

The Computational Structure of Unintentional Meaning

Mark K. Ho (mho@princeton.edu)

Department of Psychology, Princeton University
Princeton, NJ 08540

Joanna Korman*(jkorman@mitre.org)

The MITRE Corporation
Bedford, MA 01730

Thomas L. Griffiths (tomg@princeton.edu)

Department of Psychology, Princeton University
Princeton, NJ 08540

Abstract

Speech-acts can have literal meaning as well as pragmatic meaning, but these both involve consequences typically *intended* by a speaker. Speech-acts can also have *unintentional meaning*, in which what is conveyed goes above and beyond what was intended. Here, we present a Bayesian analysis of how, to a listener, the meaning of an utterance can significantly differ from a speaker's intended meaning. Our model emphasizes how comprehending the intentional *and* unintentional meaning of speech-acts requires listeners to engage in sophisticated model-based perspective-taking and reasoning about the *history* of the state of the world, each other's actions, and each other's observations. To test our model, we have human participants make judgments about vignettes where speakers make utterances that could be interpreted as intentional insults or unintentional *faux pas*. In elucidating the mechanics of speech-acts with unintentional meanings, our account provides insight into how communication both functions and malfunctions.

Keywords: Bayesian modeling, social cognition, common ground, speech-act theory, faux pas, theory of mind

Introduction

People sometimes communicate things that they did not intend or expect. Consider the following vignette, adapted from Baron-Cohen et al. (1999):

Curtains Paul had just moved into a new apartment. Paul went shopping and bought some new curtains for his bedroom. After he returned from shopping and had put up the new curtains in the bedroom, his best friend, Lisa, came over. Paul gave her a tour of the apartment and asked, "How do you like my bedroom?"

"Those curtains are horrible," Lisa said. "I hope you're going to get some new ones!"

Clearly, Lisa committed a social blunder or *faux pas* with her remark. What happened here? When Lisa says, "Those curtains look horrible," she is merely stating her private aesthetic experience of the curtains. The *literal* meaning is straightforward: The curtains look bad. And the *intended* or *expected* meaning of her utterance is largely captured by this literal meaning. However, to Paul, the utterance means more. Specifically, what Lisa is *really* saying is that *he* chose horrible curtains. Of course, Lisa did not "really" say that Paul's

choice in curtains was horrible—she had no intention of conveying such an idea. Paul might even realize this. Nonetheless, the remark stings. Why? Lisa and Paul each possess a piece of a puzzle, and when put together, they entail that Paul has awful taste in curtains. At the outset, neither one knew that they each had a piece of a puzzle. But once Lisa makes her remark, she inadvertently completes the puzzle, at least from Paul's perspective.

Standard models of communication (Grice, 1957; Sperber & Wilson, 1986) tend to focus on how people use language successfully. For example, people can imply more than they literally mean (Carston, 2002), convey subtle distinctions via metaphor (Tendahl & Gibbs Jr, 2008), and manage their own and others' public face using politeness (Levinson, Brown, Levinson, & Levinson, 1987; Yoon, Frank, Tessler, & Goodman, 2018). But things do not always go smoothly, as Paul and Lisa's situation indicates. Sometimes people find themselves having inadvertently stepped on conversational landmines, meaning things that they never anticipated meaning. Notably, because such situations present complex dilemmas of mutual perspective-taking against a backdrop of divergent knowledge, they can serve as advanced tests of theory of mind (Baron-Cohen et al., 1999; Zalla, Sav, Stopin, Ahade, & Leboyer, 2009; Korman, Zalla, & Malle, 2017). But how do people reason about such dilemmas? And how can this be understood computationally? Disentangling unintentional meaning can shed light on how communication works in a broader social context as well as inform the design of artificial intelligences that interact with people.

Here, we develop a rational, cognitive account of interpreting unintentional speech-acts that builds on existing Bayesian models of language (e.g., Rational Speech Act [RSA] models [Goodman & Frank, 2016]). To do this, we analyze the general epistemic structure of social interactions such as the one described above and model listeners engaging in *model-based perspective-taking*. In particular, our model explains how the same utterance could be interpreted as either an (unintentional) *faux pas* or an *intentional insult* depending on the context of a listener and speaker's interaction. We then test several model predictions in an experiment with human participants. In the following sections, we outline our computational model, experimental results, and their implications.

*The author's affiliation with The MITRE Corporation is provided for identification purposes only, and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions, or viewpoints expressed by the author.

A Bayesian Account of Unintentional Meaning

During social interactions, people reason about the world as well as each other’s perspective on the world (Brown-Schmidt & Heller, 2018). Thus, our account has two components, which we formulate as probabilistic models. First, we specify a *world model* that captures common-sense relationships between world states, actions, and events. Second, we define *agent models* of a speaker and listener reasoning about the world and one another.

