

Conversational Markers of Constructive Discussions

Vlad Niculae and Cristian Danescu-Niculescu-Mizil

{vlad|cristian}@cs.cornell.edu

Cornell University

Abstract

Group discussions are essential for organizing every aspect of modern life, from faculty meetings to senate debates, from grant review panels to papal conclaves. While costly in terms of time and organization effort, group discussions are commonly seen as a way of reaching better decisions compared to solutions that do not require coordination between the individuals (e.g. voting)—through discussion, the sum becomes greater than the parts. However, this assumption is not irrefutable: anecdotal evidence of wasteful discussions abounds, and in our own experiments we find that over 30% of discussions are unproductive.

We propose a framework for analyzing conversational dynamics in order to determine whether a given task-oriented discussion is worth having or not. We exploit conversational patterns reflecting the *flow of ideas* and the *balance* between the participants, as well as their linguistic choices. We apply this framework to conversations naturally occurring in an online collaborative world exploration game developed and deployed to support this research. Using this setting, we show that linguistic cues and conversational patterns extracted from the first 20 seconds of a team discussion are predictive of whether it will be a wasteful or a productive one.

1 Introduction

Working in teams is a common strategy for decision making and problem solving, as building on effective social interaction and on the abilities of each member can enable a team to outperform lone individuals. Evidence shows that teams often perform better than individuals (Williams and Sternberg, 1988) and even have high chances of reaching

correct answers when all team members were previously wrong (Laughlin and Adamopoulos, 1980). Furthermore, team performance is not a factor of individual intelligence, but of *collective intelligence* (Woolley et al., 2010), with interpersonal interactions and emotional intelligence playing an important role (Jordan et al., 2002).

Yet, as most people can attest from experience, team interaction is not always smooth, and poor coordination can lead to unproductive meetings and wasted time. In fact, Romano Jr and Nunamaker Jr (2001) report that one third of work-related meetings in the U. S. are considered unproductive, while a 2005 Microsoft employee survey reports that 69% of meetings are ineffective.¹ As such, many grow cynical of meetings.

Computational methods with the ability to reliably recognize unproductive discussions could have an important impact on our society. Ideally, such a system could provide actionable information as a discussion progresses, indicating whether it is likely to turn out to be productive, rather than a waste of time. In this paper we focus on the conversational aspects of productive interactions and take the following steps:

- introduce a *constructiveness* framework that allows us to characterize teams where discussion enables better performance than the individuals could reach, and, conversely, teams better off not having discussed at all (Section 3);
- create a setting that is conducive to decision-making discussions, where all steps of the process (e.g., individual answers, intermediate guesses) are observable to researchers: the *StreetCrowd* game (Sections 4–5);

¹money.cnn.com/2005/03/16/technology/survey/

- develop a novel framework for conversational analysis in small group discussions, studying aspects such as the flow of ideas, conversational dynamics, and group balance (Sections 6–7).

We reveal differences in the collective decision process characteristic of productive and unproductive teams, and show that these differences are reflected in their conversational patterns. For example, the language used when new ideas are introduced and adopted encodes important discriminative cues. Measures of interactional balance and language matching (Niederhoffer and Pennebaker, 2002; Danescu-Niculescu-Mizil et al., 2011) also prove to be informative, suggesting that more balanced discussions are most productive. Our results underline the potential held by computational approaches to conversational dynamics. To encourage further work in this direction, we render our dataset of task-oriented discussions and our feature-extraction code publicly available.²

2 Related Work

Existing computational work on task-oriented group interaction is largely focused on how well the team performs. Coetzee et al. (2015) deployed and studied the impact of a chat-based team interaction platform in massive open online courses, finding that teams reach more correct answers than individuals, and that the experience is more enjoyable. One often studied experimental setting is the HCRC Map Task Corpus (Anderson et al., 1991), consisting of 128 conversations between pairs of people, where a designated one gives directions to the other. This simplified setting avoids issues like role establishment and leadership. Reitter and Moore (2007) find that successful dialogs are characterized by long-term adaptation and alignment of linguistic structures at syntactic, lexical and character level. A notable feature of this work is the *success prediction* task attempted using only the first 5 minutes of conversation. Other attempts use authority level features inspired from negotiation theory, experimental meta-features, task-specific features (Mayfield et al., 2011), and sociolinguistic spelling differences (Mayfield et al., 2012). Another research path

uses negotiation tasks from the Inspire dataset (Kersten and Zhang, 2003), a collection of 1525 online bilateral negotiations where roles are fixed (buyer and seller) and success is defined by the sale going through. Sokolova et al. (2008) use a bag-of-words model and investigate the importance of temporal aspects. Sokolova and Lapalme (2012) measure informativeness, quantified by lexical sets of degrees, scalars and comparatives.

