

# The Pragmatics of Indirect Commands in Collaborative Discourse

Matthew Lamm\*  
Stanford Linguistics  
Stanford NLP Group  
mlamm@stanford.edu

Mihail Eric\*  
Stanford Computer Science  
Stanford NLP Group  
meric@cs.stanford.edu

## Abstract

Today’s artificial assistants are typically prompted to perform tasks through direct, imperative commands such as *Set a timer* or *Pick up the box*. However, to progress toward more natural exchanges between humans and these assistants, it is important to understand the way non-imperative utterances can indirectly elicit action of an addressee. In this paper, we investigate command types in the setting of a grounded, collaborative game. We focus on a less understood family of utterances for eliciting agent action, locatives like *The chair is in the other room*, and demonstrate how these utterances indirectly command in specific game state contexts. Our work shows that models with domain-specific grounding can effectively realize the pragmatic reasoning that is necessary for more robust natural language interaction.

## 1 Introduction

A major goal of computational linguistics research is to enable organic, language-mediated interaction between humans and artificial agents. In a common scenario of such interaction, a human issues a command in the imperative mood—e.g. *Put that there* or *Pick up the box*—and a robot acts in turn (Bolt, 1980; Tellex et al., 2011; Walter et al., 2015). While this utterance-action paradigm presents its own set of challenges (Tellex et al., 2012), it greatly simplifies the diversity of ways in which natural language can be used to elicit action of an agent, be it human or artificial (Clark, 1996; Portner, 2007; Kaufmann and Schwager, 2009; Condoravdi and Lauer, 2012; Kaufmann, 2016). Most clause types, even vanilla declaratives, instantiate as performative requests in certain contexts (Austin, 1975; Searle, 1989; Perrault and Allen, 1980).

In this work, we employ machine learning to study the use of performative commands in the Cards corpus, a set of transcripts from a web-based game that is designed to elicit a high degree of linguistic and strategic collaboration (Djalali et al., 2011, 2012; Potts, 2012). For example, players are tasked with navigating a maze-like gameboard in search of six cards of the same suit, but since a player can hold at most three cards at a time, they must coordinate their efforts to win the game.

We focus on a subclass of performative commands that are ubiquitous in the Cards corpus: Non-agentive declaratives about the locations of objects, e.g. “The five of hearts is in the top left corner,” hereafter referred to as *locatives*. Despite that their semantics makes no reference to either an agent or an action—thus distinguishing them from conventional imperatives (Condoravdi and Lauer, 2012)—locatives can be interpreted as commands when embedded in particular discourse contexts. In the Cards game, it is frequently the case that an addressee will respond to such an utterance by fetching the card mentioned.

Following work on the context-driven interpretation of declaratives as questions (Beun, 2000), we hypothesize that the illocutionary effect of a locative utterance is a function of contextual features that variably constrain the actions of discourse participants. To test this idea, we identify a set of 94 locative utterances in the Cards transcripts that we deem to be truly ambiguous, out of context, between

---

\* Authors contributed equally

informative and command readings. We then annotate their respective transcripts for a simplified representation of the tabular common ground model of Malamud and Stephenson (2015). Here, we identify the common ground with the state of a game as reflected by the utterances made by both players up to a specific point in time. Finally, we train machine learning classifiers on features of the common ground to predict whether or not the addressee will act in response to the utterance in question. Through these experiments we discover a few very powerful contextual features that predict when a locative utterance will be interpreted as a command.

## 2 Related Work

The subject of indirect commands, of which the locative utterances we study are an example, has been extensively analyzed in terms of speech act and decision theory (Austin, 1975; Clark, 1979; Perrault and Allen, 1980; Allen and Perrault, 1980; Searle, 1989). In Portner’s (2007) formal model, imperatives are utterances whose conventional effect updates an abstract “to-do list” of an addressee. More recent debate has asked whether this effect is in fact built into the semantics of imperatives, or if their directive force is resolved by pragmatic reasoning in context (Condoravdi and Lauer, 2012; Kaufmann, 2016).

