical categories [66]. For our task, we reduce the total category set to a subset of those that are cognitive- or function-oriented and those related to affect (see Table 1), culminating to 19 categories. Additionally, because we expect conversational conflict to influence team viability, we include as one feature the total count for the presence of any 17 discourse markers for recognizing disagreement from the ARGUE corpus [1, 2]. Prior work on chat-based team dialogues [59] motivated us to incorporate the number of references to other team members as another feature. As a higher level representation of word choice, we also calculate the minimum, maximum, mean, and standard deviation of the Dale-Chall readability scores for each members' chat contribution.

**3.2.2 Human-labeled Features.** Furthermore, we hypothesized that certain aspects of teamwork would be discernible to humans but not to algorithms. We have drawn 20 features from organizational behavior (Table 3) to be hand-labeled. We used the Amazon Mechanical Turk platform to recruit 366 human annotators who had completed at least 500 tasks to annotate teams' chats. For a given chat, the human annotators filled out a survey of 20 questions corresponding to the 20 features, answering each question using a 5-point Likert scale.

More specifically, our 366 human annotators or "reviewers" must have completed at least 1000 tasks and must be based in the USA. The recruiting HIT also excluded all reviewers who had participated in the original studies. Workers were paid at a rate of \$15 per hour, in accordance with existing literature [71]. We then built a platform that displays one chat per annotation task along with instructions given to the original participants in the task. With all pseudo-names of the participants replaced by a label of "Person\_<number>" (e.g. Person1, Person2, ...) format, the reviewers were asked to rate their perception of the chat interaction on questions measuring selected features on a 5-point Likert scale and to predict whether a team would fracture. We included an attention check question asking for the number of participants in the chat and excluded responses that were not consistent with the truth from the analysis. The order of questions were randomized, and reviewers were allowed to participate in the task multiple times.

At least 3 reviewers annotated each chat, and the median response was selected as the feature value for analysis. The average standard deviation of scores across all raters and chats for a given

<sup>2</sup><https://textblob.readthedocs.io>

Table 3. Human-labeled features drawn from previous literature and the statements used to ask Amazon Mechanical Turk workers to rate them. The mean standard deviation ( $\bar{\sigma}$ ) is for each feature across all included raters and across interactions is also shown.

Feature	Survey Question	$\bar{\sigma}$
Generating Many Ideas [53]	There were many ideas/suggestions/input generated from this group.	0.92
Reaching a Conclusion [21]	The group reached a decision in the end (there was a sense of consensus).	0.88
Progressing Slowly [8]	It took a long time for the team to reach a conclusion.	1.15
Thoughtful Response [6, 45]	The final answer of the team is reasonable and well-thought-out (e.g: a creative answer for ad writing).	0.90
Social Loafing [53]	At least one team member was not participating as much as the other team members.	1.15
Annoying Collaborators [34, 39]	At least one team member was annoying (used all caps, or sarcasm, or was obnoxious, etc).	1.32
Frustrated Collaborators [34, 39]	At least one team member expressed frustration/confusion.	1.27
Positive Interaction [7]	There were positive interactions (emojis, lol, great!, etc).	0.89
Starting Salutations [27, 65]	The team members greeted each other at the beginning of the interaction.	0.87
Ending Salutations [27]	The team members congratulated each other at the end of the interaction.	0.96
Sarcastic Interaction [50]	There were sarcastic comments.	1.28
Agreeable Interaction [72]	The group was agreeable.	0.89
Respectful Interaction [72]	The teammates were supportive and respectful as opposed to insulting.	0.85
Passive-aggressive Interaction [22]	Team members were passive-aggressive.	1.28
Dismissing Collaborators [13, 23]	At least one person in the team was dismissed at some point.	1.35
Punishing Collaborators [13, 23]	At least one person in the team was punished at some point.	1.35
Showing Embarrassment [13, 23]	At least one person in the team seemed embarrassed at some point.	1.31
Political Interaction [3, 23, 35, 72]	This interaction was Political as opposed to Non-political.	1.35
Playful Interaction [3, 23, 35, 72]	The conversation was Playful as opposed to ideologically charged.	0.96
Fact-driven Discussion [3, 23, 35, 72]	Ideas were Fact-based as opposed to Emotion-based.	0.99

feature was 1.08 prior to filtering. Responses that were more than 1.5 standard deviation away from the mean result were excluded prior to calculating the median. Mean standard deviation values for each feature across all interactions are shown in Table 3.

### 3.3 Machine learning models

We operationize our goal through a binary classification task to differentiate high and low viability teams. For our analysis, we divide all teams into two groups – either high viability or low viability – based on their viability score. Specifically, we set the cutoff viability score as a varying decision threshold, and classify all teams with viability score lower than the cutoff score as low viability teams and teams with viability score above the cutoff as high viability teams. For instance, if we set the cutoff threshold as 45.5 (the median viability score in our dataset), we divide all teams into two equal sized groups, where those teams with viability score above the median set as the high viability sample and those teams with viability lower than the median as the low viability sample. As such, we treat predicting team viability as a binary classification problem because we are interested in distinguishing teams with relatively higher viability from teams with relatively lower viability, rather than predicting the particular measure of viability.

We implemented 5 widely-used machine learning classification models from Python’s scikit-learn library [56], i.e., logistic regression, support vector classification (SVC), random forest classifier, multi-layer perceptron (MLP) classifier, and gradient boosting classifier, to train and classify low vs high team viability based on our derived features. We used the area under the curve of the receiver operating characteristics curve (AUC ROC) as a measure of prediction performance, since we formulated the viability differentiation as a binary classification problem and AUC is insensitive to an unbalanced testing set. Note that most experiments are done via 5-fold stratified cross validation, and each feature in the training and test sets are standardized.

## 4 RESULTS

### 4.1 Differentiating team viability under different split

To answer RQ1, we evaluated the five models across nine different trials where we set the viability decision threshold at every ten percentile points, (e.g 10th percentile, 20th percentile, ... , 90th percentile). The models are able to differentiate higher from lower viability teams (Table 4). Perhaps more importantly, our results suggest that there exist general mechanisms behind team viability vs. team interaction, independent of the team member composition.

When we consider the case of classifying a team’s viability relative to the median split, the best AUC score of the five models is 0.76. More interestingly, we find that as the data partition becomes such that we label fewer samples as high viability teams and focus on the teams with relatively higher viability scores, the classification accuracy becomes higher. For instance, when we label the top decile teams as high viability teams and designate the remaining 90% as low viability teams, we reach over 0.9 AUC, indicating there are salient differences between teams above the 90th percentile and the rest.

Prediction Model	Data Partition Percentile								
	10	20	30	40	50	60	70	80	90
Logistic Regression	<b>0.74</b>	0.72	0.69	0.72	<b>0.75</b>	0.74	0.78	0.81	<b>0.92</b>
SVC	<b>0.63</b>	0.70	0.67	0.69	<b>0.76</b>	0.75	0.79	0.85	<b>0.92</b>
Random Forest	<b>0.66</b>	0.72	0.69	0.72	<b>0.73</b>	0.75	0.79	0.86	<b>0.93</b>
MLP Classifier	<b>0.66</b>	0.68	0.66	0.68	<b>0.71</b>	0.71	0.74	0.80	<b>0.90</b>
Gradient Boosting Classifier	<b>0.67</b>	0.70	0.67	0.72	<b>0.75</b>	0.73	0.76	0.86	<b>0.93</b>

Table 4. Viability classification performance (measured by ROC AUC) with different division of positive and negative samples using all features.