Research on success in groups with more than two members is less common. Friedberg et al. (2012) model the grades of 27 group assignments from a class using measures of average entrainment, finding task-specific words to be a strong cue. Jung (2011) shows how the affective balance expressed in teams correlates with performance on engineering tasks, in 30 teams of up to 4 students. In a related study the balance in the first 5 minutes of an interaction is found predictive of performance (Jung et al., 2012). None of the research we are aware of controls for initial skill or potential of the team members.

In management science, network analysis reveals that certain subgraphs found in long-term, structured teams indicate better performance, as rated by senior managers (Cummings and Cross, 2003); controlled experiments show that optimal structures depend on the complexity of the task (Guetzkow and Simon, 1955; Bavelas, 1950). These studies, as well as much of the research on effective team crowdsourcing (Lasecki et al., 2012; Wang et al., 2011, *inter alia*), do not focus on linguistic and conversational factors.

3 Constructive Discussions

The first hurdle is to reliably quantify how productive group conversations are. In problem-solving, the ultimate goal is to find the correct answer, or, failing that, to come as close to it as possible. To quantify closeness to the correct answer, a *score* is often used, such that better guesses get higher scores; for example, school grades.

In contrast, our goal is to measure how productive a team’s interaction is. Scores are measures of correctness, so using them as a proxy for interaction quality is not ideal: a team of straight *A* students can manage to get an *A* on a project without exchanging ideas, while a group of *D* students getting a *B* is

²<https://vene.ro/constructive/>




Figure 1: Intuitive sketch for constructiveness. The solid green circle corresponds a team guess following a constructive discussion ($c_{\text{avg}} > 0$), the dashed green circle corresponds to the scenario of a team that outperforms its best member ($c_{\text{best}} > 0$), while the dashed red circle corresponds to a team that underperforms its worst member ($c_{\text{worst}} < 0$).

more interesting. In the latter case, the team’s improved performance is likely to come from a good discussion and an efficient exchange of complementary ideas—making the sum greater than the parts.

To capture this intuition we say a team discussion is *constructive* if it results in an improvement over the potential of the individuals. We can then quantify the degree of constructiveness c_{avg} as the improvement of the team score t over the mean of the initial scores g_i of the N individuals in the team:

$$c_{\text{avg}} = \text{score}(t) - \frac{\sum_{i=1}^N \text{score}(g_i)}{N}.$$

The higher c_{avg} is, the more the team’s answer, *after discussion*, improves upon the individuals’ average performance *before discussion*; zero constructiveness ($c_{\text{avg}} = 0$) means the team performed no better than its members did before discussing, while negative constructiveness ($c_{\text{avg}} < 0$) corresponds to non-constructive discussions.³ Figure 1 sketches the idea visually: the dark green circle corresponds to the team’s score after a constructive discussion ($c_{\text{avg}} > 0$), being above the average individual score.

Since individuals answers can sometimes vary widely, we also consider the extreme cases of teams

³From an operational perspective, a team can choose, instead of having a discussion, to aggregate individual guesses, e.g., by majority voting or averaging. Non-constructive discussions roughly correspond to cases where such an aggregate guess would actually be better than what the team discussion would accomplish.

that perform better than the best team member ($c_{\text{best}} > 0$) and worse than the worst member ($c_{\text{worst}} < 0$), where:

$$c_{\text{best}} = \text{score}(t) - \max_i \text{score}(g_i)$$

$$c_{\text{worst}} = \text{score}(t) - \min_i \text{score}(g_i).$$

One way to think of the extreme cases is to imagine a team supervisor that collects the individual answers and aggregates them, without any external information. An oracle supervisor can do no better than choosing the best answer. The discussion and interaction of teams where $c_{\text{best}} > 0$ leads to a better answer than such an oracle could achieve. (One such scenario is illustrated by the dashed light green circle in Figure 1.) Similarly, teams where $c_{\text{worst}} < 0$ waste their time completely, as simply picking one of their members’ answers at random is guaranteed to do better. (The dashed red circle in Figure 1 illustrates this scenario.)

The most important aspect of the constructiveness framework, in contrast to traditional measures of correctness or success, is that all constructiveness measures are designed to control for initial performance or potential of the team members, in order to focus on the effect of the discussion.⁴

In settings of importance, the true answer is not known a priori, and this constructiveness cannot be calculated directly. We therefore seek out to model constructiveness using observable conversational and linguistic correlates (Sections 6–7). To develop such a model, we design a large-scale experimental setting where the true answer is available to researchers, but unknown by the players (Section 4).