The present work synthesizes these intuitions from the theory of commands with recent computational work on natural language pragmatics (Vogel et al., 2013, 2014; Degen et al., 2013) and collaborative dialogue (Chai et al., 2014). We are particularly influenced by previous work demonstrating the complexity of pragmatic phenomena in the Cards corpus (Djalali et al., 2011, 2012; Potts, 2012).

## 3 The Cards corpus

The Cards corpus is a set of 1,266 transcripts from a two-player, collaborative, web-based game. The Cards corpus is well-suited to studying the pragmatics of commands because it records both utterances made as well as the actions taken during the course of a game.

At the start of the game, player locations are randomly initialized on a two-dimensional, grid-style game board. Cards from a conventional 52-card deck are scattered randomly throughout the board. Players are prompted with the following task:

Gather six consecutive cards of a particular suit (decide which suit together), or determine that this is impossible. Each of you can hold only three cards at a time, so youll have to coordinate your efforts. You can talk all you want, but you can make only a limited number of moves.

In addition to the fact that players can only hold three cards at a time, the game is further-constrained in ways that stimulate highly collaborative talk. In particular, while players can see their own location, they cannot see the locations of their partners and so must inquire about them. Players can only see cards within a small neighborhood around their respective locations, and so must explore the board to find relevant cards. Moreover, while some walls are visible, others are invisible and so lead to surprise perturbations in the course of exploring the gameboard.

## 4 Command types in the Cards Corpus

Commands in the Cards corpus can be coarsely divided into ones which make reference to an action with a second person agent, and those which do not.

The first of these categories is comprised of imperatives and a variety so-called performative commands: Utterances which act as commands in context but whose clause type is not conventionally associated with the effect of commanding (Clark, 1979; Searle, 1989; Wierzbicka, 1991). For example, with respect to picking up cards:

pick it up!  
pick up the 9  
or hell, grab the 234 of D  
ok when you get here pick up the 8H  
i think you should pick up the 3h, 5h, and 8h  
so if you can find 5S,6S,7S that would be great

With respect to dropping (or not dropping) cards:

drop the 2  
keep the 3  
no dont drop it.)  
ok drop the 7 i guess  
get rid of 6d. i found 7h  
so if you come and drop the 8h and pick up the 6h we are good

With respect to conversational actions (some of these utterances are shortened for clarity):

tell me where it is  
talk to me dude [...]  
tell me if you see 5 or 6  
don't just say "a lot of cards" [...]  
awesome let me know once you have it.

Imperatives and performatives that mention agents contrast with the lesser understood subclass of performative commands that are the focus of this work. Utterances like "The five of hearts is in the top corner" do not even encode an action with respect to the object mentioned, let alone an agent, but can nevertheless be used to elicit action of an addressee in certain contexts.

As a motivating example, consider the following exchange between two players describing their respective hands:

P1: 3h, 4h and ks  
P2: i have a queen of diamonds and ace of club  
P2: we have a mess lol

Despite Player 2's concerns, a strategy emerges shortly thereafter when Player 1 finds an additional hearts card:

P1: i have 3h,4h,6h  
P2: ok so we need to collect hearts then

At this point in the transcript, all that has been committed to the common ground is that Player 1 has a full hand of three proximal hearts cards that could be relevant to a winning strategy, and Player 2 has two non-hearts cards. This is the very next utterance in the exchange:

P1: there is a 5h in the very top left corner

Player 2 is seen immediately hereafter to navigate to the top left corner, pick up the five of hearts, and confirm:

P2: ok i got it :)

In this exchange, Player 2 appears to understand not only that the five of hearts is relevant to the winning strategy of six consecutive hearts, but also that it makes more sense for her to act on information about its location than it does for Player 1 to do so.