Prediction Model	Data Partition Percentile								
	10	20	30	40	50	60	70	80	90
Logistic Regression	<b>0.77</b>	0.76	0.73	0.75	<b>0.76</b>	0.76	0.78	0.82	<b>0.94</b>
SVC	<b>0.63</b>	0.74	0.72	0.73	<b>0.76</b>	0.75	0.79	0.84	<b>0.93</b>
Random Forest	<b>0.70</b>	0.70	0.69	0.73	<b>0.72</b>	0.73	0.77	0.84	<b>0.92</b>
MLP Classifier	<b>0.69</b>	0.71	0.71	0.71	<b>0.73</b>	0.74	0.75	0.82	<b>0.92</b>
Gradient Boosting Classifier	<b>0.66</b>	0.69	0.67	0.73	<b>0.76</b>	0.73	0.78	0.85	<b>0.92</b>

Table 5. Viability classification performance (measured by ROC AUC) with different division of positive and negative samples by computationally derived features.

Prediction Model	Data Partition Percentile								
	10	20	30	40	50	60	70	80	90
Logistic Regression	<b>0.56</b>	0.56	0.53	0.54	<b>0.61</b>	0.60	0.63	0.69	<b>0.71</b>
SVC	<b>0.46</b>	0.44	0.47	0.51	<b>0.58</b>	0.60	0.63	0.68	<b>0.66</b>
Random Forest	<b>0.58</b>	0.50	0.65	0.55	<b>0.52</b>	0.56	0.60	0.65	<b>0.68</b>
MLP Classifier	<b>0.55</b>	0.52	0.51	0.53	<b>0.55</b>	0.59	0.63	0.67	<b>0.70</b>
Gradient Boosting Classifier	<b>0.57</b>	0.48	0.53	0.55	<b>0.5</b>	0.56	0.59	0.67	<b>0.70</b>

Table 6. Viability classification performance (measured by ROC AUC) with different division of positive and negative samples by human labeled features.

In the feature set we are using, in order to capture a wide variety of team interaction behaviors, some features were automatically developed by computational means, and those that could not be, due to the limitations of current algorithms, were labeled by humans. For the practicality of creating a tool for differentiating team viability, we further performed a preliminary ablative analysis of our features. We repeat the previous procedure separately for the computationally derived and human labeled features, which is shown in Table 5 and Table 6, respectively.

We see a noticeable difference in the performance of the two methods, but regardless, we witness a similar pattern of increasing accuracy as the threshold for dividing the data increases, which can be seen in the right-side plot of Figure 3. Across all decision thresholds, the classification accuracy of the computationally derived features is at least as good as the full features set, having an AUC of 0.76 of at the median split and 0.94 at the 90th percentile split as seen in Table 5. Using our human labeled features alone results in only an AUC of 0.61 at the median split and 0.71 at the 90th percentile split as seen in Table 6. This discrepancy could suggest that 1) there can be significant signal overlap between computationally derived and human labeled features, or 2) human labeled features may be noisier than computationally derived features. This result provides evidence that computationally derived features alone may be suited to further study.

Finally, to rule out the possibility that team performance may be a confounder, we carried out a robustness check by measuring the correlation between team performance (as measured by actual task performance, e.g. click through rate of advertisement text in the creative task) and team viability. In our dataset, 175 teams have team performance information. We observe weak correlation between the two ( $R^2 = 0.10$ ). As another test, we add team performance as a feature to the classification model and compare it with the classification model without team performance on these 175 teams. We find no change in classification AUC: under 10% and 50% splits, the AUC is

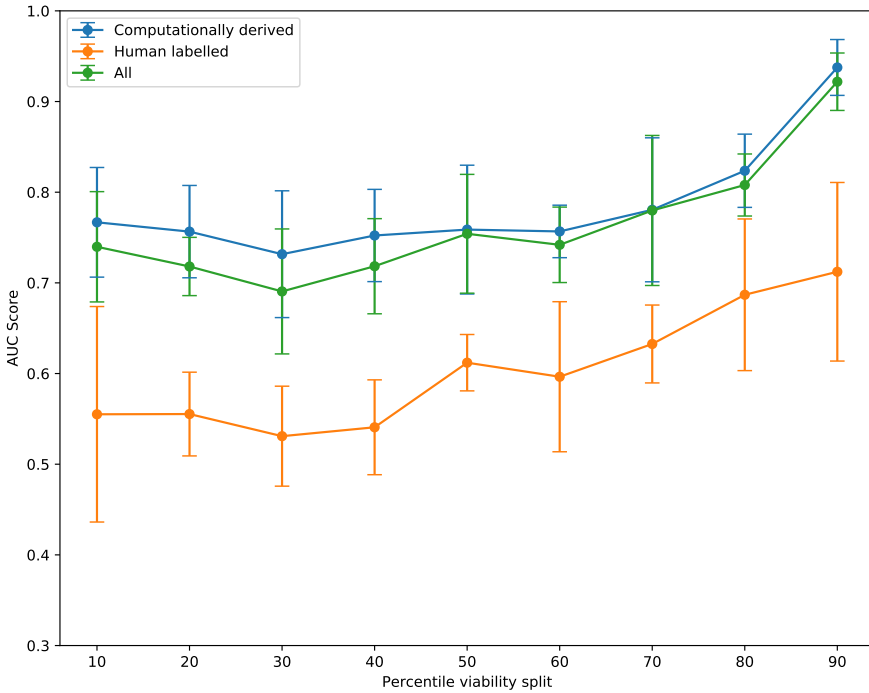


Fig. 3. Performance of the logistic regression model in AUC score with computationally derived features, human labeled features, and all features across different viability percentile thresholds. Errors bars show the 95% confidence interval across 5 folds.

stable at 0.78, and under 90% split, the AUC is stable at 0.96 under both conditions, which further verifies that team viability cannot be explained by team performance.

For the rest of our analysis, we will focus on using our logistic regression model. It achieves the best results in most viability splits and as a model, it is computationally simpler and more interpretable than the other models we tested. Additionally, we report other performance metrics specifically for logistic regression under various conditions of decision threshold and used features in Table 7.

## 4.2 Feature Analysis

To answer RQ2, we proceed to investigate the role of different features in predicting team viability. Specifically, we run penalized logistic regression with L1 regularization (Lasso regression) for our model using threshold scores of 34.95 (bottom decile split), 45.5 (median split), and 58.25 (top decile split), and examine the coefficients of the features to understand whether certain features contribute positively or negatively to team viability. We choose penalized logistic regression since it helps prevent collinearity and sparsity with the effect that the coefficients' weights are moved to the most predictive features.

We present coefficients and their confidence intervals for all the significant features in Fig. 4. As can be observed, two features stand out as significant features across all three splits. Use of exclusive language (as defined by LIWC), such as 'but', 'either', 'except', and addressing each other in second person (e.g. 'you') are significantly and positively correlated with the teams with relatively



Features	Partition				
	Percentile	F1 Score	Precision	Recall	ROC AUC
Computationally Derived	10	0.10	0.13	0.07	0.77
	50	0.67	0.70	0.64	0.76
	90	0.57	0.69	0.52	0.94
Human Labelled	10	0.03	0.20	0.02	0.56
	50	0.57	0.59	0.56	0.61
	90	0.08	0.50	0.05	0.71
All Features	10	0.20	0.32	0.15	0.74
	50	0.66	0.67	0.64	0.75
	90	0.56	0.62	0.52	0.92

Table 7. Additional metrics of viability classification for Logistic Regression when using a decision threshold at the 10th, 50th, and 90th percentile of viability score. Note that AUC is a better metric to evaluate viability classification performance in our case since the metric based on ranking, while F1 score, precision and recall are based on threshold to divide positive and negative labels. The high AUC for our model indicates the ability of our model to correctly rank high viability team ahead of low viability team, while the relatively low F1 score only reflects the cutoff threshold is not ideal.

higher viability across all three splits, which support findings that suggest teams whose members actively respond to and engage with the ideas of others are more likely to be viable [9]. In one of the top decile teams, the team members used exclusive language to consider a variety of different viewpoints and arguments when discussing their solution,

**Person3:** I believe it would be fair to allocate most, if not all of the funds to project 3. Neither of the other two projects are vital, in my opinion.