World Model

We model the interaction as a partially observable stochastic game (POSG), a generalization of Markov Decision Processes (MDPs) with multiple agents with private observations (Kuhn, 1953). Formally, a world model $\mathcal{W} = \langle I, \mathcal{S}, \mathcal{A}, \mathcal{Z}, \mathcal{T} \rangle$ where:

- I is a set of n agents indexed $1, \dots, n$;
- \mathcal{S} is a set of possible states of the world, where each state $s \in \mathcal{S}$ is an assignment to k variables, $s = (x_0, x_1, \dots, x_k)$;
- $\mathcal{A} = \times_{i \in I} \mathcal{A}^i$ is the set of joint actions, i.e., every combination of each agent i ’s actions, \mathcal{A}^i (including utterances);
- $\mathcal{Z} = \times_{i \in I} \mathcal{Z}^i$ is the set of joint private observations, which is every possible combination of each individual agent i ’s private observation set, \mathcal{Z}^i ; and
- $\mathcal{T} = P(z, s' | s, a)$ is a transition function representing the probability of a joint observation z and next state s' given a previous state $s \in \mathcal{S}$ and joint action $a \in \mathcal{A}$ was taken.

In **Curtains**, the initial state, s_0 , includes Paul with the old curtains in the apartment and Lisa elsewhere. There is also a latent state feature of interest: whether Paul has good or bad taste. At $t = 0$, Paul’s action, a_0^{Paul} , is choosing new curtains, while Lisa’s action, a_0^{Lisa} , is going to the apartment. The joint action, $a_0 = (a_0^{\text{Paul}}, a_0^{\text{Lisa}})$, results in a new state, s_1 , with them both in the apartment, the curtains either good or bad, and Paul’s taste. Paul’s observation, z_0^{Paul} , but not Lisa’s, z_0^{Lisa} , includes Paul having put up the curtains. These relationships between world states (e.g. Paul and Lisa’s locations), actions (e.g. Lisa walking to Paul’s apartment), and observations (e.g. Paul observing himself put up the curtains) are formally encoded in the transition function \mathcal{T} . The sequence of states, joint actions and observations resulting from such interactions constitute the *history* up to a point t , $\vec{h}_t = (s_0, a_0, z_0, \dots, s_{t-1}, a_{t-1}, z_{t-1}, s_t)$.

Agent Models

Agents are modeled as Bayesian decision-makers (Bernardo & Smith, 1994) who can reason about the world and other agents as well as take actions—including making utterances.

Interactive Belief State Agents’ beliefs are probability distributions over variables that represent aspects of the current state, previous states, or each other’s beliefs. The configuration of these first- and higher-order, recursive beliefs constitute their *interactive belief state* (Gmytrasiewicz & Doshi, 2005). We refer to an agent i ’s beliefs as b^i . For example, if we denote Paul’s taste as the variable T^{Paul} , then Paul’s

belief that his taste is good is $b^{\text{Paul}}(T^{\text{Paul}} = \text{Good})$. Higher-order beliefs can also be represented. For instance, we can calculate Paul’s *expectation* of Lisa’s belief in his taste as $\mathbb{E}_{b^{\text{Paul}}}[b^{\text{Lisa}}](T^{\text{Paul}}) = \sum_{b^{\text{Lisa}}} b^{\text{Paul}}(b^{\text{Lisa}}(T^{\text{Paul}}))$.

An agent i ’s beliefs are a function of their prior, model of the world, model of other agents, and *observation history* up to time t , \vec{z}_t^i . Note that \vec{z}_t^i can include observations that are completely private to i (e.g., Lisa’s personal aesthetic experience) as well as public actions and utterances (e.g., Lisa’s remark to Paul). Thus, we denote Paul’s belief about his taste at a time t as $b_t^{\text{Paul}}(T^{\text{Paul}}) = b^{\text{Paul}}(T^{\text{Paul}} | \vec{z}_t^{\text{Paul}})$. Given a sequence of observations, \vec{z}_t^i , posterior beliefs about a variable X are updated via Bayes’ rule:

$$b(X | \vec{z}_t^i) \propto b(\vec{z}_t^i | X)b(X) \quad (1)$$

$$= \sum_{\vec{h}_t} b(\vec{z}_t^i | \vec{h}_t)b(\vec{h}_t, X) \quad (2)$$

The capacity to reason about higher-order beliefs (e.g., Paul’s beliefs about Lisa’s belief in his taste), along with Equation 2 express agents’ joint inferences about events and model-based perspective-taking.