4 Experimental setting

4.1 StreetCrowd

In order to study the constructiveness of task-oriented group discussion, we need a setting that is conducive to decision-making discussions, where all steps of the process (individual answers, intermediate guesses, group discussions and decisions) are observable. Furthermore, to study at scale, we need to

⁴Due to its relative nature, constructiveness also accounts for variation in task difficulty in most scenarios. For example, in terms of c_{worst} , when a team cannot even match its worst performing member, this is a sign of poor team interaction even if the task is particularly challenging.

find a class of complex tasks with known solutions that can be automatically generated, but that cannot be easily solved by simply querying search engines.

With these constraints in mind, we built *StreetCrowd*, an online multi-player world exploration game.⁵ *StreetCrowd* is played in teams of at least two players and is built around a geographic puzzle: determining your location based on first-person images from the ground level.⁶ Each location generates a new puzzle.

Solo phase. Each player has 3 minutes to navigate the surroundings, explore, and try to find clues. This happens independently and without communicating. At the end, the player is asked to make a guess by placing a marker on the world’s map, and is prompted for an explanation and for a confidence level. The answer is not yet revealed.

Team phase. The team must then decide on a single, common guess. To accomplish this, all teammates are placed in a chatroom and are provided with a map and a shared marker. Any player can move the marker at any point during the discussion. The game ends when all players agree on the answer, or when the time limit is reached. An example discussion is given in Figure 4.

Guesses are scored according to their distance to the true location using the spherical law of cosines:

$$\text{score}(\text{guess}, \text{true}) = -Rd(\text{guess}, \text{true})$$

where d is the arc distance on a sphere, and R denotes the radius of the earth, assumed spherical. The score is given by the negative distance in kilometers, such that higher means better. To motivate players and emphasize collaboration, the main *StreetCrowd* page displays a leaderboard consisting of the best team players.

The key aspects of the *StreetCrowd* design are:

- The puzzles are complex and can be generated automatically in large numbers;
- The true answers are known to researchers, but hard to obtain without solving the puzzle, allowing for objective evaluation of both individual and group performance;

⁵<http://streetcrowd.us/start> (the experiment was approved by the IRB).

⁶We embed Google Street View data.

- Scoring is continuous rather than discrete, allowing us to quantify degrees of improvement and capture incremental effects;
- Each teammate has a different *solo phase* experience and background knowledge, making it possible for the group discussion to shed light on new ideas;
- The puzzles are engaging and naturally conducive to collaboration, avoiding the use of monetary incentives that can bias behavior.

4.2 Preprocessing

In the first 8 months, over 1400 distinct players participated in over 2800 *StreetCrowd* games. We tokenize and part-of-speech tag the conversations.⁷ Before analysis, due to the public nature of the game, we perform several filtering and quality check steps.

Discarding trivial games. We remove all games that the developers took part in. We filter games where the team fails to provide a guess, where fewer than two team members engage in the team chat, and puzzles with insufficient samples.

Preventing and detecting cheating. The *StreetCrowd* tutorial asks players to avoid using external resources to look up clues and get an unfair advantage. To prevent cheating, we detect and block chat messages that link to websites, and we employ cookies and user accounts to prevent people from playing the same puzzle multiple times. To identify games that slip through this net, we flag cases where the team, or any individual player, guesses within 10 km of the correct answer, and leaves the window while playing. We further remove a small set of games where the players confess to cheating in the chat.

After filtering, our dataset consists of 1450 games on 70 different puzzles, with an average of 3.9 games per unique player, and 12.1 messages and 64.5 words in an average conversation.

5 Constructiveness in StreetCrowd

We find that, indeed, most of the games are constructive. There are, however, 32% non-constructive

⁷We use the *TweetNLP* toolkit (Owoputi et al., 2013) with a tag set developed for Twitter data. Manual examination reveals this approach to be well suited for online chat data.





Figure 2: Distribution of team constructiveness.


games ($c_{avg} < 0$); this reflects very closely the survey by Romano Jr and Nunamaker Jr (2001). Interestingly, in 36% of games, the team arrives at a better answer than any of the individual guesses ($c_{best} > 0$). The flip side is also remarkably common, with 17% of teams performing even worse than the worst individual ($c_{worst} < 0$). The distribution of constructiveness is shown in Figure 2: the fat tails indicate that cases of large improvements and large deterioration are not uncommon.

Collective decision process. Due to the full instrumentation of the game interface, we can investigate how constructiveness emerges out of the team’s interaction. The team’s intermediate guesses during discussion confirm that a meaningful process leads to the final team decision: guesses get closer and closer to the final submitted guess (Figure 3a); in other words, the team *converges* to their final guess.