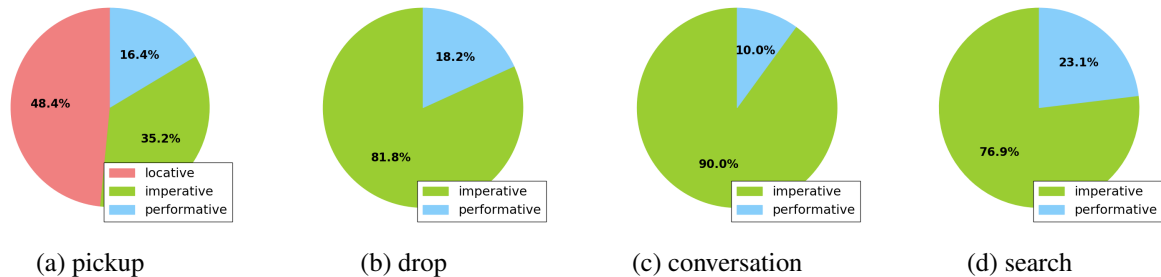


Figure 1: For each action, the distribution of command types

This discourse encapsulates the collaborative reasoning pattern described by Perrault and Allen (1980). The speaker assumes that the addressee is a cooperative agent. Thus, the addressee shares in the goal of attaining a winning game state and will act in a way so as to realize that goal. Recognizing the fact that the speaker would have to drop cards relevant to the goal to pick up the card at issue, a cooperative addressee will infer that she *should* act by picking it up instead. In this way, locative utterances can be indirectly used as commands.

The distribution of command types for a subset of actions (*pickup*, *drop*, *conversation*, and *search*) is displayed in Figure 1. As depicted, for the majority of actions compelled by a speaker and taken by an addressee, the imperative is the predominant command strategy, followed by non-locative performatives. However locative commands appear to be the dominant strategy for eliciting card pickups in the corpus, constituting nearly half of all such commands observed.

This pattern demonstrates that for certain kinds of actions, it is quite natural to use the least direct, most context-dependent command strategy to elicit action of an addressee.

## 5 Common Ground Effects on the Directive Force of Locative Utterances

We seek to understand how the discourse context of a game can influence the interpretation of locative utterances as commands. We therefore construct a binary classification task whereby we test how the role of a locative utterance can be resolved in context, evidenced by the actions that are taken as follow-ups to the utterance. In our task, one label denotes addressee follow-up in the form of acting to pick up the card in question, signaling her intention to act, or asking a clarifying question about its whereabouts. The second label denotes that either the speaker acts on their own utterance or neither agent does.

### 5.1 Annotation Details

Using a random sample of 200 transcripts from the corpus, we identify instances where a locative utterance is made and we annotate the common ground up to this utterance. This yields 55 distinct transcripts constituting 94 utterances with this particular phenomenon.

Our common ground annotations include the following information about the game state as indicated by players' utterances: cards in the players hands, player location, known information about the existence or location of cards, strategic statements made by players about needed cards, and whether a player is able to act with respect to an at-issue card.

### 5.2 Experiments

Our aim in devising this task is to investigate connection between common ground knowledge and the illocutionary effects of locative utterances. We therefore train a standard logistic regression classifier and experiment with a few carefully designed features that encode constraints on player action, and which should hypothetically trigger the interpretation of locative utterances as indirect commands. We experiment with the following features:

Model	$F_1$
<i>Random</i>	23.5
<i>Bigram</i>	58.9
Edit Distance	62.5
Explicit Goal	76.2
Full Hand	<b>82.3</b>
Explicit Goal + Full Hand	77.7

Table 1:  $F_1$  performance as reported on the test set. Note our baselines are italicized.