Speaker Model Speakers have beliefs and goals. When choosing what to say, they may have beliefs and goals with respect to the listener’s beliefs and goals. In our example, Lisa may care about being informative about how she sees the curtains, but may also think Paul cares about having good taste in curtains and care whether she hurts his feelings. Following previous work (e.g., Franke, 2009), we model speakers as reasoning about changes in *belief states*. Here, we are interested in how a speaker can intend to mean one thing but inadvertently mean another. Thus, we distinguish between state variables that the speaker wants to be *informative* about, X^{Info} (e.g., how Lisa sees the curtains), and *evaluative* variables, X^{Eval} , that the listener wants to take on a specific value $x^{\text{Eval}*}$ (e.g., Paul’s taste being good). The speaker then cares about the changes in those quantities. Formally:

$$\Delta_t^{\text{L-Info}} = b_{t+1}^{\text{L}}(X^{\text{Info}} = x^{\text{Info}}) - b_t^{\text{L}}(X^{\text{Info}} = x^{\text{Info}}), \quad (3)$$

where x^{Info} is given by \vec{h}_t ; and,

$$\Delta_t^{\text{L-Eval}} = b_{t+1}^{\text{L}}(X^{\text{Eval}} = x^{\text{Eval}*}) - b_t^{\text{L}}(X^{\text{Eval}} = x^{\text{Eval}*}). \quad (4)$$

A speaker who is interested in what the listener thinks about X^{Info} and X^{Eval} will, at a minimum, anticipate how their utterances will influence $\Delta_t^{\text{L-Info}}$ and $\Delta_t^{\text{L-Eval}}$. A speaker would then have a reward function defined as:

$$R^S(a_t^S, \vec{z}_{t+1}^L) = \theta^{\text{L-Info}} \Delta_t^{\text{L-Info}} + \theta^{\text{L-Eval}} \Delta_t^{\text{L-Eval}} \quad (5)$$

where the θ terms correspond to how the speaker values certain outcomes in the listener’s mental state. For instance, if $\theta^{\text{L-Eval}} < 0$, the speaker *wants* to insult the speaker.

Given Equation 5, a speaker can take utterances based on expected future utility/rewards (or *value* [Sutton & Barto,

