A Pragmatic Model of Spatial Language

Midas Vanooteghem (midas.vanooteghem@kuleuven.be)
Department of Psychology, University of Leuven, Belgium, Tiensestraat 102
B-3000 Leuven, Belgium

Walter Schaeken (walter.schaeken@kuleuven.be)
Department of Psychology, University of Leuven, Belgium, Tiensestraat 102
B-3000 Leuven, Belgium

Abstract

Spatial language is an integral part of everyday communication, but existing theories fail to explain the cognitive processing underlying the interpretation of spatial descriptions. The traditional spatial template theory does not attempt to provide a mechanism behind the computation of acceptability ratings, nor does it adequately address the effect of distractors or the conversational context. In this study, we propose a new model of *spatial pragmatics*, based on the Rational Speech Act (RSA) framework. This Bayesian model shows how the patterns of spatial language understanding follow directly from the principles of pragmatics. Our model accounts for previously found effects of angle, distance, and distractors on acceptability ratings and indication behavior. We test these predictions in an experiment consisting of various tasks. The results largely support our model's predictions, suggesting that pragmatic reasoning might play a role in spatial language use.

Keywords: Bayesian modeling; pragmatics; spatial language; acceptability.

Introduction

In everyday communication, we frequently describe the location of one object relative to another. For instance, at a party, one might tell a friend, "There's a tasty snack to the left of the bottle of soda." Here, the snack is the *located object* and the bottle of soda is the *reference object* (Carlson & Logan, 2001).

In some cases, the location of the located object can be described using multiple different directional terms. For example, the snack could be described as "to the left of," "behind," or even "to the left of and behind" the bottle of soda. Research has shown that the acceptability of different spatial descriptions depends on *angular deviation*: the angle between the line connecting the located object with the reference object and the ideal direction implied by the direction word (e.g. horizontally to the left for "to the left"). Larger deviations from the ideal angle consistently lead to lower acceptability ratings (Crawford, Regier, & Huttenlocher, 2000; Gapp, 1995; Hayward & Tarr, 1995; Logan & Sadler, 1996; Regier & Carlson, 2001). Interestingly, when objects are arranged in grids, acceptability ratings for objects with an ideal angle are higher than what would be expected based on a linear effect of angle alone. Similar patterns were found in tasks where participants judged spatial utterances as either "true" or "false", with smaller angles corresponding to more "true" responses and faster reaction times.

In most everyday scenes, more than two objects are present. Objects that are near the reference or located object but not mentioned in the spatial description are referred to as *distractors*. Carlson and Logan (2001) found that the distractors lower acceptability ratings and slow reaction times. Without distractors, the distance between the located and reference object generally does not affect acceptability (Gapp, 1995; Kojima & Kusumi, 2006), but in the presence of distractors, greater distances lead to lower ratings (Carlson & Logan, 2001; Kojima & Kusumi, 2006). Finally, when participants were asked to predict an object's location based on a spatial description, they tended to choose positions that were both close to the reference object and aligned with the ideal angle implied by the directional term (Logan & Sadler, 1996).

The prevailing explanation for patterns in acceptability ratings could be called the *spatial template theory*, which posits that each spatial term has its own template. These templates are typically divided into three regions: *good* (ideal angle), *acceptable* (angles between 0° and 90°), and *bad* (angles greater than 90°; Kojima & Kusumi, 2006; Logan & Compton, 1996; Logan & Sadler, 1996; Regier & Carlson, 2001). According to this theory, people evaluate acceptability, by aligning the spatial template with the reference object and determining in which region the located object lies.

Logan and Compton (1996) added further nuance to this theory to account for the heterogeneity of their data. They proposed that the boundary between the good and bad regions is sharper than the boundary between the good and acceptable regions and that even within the acceptable region, smaller angular deviations are rated as more acceptable than larger ones.

Building on Herskovits (1985), Logan and Sadler (1996) extended the spatial template theory to account for distractor effects. They argued that when a distractor is present, people attempt to fit both the located object and the distractor into the same spatial template centered on the reference object. Acceptability ratings decrease when the distractor provides a better example of the spatial relation than the located object.

We argue that the spatial template theory has several important limitations. First, it does not explain the origin or rationale for the template's shape. For instance, the sharp boundary between the regions appears arbitrary.

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

Second, although the theory assumes that more acceptable locations lead to higher ratings, it does not specify how acceptability or production probabilities are computed from the spatial template.