Notably, when considering how *correct* the intermediate guesses are, we notice an important difference between the way constructive and non-constructive teams converge to their final guess (Figure 3b). During their collaborative decision process, constructive teams make guesses that get closer and closer to the correct answer; in contrast, non-constructive teams make guesses that take them farther from the correct answer. This observation has two important consequences. First, it shows that the two types of teams behave differently throughout, suggesting we could potentially detect non-constructive discussions early on, using interaction patterns. Second, it emphasizes the potential practical value of such a task: stopping a non-constructive team early could lead to a better answer than if they would carry on.



(a) Distance between the last three intermediate guesses and the final guess.



(b) Score of the first three intermediate guesses; the mean score of the initial individual guesses and the score of the final team guess are shown for reference.

Figure 3: Intermediate guesses offer a glimpse at the decision process: (a) guesses converge rather than zig-zag (b) in constructive games, guesses get incrementally better than the mean individual score, while in non-constructive games, they get worse. (Games with ≥ 3 intermediate guesses.)

6 Conversation analysis

The process of team convergence revealed in the previous section suggests a relation between the interaction leading to the final group decision and the relative quality of the outcome. In this section, we develop a conversation analysis framework aimed at characterizing this relation. This framework relies on conversational patterns and linguistic features, while steering away from lexicalized cues that might not generalize well beyond our experimental setting. To enable reproducibility, we make available the feature extraction code and the hand-crafted resources on which it relies.⁸

⁸<https://vene.ro/constructive/>

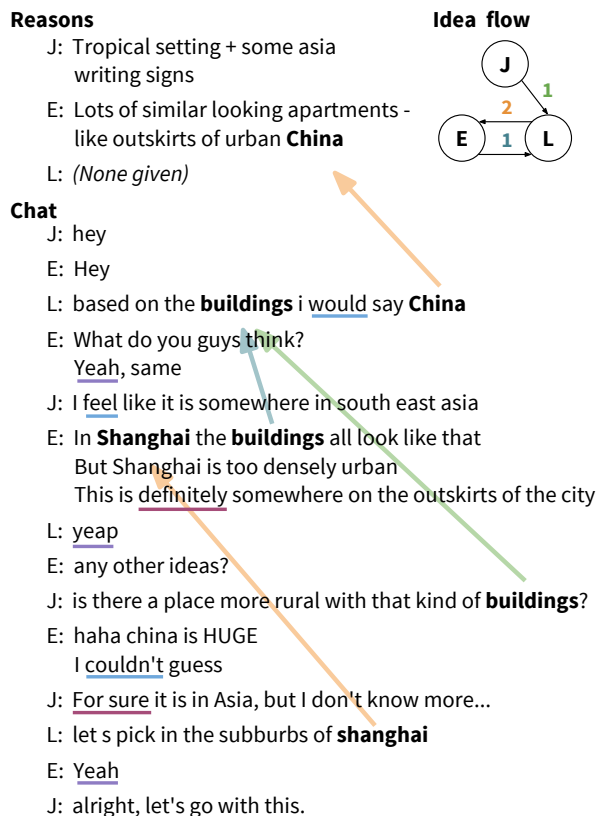


Figure 4: Example (constructive) conversation and the corresponding flow of ideas. Idea mentions are in bold, and relevant word classes are underlined. Arrow colors map to introducer–adopter pairs, matching the edges in the top-right graph.

6.1 Idea flow

Task-oriented discussions are the primary way of exchanging ideas and opinions between the group members; some are quickly discarded while others prove useful to the final guess. The arrows in Figure 4 show how ideas are introduced and discussed in that example conversation. We attempt to capture the shape in which the ideas flow in the discussion. In particular, we are interested in how many ideas are discussed, how widely they are adopted, who tends to introduce them, and how.

We consider as candidate ideas all nouns, proper nouns, adjectives and verbs that are not stopwords. As soon as a candidate idea introduced by a player is *adopted* by another, we count it. Henceforth, we'll refer to such adopted ideas simply as *ideas*. In general chat domains, state-of-the-art models of conver-

sation structure use unsupervised probabilistic models (Ritter et al., 2010; Elsner and Charniak, 2010). Since *StreetCrowd* conversations are short and focused, the *adoption* filter is sufficient to accurately capture what ideas are being discussed; a manual examination of the ideas reveals almost exclusively place names and words such as *flag*, *sign*, *road*—highly relevant clues in the context of *StreetCrowd*. In Figure 4, three ideas are adopted: *China*, *buildings* and *Shanghai*. The only idea adopted by all players is *buildings*, a good signal that this was the most important clue. A notable limitation is that this approach cannot capture the connections between *Shanghai* and *China*, or *buildings* and *apartments*. Further work is needed to robustly capture such variations in idea flow, as they could reveal trajectories (discussion getting more specific or more vague) or lexical choice disagreement.