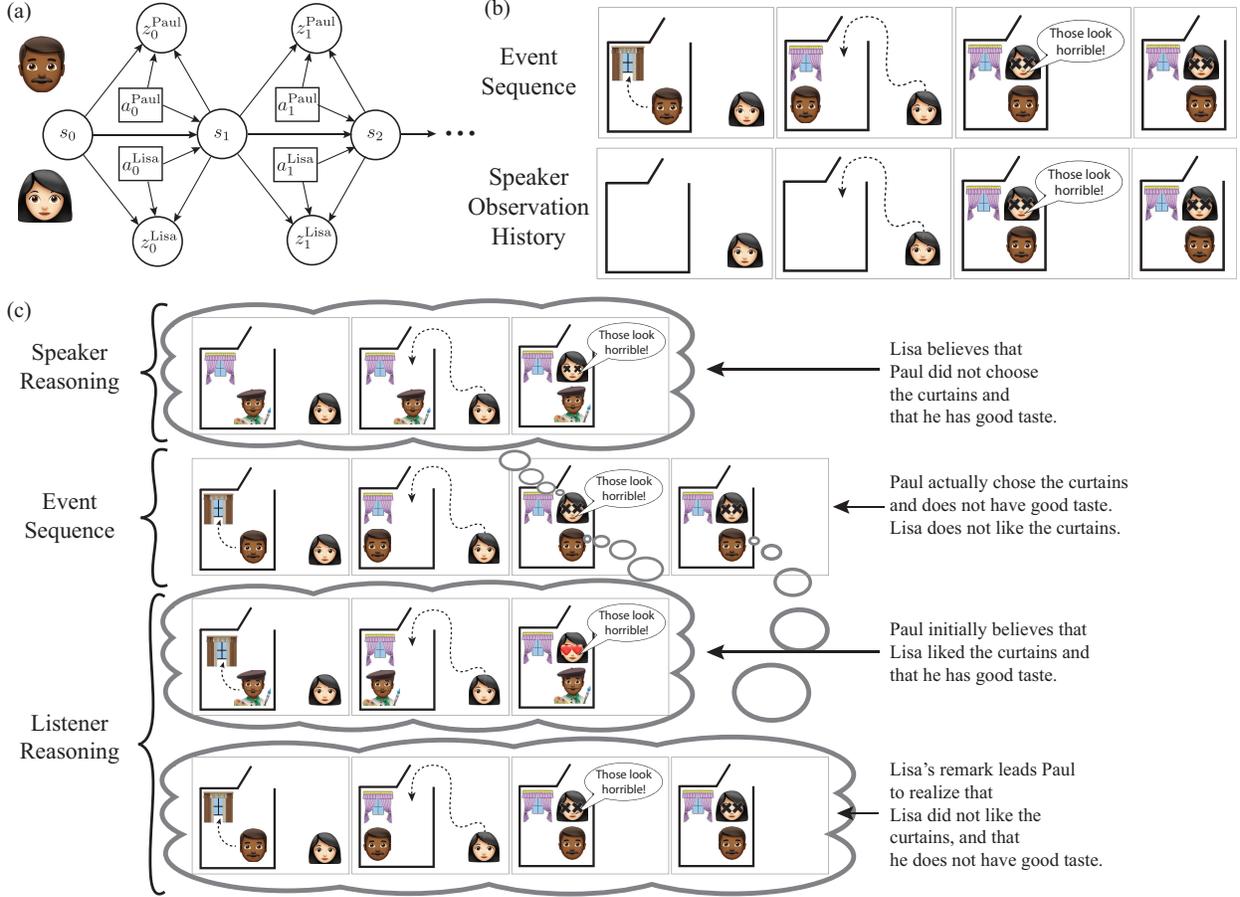


Figure 1: Model and example of unintentional meaning. (a) Influence diagram with state, action, and observation dependencies. Circles correspond to world state (e.g., s_t) and observation (e.g., z_t^i) variables; squares correspond to agent action variables (including utterances) (e.g., a_t^i). (b) Event sequence in **Curtains** (top) and speaker observation history (bottom). Lisa does not observe Paul choose the curtains. Only Lisa experiences whether the curtains look good or bad and comments on this experience. (c) Diagram of interactive belief state over time in **Curtains**.

1998]), where the expectation is taken with respect to the speaker’s beliefs, b_t^S . That is, given observations \vec{z}_t^S , the value of a_t^S is $V^S(a_t^S; \vec{z}_t^S) = \mathbb{E}_{b_t^S}[R^S(a_t^S, \vec{z}_{t+1}^L)]$, and an action is chosen using a Luce choice rule (Luce, 1959).

Listener Inference Our goal is to characterize how a listener’s interpretation of an utterance can differ from a speaker’s intended meaning, which requires specifying listener inferences. We start with a simple listener that understands the literal meanings of words when spoken. Following previous models (Franke, 2009; Goodman & Frank, 2016), the literal meaning of an utterance a^S is determined by its truth-functional denotation, which maps histories to Boolean truth values, $[[a^S]] : \vec{h}_t \mapsto y, y \in \{\text{True}, \text{False}\}$. A literal listener’s model of speaker utterances is:

$$b(a^S | \vec{h}_t) \propto \begin{cases} 1 - \epsilon & \text{if } [[a^S]](\vec{h}_t) \\ \epsilon & \text{if } \neg [[a^S]](\vec{h}_t) \end{cases}$$

where ϵ is a small probability of a^S being said even if it happens to be false.

We can also posit a more sophisticated listener who, rather than assuming utterances literally reflect reality, reason about how a speaker’s beliefs and goals mediate their use of language. This type of listener draws inferences based on an *intentional* model of a speaker that track the quantities in Equations 3 and 4 as well as maximize the expected rewards. These inferences, however, occur while the listener is also reasoning about the actual sequence of events \vec{h}_t , making it possible to draw inferences based on utterances that the speaker did not anticipate.

Model Simulations

In the original **Curtains** scenario, Lisa was not present when Paul put up the curtains. As a result, Lisa’s comment (“Those curtains are horrible”) is interpreted in a *diverging* observation history context. But what if Lisa had been present when Paul put up the curtains and made the same utterance? Given a *shared* observation history, Lisa’s utterance is still offensive, but now Lisa has all the information needed to realize it would be offensive. Put simply, in the *diverging history* con-