**Person4:** great idea but sooo many people are hurting

**Person2:** #2 is the artsy option that we might get donations from the art community or art enthusiasts to help

On the other other hand, we observe that there are quite a few non-overlapping features that are significant when differentiating high and low viability teams at different thresholds, indicating that bottom decile viability teams, above and below median teams and top decile viability teams demonstrate different interaction behavior.

What distinguishes top decile teams from other teams? The use of exclusive language and second person are very strongly correlated with top decile teams, even more so than the other two splits for which these two features are also positive indicators. Other features such as articles and adverbs support this overarching trend of active engagement with other team members. As such, more active engagement appears to be an extremely important trend for top decile teams. A higher average reading level for the chat interaction is also positively correlated with membership in the top decile, suggesting that top decile teams often have the willingness and ability to communicate complex topics. In addition, the top decile has two interesting indicators which seem unintuitive upon first glance. The human-labeled presence of sarcastic interaction is positively correlated with viability scores above the 90th percentile. Upon further investigation, we found that while sarcasm has a negative connotation generally, the instances of human-labeled sarcasm which were present in top decile team interactions were of the humorous variety.

**Person2:** 5) it's in the 120 range maybe? anyone live in Death Valley? ;)

In addition, the computationally-derived feature of Anxiety Words was also a positive indicator of top decile teams, but further investigation reveals that team members were using words which the algorithm identified as anxious in order to talk about their task, e.g., that "veterans were *struggling*".

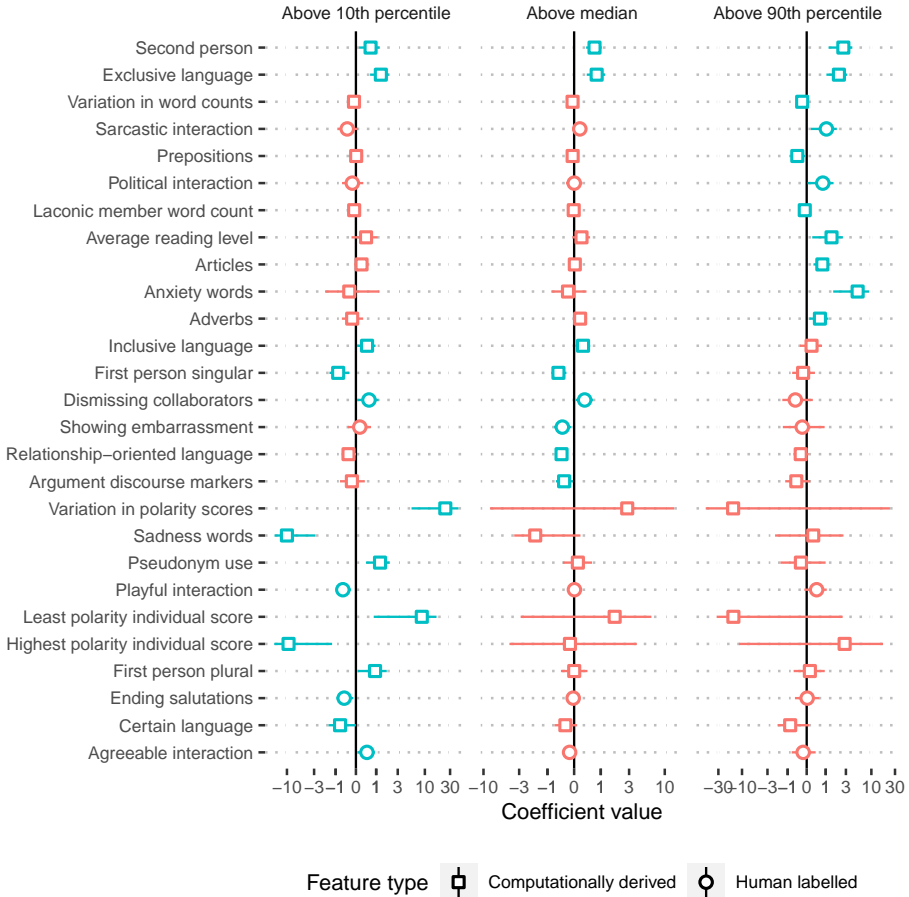


Fig. 4. Logistic regression coefficients for features. Features with confidence interval distinguishable from zero (significant) are shown in blue, while features with confidence interval overlapping zero (insignificant) are shown in red. Features are ordered from most to least useful in classification, and they are clustered by viability split. Values are shown on a hyperbolic arcsine scale to cope with variance.

What makes a team more likely to be a bottom decile team? The significant features that stand out make intuitive and theoretical sense. More exhibits of sadness, as well as the presence of a highly polarized individual – an individual who outputs very positive or very negative chat – are associated with the bottom decile of teams. These findings are consistent with the organizational behavior theory that displays of emotion, especially negative ones, will negatively impact the effectiveness of the team. [43] Interestingly, more use of first person singular terms (e.g., I, me, mine) and less use of first person plural (e.g., we, our, us) and second person words (e.g. you, your) is also associated with a higher chance of being in the bottom decile teams. This trend contributes to our conclusion that more active engagement with other team members correlates with higher viability, and less active engagement with other members correlates with lower viability. Finally, playful interactions were correlated with the bottom decile teams as well. Further investigation shows that overly playful teams had difficulty actually resolving the task at hand, ultimately leading

to lower viability scores. For example, one team had a hard time deciding on a slogan for a chair after joking about a pun,

**Person3:** What about Bar Stool for American's Behind

**Person4:** that's America's Ass!

**Person4:** it would get meme points lol

**Person3:** A Stool from America's Ass

...

**Person2:** we have to decide ppl

When distinguishing teams above and below the median viability score, the general trends we discover in the more extreme splits hold true. More use of exclusive language and second person pronouns, which directly engage with other team members, are positive indicators of team viability greater than the median. The presence of argument discourse markers, and exhibitions of embarrassment, are correlated with team viability lower than the median. Two unusual features are specifically significant for the median viability split. Dismissing collaborators is a positive indicator for teams above median viability, which further corroborates the overarching trend that directly engaging with other team members — even when such engagement is active disagreement — is correlated with higher viability, since direct engagement allows team members to feel that the team actively communicates on issues. Finally, relationship-oriented language is interestingly correlated with teams below the median viability. We note that this LIWC category contained words that were colloquial and could easily be interpreted negatively, e.g. "give advice" and "bro". While these words could be used positively in some situations, we expect that their negative interpretation may be disproportionately more impactful on team viability, while their positive interpretation may have little to no impact.

### 4.3 Differentiating team viability over time

To answer RQ3, we further investigate whether our model can be used to classify team viability over time. Specifically, we are curious about whether mechanisms behind team viability are time dependent or generalizable over time, i.e. is a team interaction - viability model learnt from team interactions at a certain stage applicable to classification of team viability at another stage?

The dataset we use is well-suited to study this research question since in the experiment design, people are asked to engage in multiple rounds of tasks (round 1 to round 4), and for each round they join newly-convened teams with teammates they do not recognize.