- **Edit Distance:** We use the minimal number of edits for an optimal solution as a feature. Given the cards in the players’ hands at a given point in the game, we can define an optimal solution based on the number of edits that must be made to the hands to achieve that optimal solution. An edit is defined as either picking up or dropping a card, and each such action has a cost of 1. An optimal solution is defined as the one that requires the minimal number of edits given the current hands. For example, if player 1 has a 2H, 3H, and 4H and player 2 has a 6H and a 7H, the optimal solution is the 2H, 3H, 4H, 5H, 6H, and 7H. Such a solution requires a single edit because player 2 simply has to pick up a 5H. This feature seeks to capture the intuition that an addressee should tend to act with respect to a card when the edit distance is not particularly high and hence the game is near a winning state.
- **Explicit Goal:** This binary feature is triggered in two cases: 1) When the suit of card mentioned matches the agreed-upon suit strategy in the common ground and 2) When the card mentioned appears in the set of cards the addressee claims to need. This models the prediction that locative utterances are more likely to be indirect commands when they are relevant to a well-defined goal.
- **Full Hands:** This binary feature is triggered when the speaker has three cards of the same suit as the card mentioned, and which are associated with some winning six-card straight, but the addressee does not. This models the prediction that locative utterances are likely to be indirect commands when they provide information relevant to winning, but only the addressee can act as such.

Single-feature classifiers are compared against a number of baselines to help benchmark our predictive task. Our first baseline, which is context-agnostic, seeks to capture the intuition that the role of a locative utterance is entirely ambiguous when considered in isolation. This baseline predicts the agent follow-up using a Bernoulli distribution weighted according to the class priors of the training data.

The second baseline incorporates surface-level dialogue context via bigram features of all the utterance exchanged between players up to and including the locative utterance. We also experimented with a unigram baseline but found that its performance was inferior to that of the bigram.

### 5.3 Results

We test our common-ground features one at a time with our logistic regression model, as we are interested in seeing how successfully they encode agents’ pragmatic inferences. We also combine the two best-performing common-ground features. We report the results of our experiments using an  $F_1$  measure and a 0.8/0.2 train/test split of our data in Table 1.

We see that of our two baselines, the bigram model performs better. This bigram model also uses 2,916 distinct lexical features which makes it a highly overspecified model for our moderate data size.

We find that our single-feature context-sensitive models both significantly outperform our baselines. Our Explicit Goal feature outperforms the Edit Distance feature by about 14%, which indicates that locative utterances are often interpreted as commands in the presence of an explicit, common goal. The Full Hands feature outperforms the Explicit Goal feature by about 6%. This strongly suggests that constraints on speaker action play a role in determining the illocutionary effect of a locative utterance. An addressee

of such an utterance will tend to act accordingly when their partner cannot pick up the card mentioned, and when the card in question brings them closer to winning the game. We find that combining the Explicit Goal and Full Hands features improved performance over only using the Explicit Goal feature but reduced overall performance. This could be because the two features encode some common information about the agents' pragmatic implicatures during the game, and hence their correlative effects tend to degrade the combined model performance.

## 6 Conclusion

In this work, we have performed an extensive study of command types as present in the Cards corpus. Using the corpus as a test bed for grounded natural language interaction among agents with a shared goal, we describe a variety of utterances that may function as indirect commands when regarded in context. In particular, locative utterances, which are not conventionally associated with command interpretations, are shown to operate as commands when considered in relation to situational constraints in the course of collaborative interaction. We develop a predictive task to show that models with carefully-designed features incorporating game state information can help agents effectively perform such pragmatic inferences.

## 7 Acknowledgments

The authors would like to thank Christopher Potts and all of the anonymous reviewers for their valuable insights and feedback.