Balance in idea contributions between the team members is a good indicator of productive discussions. In particular, in the best teams (the ones that outperform the best player, i.e., $c_{\text{best}} > 0$) the most idea-prolific player introduces fewer ideas, on average, than in the rest of the games (Figure 5a, $p = 0.01$).⁹ In Figure 4, *E* is the most prolific player and only introduces two ideas. To further capture the balance in contribution between the team members, we use the entropy of the number ideas introduced by each player. We also count the number of ideas adopted unanimously as an indicator of convergence in the conversation.

In terms of the overall number of ideas discussed, both the best teams (the ones that outperform the best player) and the worst teams (the ones that perform worse than the worst player) discuss fewer ideas than the rest (Figure 5b, $p = 0.006$). Indeed, an ideal interaction would avoid distracting ideas, but in teams with communication breakdowns, members might fail to adequately discuss the ideas that led them to their individual guesses.

The language used to introduce new ideas can indicate confidence or hesitation; in Figure 4, a hedge (*would*) is used when introducing the *buildings* cue. We find that, in teams that outperform the best player, ideas are less likely to be accom-

⁹All p-values reported in this section are based on one-sided Mann-Whitney rank statistical significance tests.

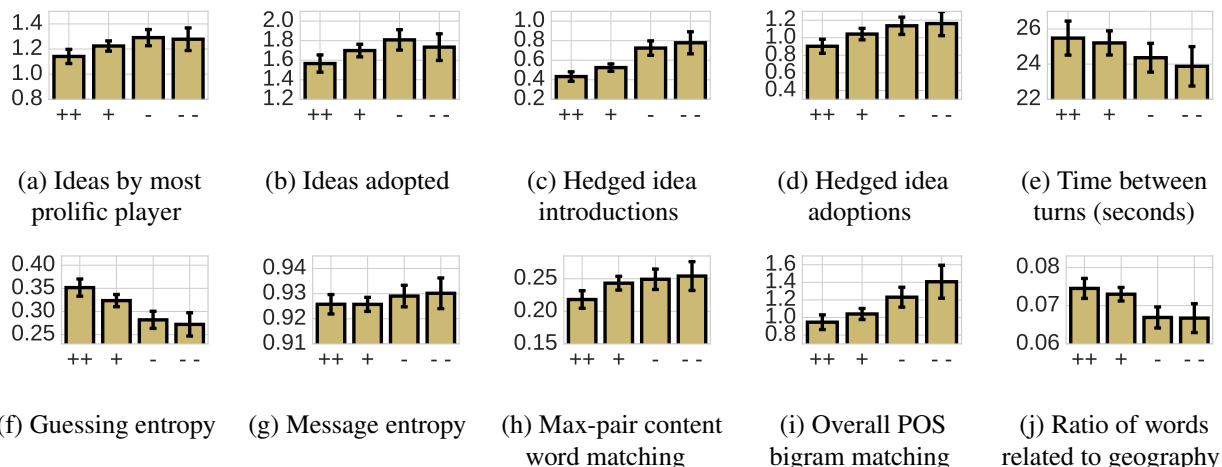


Figure 5: Averages for some of the predictive features, based on idea flow (a-d), balance (e-i), and lexicons (j). Error bars denote standard errors. Legend: (++): teams that do better than their best member ($N = 525$), (+): constructive ($N = 986$), (-): non-constructive ($N = 464$), (--): worse than worst member ($N = 248$).

panied by hedge words when introduced (Figure 5c, $p < 10^{-4}$), showing less hesitation. Furthermore, the level of confidence used when players adopt others’ ideas is also informative (Figure 5d). Interestingly, overall occurrences of certainty and hedge words (detailed in Section 6.3) are not predictive, suggesting that ideas are good selectors for important discussion segments.

6.2 Interaction dynamics

Balance. Interpersonal balance has been shown to be predictive of team performance (Jung, 2011; Jung et al., 2012) and, similarly, forms of linguistic balance have been shown to characterize stable relationships (Niculae et al., 2015). Here we focus on balance in contributions to the discussion and the decision process. In search of measures applicable to teams of arbitrary sizes, we use binary indicators of whether all players participate in the discussion and in moving the marker, as well as whether at least two players move the marker. To measure team balance with respect to continuous user-level features, we use the entropy of these features:

$$\text{balance}(S) = - \sum_{\bar{s} \in S} \bar{s} \log_{|S|} \bar{s},$$

where, for a given feature, S is the set of its values for each user, normalized to sum to 1. For instance, the chat message entropy is 1 if everybody chats equally, and decreases toward 0 as one or more

players dominate. We use the entropy of the number of messages, words per message, and number of intermediate guesses. In teams that outperform the best player, users take turns controlling the marker more uniformly (Figure 5f, $p = 0.006$), adding further evidence that well-balanced teams perform best. **Language matching.** We investigate matching at stopword, content word, and POS tag bigram level: the stopword matching at a turn is given by the number of stopwords from the earlier message repeated in the reply, divided by the total number of distinct stopwords to choose from; similarly for the rest. We micro-average over the conversation:

$$\text{match} = \frac{\sum_{(\text{msg}, \text{reply}) \in \text{Turns}} |\text{msg} \cap \text{reply}|}{\sum_{(\text{msg}, \text{reply}) \in \text{Turns}} |\text{msg}|}.$$

