

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

No Such Thing as the Average Listener: Belief-driven versus Action-driven Strategies in Signaling

Permalink

<https://escholarship.org/uc/item/9d42m3tx>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Stacy, Stephanie

Yun, Yiling

Potter, Max

et al.

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

No Such Thing as the Average Listener: Belief-driven versus Action-driven Strategies in Signaling

Stephanie Stacy¹ Yiling Yun¹ Max Potter¹
stephaniestacy@g.ucla.edu yiling.yun@g.ucla.edu mpotter2001@g.ucla.edu
Naomi Moskowitz³ Tao Gao^{1,2}
nmoskowitz22@g.ucla.edu tao.gao@stat.ucla.edu

¹ Department of Statistics, UCLA ² Department of Communication, UCLA ³ Department of Psychology, UCLA

Abstract

Resolving overloading in communication requires attention to context. Previous research has found that the mutual assumption of cooperation during communication can act as a powerful constraint, allowing successful resolution under ambiguity. In this study, we investigate two specific types of cooperative context used in a communicative task which arise from different sources: beliefs and actions. In belief-driven communication, signals are interpreted in context of what else a speaker could have said about the world. Here communicators assume that the speaker aims to change the listener's beliefs by providing the most straightforward signal. In action-driven communication, signals are considered in terms of what a speaker can reasonably ask others to do. Signaling can be sensitive to utility considerations of acting and interacting in the physical world. Through a communication game, we tested how listeners would interpret an ambiguous signal using belief-driven or action-driven strategies. We find that while no one strategy is dominant overall, individuals are highly consistent in which strategy they employ when forced to decide.

Keywords: Cooperation; communication; pragmatics; speech acts; joint utility

Introduction

“It is impossible to speak in such a way that you cannot be misunderstood” - Karl Popper

Everyday communication can be incredibly overloaded: a single word can have many potential meanings. Despite this, even a sparse, ambiguous signal can often be enough to communicate successfully (Liszkowski, Carpenter, & Tomasello, 2008). This success stems from how attentive humans are to the context in which exchanges are framed (Sperber & Wilson, 1986). In addition, cooperation in general has been viewed as important for communication: this can explain why someone might ask a stranger for directions or grab a coat after being told it is cold outside. We focus on building communication under this cooperative default. Furthermore, we make the distinction between two separate types of rational cooperative logic communicators employ: speech acts and joint planning. Speech acts involve reasoning over all possible utterances under the assumption that the signal received is chosen deliberately to convey specific, relevant information. Joint planning involves communicating under the assumption that the shared environment helps cooperators decide on fair actions. These discrete but complementary views of cooperation in communication offer distinct mechanisms for resolving ambiguity and thus become most useful in different contexts. Moreover, these perspectives act as powerful

heuristics to guide communication because of the strong constraints they offer on how to send and interpret signals.

Both cooperative aspects of communication have previously been explored; however, they have been typically viewed separately, from the perspectives of different contexts. In the present study we incorporate them in the same behavioral task to explore whether humans can flexibly employ these two cooperative heuristics for disambiguation based on the context they are in. In addition, when both strategies can be used to solve the task but provide conflicting answers, we examine whether a strategy is dominant, both across participants and within an individual.

Context of Beliefs: Cooperative Speech Acts

The first type of cooperative logic employed during communication is speech acts. Speech acts fall under the umbrella of language pragmatics – the branch of linguistics which focuses specifically on the context sensitive interpretation of utterances. Grice's insights in developing a cooperative framework for communication have been highly influential in guiding a formalization of pragmatics. Specifically, Gricean cooperative logic treats communication as a truthful, concise, relevant, and straightforward exchange (Grice, 1975). To be considered cooperative, a signal should be straightforward, maximally efficient, and predicted to be interpretable by the receiver. In order to determine what is straightforward or efficient, communicators must engage in social reasoning about their partners. Although the signaler must ultimately decide on a signal, this process implicitly requires considering the context of all available – but not chosen – options. As a result, a signal with multiple literal meanings may now have a clear pragmatic interpretation which can be inferred using the situational context of the utterance.