Third, Carlson and Logan (2001) found the theory's explanation of distractor effects—based on comparing the located object and distractor—inadequate. As an alternative explanation, we suggest instead that distractors lower acceptability because they serve as better reference objects than the reference object in the sentence.

Fourth, spatial template theory does not explain how listeners use spatial descriptions to infer object locations.

Finally, the theory overlooks the conversational context in which spatial language is used, including the reasoning that speakers and listeners engage in regarding each other's intentions and interpretations. This aspect of communication, known as pragmatics, is central to the model we propose in this paper.

In this paper, we propose a model of spatial language grounded in the Rational Speech Act (RSA) framework (Frank & Goodman, 2012) and test its predictions in an online experiment. RSA is a Bayesian theory of pragmatics that models how speakers and listeners reason about each other's intentions when producing and interpreting utterances. While RSA has previously been applied to spatial terms like 'near' (Ullman, Xu, & Goodman, 2016) and 'inside' (Carstensen, Kon, & Regier, 2014), it has not yet been applied to directional terms.

Our model addresses several limitations of the spatial template theory. It explains how spatial templates are derived, accounts for the sharp boundary between the good and acceptable regions, derives acceptability ratings from production probabilities, models listeners' inference, and explains why distractors affect ratings.

In the following sections, we introduce our experimental design, present the model, derive hypotheses, and evaluate these hypotheses using data from our experiment by comparing multiple linear regression models.

Experiment

To evaluate our model, we designed a novel experiment comprising four distinct tasks: (1) the Acceptability Rating Task, in which participants provided acceptability ratings for spatial descriptions, (2) the Verification Task, in which participants verified spatial descriptions against given scenarios, (3) the Simple Indication Task, wherein participants identified the most likely location of an object based on a spatial description, and (4) the Multiple Indication Task, wherein participants assessed the likelihood of various locations being the located object's position. We hypothesize that this experiment will replicate findings from the literature regarding the effects of angle, ideal angle, distance, and distractors on acceptability judgments, reaction times, proportions of true responses, and location-indication behaviors.

Participants

Thirty participants (10 for prior elicitation and 20 for hypothesis testing and model fitting) were recruited for this experiment. Participants were drafted from the native English-speaking population with an age between 18 and 35 years on the recruitment platform Prolific. The mean age in our sample was 28.63 years. We used quota sampling to ensure that we had an equal number of men and women in our sample. All participants provided informed consent prior to participation and were compensated £7 for their time.

Materials

The experiment used an invisible 7x7 grid on a grey background as the visual stimulus, with the reference object positioned in the center of the grid. Objects were represented as black silhouettes in white circles. The scenes varied in angle, distance, and the presence or absence of a distractor. Spatial relations between objects were described using spatial descriptions that varied systematically by direction (2 horizontal, 2 vertical, and 4 diagonal directions) and objects.[1]

Design

The study employed a within-subject design, with each participant completing all four tasks. In the Acceptability Rating Task (236 trials), participants rated the applicability of a spatial description on a continuous scale (from 0 to 1). In the Verification Task (236 trials), participants judged the truth of a proposition using a binary scale ('true'/'false'). In the Simple Indication Task (16 trials), participants indicated the most likely location of the located object based on the spatial description by clicking on the spot. In the Multiple Indication Task (16 trials), participants indicated the likelihood of each grid square being the object's location by clicking each square a number of times proportional to their estimation of the likelihood. With each click, the square became a darker shade of gray, indicating a larger likelihood. Each task included trials with and without a distractor placed in the ideal location for the described spatial relation. The task order was partially counterbalanced: the two indication tasks were always presented first, followed by the verification task and the acceptability rating task.

Procedure

The experiment had a median completion time of 1 hour and 2 minutes. Each trial began with a 1-second fixation cross, followed by a 2-second presentation of a proposition describing the spatial relation between the located and reference object. Afterward, participants responded on a screen that remained visible until they provided an answer. If participants forgot the spatial description, they could request to view it again.

[1] Instructions and examples of the stimuli are available on OSF: https://osf.io/s9z8f/

To maintain concentration, breaks were offered after every 100 trials. In the verification task, a prompt appeared if participants responded the same way in over 95% of trials for at least 20 consecutive trials, asking whether they intended to give identical responses for similar trials (what constituted similar trials was shown on the screen). If confirmed, such trials were skipped.