We also micro-average at the player-pair level, and use the maximum pair value as a feature. This gives an indication of how cohesive the closest pair is, which can be a sign of the level of power imbalance between the two (Danescu-Niculescu-Mizil et al., 2012). Figure 5h shows that in teams that outperform the best individual the most cohesive pair matches fewer content words ($p = 0.023$). Overall matching is also significant, notably in terms of part-of-speech bigrams; in teams that outperform the best individual there is less overall matching (Figure 5i, $p = 0.007$). These results suggest that in constructive teams the relationships between the members are less subordinate.

Agreement and confidence. We capture the amount of agreement and disagreement using high-precision keywords and filters validated on a subset of the data. (For instance, the word *sure* marks agreement if found at the beginning of a message, but not otherwise.) In Figure 4, agreement signals are underlined with purple; the team exhibits no disagreement.

The relative position of successive guesses made can also indicate whether the team is refining a guess or contradicting each other. We measure the median distance between intermediate guesses, as well as between guesses made by different players; in constructive teams, the jumps between different player guesses are smaller ($p < 10^{-16}$).

Before the discussion starts, players are asked to self-evaluate their confidence in their individual guesses. Constructive teams have more confident members on average ($p < 10^{-5}$).

6.3 Other linguistic features

Length and variation. We measure the average number of words per message, the total number of words used to express the *solo phase* reasons, and the overall type/token ratio of the conversation. We also measure responsiveness in terms of the mean time between turns and the total number of turns.

Psycholinguistic lexicons. We use hand-crafted lexicons inspired from LIWC (Tausczik and Pennebaker, 2010) to capture certainty and pronoun use. For example, the conversation in Figure 4 has two confident phrases, underlined in red. We also use a custom hedging lexicon adapted from Hyland (2005) for conversational data; hedging words are underlined in blue in Figure 4. To estimate how grounded the conversation is, we measure the average *concreteness* of all content nouns, adjectives, adverbs and verbs, using scalar word and bigram ratings from Brysbaert et al. (2014).¹⁰ Concreteness reflects the degree to which a word denotes something perceptible, as opposed to ideas and concepts. Words like *soil* and *coconut* are highly concrete, while words like *trust* have low concreteness.

Game-specific words. We put together a lexicon of geography terms and place names, to capture task-

¹⁰We scale the ratings to lie in $[0, 1]$. We extrapolate to out-of-vocabulary words by regressing on dependency-based word embeddings (Levy and Goldberg, 2014); this approach is highly accurate (median absolute error of about 0.1).

Features	Full conversation			First 20s		
	(++)	(+)	(--)	(++)	(+)	(--)
Baseline	.51	.52	.55	.52	.50	.54
Linguistic	.54	.52	.50	.50	.51	.50
Interaction	.55 [†]	.56 [†]	.53	.55 [†]	.57 [*]	.56
POS	.55 [†]	.59 [*]	.55	.54	.54	.53
All	.56 [*]	.60 [*]	.56	.56 [*]	.57 [*]	.57 [†]

Table 1: Cross-validation AUC scores. Significantly better than chance scores after 5000 permutations denoted with $*$ ($p < 0.05$) and $†$ ($p < 0.1$).

specific discussion. We use a small set of words specific to the *StreetCrowd* interface, such as *map*, *marker*, and *game*, to capture phatic communication. Figure 5j shows that constructive teams tend to use more geography terms ($p = 0.008$), possibly because of more on-topic discussion and a more focused vocabulary.

Part-of-speech patterns. We use n-grams of coarse part-of-speech tags as a general way of capturing common syntactic patterns.

7 Predicting constructiveness

7.1 Experimental setup

So far we have characterized the relation between a team’s interaction patterns and its level of productivity. This opens the door towards recognizing constructive and non-constructive interactions in realistic settings where the true answer is not known. Ideally, such an automatic system could prompt unproductive teams to reconsider their approach, or to aggregate their individual answers instead. With *early detection*, non-constructive discussions could be stopped or steered on the right track. In order to assess the feasibility of such a challenging task and to compare the predictive power of our features, we consider three classification objectives:

(++): Team outperforms its best member ($c_{\text{best}} > 0$)?