While Grice's maxims are intuitively important for communication, alone, they are not enough to solve uncertainty in communication. Instead they must be combined with the insight that exchanges center around the *use* of language. This is useful because viewing communication through its use ties signals to communicative goals, making their utilities easier to define (Allen & Perrault, 1980; Goodman & Frank, 2016). Under this formulation, communication is a type of rational action: a speech act (Austin, 1975; Clark, 1996; Grice, 1975). When viewed as such, signals have the communicative goal of conveying information about a referent or state of the world

to a listener given the decision context (Van Rooy, 2003). A rational, utility driven signaler chooses a signal by evaluating all possible things she could say and picking a good option. Having a communicative goal provides the mechanism for that evaluation of what is good: a signal's value comes from how it is expected to change the listener's beliefs to reflect the intended referent. In turn, under these same assumptions, the listener can use these cooperative constraints to infer the intended pragmatic meaning of the signal.

Empirical evidence also supports a cooperative pragmatic account of communication in adults. Referential language games provide a controlled environment well suited for studying pragmatic reasoning (Lewis, 1969; Wittgenstein, 1953). In these games, a set of items with different features (e.g. shape, color) act as context, and a listener aims to understand which referent a speaker is indicating from a potentially ambiguous signal. In one game, listeners were asked to bet on which item they believed the signal referred to by distributing money across the different possibilities (Frank & Goodman, 2012). The listeners' bets (combined with empirical ratings of feature salience) agreed highly with how informative the speaker's signal was for disambiguating the item. These effects were also replicated in a forced choice task with a similar setup (Qing & Franke, 2015) as well as in a setting with more complex stimuli depicting ambiguous spatial relations, albeit with more noise (Carstensen, Kon, & Regier, 2014).

Context of Actions: Cooperative Joint Planning

The second type of cooperation we focus on is the context joint planning provides in a shared task. Much of communication occurs face-to-face where perceptual cues in the environment provide important context for framing an exchange. From this perspective, communication is simply a social tool which can enable individuals to coordinate and get things done together more effectively (Bruner, 1985; Tomasello, 2000; Vygotsky, n.d.). Again, communication is framed in terms of use, but this time studied using commonsense knowledge outside of language. Instead, this knowledge lies in considering consequences in the physical world through action planning and in others' mental world which provide the beliefs and desires to create a plan.

We motivate our emphasis on non-linguistic context by examining how even young children who do not yet have the capacity for fully-developed language can intelligently and flexibly reason using sparse communication. Before they have mastered language, toddlers can use visual communication to monitor and regulate their partner through protesting or attempting to re-engage them when they break a joint commitment formed through verbal acknowledgement (Gräfenhain, Behne, Carpenter, & Tomasello, 2009; Warneken, Chen, & Tomasello, 2006). At as young as four years old, children already exhibit sensitivity to minimal communication: they establish commitment using simple cues such as joint attention to help offset risks of cooperating in a stag hunting paradigm (Wyman, Rakoczy, & Tomasello, 2013). Moreover, slightly older children protest when their partner does not cooper-

ate, even when eye contact was the only established form of joint commitment (Siposova, Tomasello, & Carpenter, 2018). These findings indicate that communication can help establish strong joint goals in the context of cooperation.

One of the early non-verbal uses of communication is also demonstrated in the context of fairness. Children, but not chimpanzees are able to split rewards fairly in a collaborative task where it is easy for one party to monopolize rewards (Hamann, Warneken, Greenberg, & Tomasello, 2011; Warneken, Lohse, Melis, & Tomasello, 2011). In the few cases where one child tries to take more than is fair, sparse communication (e.g. "Hey!") quickly and decisively resolves disputes. Here "Hey" is overloaded, and this overloading is not solved by considering alternative protests or signals as in pragmatic reasoning. Instead, it is solved by considering the context of the task – where the principle of fairness is being violated. This protest comes from not only a preference for equality but also a resentment at being treated unfairly (Engelmann & Tomasello, 2019). These developmental studies demonstrate the importance of task-based cooperation in communication stripped down to its most fundamental form, without syntax or grammar.

These cooperative properties of commitment and fairness can be realized through utility driven joint planning: cooperators act under a rational plan that apportions fair costs and rewards to all parties given a joint goal. Empirical evidence has also shown that adults engage in joint utility planning for cooperative tasks, preferring co-efficient actions which prioritized the group utility over the utility of any individual (Török, Pomiechowska, Csibra, & Sebanz, 2019). From a utility driven standpoint, even toddlers understand the cooperative logic of ambiguous requests from a joint utility perspective (Grosse, Moll, & Tomasello, 2010). In this experiment, two equivalent items are equidistant from the toddler, but near and far relative to the speaker. When the speaker makes an ambiguous request for the item, children are able to use cooperative logic to reason over the *joint* utility dynamics of the environment in the context of the speaker's capabilities: reaching the far item more often when the speaker had their hands free than occupied. These studies support how communication should be taken in context of committing to achieve a shared goal fairly and respectfully. In both children and adults, joint planning ultimately makes it irrational to ask a collaborator to do something more easily accomplished by oneself.