To shed light on this question, we ran a series of ablative experiments where for  $j \in [1, 4]$ , we trained a logistic regression model using team interactions from all rounds except round  $j$  and tested the model on round  $j$ , i.e. can we use interaction patterns mined from rounds other than  $j$  to predict team viability at round  $j$ ? Under this setting, we expect to observe rather consistent viability classification performance should the relationship between viability and team dynamics be static, otherwise there will be a significant performance drop since signals from round  $j$  are not encoded in the trained model.

In other words, if the behaviors that are associated with high or low viability are stable, we would expect the model trained on a team's first meeting to still function well on their fourth meeting; if the behaviors are not stable, we would expect the model to have weaker performance when run on a different meeting than its training set. We ran separate experiments for all features, computationally derived features and human labeled features with varying decision thresholds of viability scores. Additionally, we ensure there are equal numbers of training/testing samples (300 training, 60 testing) across trials to provide a fair comparison.

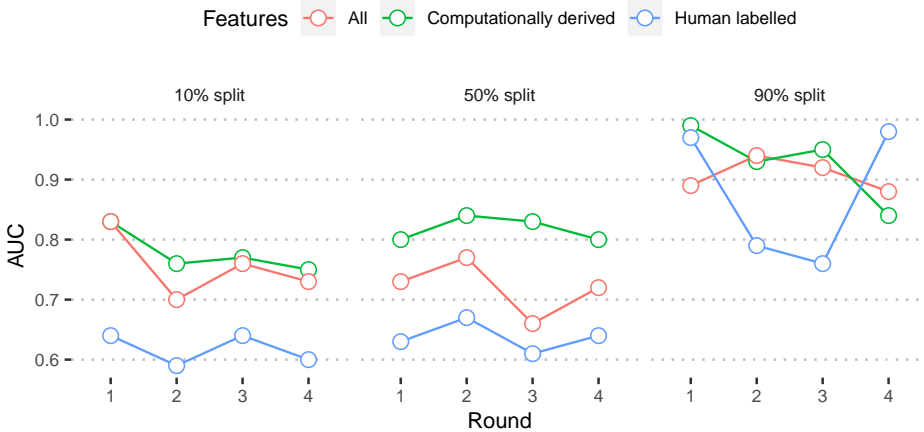


Fig. 5. Predictive accuracy across rounds appears to not show any major trend. Notably, no confidence intervals are provided because these could not be calculated with cross validation — these results are exact given the data studied.

We present the results in Fig 5. As can be observed, in the case of the 10% and 50% split, there is not much deviation in the model’s performance between different rounds used to train the model whether all features, computationally derived or human labeled features are used. Classification performance is stable around AUC of 0.8 for all features and algorithm derived features while around AUC of 0.65 for human labeled features. For the 90% viability split, we do observe a large drop from AUC close to 1.0 to AUC 0.8 for human labeled features at rounds 2 and 3 and for computationally derived features at round 4, yet there is no clear monotonic trend between performance and testing round indicating the variation is likely due to noise. Moreover, classification performance for all features remains relatively stable around 0.9 under the 90% split. Thus we conclude that our proposed model is rather robust regardless of the length of time team members have engaged in the work, which provides evidence that the relationship between viability and team interaction is static over time.

#### 4.4 Differentiating Initial and Reconvened Team Viability

As mentioned in Section 3.1, the dataset we used has a unique property of reconvened teams [69, 70] in addition to the independent rounds we have just discussed. In reconvened rounds, team members were aware that they were collaborating with the same team members for more than one separate round. These reconvened teams have not appeared in our training or testing sets so far to avoid any potential confounds, however, here we are interested to investigate whether our model will perform differently under the reconvened scenario, so as to answer RQ3 from another perspective. Specifically, we trained the viability classification model using all the masked interactions (300 interactions), and test on unmasked initial interaction (133 rounds) and reconvened interaction (132 rounds), for the three feature groups under the 10%, 50% and 90% division of viability. The result is reported in Figure 6.

We find that when using all features and computationally derived features, classifying team viability for reconvened interactions achieves much lower performance compared to classifying team viability for the initial interactions when we use the 10% and 50% division: for example, for the 10% split, computational features achieve 0.81 AUC when classifying the initial interaction

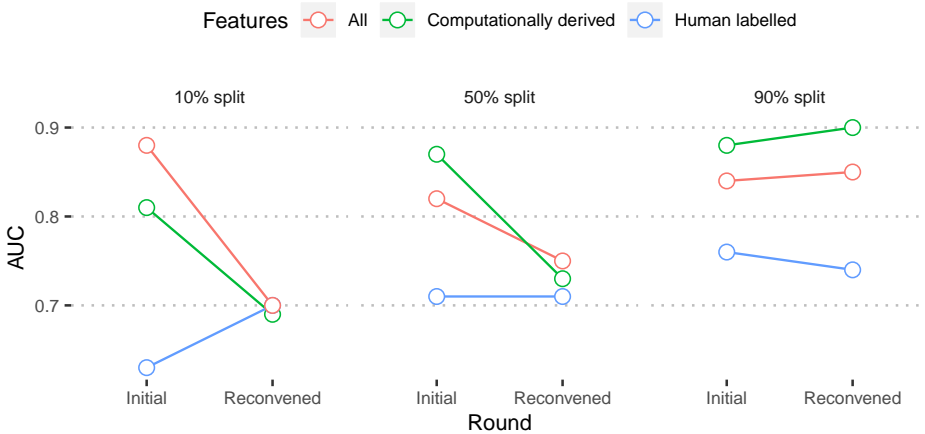


Fig. 6. Predictive accuracy in reconvened rounds sees a substantial dip for algorithmic features in lower viability split cases, while human features appear less affected. Notably, no confidence intervals are provided because these could not be calculated with cross validation – these results are exact given the data studied.

but only 0.69 AUC when reconvened. Given prior analysis where we have shown that patterns behind team viability is rather consistent over time for 10% and 50% division, we conclude that the result indicates that computationally derived features are not so consistent at differentiating team viability in the reconvened condition, which may be a result of the model failing to capture signals embedded in the fact that team members know they are working with the same group of people for the reconvened rounds. Meanwhile, at the 90th percentile split, the viability classification performance remains relatively stable, indicating that the top decile viable teams have behavior patterns that remain consistent despite knowing they are working with the same group.

Interestingly, we find that for human labeled features, the viability classification performances for the initial and reconvened interactions are identical or even increase, indicating that human labeled features are perhaps more robust at classifying reconvened interactions.

#### 4.5 Differentiating viability based on different amounts of interaction

Finally, to answer RQ4, we performed an additional analysis using our computationally derived features to test how little of the chat discourse is necessary to effectively differentiate teams' viabilities. Because we will need to consider smaller windows of the chat interactions, we can only apply the features that were derived computationally; the method with which we produced human labeled features was only performed once at the end of the whole chat interaction. We expect that the group interactions are dynamic and may change over the course of completing the task. For this study, since the chat histories we are analyzing are about 10 minutes long, we consider windows of data at 10 second increments, totaling to 60 trials. We start by calculating our features on the first ten seconds and evaluating the performance of our model; next, we take the data within the first 20 second window, calculate features, and evaluate the AUC score; this continues until the full 10 minutes of data is used. All trials were performed only with our Logistic Regression model with three different viability threshold splits: 10%, 50%, and 90%.

In Figure 7, we see that when using the 10th and 50th percentile splits, the rate of improvement is gradual, and reaches consistent performance when at least 200 seconds of chat data is used. In these two cases, our approach obtains 95% of the full accuracy with at least 35% of the entire interaction.