(+): Team is constructive ($c_{\text{avg}} > 0$)?

(--): Team underperforms its worst member ($c_{\text{worst}} < 0$)?

To investigate *early detection*, we evaluate the classification performance when using data from only the first 20 seconds of the team’s interaction.¹¹

¹¹Measured from the first chat message or guess. For this evaluation, we remove teams where the first 20 seconds contain over 75% of the interaction, to avoid distorting the results with

Since all three objectives are imbalanced (Figure 2), we use the area under the ROC curve (AUC) as the performance metric, and we use logistic regression models. We perform 20 iterations of puzzle-aware shuffled train-validation splitting, followed by 5000 iterations on the best models, to estimate variance. This ensures that the models don't learn to overfit puzzle-specific signals. The combined model uses weighted model averaging. We use grid search for regularization parameters, feature extraction parameters, and combination weights.

7.2 Discussion of the results (Table 1)

We compare to a baseline consisting of the team size, average number of messages per player, and conversation duration. For comparison, a bag-of-words classifier does no better than chance and is on par with the baseline. We refer to idea flow and interaction dynamics features (Section 6.2) as *Interaction*, and to linguistic and lexical features (Section 6.3) as *Linguistic*. The combination model including baseline, interaction, linguistic and part-of-speech n-gram features, is consistently the best and significantly outperforms random guessing (AUC .50) in nearly all settings. While overall scores are modest, the results confirm that our conversational analysis framework has predictive power, and that the high-stakes task of *early prediction* is feasible. The language used when introducing and adopting ideas, together with balance and language matching features, are selected in nearly all settings. The least represented class (-) has the highest variance in prediction, suggesting that more data collection is needed to successfully capture extreme cases. Useful POS patterns capture the amount of proper nouns and their contexts: proper nouns at the end of messages are indicative of constructiveness, while proper nouns followed by verbs are a negative feature. (The constructive discussion shown in Figure 4 has most proper nouns at the end of messages.)

A manual error analysis of the false positives and false negatives where our best model is most confident points to games with very short conversations and spelling mistakes, confirming that the noisy data problem causes learning and modeling difficulties.

teams who make their decision early, but take longer to submit. The 20 second threshold was chosen as a trade-off in terms of how much interaction it covers in the games.

8 Conclusions and Future Work

We developed a framework based on conversational dynamics in order to distinguish between productive and unproductive task-oriented discussions. By applying it to an online collaborative game we designed for this study, we reveal new interactions with conversational patterns. Constructive teams are generally well-balanced on multiple aspects, with team-members participating equally in proposing ideas and making guesses and showing little asymmetry in language matching. Also, the flow of ideas between teammates marks predictive linguistic cues, with the most constructive teams using fewer hedges when introducing and adopting ideas.

We show that such cues have predictive power even when extracted from the first 20 seconds of the conversations. In future work, improved classifiers could lead to a system that can intervene in non-constructive discussions early on, steering them on track and preventing wasted time.

Further improving classification performance on such a difficult task will hinge on better conversation processing tools, adequate for the domain and robust to the informal language style. In particular, we plan to develop and evaluate models for idea flow and (dis)agreement, using more advanced features (e.g., from dependency relations and knowledge graphs).

The *StreetCrowd* game is continuously accumulating more data, enabling further development on conversation analysis. Our full control over the game permits manipulation and intervention experiments that can further advance research on teamwork. In future work, we envision applying our framework to settings where teamwork takes place online, such as open-source software development, Wikipedia editing, or massive open online courses.

Acknowledgements We are particularly grateful to Bob West for the poolside chat that inspired the design of the game, to Daniel Garay, Jinjing Liang and Neil Parker for participating in its development, and to the numerous passionate players. We are also grateful to Natalya Bazarova, David Byrne, Mala Gaonkar, Lillian Lee, Sendhil Mullainathan, Andreas Veit, Connie Yuan, Justine Zhang and the anonymous reviewers for their insightful suggestions. This work was supported in part by a Google Faculty Research Award.