Methods

This task combined feature overloading enriched by a spatial scene, which combined abilities to disambiguate signals both using the belief-driven context of words and the action-driven context of utility dynamics. In the grid-world environment, participants played a referential communication game where they were told the goal was to cooperate with their partner to reach a target item in the fewest steps. During the game, the participant always played the role of a receiver who could

observe the entire environment but did not know which item was the intended target. Participants were told they were collaborating with a cooperative signaler who had a full view of the grid and knew the target.

Participants

Thirty-four undergraduate students in the Department of Communication at University of California, Los Angeles (UCLA) participated in this online study for class credits. We analyzed the data of 27 participants after excluding one participant for not finishing the experimental trials, two participants for failing the comprehension quiz more than twice, and four participants for self-reporting not being serious in the experiment. The experiment was performed in accordance with guidelines and regulations approved by the UCLA institutional review board IRB#19-001990.

Stimuli and Task

Participants were able to access the experiment on their personal computer or laptop. On each trial, a 9 by 10 grid layout was presented to participants. Each grid square was 50 px \times 50 px and, three items were placed in the grid. Each item had two features of color (orange, purple, or green) and shape (triangle, circle, or square) for a total of nine distinct items. An icon representing the participant was located at grid location (4, 6) while their partner was located at (4, 0).

Design

The experiment followed a within-subject design with four conditions: pragmatic, utility, conflict, and signaler-walk (see Fig. 1). Participants played a total of 80 trials, presented in a random order with 20 trials per condition. The main dependent variable was the strategy the participant employed to solve each condition, reflected by the item they chose as the target. The receiver's response time from when they received a signal to when they selected an item was recorded. In addition, participants rated their own confidence after each decision. Although participants were told their partner was cooperative and intelligent, signals were pre-programmed. The signaler's decision depended on the condition and consisted of either an ambiguous signal – consistent with multiple potential items in the trial – or walking to an item when that item was closer to the signaler than the receiver.

The pragmatic condition coincided with the example from Frank & Goodman (2012) forcing participants to only used contextual information of item features, but was spread spatially in a visual display. This condition consisted of two items which had one unique feature and one feature shared with each of the features in a third item (see Figure 1a). The signaler always chose one of the shared features, making two items consistent with the signal. These two consistent items were equidistant from the receiver and all items were closer to the receiver than the signaler, so that utility dynamics could not influence the receiver's decision. Given the signal, receivers could select an item that was irrational: inconsistent with the signal, literal: consistent but could be indicated with




Figure 1: Example trials: Each condition where the signaler says “green” (a-c) has two literally consistent responses given the signal. However, the items can be distinguished using the context of other shapes and utility dynamics. 1a: green square is pragmatic; green triangle is literal. 1b: green circle at (4,2) is closer by individual utility (than the green circle at (0,5)) but jointly inefficient; green circle at (0,5) is jointly efficient for the receiver but farther by individual utility. 1c: green triangle is pragmatic but jointly inefficient; green square is jointly efficient for the receiver but not pragmatic. 1d: The signaler walks to the purple circle herself, ending the trial without a receiver action.

a more straightforward signal, or pragmatic: consistent and most straightforward because both features were overloaded.

The utility condition contained two items with identical features (and one irrelevant distinct one), which forced participants to make a purely utility-based decision. The observed signal was a feature shared by the identical items and because the items were identical, considering the context of language pragmatics was unable to help (see Figure 1b). This setup reflected the dynamic in Grosse et al. (2010), but with a stronger individual utility component. One identical item (A) was always closer to the signaler than the receiver, and the other (B) was closer to the receiver than signaler. However, from the receiver's perspective, item A was closer than item B. In response to the signal, receivers could select an item that was irrational: the non-identical inconsistent one; individual: item A, closer than item B from an individual utility perspective; or joint: item B, closer to the receiver than signaler from a joint utility perspective.