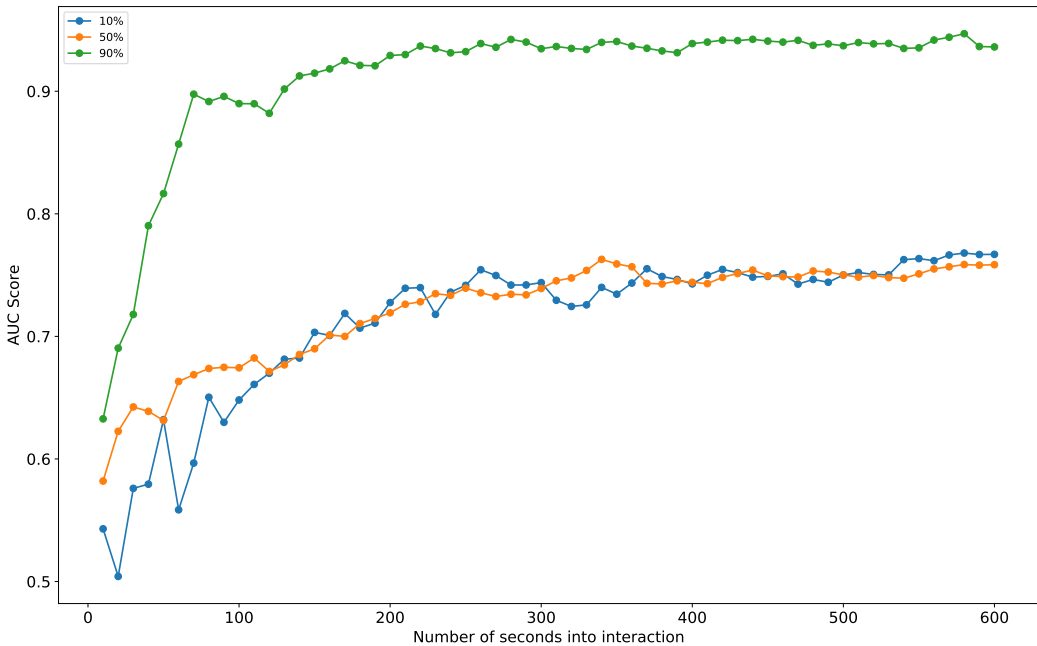


Fig. 7. Results of training the Logistic Regression model with data from different numbers of seconds of the chat interaction when low v. high viability is split at the 10, 50, and 90 percentile viability.

The performance trend when analyzing the top decile teams (90% split) is notably different; as the number of seconds of data increases, the AUC score rapidly increases and converges to steady performance as soon as when the training procedure considers at about 70 seconds of chat data. With only 12% of the chat interaction, we achieve 95% of the full accuracy when differentiating the top decile teams. This could suggest that the teams in the top decile of team viability exhibit some behavior earlier on that strongly contributes to predicting their final team viability. Regardless of the viability threshold used to divide the teams, we see that our model can consistently predict low v. high team viability much earlier in the process than at the full 10 minute mark.

To understand why this small window of chat data early in the collaboration can be predictive of the team viability by the end of the interaction, we compare the performance of consecutive disjoint windows of data to capture behaviors at different times in the interaction. Because the model achieved 95% of the full accuracy in about 200 seconds for the 10th and median viability split, we divide the ten minute interaction into three 200 second windows. Whereas the previous thin slicing analysis was done by considering cumulative windows of data, this time, none of the windows of data for each trial share any data. Looking at Figure 8, we see that the performance for all three viability splits remains relatively consistent over the three windows. This suggests that a small slice of data in the initial stages of the interaction is effective because by 200 seconds, team members start to exhibit consistent signals of behaviors that influence viability. Latter windows in the middle and end of the interaction continue to have the presence of these behaviors, hence the relatively consistent performance scores.



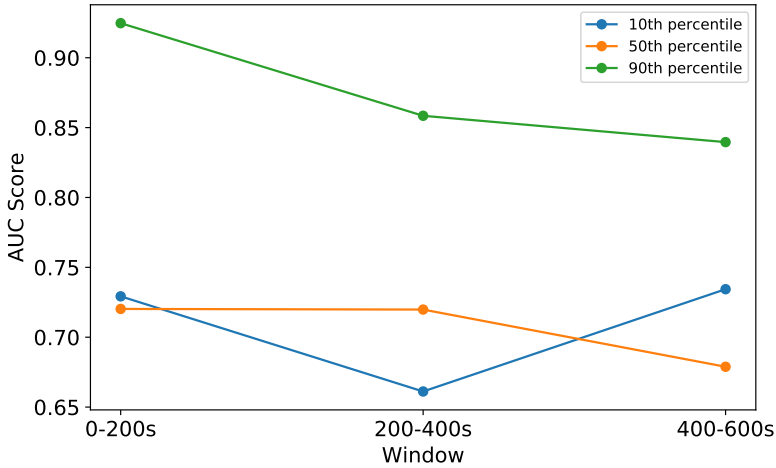


Fig. 8. Performance of Logistic Regression model trained on disjoint windows of 200 seconds at different percentile splits (10th, 50th, 90th).

## 5 DISCUSSION

### 5.1 Research Implications

The strong performance of our models in classifying teams by viability suggests that team viability largely is driven by underlying mechanisms that are common across different teams and repeated interactions of the same team. In particular, we find that work patterns and the language usage, which can be computationally derived, are effective in differentiating high and low viability teams with high accuracy of  $AUC = 0.76$  when equally splitting the data, and  $AUC \geq 0.9$  when differentiating the top decile of high viability teams from the rest. Furthermore, the considerably high accuracy for classifying the teams with the top decile viability score and that only two features are shared as being significant across all decision thresholds of viability score suggest that the behavior of the most viable teams are notably different from that of the rest of the teams, with lowest to moderate viability. The relative decrease in the performance when classifying relative to the median or 10th percentile viability score implies that differentiating teams with an average level of viability and those with the lowest viability is more challenging. While all teams with the highest viability in our dataset are viable for the same reasons, indicated by the common signals of active contribution, equal division of labor, more accessible, objective, and positive language, less viable teams appear to reach that state for differing reasons, not as consistently as the former.

We also note that a majority of our significant features in all three classification viability thresholds (10th, 50th, 90th) were computationally derived, so they can be evaluated without exposing the interaction text to another party or person and remain accurate given a fraction of each interaction. On the other hand, human-labelled features were more robust against changes in context, suggested by the consistency in differentiating team viability in the initial interactions and the reconvened rounds. Although our features are not exhaustive by any means, we suggest that automated features, which offer a consistent comparison of all teams, identify distinctive characteristics of teams depending on their viability, while the human-labelled features reflect human judgment that can detect nuanced signals of team dynamics, such as sarcastic humor, more efficiently.

## 5.2 Design Implications

Our work also has several design implications for building computer-supported collaboration platforms. For instance, in online crowd-sourced collaboration systems, we show that it is possible to differentiate high and low viability teams with even a few instances of member interaction, providing opportunities for timely intervention if necessary. Designers could potentially leverage our detection pipelines to help teams monitor their viability and build intervention tools that respond to warning signals to promote a successful collaboration experience. To illustrate this, if the system recognizes signals associated with low team viability (e.g., overuse of negation, absence of second person pronouns), the system can remind the teams to take steps for positive behavior change.