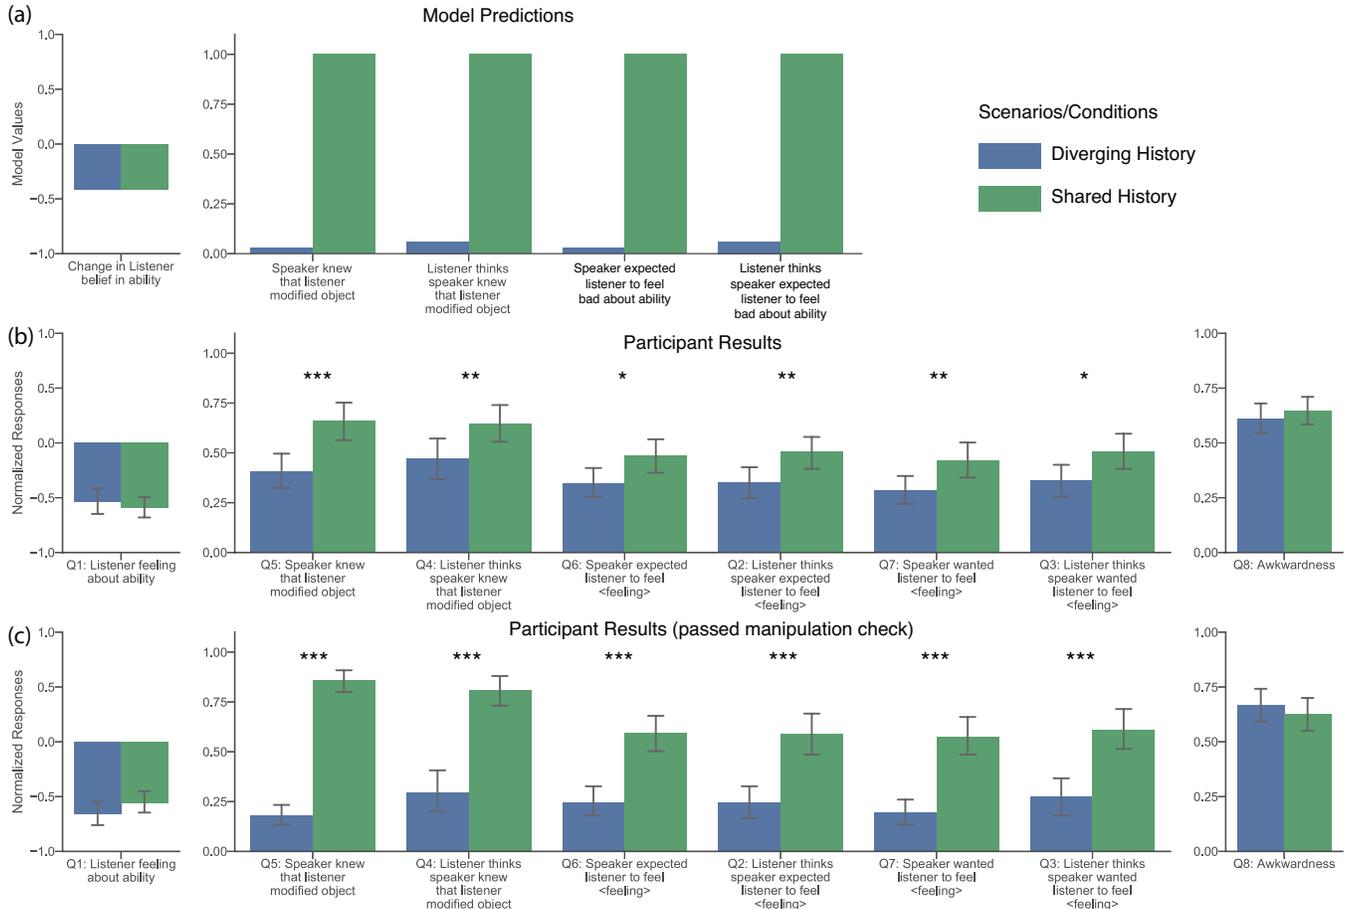


Figure 2: (a) Model predictions. The model predicts that the listener’s change in belief in the evaluative variable (Δ_t^{L-Eval}) is equally negative in the diverging and shared history scenarios. However, whether the speaker anticipated the offensiveness of their comment differs between the two scenarios, as do the listener’s beliefs about the speaker’s anticipation. (b) Judgments from all participants by question. Responses were normalized depending on whether response scales were valanced (Q1), likelihood (Q2-Q7), or qualitative (Q8). (c) Judgments from participants who correctly identified whether the speaker knew the listener modified the object. *: $p < .05$, **: $p < .01$, ***: $p < .001$.

text, the utterance is a faux pas, whereas in the *shared history* context, it is an intentional insult.

In this section, we discuss how our model can be used to make these intuitive predictions precise and explain how they arise from agents’ interactions and model-based perspective-taking within a shared environment. We implemented our model in WebPPL (Goodman & Stuhlmüller, 2014), a programming language that can express stochastic processes like POSGs as well as Bayesian inference.

Generative Model

To model a scenario like **Curtains**, we define agents, objects, and features assigned to them. These are the curtains, which have a location (inside Paul’s apartment); the speaker (Lisa), who has a location (inside or outside Paul’s apartment) and a perception of the curtains (good or bad); and the listener (Paul), who has a location (inside or outside) and ability to choose curtains (high or low). Additionally, the listener can either act on the curtains or not, while the speaker can enter

the apartment and make an utterance about the curtains (“the curtains look good”, “the curtains look bad”, or <nothing>). The truth-conditional semantics of the utterances map onto world features in a standard manner, and we set $\epsilon = .05$.

Observations depend on whether agents and objects are co-located and are defined as subsets of state and action variables. For instance, if Paul and Lisa are both inside the house and Paul modifies the curtains, they both observe that Paul acted on the curtains, but only Lisa directly knows whether they look good to her. Finally, we define a state and action prior for both agents such that the listener’s ability is initially high ($p = 0.90$), the speaker’s perception of the object is initially random ($p = 0.50$), and the listener has a low probability of modifying the object ($p = 0.05$).