The conflict condition was designed to force participants to choose between a joint utility and pragmatic strategy. It was identical to the pragmatic condition in terms of item feature structure and signal. Additionally, the two items consistent with the signal were equidistant from the receiver. However, instead of all items being closer to the receiver than signaler from a joint perspective, the pragmatic item was closer to the signaler than to the receiver (see Figure 1c). Receivers could still select an irrational item, but now had two previous heuristics in conflict and could select either a pragmatic interpretation inconsistent with joint utility (pragmatic) or a joint utility interpretation that was literally true, but with an alternative signal that would be more straightforward (joint).

Finally, in the signaler-walk condition, the signaler walked to an item, and participants did not make a decision. In all cases, the item walked to would be closer to the signaler than receiver from a joint perspective. This control condition was to establish that the signaler was making rational decisions from a cooperative joint utility perspective.

Items and signals were counterbalanced to account for preference of feature or feature value. In addition, items were separated by a minimum distance of two grid units to reduce potential perceptual chunking. Items always were always at least two grid units farther from one agent than the other in order to ensure clear joint utility judgments. Finally, item locations within a condition were sampled randomly without replacement, subject to the utility constraints defined by the condition and aforementioned restrictions.

Procedure

Participants entered the experiment by opening the link to the experiment on their own device. They started with an instruction tutorial which established the rules and cooperative context of the task then completed a comprehension quiz that tested them on the goal and set-up of the experiment. Before beginning the actual task, participants also completed eight practice trials to familiarize them with the task and types of decisions the signaler could make. Practice consisted of two trials in each condition presented in a random order.

In each trial, participants first waited for the signaler to make a decision. The signaler either walked to an item herself or sent a signal to the participant describing a single feature (e.g. “circle”). If the signaler sent a signal, the participant then had a chance to walk to the item they believed was the target by clicking on it. Before they made a decision, hovering the cursor over any item in the grid would display the distance each agent was from that item: the cost of traveling to that location. If the signaler moved to the target herself, participants observed the signaler walking to the item. Both agents traveled along the grid taking steps in the four cardinal directions. The trial ended when either agent reached an item. Once either agent reached an item, a review box would pop up, showing who took how many steps to reach which item. Participants were asked to rate their confidence in their selection from one (least confident) to five (most confident). Participants then proceeded to the next trial. After all experi-

mental trials, participants took an exit survey for self-reports on how serious they were throughout the experiment, strategies they used, and performance of their partner.

Results

We analyzed the strategy, response time, and confidence rating on 540 trials. Across all conditions, only nine trials had irrational responses (Pragmatic: 5, Utility: 3, Conflict: 1), thus we restricted our analyses to focus on the major strategies employed – for the Pragmatic condition, pragmatic/literal; for the Utility condition, joint/individual; and for the Conflict condition, joint/pragmatic.

Strategy Preferences: Population versus Individual

Our first research question of interest was whether any clear strategy preferences emerged across the population. For each condition, we averaged subject-specific trials and used a two tailed z-test under the hypothesis $H_0 : \mu = .5$, which tested for a strategy preference across the sampled individuals. For the Pragmatic condition, we found a significant preference for pragmatic signal interpretations over literal ones ($\bar{x}_{pragmatic} = .717$, $p = 1.18 \times 10^{-7}$). For the Utility condition, there was no preference for one type of utility reasoning over the other ($\bar{x}_{joint} = .567$, $p = .436$). Similarly, there was no dominant strategy in the Conflict condition ($\bar{x}_{joint} = .539$, $p = .630$). Over the three conditions, the only strategy consistently dominant across individuals was the preference for pragmatic over literal interpretations of an ambiguous signal.

While there seems to be no clear preference across individuals for one strategy over another, the picture becomes quite different when examined from the individual level. Looking at the individual-level strategy breakdown suggests people were highly consistent in choosing a strategy (see Figure 2). We focused on the Conflict condition in particular to explore this idea. Instead of looking whether people employed a dominant strategy overall, we tested the hypothesis of whether an *individual* employed a dominant strategy, adjusting for multiple comparisons using the Benjamini-Hochberg criteria. In 23 out of 27 cases, participants adopt a dominant strategy (all $p_{adj} < .05$ except subjects 27: $p = .503$, 32: $p = .273$, 41: $p = .125$, 43: $p = .125$).