Such systems can help improve the sustainability and longevity of teams in the workplace. Today, many teams rely on human leadership (i.e. a manager or a supervisor) to resolve team viability issues. However, such human figures can only detect and intervene after an incident has already occurred. Often times, after an incident like a heated dispute occurs, a team's dynamics are beyond remedy. Meanwhile, the aforementioned systems can detect team viability issues in real time, allowing leadership to intervene before a major conflict erupts. This system affords teams a chance to resolve any growing issues and stay together before an irrevocable episode materializes.

The features used by these systems can also shed light on why a team's viability is suffering and recommend tailored intervention depending on the factors for low viability. For example, if a team's low viability can be explained by highly unequal participation, managers can subsequently address this specific issue and improve the team's work patterns. On the other hand, if low viability of a team can be attributed to overuse of negation, a nudging reminder to be more positive may be more effective. It is important to note that a combination of behaviors may be affecting a team's viability. Exploring different combinations may allow systems to learn the behavior patterns of a team and make more personalized recommendations for intervention. Furthermore, understanding how more viable teams behave can help provide areas for improvement for teams that are not necessarily considered not viable but simply want to improve how they work together. To illustrate, behavior associated with highest viability teams can be referenced when a team establishes a team norm.

Despite the previously mentioned beneficial use cases, these systems may also result in undesirable outcomes. To illustrate, a faulty prediction of the system can lead to the deterioration of viability going unnoticed until a major episode occurs or costly and unnecessary interventions. A more concerning implication of an automated system assessing viability is the possibility of the system adopting biases against certain communities based on distinctive language patterns, such as culturally-dependent colloquialism or accents. If the system is more sensitive to the language of certain groups of people, it may systematically and wrongly flag them as being problematic. Such an algorithmic bias would lead to discriminatory practices and inequitable treatment based on one's cultural identity, which is both detrimental to team viability and unacceptable. Finally, the system may be hijacked for malicious or inappropriate uses. For example, the sheer amount of data that would be collected regarding individuals' dialogue behavior raises concern for protecting one's privacy against surveillance. Furthermore, the system could also be inappropriately applied to predict characteristics other than viability (e.g. performance), leading to flawed evaluations and perceptions of teams.

These issues motivate a variety of design principles and guidelines for usage in order to avoid unnecessary cost, inefficacy, injustices, and malicious usage. First, signals used to assess viability should be made public and should be readily audited. Further, indications of viability should not be made for individuals and attempts to measure individual viability potential should be avoided.

Since viability is inherently a team property, extrapolation of assessment to individuals will be both flimsy and potentially damaging. Although it is tempting to develop recommendations with high precision, the trade-off between maximizing precision and minimizing potential abuse of targeted recommendations should be considered seriously. To resolve these issues in our own work, we are making available a version of our algorithm in which the features and their coefficients are transparently reported, and additionally, where a suitable amount of interaction material will be required to make assessments. Further discussion of ideas to respect privacy and ensure algorithmic fairness should be continued throughout the process of design, implementation, and deployment.

### 5.3 Limitations & Future Work

Considering the methods with which our dataset was gathered, there are a few limitations to the generalizability of the insights we draw from our analysis. First, the teams we analyze are relatively small in size and cannot represent the diversity and complexity in team roles and hierarchies in large organizations. Second, there may be important differences between remote collaboration and collocated work (even though the two forms demonstrate considerable characteristic similarities [14, 17]), particularly around finding common ground [54]. Prior research suggests that conflict of all types (task, affective, and process) will be detrimental to the performance of remote teams but may not have the same effects on traditional teams [31]. Thus, our approach may not fully capture interactions in collocated teams. Third, the chat interactions themselves are short in duration (10 mins for each round) and involve well-structured tasks. Coordination and problem defining in more complex, open-ended tasks may follow different dynamics than what participants did in our study, which could affect a team's viability.

In future work, it would be interesting to extend this approach to real world cases like software engineering tasks, which is a prolonged collaboration and involves complex interactions between people and the work itself. This work focuses on remote synchronous collaboration, but whether the same insights apply to face-to-face interactions or what differences there may be are both useful areas to explore. Also, due to the size of the collected dataset, we are unable to do fine-grained analysis comparing our model under different task categories. In the future, it would be interesting to extend our analysis to different task categories.

Nevertheless, we believe our derived insights do generalize and can be of great value to understand team viability under a remote collaboration/interaction setting (which is becoming increasingly important as more collaborations are shifted online after the COVID-19 pandemic), especially for short-term team interactions where members do not know each other's identities. Such flash interactions indeed happen often in online communities, e.g., in online forums, discussions under social media posts, and news reviews, and are also becoming common in crowdsourcing tasks. As such interactions are essentially similar in property to our collected data, we argue similar patterns are very likely to hold. Our findings and implications on team viability can help better understand user engagement in these settings.

Another possible line of analysis is delving into the cultural context of the teams, the human raters, and the computationally judged features. In the current work, all of the data are collected from US Amazon Mechanical Turk workers with moderate experience, and this sampling is likely to influence aspects of our result. Notably, the researchers attempted to code chats using the same prompts as the human raters did and found no notable difference in accuracy or inter rater reliability. Further, no stratification of the human rating data implicate sub populations within our analysis who systematically performed better. With this said, we expect these results may differ when other populations or cultural backgrounds are engaged. Finally, it would be useful to include other factors of interpersonal communications (e.g., linguistic accommodation, between individual personality/dissimilarity, etc.) as future work to gain a deeper understanding of team viability.

## 6 CONCLUSION

Recognizing the importance of team viability — rather than performance — in effective collaboration, we study signals of team viability in text interaction data of remote teams. Our machine learning models using features motivated by previous literature can reliably differentiate teams that belong to the top decile, above median, and the bottom decile of viability scores. Further investigation of the models suggest that while the use of exclusive language and second person pronouns are associated with higher viability teams across thresholds, other behaviors are unique to distinguishing the top decile, such as higher average reading level, and the bottom decile, such as sadness and displays of extreme emotions. Furthermore, we demonstrate that the relationship between viability and team interaction is consistent over time, yet the pattern differs by de novo teams and reconvened teams. Finally, we find that signals from the initial moments of interaction are sufficient for successful prediction of the whole interaction, which suggests that viability can be evaluated and acted upon early. The strength of our models using automated features indicate opportunities for real time assessment of viability, which lends further space for intervention.

## ACKNOWLEDGMENTS

We are grateful for the efforts of numerous research assistants, Amazon Mechanical Turk worker participants and Upwork research engineers — without them we would not have been able to collect this data. We thank Daniel A. McFarland for helpful discussions and suggestions. We also thank the associate editor and reviewers for their valuable comments. The data collection was supported by the Stanford Data Science Initiative, RISE Thailand Consortium, the Hasso Plattner Design Thinking Research Program, the Office of Naval Research (N00014-16-1-2894) and a National Science Foundation award IIS-1351131.