A practice trial was administered prior to each task to familiarize participants with the procedure. At the end of the experiment, participants were asked follow-up questions regarding their responses (e.g., whether they answered automatically) and their sense of orientation. Participants were also encouraged to provide any additional remarks that might help interpret their responses.

The Spatial RSA Model

We propose that a pragmatic reasoning process underlies the participants' responses in our experiment. Specifically, we suggest that when judging the acceptability of an utterance or verifying its truth, participants take the perspective of a *pragmatic speaker*, evaluating the utterance based on how likely they would be to produce it themselves. In contrast, when completing the indication tasks, participants adopt the role of a *pragmatic listener*, reasoning about the intended meaning behind the utterance.

According to the RSA framework, which we extend in our model, meaning is shaped by semantics as well as reasoning (Scontras, Tessler, & Franke, 2021). When a pragmatic speaker chooses an utterance to communicate a state of the world, they consider what a *literal listener* is likely to infer. The literal listener's understanding is influenced by the semantic meaning of the utterance as well as the context, including knowledge about which states of the world are more likely. Conversely, a pragmatic listener could in turn reason about how the pragmatic speaker likely would have expected that the literal listener would interpret the utterance. This recursive reasoning process between speakers and listeners is formalized within the Bayesian RSA framework, which defines the probabilities of a speaker selecting a particular utterance and the probabilities of a listener interpreting it as a specific state of the world.

The Spatial Template

Our model maintains the assumption of the spatial template theory that spatial templates underlie all aspects of spatial language use. However, we propose that these templates are not directly used to generate responses. Instead, a series of intermediate processing steps occur before responses are generated.

Moreover, our model differs from the traditional spatial template theory in two important ways: the template's semantic interpretation and its shape. Specifically, we take inspiration from the RSA model of quantifiers by van Tiel, Franke, and Sauerland (2021) and argue that the spatial template represents a prototype semantics, where the degree of truth of a spatial description depends on its similarity to a prototype. For instance, in the case of directional terms, the prototype corresponds to a spatial relation where the located object forms an ideal angle with the reference object—e.g., a horizontally left position would be the prototype for the term 'to the left'. The closer a scene aligns with the prototype (i.e., the smaller the angular deviation), the higher its degree of truth.

We formalize this relation between the degree of truth $[[u]](s)$ of a spatial description (or *utterance*) u and a location of the located objects (or *state*) s, and an angle $\angle(s;u)$ using a Gaussian function as proposed by van Tiel et al. (2021):

$$[[u]](s) = \exp\left(-\left(\frac{\angle(s;u)}{\beta_i}\right)^2\right) \qquad (1)$$

Here, β_i is a participant-specific parameter with values ranging from 0 to $+\infty$, which determines the slope of the relation. Lower values of β_i correspond to steeper slopes, indicating a sharper decline in the truth value as the angle deviates from the prototype. This Gaussian function differs significantly from the sharp boundary between "acceptable" and "good" regions proposed in traditional spatial template theory. Rather than discrete zones, the function provides a gradual variation in truth value based on proximity to the prototype (Figure 1, left).

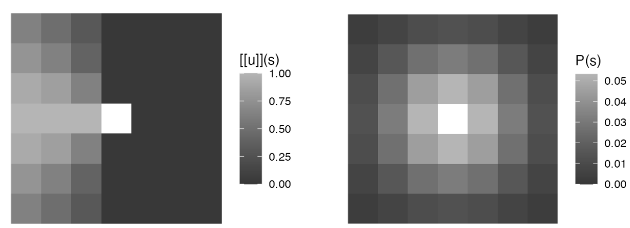

Figure 1: The spatial template for 'to the left' (left) and the state prior for a reference object in the middle of the display (right).

The State Prior

A core principle of the Rational Speech Act (RSA) framework is that the likelihood of the located object's position (referred to as *states* in RSA terminology) is influenced by prior expectations, which listeners consider during interpretation. This likelihood is represented by the **state prior**. Under certain assumptions, it can be demonstrated that objects are more likely to be located near the reference object rather than farther away. If an object were farther, the speaker would likely have chosen a closer reference object to enhance informativity.[2] This expectation

[2] This can be proven in the RSA framework with a uniform state prior and possible reference objects in all locations of the grid. See OSF.

is captured by the state prior $P(s)$, which is modeled as a Gaussian function of the distance from the reference object $\delta(s)$:

$$P(s) \propto \exp\left(-\left(\frac{\delta(s)}{\gamma_i}\right)^2\right) \tag{2}$$

Here, γ_i is a participant-specific parameter determining the steepness of the function. The resulting values are normalized so that their sum equals 1. As illustrated in Figure 1 (right), the probabilities for states where the located object is closer to the reference object are higher.