Moreover, we investigated whether participants' strategies correlated between conditions. Pairwise correlation analyses indicated a strong positive relationship between an individual's strategy in the Utility and Conflict condition (Spearman's $\rho = .91$, $p = 1.20 \times 10^{-11}$). That is, individuals who chose a joint utility strategy in the Utility condition were also likely to choose a joint utility strategy when pragmatic reasoning and utility reasoning were in conflict (see Figure 3). This effect was not observed for the Pragmatic and Conflict ($\rho = .22$, $p = .267$) or Pragmatic and Utility conditions ($\rho = .27$, $p = .173$).

Strategy Difficulty: Decision Time and Confidence

In this task, we examined decision time which can act as a rough proxy for the cognitive difficulty involved in employ-




Figure 2: Strategy breakdown for all participants across conditions, ordered by a participant's preference in the Conflict condition. We observe high consistency in joint utility reasoning within individuals across conditions.




Figure 3: Strategy correlations across conditions. Correlation coefficients (upper triangle), corresponding to the individual responses (lower triangle). Histogram describing distribution of strategy preference (on diagonal). Data concentrated at the extremes of the histogram indicate the strong, divergent preferences seen in the Utility and Conflict conditions.

ing that strategy (Townsend, 1992). Because participants took the study on their personal device instead of a controlled laboratory setting, the data included extreme decision times which could not reasonably be attributed to deliberation on the task.

While it is common practice to remove reaction times above a certain threshold, typically three Z-scores away (Tabachnick, Fidell, & Ullman, 2007), this can substantively inflate Type I error rate (Bakker & Wicherts, 2014). Because we had no strong literature-based intuition for a decision time cut-off to indicate when subjects were no longer paying attention, we relied on nonparametric testing which is robust to outliers and skew inherent in reaction time data. In the Pragmatic condition, we found participants to take more time when employing pragmatic reasoning than literal reasoning ($\bar{x}_{pragmatic} = 5.74$ sec, $\bar{x}_{literal} = 4.14$ sec, Mann-Whitney-Wilcoxon test (MWW); 95% CI of median difference [.451, 1.951], $p = 6.15 \times 10^{-4}$). In the Utility condition, participants took more time to respond when employing joint utility reasoning ($\bar{x}_{joint} = 3.42$ sec, $\bar{x}_{individual} = 2.54$ sec, MWW; 95% CI of median difference [.122, .822], $p = 4.36 \times 10^{-3}$). Finally, in the Conflict condition, participants spent longer to make a decision when employing pragmatics as opposed to joint utility reasoning ($\bar{x}_{pragmatic} = 3.79$ sec, $\bar{x}_{joint} = 3.07$ sec, MWW; 95% CI of median difference [.002, .886], $p = 0.024$).




Figure 4: Distribution of decision times for each strategy employed in each condition. Responses above 45 seconds ($n_{pragmatic} = 13$, $n_{utility} = 4$, $n_{conflict} = 5$) are included in analyses but not shown here for legibility.

In addition to decision time, we also examined self-reported confidence as a function of decision strategy. In the Pragmatic condition, participants were significantly more confident when choosing pragmatic items than when choosing literal ones ($\bar{x}_{pragmatic} = 3.48$, $\bar{x}_{literal} = 3.03$, $p = 1.80 \times 10^{-5}$ under Welch's t-test). In the Utility condition, participants were significantly more confident when choosing items that maximized joint utility than when they chose the items that maximized individual utility ($\bar{x}_{joint} = 3.99$, $\bar{x}_{individual} = 3.65$, $p = 6.80 \times 10^{-5}$ under Welch's t-test). Finally, in the Conflict condition, participants were significantly more confident when choosing the joint utility items than when choosing the pragmatic ones ($\bar{x}_{joint} = 3.90$, $\bar{x}_{pragmatic} = 3.63$,

$p = 0.00208$ under Welch's t-test).

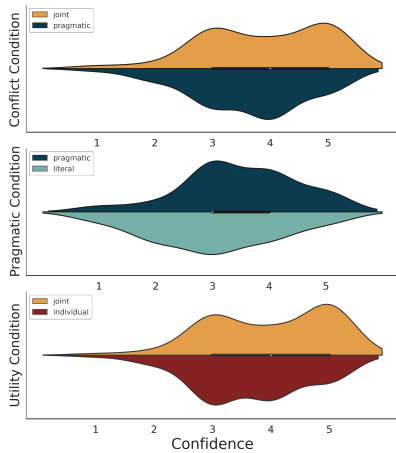


Figure 5: Self-rated confidence split by strategy decisions.