## REFERENCES

- [1] Rob Abbott, Marilyn Walker, Pranav Anand, Jean E Fox Tree, Robeson Bowmani, and Joseph King. 2011. How can you say such things!?: Recognizing disagreement in informal political argument. In *Proceedings of the Workshop on Languages in Social Media*. Association for Computational Linguistics, 2–11.
- [2] Kelsey Allen, Giuseppe Carenini, and Raymond Ng. 2014. Detecting disagreement in conversations using pseudo-monologic rhetorical structure. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1169–1180.
- [3] Pranav Anand, Marilyn Walker, Rob Abbott, Jean E. Fox Tree, Robeson Bowmani, and Michael Minor. 2011. Cats Rule and Dogs Drool!: Classifying Stance in Online Debate. In *Proceedings of the 2Nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (Portland, Oregon) (WASSA '11)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 1–9. <http://dl.acm.org/citation.cfm?id=2107653.2107654>
- [4] Caroline Aube and Vincent Rousseau. 2005. Team goal commitment and team effectiveness: the role of task interdependence and supportive behaviors. *Group Dynamics: Theory, Research, and Practice* 9, 3 (2005), 189.
- [5] Murray R Barrick, Greg L Stewart, Mitchell J Neubert, and Michael K Mount. 1998. Relating member ability and personality to work-team processes and team effectiveness. *Journal of applied psychology* 83, 3 (1998), 377.
- [6] Murray R Barrick, Greg L Stewart, Mitchell J Neubert, and Michael K Mount. 1998. Relating member ability and personality to work-team processes and team effectiveness. *Journal of applied psychology* 83, 3 (1998), 377.
- [7] Sigal G. Barsade. 2002. The Ripple Effect: Emotional Contagion and Its Influence on Group Behavior. *Administrative Science Quarterly* (2002), 644–675. <https://doi.org/10.2307/3094912>
- [8] Daniel J Beal, Robin R Cohen, Michael J Burke, and Christy L McLendon. 2003. Cohesion and performance in groups: A meta-analytic clarification of construct relations. *Journal of applied psychology* 88, 6 (2003), 989.
- [9] Kristin J Behfar, Randall S Peterson, Elizabeth A Mannix, and William MK Trochim. 2008. The critical role of conflict resolution in teams: A close look at the links between conflict type, conflict management strategies, and team outcomes. *Journal of applied psychology* 93, 1 (2008), 170.
- [10] Suzanne Bell and Brian Marentette. 2011. Team viability for long-term and ongoing organizational teams. *Organizational Psychology Review* 1 (11 2011), 275–292. <https://doi.org/10.1177/2041386611405876>
- [11] Suzanne T Bell. 2007. Deep-level composition variables as predictors of team performance: a meta-analysis. *Journal of applied psychology* 92, 3 (2007), 595.

- [12] Suzanne T Bell and Brian J Marentette. 2011. Team viability for long-term and ongoing organizational teams. *Organizational Psychology Review* 1, 4 (2011), 275–292.
- [13] Bret H Bradley, Bennett E Postlethwaite, Anthony C Klotz, Maria R Hamdani, and Kenneth G Brown. 2012. Reaping the benefits of task conflict in teams: The critical role of team psychological safety climate. *Journal of Applied Psychology* 97, 1 (2012), 151.
- [14] Hancheng Cao, Zhilong Chen, Fengli Xu, Tao Wang, Yujian Xu, Lianglun Zhang, and Yong Li. 2020. When Your Friends Become Sellers: An Empirical Study of Social Commerce Site Beidian. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 83–94.
- [15] Sybil Carrère and John Mordechai Gottman. 1999. Predicting divorce among newlyweds from the first three minutes of a marital conflict discussion. *Family process* 38, 3 (1999), 293–301.
- [16] Wayne Cascio and Ramiro Montealegre. 2016. How Technology Is Changing Work and Organizations. *Annual Review of Organizational Psychology and Organizational Behavior* 3 (03 2016), 349–375. <https://doi.org/10.1146/annurev-orgpsych-041015-062352>
- [17] Zhilong Chen, Hancheng Cao, Fengli Xu, Mengjie Cheng, Tao Wang, and Yong Li. 2020. Understanding the Role of Intermediaries in Online Social E-commerce: An Exploratory Study of Beidian. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–24.
- [18] Jessica Nicole Cooperstein. 2017. Initial Development of a Team Viability Measure. (2017).
- [19] Glen Coppersmith, Mark Dredze, and Craig Harman. 2014. Quantifying mental health signals in Twitter. In *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic signal to clinical reality*. 51–60.
- [20] Tugrul U Daim, Anita Ha, Shawn Reutiman, Brennan Hughes, Ujjal Pathak, Wayne Bynum, and Ashok Bhatla. 2012. Exploring the communication breakdown in global virtual teams. *International Journal of Project Management* 30, 2 (2012), 199–212.
- [21] Shelley D. Dionne, Hiroki Sayama, Chanyu Hao, and Benjamin J. Benjamin. 2010. The role of leadership in shared mental model convergence and team performance improvement: An agent-based computational model. (2010), 1035–1049. <https://doi.org/10.1016/j.leaqua.2010.10.007>
- [22] Discourse. [n.d.]. This is a Civilized Place for Public Discussion. ([n. d.]). <http://bit.ly/1TM8K5x>
- [23] Maeve Duggan. 2017. Online Harassment 2017. *Pew Research Center* (2017).
- [24] Amy Edmondson. 1999. Psychological safety and learning behavior in work teams. *Administrative science quarterly* 44, 2 (1999), 350–383.
- [25] R. M. Krauss, Fussell, S. R. 1992. Coordination of knowledge in communication: Effects of speakers’ assumptions about what others know. *Journal of Personality and Social Psychology* 62 (1992).
- [26] John Mordechai Gottman and Robert Wayne Levenson. 2000. The timing of divorce: Predicting when a couple will divorce over a 14-year period. *Journal of Marriage and Family* 62, 3 (2000), 737–745.
- [27] Penelope S. Greenberg, Ralph H. Greenberg, and Yvonne L. Antonucci. 2007. Creating and sustaining trust in virtual teams. (2007), 325–333. <https://doi.org/10.1016/j.bushor.2007.02.005>
- [28] J Richard Hackman. 1980. Work redesign and motivation. *Professional Psychology* 11, 3 (1980), 445.
- [29] J Richard Hackman. 2011. *Collaborative intelligence: Using teams to solve hard problems*. Berrett-Koehler Publishers.
- [30] Pamela Hinds. 2003. Out of Sight, Out of Sync: Understanding Conflict in Distributed Teams. *Organization Science* 14 (11 2003), 615–632. <https://doi.org/10.1287/orsc.14.6.615.24872>
- [31] Pamela J Hinds and Diane E Bailey. 2003. Out of sight, out of sync: Understanding conflict in distributed teams. *Organization science* 14, 6 (2003), 615–632.
- [32] Helene V Hjalmarson and Margaretha Strandmark. 2012. Forming a learning culture to promote fracture prevention activities. *Health Education* (2012).
- [33] Malte F Jung. 2016. Coupling interactions and performance: Predicting team performance from thin slices of conflict. *ACM Transactions on Computer-Human Interaction (TOCHI)* 23, 3 (2016), 1–32.
- [34] Ben Kirman, Conor Lineham, and Shaun Lawson. 2012. Exploring mischief and mayhem in social computing or: how we learned to stop worrying and love the trolls. In *Extended Abstracts on Human Factors in Computing Systems* (Austin, Texas, USA) (*CHI ’12*). ACM, New York, NY, USA, 121–130. <https://doi.org/10.1145/2212776.2212790>
- [35] Aniket Kittur, Bongwon Suh, and Ed H. Chi. 2009. What’s in Wikipedia? Mapping Topics and Conflict Using Socially Annotated Category Structure. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) (*CHI ’09*). ACM, New York, NY, USA, 1509–1512. <https://doi.org/10.1145/1518701.1518930>
- [36] Steve WJ Kozlowski. 1998. Training and developing adaptive teams: Theory, principles, and research. (1998).
- [37] Steve WJ Kozlowski and Bradford S Bell. 2007. Team learning, development, and adaptation. In *Work group learning*. Psychology Press, 39–68.
- [38] Steve WJ Kozlowski, Stanley M Gully, Earl R Nason, and Eleanor M Smith. 1999. Developing adaptive teams: A theory of compilation and performance across levels and time. *Pulakos (Eds.), The changing nature of work performance: Implications for staffing, personnel actions, and development* 240 (1999), 292.