Model Predictions

Given the generative model, we can provide scenarios and calculate aspects of the resulting interactive belief state (the listener and speaker’s beliefs about the world and each other’s

beliefs). In particular, we compare the results of a *shared history* with those of a *diverging history*. In the shared history, the speaker and listener are both present when the listener modifies the object, whereas in the diverging history, the speaker is not present when the listener acts on the object. Otherwise, the two scenarios are the same and the speaker comments on the curtains being bad. Figure 2a displays the results of the simulation when given each of the two histories. In both histories, the listener learns that their ability when modifying the object, X^{Eval} , is low (i.e., $\Delta_T^{\text{L-Eval}} < 0$). They also learn about the informative variable (i.e., $\Delta_T^{\text{L-Info}} > 0$).

However, the resulting interactive belief states differ in important ways. For example, in the diverging history, although the listener concludes that the evaluative variable is *low*, the speaker thinks the evaluative variable is *high*. Relatedly, the speaker thinks the utterance was informative ($\mathbb{E}_{\beta_S}[\Delta^{\text{L-Info}}] > 0$) but not offensive ($\mathbb{E}_{\beta_S}[\Delta^{\text{L-Eval}}] = 0$). Moreover, the listener knows the speaker believes that their comment was expected to be informative and not offensive. In the shared history, this is not the case: The listener and speaker both believe the evaluative variable is low, and they both know the resulting informational and evaluative effects. Because they were both present when the listener modified the object, they share expectations about the utterance’s meaning.

Put intuitively, whereas the shared history leads to an *expected insult*, the diverging history leads to a *faux pas*. Our model explains this difference in terms of differential transformations of the listener and speaker’s interactive belief state.

Experiment

Our model explains how different observation histories result in interactive belief states, which can produce unintentional meaning. To test whether this accurately describes people’s capacity to reason about unintentional meaning, we had people read vignettes that described scenarios involving shared or diverging observation histories. The underlying logical structure of all the vignettes mirrored that of **Curtains**, and so the model predictions described in the previous section apply to all of them. Participants then provided judgments corresponding to predicted differences in listener/speaker beliefs. The study’s main hypotheses were preregistered on the Open Science Framework platform (<https://osf.io/84wqn>). Overall, we find that our model captures key qualitative features of people’s inferences.

Materials

We developed a set of vignettes that included interactions in different contexts as well as different histories of interaction. Each vignette involved a listener (e.g., Paul) who could potentially interact with an object (e.g., curtains) as well as a speaker (e.g., Lisa) who makes an utterance about their negative aesthetic experience of the object (e.g., “The curtains look horrible”). In the shared history versions of the vignettes, the two agents were described as being both present when the listener acted on an object. In the diverging history

versions of the vignettes, the speaker was not present when the listener interacted with the object. Each vignette involved one of five contexts: *Curtain*, *Story-Prize*, *Wine-bottle*, *Cupcakes*, and *Parking*. Thus there were a total of ten items (Diverging/Shared history \times 5 contexts). All items used in the experiment are available on the primary author’s website.

Procedure

One-hundred participants were recruited via MTurk to participate in our experiment using PsiTurk (Gureckis et al., 2016). Each participant read one of the ten context-history items, and then answered the following questions in order:

- Q1: At this point, how does <listener> feel about their ability to <action>? [6 point scale ranging “Very Bad” to “Very Good” with no neutral option]
- Q2: <listener> thinks that <speaker> expected that their remark would make them feel <Q1_response>.
- Q3: <listener> thinks that in making the remark, <speaker> wanted to make them feel <Q1_response>.
- Q4: <listener> thinks that <speaker> thinks that <listener> <action>.
- Q5: <speaker> knew that <listener> <action>.
- Q6: In making the remark, <speaker> expected <listener> to feel <Q1_response>.
- Q7: In making the remark, <speaker> wanted <listener> to feel <Q1_response>.
- Q8: How awkward is this situation? [5 point scale ranging “Not at all” to “Extremely”]

The values for <listener>, <speaker>, and <action> were specified parametrically based on the context, while the value for <Q1_response> was filled in based on the answer to the first question. The response scale for questions 2-7 was a six-point scale ranging from “Definitely Not” to “Definitely”, with no neutral point. We included question 8 because previous work studying faux pas have focused on this question (Zalla et al., 2009). Participants were also given free response boxes to elaborate on their interpretation of the situation and answered demographic questions.