Discussion

At first glance, overall results on strategy preference may appear inconclusive. A preference for pragmatic reasoning in the Pragmatic condition supports previous empirical findings in referential language games (Frank & Goodman, 2012; Qing & Franke, 2015), replicating this phenomena in our visual task. However, we found no general preference for either joint vs. individual utility or, when put in conflict, pragmatics vs. joint utility reasoning. However, when we consider the Conflict condition at the individual level, we see that people are exceptionally strategic in their decisions. In reality, there is no “average” communicative reasoner employing the dominant strategy but rather groups of highly consistent decision-makers who have overwhelming preferences for *different* strategies. In this task, some individuals took a belief-driven perspective to signal understanding: signals influence beliefs about the visual features of the referent item based on the speaker’s intention to be straightforward. Other individuals took an action-driven perspective to signal understanding: signals should be interpreted in terms of utility to help make actions efficient under the joint task.

When we consider individual preferences of the Conflict condition in conjunction with the other conditions, an interesting pattern emerges. First an individual’s preference for pragmatic reasoning was not indicative of their dominant strategy in the Conflict condition, though the average strategy analysis indicated that people were generally capable of adopting a pragmatic approach. However, preference for a joint over individual utility approach in the Utility condition had a strong positive correlation with the individual’s dominant strategy in the Conflict condition. This suggests that while only a subset of individuals used a joint utility based strategy, it was an incredibly powerful heuristic that could generalize across contexts for those people, a phenomenon not observed for pragmatic reasoning.

Counter to our initial expectations around cooperative

planning, in the Utility condition, many individuals preferred an individual utility strategy. One potential explanation for this is the lack of interaction between partners in the task. Although framed as cooperative, in reality the signaler’s responses were pre-programmed and there was no regulation or feedback between partners. These findings leave room for interesting potential future work in a truly interactive version of this task, where the role of communicator and listener are not fixed, and could lead to much stronger preferences for fairness and cooperation.

Another explanation for the prevalence of individual utility reasoning is that being cooperative requires effort. It is more intuitive to reason from an egocentric perspective (Epley, Keysar, Van Boven, & Gilovich, 2004), and participants who have the capacity to plan jointly may not have had high enough motivation in the task to surpass the effort threshold required to engage it. Empirical work points to the idea that when interpreting referring expressions, individuals weigh both egocentric and joint perspectives depending on context (Heller, Parisien, & Stevenson, 2016) leading to a division of labor in communication. One factor that could contribute to this division is an estimation of the degree of effort one’s partner is exerting (Hawkins, Gweon, & Goodman, 2021). Some evidence that could align with this explanation comes from analyzing the decision time and confidence ratings of participants in conjunction. We see that people were faster at making individual utility based decisions than joint utility based ones. At the same time, confidence ratings were higher on trials where people employed joint utility. We see a similar pattern when comparing decisions in the Pragmatic condition. On average, people took longer to make a pragmatic decision than a literal one, which is highly consistent with the computational models of pragmatics. In order to come up with a pragmatic interpretation of a signal, a listener must first reason over literal interpretations (Goodman & Frank, 2016). At the same time, people were more confident about pragmatic selections than literal ones. Both these conditions suggest that cooperative communication takes work: a listener must do their share of the heavy-lifting in language to reason flexibly under ambiguity, using a variety of contextual cues.

While these three experimental conditions have shown how individuals employ pragmatic and utility-based reasoning independently as well as how strategies diverge when pitted against each other, future research should address how these heuristics interact with each other, which has been demonstrated to be a theoretically promising approach to communication (Stacy et al., 2021). These communicative strategies are not necessarily incompatible with each other. In fact, context – and the constraints it provides – likely accumulate evidence to resolve ambiguity in linguistic communication (Roy & Mukherjee, 2005). Integration of many simpler contextual heuristics may be a key to fast, flexible, and sparse signaling.

Acknowledgments

We thank our anonymous reviewers for their thoughtful feedback. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1650604 to SS and by DARPA PA 19-03-01 and ONR MURI project N00014-16-1-2007 to TG.