The Literal Listener

The model defines a **literal listener** L_0, who interprets utterances based only on their literal meanings and the prior probabilities of states. For this listener, the probability of a state s given an utterance u, $P_{L_0}(s \mid u)$, is proportional to the product of the spatial template of the utterance $[[u]](s)$ and the state prior $P(s)$:

$$P_{L_0}(s \mid u) \propto [[u]](s) \cdot P(s) \tag{3}$$

In Figure 2 (upper left), we observe that the highest probability is associated with the ideal spot, directly to the left of the reference object, with probabilities decreasing as the distance from this spot increases. This pattern arises from the combined effect of angle, as determined by the spatial template, and the distance from the reference object, as determined by the state prior.

The Pragmatic Speaker

The model incorporates a **pragmatic speaker** S_1, who behaves as an approximately rational actor. This means that the probability of the speaker using a particular utterance is proportional to the utility of that utterance $U_{S_1}(u;s)$ in comparison to other possible utterances. Specifically, the probability of choosing an utterance given a state, $P_{S_1}(u \mid s)$, is determined by a softmax function of the utterance's scaled utility:

$$P_{S_1}(u \mid s) \propto \exp\left(\alpha \cdot U_{S_1}(u;s)\right) \tag{4}$$

In this equation, α is the **rationality parameter**. When $\alpha = 0$, the pragmatic speaker selects an utterance at random, without regard for utility. As α approaches infinity, the speaker always chooses the utterance with the highest utility. Importantly, the probabilities of all utterances given a state sum to one.

The **utility** of an utterance in a specific state, $U_{S_1}(u;s)$, is defined as the difference between the logarithm of the literal listener's probability for the state given the utterance and the cost of the utterance $C(u)$:

$$U_{S_1}(u;s) = \log\left(P_{L_0}(s \mid u)\right) - C(u) \tag{5}$$

This formulation reflects two important factors: utterances that are longer (i.e. more costly) or those that are less likely to be understood by the literal listener (i.e. having a lower $P_{L_0}(s \mid u)$), are less likely to be selected by the pragmatic speaker.

In the lower-left panel of Figure 2, the pragmatic speaker probabilities are shown for the square with coordinates [-3,2]. This square is located directly below the square in the upper-left corner of the grid. The chart demonstrates that the utterance 'above and to the left' has the highest probability, followed by 'to the left' and then 'above'. This pattern makes intuitive sense: the square is closest to the diagonal prototypical angle and closer to the left prototypical angle than to the above prototypical angle.

The upper-right panel of Figure 2 visualizes the pragmatic speaker probabilities for the utterance 'to the left' across all states. For squares along the horizontal line directly to the left of the reference square, the probability of saying 'to the left' is 1, as this is the only valid utterance. However, as the angle deviates from the prototypical angle, the probability decreases.

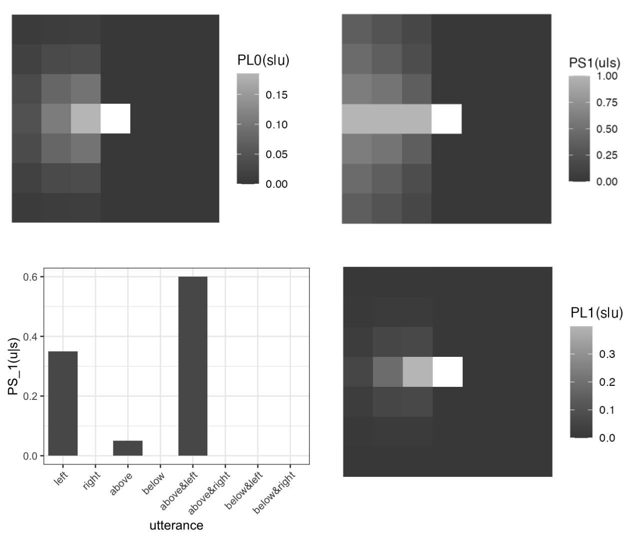

Figure 2: (Upper left) Literal listener probabilities for 'to the left' for all states. (Upper right) Pragmatic speaker probabilities for 'to the left' for all states. (Lower left) Bar chart of pragmatic speaker probabilities for different utterances for the square with coordinates [−3,2]. (Lower right) Pragmatic listener probabilities for 'to the left' for all states.