Question	β	S.E.	<i>df</i>	<i>t</i>	<i>p</i>
Q1	-0.06	0.07	94.0	-0.77	
Q2	0.15	0.06	94.0	2.65	**
Q3	0.15	0.06	94.0	2.50	*
Q4	0.18	0.06	94.0	2.78	**
Q5	0.25	0.06	94.0	4.34	***
Q6	0.14	0.06	94.0	2.53	*
Q7	0.15	0.06	94.0	2.64	**
Q8	0.04	0.05	94.0	0.78	

Table 1: Tests for Diverging/Shared history factor.

Experimental Results

Manipulation check To assess whether the Diverging/Shared history manipulation worked, we examined responses to Q5 (whether the speaker knew the listener acted on the object). A comparison in which the responses were coded

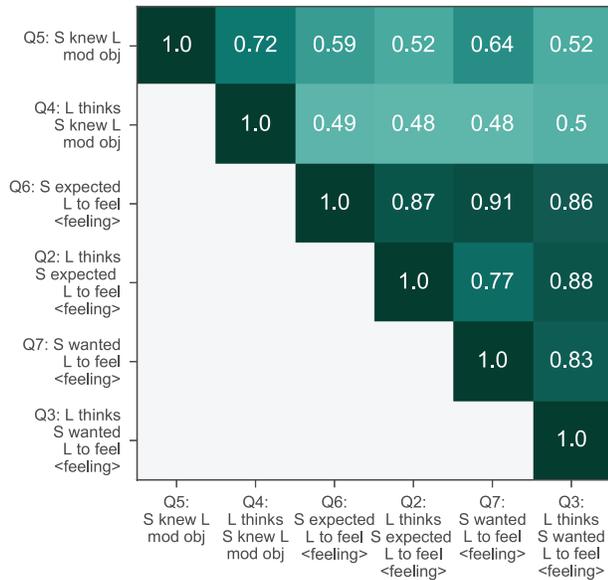


Figure 3: Judgment correlations (Pearson’s r).

as Yes or No (i.e., above or below the middle of the response scale) showed that it was effective ($\chi^2(1) = 7.92, p < .01$). However, a number of participants (15 of 50 in Shared; 20 of 50 in Diverging) did not pass this manipulation check and gave opposite answers than implied by the stories. Whether their responses are included does not affect our qualitative results, and in our analyses we use the full data set. Figure 2c plots the results for those who passed this check.

Judgment differences Responses paralleled the model predictions for the Shared versus Diverging history versions of the vignettes (Figure 2b). For each judgment, we fit mixed-effects linear models with context intercepts as a random effect and history as a fixed effect. Table shows tests of significance on the Diverging/Shared history parameters. Judgments about the listener’s feelings (Q1) were negative and not significantly different, indicating that people perceived the psychological impact (at least with respect to ability) of the utterance as roughly equivalent. In contrast, questions about the interactive belief state—the listener and speaker’s beliefs about the world and each other’s beliefs (Q2-Q7)—differed as predicted by the model. In particular, participants thought that the speaker neither expected that their utterance would hurt the listener’s feelings, nor that they wanted to do so. Participants judged that the listener recognized this as well.

Judgment correlations Judgments among questions about higher order mental states were strongly correlated, while those between the higher order mental states and the listener’s action were weaker (Figure 3). Specifically, those about speaker mental states (Q6, Q7) and listener beliefs about speaker mental states (Q2, Q3) were all highly correlated (all $r \in [0.77, 0.91], p < .001$). In contrast, questions about knowledge of the object being modified (Q4, Q5) were only moderately correlated with those about anticipated effects (Q2, Q3, Q6, Q7) (all $r \in [0.48, 0.64], p < .001$).

Discussion

People’s actions can have unexpected consequences, and speech-acts are no different. To understand unintentional meaning though, we need to characterize how a communicative act can lead to unanticipated epistemic consequences. Sometimes, a listener can learn something from an utterance that a speaker did not intend to convey or may not even believe (e.g., as in *Curtains*). Here, we have presented a Bayesian model and experiments testing how people reason about scenarios involving unintentional speech acts. Specifically, our account treats speech-acts as actions taken by a speaker that influence a shared interactive belief state—the beliefs each agent has about the world and each other’s beliefs. In doing so, we can capture the inferences that underlie unintentional meaning.

The current work raises important empirical and theoretical questions about how people reason about interactive beliefs and unintentional meaning. For instance, our experiments focus on third-party judgments about how a listener interprets the unintended meanings of utterances, but further work would be needed to assess how listeners do this (e.g., when the victim of an offhand comment) or even how speakers can recognize this (e.g., realizing one has put their foot in their mouth). Additionally, we have presented a Bayesian account of unintentional meaning in which agents reason about a large but finite set of possible histories of interaction. In everyday conversation, the space of possible histories can be much larger or even infinite. It is thus an open question how people can approximate the recursive inferences needed to make sense of unintentional meaning.