References

- Allen, J. F., & Perrault, C. R. (1980). Analyzing intention in utterances. *Artificial Intelligence*, 15(3), 143–178.
- Austin, J. L. (1975). *How to do things with words*. Oxford University Press.
- Bakker, M., & Wicherts, J. M. (2014). Outlier removal, sum scores, and the inflation of the type I error rate in independent samples t tests: The power of alternatives and recommendations. *Psychological Methods*, 19(3), 409.
- Bruner, J. (1985). Child's talk: Learning to use language. *Child Language Teaching and Therapy*, 1(1), 111–114.
- Carstensen, A., Kon, E., & Regier, T. (2014). Testing a rational account of pragmatic reasoning: The case of spatial language. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th annual conference of the cognitive science society* (p. 2009–2013). Austin, TX: Cognitive Science Society.
- Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Engelmann, J. M., & Tomasello, M. (2019). Children's sense of fairness as equal respect. *Trends in Cognitive Sciences*, 23(6), 454–463.
- Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of personality and social psychology*, 87(3), 327.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818–829.
- Gräfenhain, M., Behne, T., Carpenter, M., & Tomasello, M. (2009). Young children's understanding of joint commitments. *Developmental Psychology*, 45(5), 1430–1443.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & M. J. (Eds.), *Syntax and semantics* (Vol. 3, p. 41–58). New York: Academic Press.
- Grosse, G., Moll, H., & Tomasello, M. (2010). 21-month-olds understand the cooperative logic of requests. *Journal of Pragmatics*, 42(12), 3377–3383.
- Hamann, K., Warneken, F., Greenberg, J. R., & Tomasello, M. (2011). Collaboration encourages equal sharing in children but not in chimpanzees. *Nature*, 476(7360), 328–331.
- Hawkins, R. D., Gweon, H., & Goodman, N. D. (2021). The division of labor in communication: Speakers help listeners account for asymmetries in visual perspective. *Cognitive science*, 45(3).
- Heller, D., Parisien, C., & Stevenson, S. (2016). Perspective-taking behavior as the probabilistic weighing of multiple domains. *Cognition*, 149, 104–120.
- Lewis, D. (1969). *Convention: A philosophical study*. John Wiley & Sons.
- Liszkowski, U., Carpenter, M., & Tomasello, M. (2008). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition*, 108(3), 732–739.
- Qing, C., & Franke, M. (2015). Variations on a bayesian theme: Comparing bayesian models of referential reasoning. In H. Zeevat & H.-C. Schmitz (Eds.), *Bayesian natural language semantics and pragmatics* (pp. 201–220). Springer.
- Roy, D., & Mukherjee, N. (2005). Towards situated speech understanding: Visual context priming of language models. *Computer Speech & Language*, 19(2), 227–248.
- Siposova, B., Tomasello, M., & Carpenter, M. (2018). Communicative eye contact signals a commitment to cooperate for young children. *Cognition*, 179, 192–201.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition* (Vol. 142). Cambridge: Harvard University Press.
- Stacy, S., Li, C., Zhao, M., Yun, Y., Zhao, Q., Kleiman-Weiner, M., & Gao, T. (2021). Modeling communication to coordinate perspectives in cooperation. In *Proceedings of the 43rd annual conference of the cognitive science society*. cognitive science society.
- Tabachnick, B. G., Fidell, L. S., & Ullman, J. B. (2007). *Using multivariate statistics* (Vol. 5). Boston, MA: Pearson.
- Tomasello, M. (2000). The social-pragmatic theory of word learning. *Pragmatics*, 10(4), 401–414.
- Török, G., Pomiechowska, B., Csibra, G., & Sebanz, N. (2019). Rationality in joint action: Maximizing efficiency in coordination. *Psychological Science*, 30(6), 930–941.
- Townsend, J. T. (1992). On the proper scales for reaction time. *Cognition, information processing, and psychophysics: Basic issues*, 105–120.
- Van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, 26(6), 727–763.
- Vygotsky, L. S. (n.d.). *Mind in society: The development of higher psychological processes*.
- Warneken, F., Chen, F., & Tomasello, M. (2006). Cooperative activities in young children and chimpanzees. *Child Development*, 77(3), 640–663.
- Warneken, F., Lohse, K., Melis, A. P., & Tomasello, M. (2011). Young children share the spoils after collaboration. *Psychological Science*, 22(2), 267–273.
- Wittgenstein, L. (1953). *Philosophical investigations*.
- Wyman, E., Rakoczy, H., & Tomasello, M. (2013). Non-verbal communication enables children's coordination in a “stag hunt” game. *European Journal of Developmental Psychology*, 10(5), 597–610.