## References

- Allen, J. F. and C. R. Perrault (1980). Analyzing intention in utterances. *Artificial intelligence* 15(3), 143–178.
- Austin, J. L. (1975). *How to do things with words*. Oxford University Press.
- Beun, R.-J. (2000). Context and form: Declarative or interrogative, that is the question. *Abduction, Belief, and Context in Dialogue: Studies in Computational Pragmatics 1*, 311–325.
- Bolt, R. A. (1980). *Put-that-there: Voice and gesture at the graphics interface*, Volume 14. ACM.
- Chai, J. Y., L. She, R. Fang, S. Ottarson, C. Littley, C. Liu, and K. Hanson (2014). Collaborative effort towards common ground in situated human-robot dialogue. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, pp. 33–40. ACM.
- Clark, H. H. (1979). Responding to indirect speech acts. *Cognitive psychology* 11(4), 430–477.
- Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Condoravdi, C. and S. Lauer (2012). Imperatives: Meaning and illocutionary force. *Empirical issues in syntax and semantics* 9, 37–58.
- Degen, J., M. Franke, and G. Jäger (2013). Cost-based pragmatic inference about referential expressions. In *CogSci*.
- Djalali, A., D. Clausen, S. Lauer, K. Schultz, and C. Potts (2011, November). Modeling expert effects and common ground using Questions Under Discussion. In *Proceedings of the AAAI Workshop on Building Representations of Common Ground with Intelligent Agents*, Washington, DC. Association for the Advancement of Artificial Intelligence.

- Djalali, A., S. Lauer, and C. Potts (2012). Corpus evidence for preference-driven interpretation. In M. Aloni, V. Kimmelman, F. Roelofsen, G. W. Sassoon, K. Schulz, and M. Westera (Eds.), *Proceedings of the 18th Amsterdam Colloquium: Revised Selected Papers*, Berlin, pp. 150–159. Springer.
- Kaufmann, M. (2016). Fine-tuning natural language imperatives. *Journal of Logic and Computation*, exw009.
- Kaufmann, S. and M. Schwager (2009). A unified analysis of conditional imperatives. In *Semantics and Linguistic Theory*, Volume 19, pp. 239–256.
- Malamud, S. A. and T. Stephenson (2015). Three ways to avoid commitments: Declarative force modifiers in the conversational scoreboard. *Journal of Semantics* 32(2), 275–311.
- Perrault, C. R. and J. F. Allen (1980). A plan-based analysis of indirect speech acts. *Computational Linguistics* 6(3-4), 167–182.
- Portner, P. (2007). Imperatives and modals. *Natural Language Semantics* 15(4), 351–383.
- Potts, C. (2012). Goal-driven answers in the Cards dialogue corpus. In N. Arnett and R. Bennett (Eds.), *Proceedings of the 30th West Coast Conference on Formal Linguistics*, Somerville, MA, pp. 1–20. Cascadilla Press.
- Searle, J. R. (1989). How performatives work. *Linguistics and philosophy* 12(5), 535–558.
- Tellex, S., P. Thaker, R. Deits, T. Kollar, and N. Roy (2012). Toward information theoretic human-robot dialog. In *Robotics: Science and Systems*, Volume 2, pp. 3.
- Tellex, S. A., T. F. Kollar, S. R. Dickerson, M. R. Walter, A. Banerjee, S. Teller, and N. Roy (2011). Understanding natural language commands for robotic navigation and mobile manipulation.
- Vogel, A., A. Gómez Emilsson, M. C. Frank, D. Jurafsky, and C. Potts (2014, July). Learning to reason pragmatically with cognitive limitations. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, Wheat Ridge, CO, pp. 3055–3060. Cognitive Science Society.
- Vogel, A., C. Potts, and D. Jurafsky (2013, August). Implicatures and nested beliefs in approximate Decentralized-POMDPs. In *Proceedings of the 2013 Annual Conference of the Association for Computational Linguistics*, Stroudsburg, PA, pp. 74–80. Association for Computational Linguistics.
- Walter, M. R., M. Antone, E. Chuangsuwanich, A. Correa, R. Davis, L. Fletcher, E. Frazzoli, Y. Friedman, J. Glass, and J. P. How (2015). A situationally aware voice-commandable robotic forklift working alongside people in unstructured outdoor environments. *Journal of Field Robotics* 32(4), 590–628.
- Wierzbicka, A. (1991). *Cross-cultural pragmatics: the semantics of human interaction*. Berlin: Mouton de Gruyter.