- [39] Travis Kriplean, Michael Toomim, Jonathan Morgan, Alan Borning, and Andrew Ko. 2012. Is This What You Meant? Promoting Listening on the Web with Reflect. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). ACM, New York, NY, USA, 1559–1568. <https://doi.org/10.1145/2207676.2208621>
- [40] Christina N Lacerenza, Shannon L Marlow, Scott I Tannenbaum, and Eduardo Salas. 2018. Team development interventions: Evidence-based approaches for improving teamwork. *American Psychologist* 73, 4 (2018), 517.
- [41] Claus W Langfred. 2007. The downside of self-management: A longitudinal study of the effects of conflict on trust, autonomy, and task interdependence in self-managing teams. *Academy of management journal* 50, 4 (2007), 885–900.
- [42] Lindsay Larson, Harrison Wojcik, Ilya Gokhman, Leslie DeChurch, Suzanne Bell, and Noshir Contractor. 2019. Team performance in space crews: Houston, we have a teamwork problem. *Acta Astronautica* 161 (2019), 108–114.
- [43] Kristi M Lewis. 2000. When leaders display emotion: How followers respond to negative emotional expression of male and female leaders. *Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior* 21, 2 (2000), 221–234.
- [44] M. Lewis. 2017. *The Undoing Project: A Friendship that Changed the World*. Penguin Books, Limited. <https://books.google.ca/books?id=-ltNvgAACAAJ>
- [45] Merce Mach, Simon Dolan, and Shay Tzafir. 2010. The differential effect of team members' trust on team performance: The mediation role of team cohesion. *Journal of Occupational and Organizational Psychology* 83, 3 (2010), 771–794.
- [46] Christina Maslach, Michael P Leiter, and Susan E Jackson. 2012. Making a significant difference with burnout interventions: Researcher and practitioner collaboration. *Journal of Organizational Behavior* 33, 2 (2012), 296–300.
- [47] John Mathieu, M Travis Maynard, Tammy Rapp, and Lucy Gilson. 2008. Team effectiveness 1997-2007: A review of recent advancements and a glimpse into the future. *Journal of management* 34, 3 (2008), 410–476.
- [48] Jessica R Mesmer-Magnus and Leslie A DeChurch. 2009. Information sharing and team performance: A meta-analysis. *Journal of applied psychology* 94, 2 (2009), 535.
- [49] Frank Mueller. 1994. Societal effect, organizational effect and globalization. *Organization Studies* 15, 3 (1994), 407–428.
- [50] Smaranda Muresan, Roberto Gonzalez-Ibanez, Debanjan Ghosh, and Nina Wacholder. 2016. Identification of nonliteral language in social media: A case study on sarcasm. *Journal of the Association for Information Science and Technology* (2016), 2725–2737. <https://doi.org/10.1002/asi.23624>
- [51] Thien Hai Nguyen, Kiyooki Shirai, and Julien Velcin. 2015. Sentiment analysis on social media for stock movement prediction. *Expert Systems with Applications* 42, 24 (2015), 9603–9611.
- [52] Vlad Niculae and Cristian Danescu-Niculescu-Mizil. 2016. Conversational markers of constructive discussions. *arXiv preprint arXiv:1604.07407* (2016).
- [53] Vlad Niculae and Cristian Danescu-Niculescu-Mizil. 2016. Conversational Markers of Constructive Discussions. *CoRR abs/1604.07407* (2016). [arXiv:1604.07407](http://arxiv.org/abs/1604.07407) <http://arxiv.org/abs/1604.07407>
- [54] Gary M Olson and Judith S Olson. 2000. Distance matters. *Human-computer interaction* 15, 2-3 (2000), 139–178.
- [55] Michelle O'Daniel and Alan H Rosenstein. 2008. Professional communication and team collaboration. In *Patient safety and quality: An evidence-based handbook for nurses*. Agency for Healthcare Research and Quality (US).
- [56] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [57] P. Nye P. V. Balakrishnan Purdy, J. M. 2000. The impact of communication media on negotiation outcomes. *International Journal of Conflict Management* 11 (2000).
- [58] Rebecca L Robinson, Reanelle Navea, and William Ickes. 2013. Predicting final course performance from students' written self-introductions: A LIWC analysis. *Journal of Language and Social Psychology* 32, 4 (2013), 469–479.
- [59] Antonette Shibani, Elizabeth Koh, Vivian Lai, and Kyong Jin Shim. 2017. Assessing the language of chat for teamwork dialogue. *Journal of Educational Technology & Society* 20, 2 (2017), 224–237.
- [60] Swapna Somasundaran and Janyce Wiebe. 2010. Recognizing stances in ideological on-line debates. In *Proceedings of the NAACL HLT 2010 workshop on computational approaches to analysis and generation of emotion in text*. Association for Computational Linguistics, 116–124.
- [61] Christine A Sprigg, Paul R Jackson, and Sharon K Parker. 2000. Production teamworking: The importance of interdependence and autonomy for employee strain and satisfaction. *Human Relations* 53, 11 (2000), 1519–1543.
- [62] Greg L Stewart. 2006. A meta-analytic review of relationships between team design features and team performance. *Journal of management* 32, 1 (2006), 29–55.
- [63] J. E. McGrath Straus, S. 1994. Does the medium matter? The interaction of task type and technology on group performance and member reactions. *Journal of Applied Psychology* 79 (1994).
- [64] Susan G Straus. 1999. Testing a typology of tasks: An empirical validation of McGrath's (1984) group task circumplex. *Small Group Research* 30, 2 (1999), 166–187.
- [65] Jim Suchan and Greg Hayzak. 2001. The communication characteristics of virtual teams: a case study. *IEEE Transactions on Professional Communication* (2001), 174–186. <https://doi.org/10.1109/47.946463>



- [66] Yla R Tausczik and James W Pennebaker. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of language and social psychology* 29, 1 (2010), 24–54.
- [67] Steven Vallas. 2003. Why Teamwork Fails: Obstacles to Workplace Change in Four Manufacturing Plants. *American Sociological Review* 68 (04 2003). <https://doi.org/10.2307/1519767>
- [68] Ben Weidmann and David J Deming. 2020. *Team Players: How Social Skills Improve Group Performance*. Technical Report. National Bureau of Economic Research.
- [69] Mark E. Whiting, Allie Blaising, Chloe Barreau, Laura Fiuza, Nik Marda, Melissa Valentine, and Michael S. Bernstein. 2019. Did It Have To End This Way? Understanding The Consistency of Team Fracture. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 209 (Nov. 2019), 23 pages. <https://doi.org/10.1145/3359311>
- [70] Mark E Whiting, Irena Gao, Michelle Xing, Junior Diarrassouba N’Godjigui, Tonya Nguyen, and Michael S Bernstein. 2020. Parallel Worlds: Repeated Initializations of the Same Team To Improve Team Viability. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 67 (May 2020), 23 pages. <https://doi.org/10.1145/3392877>
- [71] Mark E Whiting, Grant Hugh, and Michael S Bernstein. 2019. Fair Work: Crowd Work Minimum Wage with One Line of Code. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 197–206.
- [72] Justine Zhang, Jonathan P Chang, Cristian Danescu-Niculescu-Mizil, Lucas Dixon, Yiqing Hua, Nithum Thain, and Dario Taraborelli. 2018. Conversations gone awry: Detecting early signs of conversational failure. *arXiv preprint arXiv:1805.05345* (2018).

Received June 2020; accepted July 2020