Similar to the distinction between the good and acceptable region in the spatial template theory, pragmatic speaker probabilities for the cardinal directions exhibit an ideal angle effect—a pattern consistent with acceptability ratings in prior studies. This effect arises because, for squares aligned with the ideal angle, only a single direction term is applicable. As a result, the pragmatic speaker assigns that utterance a probability of one. In contrast, for diagonal directions, three utterances are always applicable, and the ideal angle effect does not appear. Since Degen and Goodman (2014) hypothesized that the pragmatic speaker probabilities underlie participants' ratings in acceptability tasks, the ideal angle effect in the pragmatic speaker probabilities could be the explanation for the same effect in the acceptability ratings.

The Pragmatic Listener

The model also incorporates a **pragmatic listener** L_1, who interprets utterances by considering the probabilities assigned by the pragmatic speaker ($P_{S_1}(u|s)$) and the state prior ($P(s)$):

$$P_{L_1}(s|u) \propto P_{S_1}(u|s) \cdot P(s)$$

The pragmatic listener's probabilities, as visualized in Figure 2 (lower right)**,** are influenced by both the angle and distance between the located object and the reference object. Compared to the literal listener probabilities shown in Figure 2 (upper left), the pragmatic listener probabilities exhibit an even stronger distance effect. This occurs because the state prior $P(s)$ is applied again during the calculation of $P_{L_1}(s|u)$, which reinforces the impact of proximity to the reference object.

Previous research supports the relevance of the pragmatic listener probabilities in predicting indication behavior. For instance, Degen and Goodman (2014) suggested that participants' behavior in indication tasks is derived from the pragmatic listener probabilities. This agrees with findings from Logan and Sadler (1996), who observed that both angle and distance significantly influenced participants' behavior in their simple indication task.

The Distractor

In some trials, we included a distractor: an additional, irrelevant object placed in an ideal location—next to the reference object and at an ideal angle. We assume that in situations with a distractor, speakers may also consider utterances where the distractor functions as the reference object.

The introduction of a distractor has a straightforward effect on the literal listener probabilities. When the location that would normally have the highest probability is occupied by the distractor, the probabilities for surrounding locations increase because the total probabilities must still sum to 1.

However, the effect on the pragmatic speaker probabilities is more complex. In cases where a distractor is present, we observe an emerging effect of distance that was previously absent. As the lower panel of Figure 3 illustrates, locations with a perfect angle no longer have a probability of 1, as the utterance 'to the left' is no longer the only viable option. For instance, the probability of saying 'to the left' decreases as the distance increases (e.g., $P_{S_1}(\text{left} \mid [-3{,}0]) < P_{S_1}(\text{left} \mid [-2{,}0])$). At the same time, the probability of using the distractor as the reference object (e.g., 'left of the distractor') increases.[3] Since our model assumes that acceptability ratings are derived from pragmatic speaker probabilities, we predict that this distance effect should also appear in participant ratings. Additionally, the presence of the extra alternative utterances leads to lower overall pragmatic speaker probabilities (and thus acceptability ratings) and to the absence of an ideal angle effect.

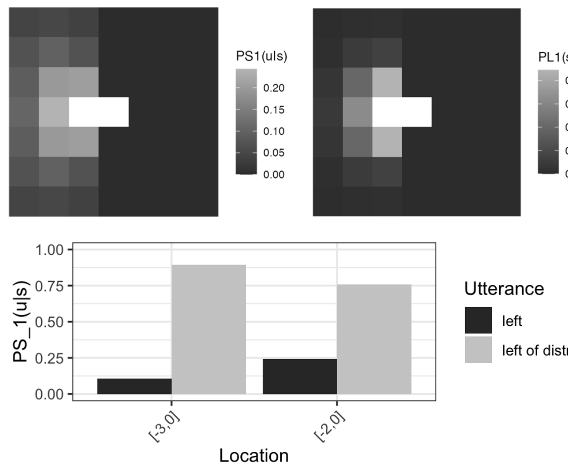

Figure 3: (Upper left) Pragmatic speaker probabilities for 'to the left' for all states when a distractor is present. (Upper right) Pragmatic listener probabilities for 'to the left' for all states when a distractor is present. (Lower panel) Pragmatic speaker probabilities for 'to the left' and 'to the left of the distractor' for the states with location [-3,0] and [-2,0].