A rigorous characterization of unintentional meaning can deepen our understanding of how communication works in a broader social context. For example, attempts to build common ground through shared experience (Clark & Marshall, 1981; McKinley, Brown-Schmidt, & Benjamin, 2017) or manage face with polite speech (Levinson et al., 1987; Yoon et al., 2018) could be understood, in part, as strategies for forestalling unintentional meaning. And given that intentionality plays a key role in judgments of blame (Baird & Astington, 2004), phenomena like plausible deniability could be understood as people leveraging the possibility of unintentional meaning to covertly accomplish communicative goals (Pinker, Nowak, & Lee, 2008). Although further investigation is needed to test the extent to which people can track and influence interactive belief states (as well as how artificial agents can do so), this work provides a point of departure for computationally investigating these social and cognitive aspects of communication.

Acknowledgments

This material is based upon work supported by the NSF under Grant No. 1544924.

References

- Baird, J. A., & Astington, J. W. (2004). The role of mental state understanding in the development of moral cognition and moral action. *New Directions for Child and Adolescent Development*, 2004(103), 37-49. doi: 10.1002/cd.96
- Baron-Cohen, S., O'Riordan, M., Stone, V., Jones, R., & Plaisted, K. (1999). Recognition of Faux Pas by Normally Developing Children and Children with Asperger Syndrome or High-Functioning Autism. *Journal of Autism and Developmental Disorders*, 29(5), 407-418. doi: 10.1023/A:1023035012436
- Bernardo, J. M., & Smith, A. F. (1994). *Bayesian theory*. John Wiley & Sons.
- Brown-Schmidt, S., & Heller, D. (2018). Perspective-taking during conversation. In G. Gaskell & S. A. Rueschemeyer (Eds.), *Oxford handbook of psycholinguistics (2nd ed.)*. Oxford: Oxford University Press Oxford.
- Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*. Blackwell Publishing.
- Clark, H. H., & Marshall, C. R. (1981). Definite reference and mutual knowledge. *Elements of discourse understanding*.
- Franke, M. (2009). *Signal to act: Game theory in pragmatics*. Institute for Logic, Language and Computation.
- Gmytrasiewicz, P. J., & Doshi, P. (2005). A framework for sequential planning in multi-agent settings. *Journal of Artificial Intelligence Research*, 24, 49-79.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, 20(11), 818-829.
- Goodman, N. D., & Stuhlmüller, A. (2014). *The Design and Implementation of Probabilistic Programming Languages*. <http://dippl.org>. (Accessed: 2018-9-12)
- Grice, H. P. (1957). Meaning. *The philosophical review*, 66(3), 377-388.
- Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., ... Chan, P. (2016). psiTurk: An open-source framework for conducting replicable behavioral experiments online. *Behavior research methods*, 48(3), 829-842.
- Korman, J., Zalla, T., & Malle, B. F. (2017). Action understanding in high-functioning autism: The faux pas task revisited. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th annual conference of the cognitive science society* (p. 2451-2456). Austin, TX: Cognitive Science Society.
- Kuhn, H. (1953). Extensive games and the problem of information. In H. Kuhn & A. Tucker (Eds.), *Contributions to the theory of games II* (pp. 193-216). Princeton University Press.
- Levinson, P., Brown, P., Levinson, S. C., & Levinson, S. C. (1987). *Politeness: Some universals in language usage* (Vol. 4). Cambridge university press.
- Luce, R. D. (1959). On the possible psychophysical laws. *Psychological review*, 66(2), 81.
- McKinley, G., Brown-Schmidt, S., & Benjamin, A. (2017). Memory for conversation and the development of common ground. *Memory and Cognition*, 45(8), 1281-1294. doi: <https://doi.org/10.3758/s13421-017-0730-3>
- Pinker, S., Nowak, M. A., & Lee, J. J. (2008). The logic of indirect speech. *Proceedings of the National Academy of Sciences*, 105(3), 833-838. doi: 10.1073/pnas.0707192105
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Cambridge, MA, USA: Harvard University Press.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT press.
- Tendahl, M., & Gibbs Jr, R. W. (2008). Complementary perspectives on metaphor: Cognitive linguistics and relevance theory. *Journal of pragmatics*, 40(11), 1823-1864.
- Yoon, E. J., Frank, M. C., Tessler, M. H., & Goodman, N. D. (2018, Dec). *Polite speech emerges from competing social goals*. PsyArXiv. Retrieved from psyarxiv.com/67ne8 doi: 10.31234/osf.io/67ne8
- Zalla, T., Sav, A.-M., Stopin, A., Ahade, S., & Leboyer, M. (2009, Feb 01). Faux pas detection and intentional action in asperger syndrome. a replication on a french sample. *Journal of Autism and Developmental Disorders*, 39(2), 373-382. doi: 10.1007/s10803-008-0634-y