References

- Anne H Anderson, Miles Bader, Ellen Gurman Bard, Elizabeth Boyle, Gwyneth Doherty, Simon Garrod, Stephen Isard, Jacqueline Kowtko, Jan McAllister, and Jim Miller. 1991. The HCRC map task corpus. *Language and speech*, 34(4):351–366.
- Alex Bavelas. 1950. Communication patterns in task-oriented groups. *Journal of the Acoustical Society of America*.
- Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. 2014. Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3):904–911.
- Derrick Coetzee, Seongtaek Lim, Armando Fox, Bjorn Hartmann, and Marti A Hearst. 2015. Structuring interactions for large-scale synchronous peer learning. In *Proceedings of CSCW*.
- Jonathon N Cummings and Rob Cross. 2003. Structural properties of work groups and their consequences for performance. *Social Networks*, 25(3):197–210.
- Cristian Danescu-Niculescu-Mizil, Michael Gamon, and Susan Dumais. 2011. Mark my words! Linguistic style accommodation in social media. In *Proceedings of WWW*.
- Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. 2012. Echoes of power: Language effects and power differences in social interaction. In *Proceedings of WWW*.
- Micha Elsner and Eugene Charniak. 2010. Disentangling chat. *Computational Linguistics*, 36(3):389–409.
- Heather Friedberg, Diane Litman, and Susannah BF Paletz. 2012. Lexical entrainment and success in student engineering groups. In *Proceedings of the Spoken Language Technology Workshop*.
- Harold Guetzkow and Herbert A Simon. 1955. The impact of certain communication nets upon organization and performance in task-oriented groups. *Management Science*, 1(3-4):233–250.
- Ken Hyland. 2005. *Metadiscourse: Exploring interaction in writing*. Continuum.
- Peter J Jordan, Neal M Ashkanasy, Charmine EJ Härtel, and Gregory S Hooper. 2002. Workgroup emotional intelligence: Scale development and relationship to team process effectiveness and goal focus. *Human Resource Management Review*, 12(2):195–214.
- Malte F Jung, Jan Chong, and Larry Leifer. 2012. Group hedonic balance and pair programming performance: affective interaction dynamics as indicators of performance. In *Proceedings of SIGCHI*.
- Malte F Jung. 2011. *Engineering team performance and emotion: Affective interaction dynamics as indicators of design team performance*. Ph.D. thesis, Stanford University.
- Gregory E Kersten and Grant Zhang. 2003. Mining in-spire data for the determinants of successful internet negotiations. *Central European Journal of Operations Research*, 11(3).
- Walter S Lasecki, Christopher D Miller, Adam Sadilek, Andrew Abumoussa, Donato Borrello, Raja Kushalnagar, and Jeffrey P Bigham. 2012. Real-time captioning by groups of non-experts. In *Proceedings of UIST*.
- Patrick R Laughlin and John Adamopoulos. 1980. Social combination processes and individual learning for six-person cooperative groups on an intellectual task. *Journal of Personality and Social Psychology*, 38(6):941.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proceedings of ACL*.
- Elijah Mayfield, Michael Garbus, David Adamson, and Carolyn Penstein Rosé. 2011. Data-driven interaction patterns: Authority and information sharing in dialogue. In *Proceedings of AAAI Fall Symposium on Building Common Ground with Intelligent Agents*.
- Elijah Mayfield, David Adamson, Alexander Rudnicky, and Carolyn Penstein Rosé. 2012. Computational representation of discourse practices across populations in task-based dialogue. In *Proceedings of ICIC*.
- Vlad Niculae, Srijan Kumar, Jordan Boyd-Graber, and Cristian Danescu-Niculescu-Mizil. 2015. Linguistic harbingers of betrayal: A case study on an online strategy game. In *Proceedings of ACL*.
- Kate G Niederhoffer and James W Pennebaker. 2002. Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, 21(4):337–360.
- Olutobi Owoputi, Brendan O’Connor, Chris Dyer, Kevin Gimpel, Nathan Schneider, and Noah A Smith. 2013. Improved part-of-speech tagging for online conversational text with word clusters. In *Proceedings of NAACL*.
- David Reitter and Johanna D Moore. 2007. Predicting success in dialogue. In *ACL*, volume 45, page 808.
- Alan Ritter, Colin Cherry, Bill Dolan, et al. 2010. Unsupervised modeling of Twitter conversations. In *Proceedings of NAACL*.
- Nicholas C Romano Jr and Jay F Nunamaker Jr. 2001. Meeting analysis: Findings from research and practice. In *Proceedings of HICSS*.
- Marina Sokolova and Guy Lapalme. 2012. How much do we say? Using informativeness of negotiation text records for early prediction of negotiation outcomes. *Group Decision and Negotiation*, 21(3):363–379.

- Marina Sokolova, Vivi Nastase, and Stan Szpakowicz. 2008. The telling tail: Signals of success in electronic negotiation texts.
- 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):24.
- Hao-Chuan Wang, Susan R Fussell, and Dan Cosley. 2011. From diversity to creativity: Stimulating group brainstorming with cultural differences and conversationally-retrieved pictures. In *Proceedings of CSCW*.
- Wendy M Williams and Robert J Sternberg. 1988. Group intelligence: Why some groups are better than others. *Intelligence*, 12(4):351–377.
- Anita Williams Woolley, Christopher F Chabris, Alex Pentland, Nada Hashmi, and Thomas W Malone. 2010. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004):686–688.