For the pragmatic listener, the presence of the distractor shifts the probabilities to certain locations, such as squares directly above or below the distractor, as shown in the upper-right panel of Figure 3. However, the shape of this pattern is highly dependent on all the parameters of the model as well as the location of the distractor.

Results

We used a Bayesian approach for data analysis, first fitting a multiple linear regression model containing parameters for all possible effects and then calculating Bayes Factors using the Savage-Dickey method (Wagenmakers et al., 2010). This approach requires specifying realistic prior distributions, which we estimated based on data collected from a separate group of 10 participants.[4]

The analysis revealed strong evidence supporting the influence of **angle** on acceptability scores across various conditions. Specifically, for the cardinal directions without a

[3] The reason for this is that due to the sigmoidal shape of the state prior, the ratio of the literal listener probabilities for 'to the left of the reference object' as compared to 'to the left of the distractor' is higher for the location [-2,0] than for [-3,0] : $\frac{P_{L_0}(\text{left} \mid [-2{,}0])}{P_{L_0}(\text{left of distr} \mid [-2{,}0])} > \frac{P_{L_0}(\text{left} \mid [-3{,}0])}{P_{L_0}(\text{left of distr} \mid [-3{,}0])}$

[4] Parameter estimates are available on OSF.

distractor, the Bayes factor indicated *extreme evidence in favor* of an effect (BF = 686.52), and it remained *extreme* with a distractor (BF = 444.41). For diagonal directions, the evidence was *very strong without a distractor* (BF = 32.04) and *moderate* with a distractor (BF = 5.83).

Regarding the effect of **distance**, the evidence provided some support for our hypothesis. For cardinal directions without a distractor, there was *moderate evidence against* an effect of distance (BF = 1/9.86), and *moderate evidence for* an effect when a distractor was present (BF = 7.98). In diagonal directions, the *evidence against the effect of distance was strong* without a distractor (BF = 1/11.38) and *anecdotal* with a distractor (BF = 1/2.67).

The results confirm our hypotheses about the effect of **ideal angle** (i.e. higher scores for prototypical directions than could be explained by a linear effect of angle). For cardinal directions without a distractor, *moderate evidence supported an effect* (BF = 5.81), whereas there was, as we expected, *moderate evidence against* it with a distractor (BF = 1/5.56). In diagonal directions, the *evidence against an effect was moderate* without a distractor (BF = 1/4.76) and *anecdotal* with a distractor (BF = 1/2.83), which confirms our hypotheses.

We found consistently strong evidence for the negative impact of the **distractor** on acceptability scores, with *extreme evidence* for cardinal directions (BF = 248.24) and *very strong evidence* for diagonal directions (BF = 80.66).

In the **verification task**, nearly all participants consistently responded with "true" on almost all trials, making further analysis impossible.

In the **simple indication task**, participants strongly preferred squares that were as close as possible to the reference object and aligned with the ideal angle. When participants selected squares that were not the closest to the reference object, they mostly adhered to the ideal angle. Conversely, when participants chose squares that were not on the ideal angle, these squares tended to be close to the reference object.

Responses in the **multiple indication task** showed significant variability across participants and even within the same participant across different trials. While certain response patterns recurred, they were difficult to interpret. This unintelligibility may stem from participants' lack of familiarity with this type of task or behavior.

Discussion

Our research provides evidence supporting the hypotheses derived from our model. We observed significant effects of angle and the presence of a distractor. More importantly, we also found effects of ideal angle for the cardinal directions without distractor and an interaction between distance and distractor. These latter effects were previously reported in the literature but lacked an explanation, which we believe our model offers.

Although these findings align with our hypotheses, the relatively limited evidence for the effects of ideal angle and the interaction between distractor and distance raises questions about whether these effects are as strong as our model predicts. Furthermore, visual inspection (plots are available on OSF*)* of the data reveals considerable individual variation in participants' response patterns. While some participants' responses closely align with our hypotheses, others did not exhibit effects of angle, ideal angle, and/or the interaction between distance and distractor. However, even these patterns might be explained within the spatial RSA model under specific parameter configurations. This does not necessarily imply that our model is so flexible that it cannot be falsified, but rather that the response patterns it predicts might not be fully captured by the simplified effects tested in this study. A more conclusive evaluation of the model's explanatory power would require comparing the fit of multiple competing models.

We also found that the evidence was consistently weaker for diagonal directions. This may be due to a smaller amount of data available for these hypotheses. While each cardinal term was tested with 21 squares, only 9 squares were available for each diagonal term. Another explanation could be that, being compositional directions, they might not have clear prototypes.

It is important to note that our model assumes that speakers consider only directional utterances. Of course, other spatial expressions—such as 'close to'—may also be applicable. However, in our experiment, we aimed to reduce the salience of such alternatives by including only utterances with directional terms in the trial spatial descriptions.

Our theory accounts for the effect of the distractor and its interaction with distance on acceptability ratings by proposing that individuals compute pragmatic speaker probabilities for the located object in reference to the distractor when interpreting the relationship between the located object and the reference object. To further investigate this, we plan a follow-up study examining whether these effects depend on the position and number of distractors, as follows from our theory. Additionally, we will ask participants to produce utterances, as these responses are more directly linked to pragmatic speaker probabilities than acceptability ratings.

Conclusion

In this paper, we introduced a pragmatic model of spatial language, which addresses some limitations of the spatial template theory. Our model can explain previously found effects of angle, ideal angle, distance, and distractors on spatial language evaluation. Additionally, hypotheses derived from our model were largely supported by the experimental data.

However, the patterns observed in our data appear more complex than can be captured by simplified model comparisons. These patterns can only be meaningfully evaluated by fitting the model directly to individual-level data, which we will do in future research. Additionally, we will extend this research in an experiment that focuses on spatial language production in displays with multiple objects.

References

Carlson, L. A., & Logan, G. D. (2001). Using spatial terms to select an object. *Memory & Cognition, 29*(6), 883–892.

Carstensen, A., Kon, E., & Regier, T. (2014). Testing a rational account of pragmatic reasoning: The case of spatial language. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society* (pp. 2009-2013).

Crawford, L. E., Regier, T., & Huttenlocher, J. (2000). Linguistic and non-linguistic spatial categorization. *Cognition, 75*(3), 209–235. https://doi.org/10.1016/S0010-0277(00)00064-0

Degen, J., & Goodman, N. (2014). Lost your marbles? The puzzle of dependent measures in experimental pragmatics. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society* (pp. 397-402).

Gapp, K. P. (1995). Angle, distance, shape, and their relationship to projective relations. *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 112–117).

Hayward, W. G., & Tarr, M. J. (1995). Spatial language and spatial representation. *Cognition, 55*(1), 39–84. https://doi.org/10.1016/0010-0277(94)00643-Y

Herskovits, A. (1985). Semantics and pragmatics of locative expressions. *Cognitive Science, 9*(3), 341–378.

Kojima, T., & Kusumi, T. (2006). The effect of an extra object on the linguistic apprehension of the spatial relationship between two objects. *Spatial Cognition & Computation, 6*(2), 145–160. https://doi.org/10.1207/s15427633scc0602_2

Logan, G. D., & Compton, B. J. (1996). Distance and distraction effects in the apprehension of spatial relations. *Journal of Experimental Psychology: Human Perception and Performance, 22*(1), 159–172.

Logan, G. D., & Sadler, D. D. (1996). A computational analysis of the apprehension of spatial relations. In P. Bloom, M. F. Garrett, L. Nadel, & M. A. Peterson (Eds.), *Language and space*. MIT Press.

Regier, T., & Carlson, L. A. (2001). Grounding spatial language in perception: An empirical and computational investigation. *Journal of Experimental Psychology: General, 130*(2), 273–298.

Scontras, G., Tessler, M. H., & Franke, M. (2021). A practical introduction to the Rational Speech Act modeling framework. *arXiv*. https://arxiv.org/abs/2105.09867

Ullman, T., Xu, Y., & Goodman, N. D. (2016). The pragmatics of spatial language. *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (pp. 1613-1618).

van Tiel, B., Franke, M., & Sauerland, U. (2021). Probabilistic pragmatics explains gradience and focality in natural language quantification. *Proceedings of the National Academy of Sciences, 118*(9). https://doi.org/10.1073/pnas.2005453118

Wagenmakers, E. J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method. *Cognitive Psychology, 60*(3), 158–189. https://doi.org/10.1016/j.cogpsych.2009.12.001