Pronominal Reference and Pragmatic Enrichment: A Bayesian Account

Andrew Kehler (akehler@ucsd.edu)
Department of Linguistics, University of California, San Diego
9500 Gilman Drive, La Jolla, CA, 92093-0108 USA

Hannah Rohde (hannah.rohde@ed.ac.uk)
Department of Linguistics and English Language, University of Edinburgh
3 Charles Street, Edinburgh, EH8 9AD, UK

Abstract

A standard assumption in linguistic, psycholinguistic, and computational research on pronoun use is that production and interpretation are guided by the same set of contextual factors. Kehler et al. (2008) and Kehler & Rohde (2013) have argued instead for a Bayesian model, one in which pronoun production is insensitive to a class of semantically- and pragmatically-driven contextual biases that have been shown to influence pronoun interpretation. Here we evaluate the model using a passage completion study that employs a subtle contextual manipulation to which traditional analyses are insensitive, specifically by varying whether or not a relative clause that modifies the direct object in the context sentence invites the inference of a cause of the event that the sentence denotes. The results support the claim that pronoun interpretation biases, but not production biases, are sensitive to this pragmatic enrichment, revealing precisely the asymmetry predicted by the Bayesian analysis. A correlation analysis further establishes that the model provides better estimates of measured pronoun interpretation biases than two competing models from the literature.

Keywords: Pronoun interpretation, discourse coherence, pragmatic enrichment, Bayesian models

Introduction

A common wisdom, one assumed in the literature on pronominal reference for decades, is that there is a unified notion of entity SALIENCE that underlies pronoun usage. This notion of salience (alternately referred to as PROMINENCE, being in PSYCHOLOGICAL FOCUS, being THE CENTER OF ATTENTION, and so forth) will determine when a speaker will choose to use a pronoun on the one hand, and hence be used by the addressee to successfully interpret the reference on the other. On this model, the task for discourse researchers is then clear: One merely needs to identify the different factors that contribute to entity salience. Many such factors have been posited, including grammatical role (Crawley, Stevenson, & Kleinman, 1990), grammatical parallelism (Smyth, 1994), thematic role (Stevenson, Crawley, & Kleinman, 1994; Arnold, 2001), information structure (Grosz, Joshi, & Weinstein, 1995; Rohde & Kehler, 2014), semantics (Koornneef & van Berkum, 2006), and world knowledge (Hobbs, 1979), among others.

A central goal of this paper is to disabuse the reader of this model. We do this by evaluating a proposal put forth by Kehler, Kertz, Rohde, and Elman (2008) and Kehler and Rohde (2013) that states that the relationship between pronoun interpretation and production is suitably cast in Bayesian terms, and further that the types of factors that condition the likelihood term (the production bias) and the prior (the bias toward entity next mention) are different. This entails the counterintuitive conclusion that a set of factors that the addressee will use in resolving the referent will not be taken into account by the speaker when deciding whether to use a pronoun. We evaluate the causal structure that underlies the proposal with a novel passage completion experiment which confirms the predictions of the analysis. We further show that the biases revealed by the experimental results are more highly correlated with the predictions of the Bayesian analysis than those of two other prominent models.

Background

The experiment described in this paper utilizes so-called IMPLICIT CAUSALITY (IC) verbs, so let us start there. IC verbs are undoubtedly the most well-studied verb class in the psycholinguistics of pronoun interpretation literature since the seminal papers of Caramazza and colleagues in the 1970s (Garvey & Caramazza, 1974; Caramazza, Grober, Garvey, & Yates, 1977; Brown & Fish, 1983; Terry Kit-fong Au, 1986; McKoon, Greene, & Ratcliff, 1993; Koornneef & van Berkum, 2006; Kehler et al., 2008, inter alia). Here we survey three facts that will be important for understanding the experimental design we employ. First, IC verbs are characteristically known for their large and divergent biases toward subsequent entity mention. If participants in a passage completion task are presented with a prompt like (1a),

(1) a. Amanda amazed Brittany because she _____
     b. Amanda detested Brittany because she _____

the large majority of completions will point to Amanda as the pronominal referent. After all, Amanda must be amazing, and hence one expects to hear why. Because causality is imputed to the subject, verbs like *amaze* are called SUBJECT-BIASED IC VERBS. If participants are given a prompt like (1b), on the other hand, the large majority of completions will point to Brittany as the pronominal referent. After all, Brittany must be detestable, and hence one expects to hear why. Because causality is imputed to the object, verbs like *detest* are called OBJECT-BIASED IC VERBS. The existence of IC biases has been replicated repeatedly, and is hence one of the bedrock results in the field.

The second fact of interest is that IC verbs are associated with another type of strong bias, in that they give rise to a greater expectation that the ensuing clause will provide a cause or reason for the described eventuality as compared to

typical non-IC verbs (Kehler et al., 2008). Whereas most previous IC studies examined pronoun biases using prompts with *because* as in (1a-b) above, Kehler et al. used free prompts as in (2) and had annotators categorize the COHERENCE RELATIONS – that is, the relevancy relations that semantically bind the clauses – in the completed passages.

(2) a. John amazed Mary. _____
    b. John detested Mary. _____
    c. John saw Mary. _____

The results showed that participants completed passages such as (2a-b) using Explanation relations – i.e., in which an event or state is followed by a cause or reason – approximately 60% of the time, compared to 24% for a control group of non-IC contexts (2c). This result accords with intuitions: Upon hearing *John amazed Mary*, it seems likely that the addressee will wonder *Why?*, and thus expect to hear an answer. On the other hand, upon hearing *John saw Mary*, it seems less likely that the addressee will wonder *Why?*, and instead expect an answer to the question *What happened next?*, for example.

The third and final fact to be highlighted is that the occurrence of a pronoun alters referential biases from what they were before the pronoun, specifically toward the entity mentioned from the subject position of the previous clause. Stevenson et al. (1994) reported on a series of passage completion experiments that investigated pronoun biases across eight distinct context types, comparing conditions in which a pronoun was or was not included in the prompt:

(3) a. Amanda detested Brittany. She _____
    b. Amanda detested Brittany. _____

Unlike for prompts like (3a), in the free prompt condition (3b) participants picked their own referring expressions for the first-mentioned entity. Stevenson et al. found two results of interest. First, across all eight context types, there were a greater number of references to the previous subject in the pronoun prompt condition than in the free prompt condition. Crucially, this did not always result in an overall pronoun bias toward the subject in the pronoun prompt condition: For instance, for object-biased IC verbs as in (3a), the overall pronoun bias was still toward the object. Instead, the key finding is that the occurrence of a pronoun always shifted the distribution of references toward the subject compared to when no pronoun was provided – that is, for prompts like (3b), the first-mention bias toward the object was even stronger.

Stevenson et al.'s second finding was that, in their free prompt conditions across all stimulus types, participants' choice of referential form for the first-mentioned entity was heavily biased towards a pronoun when the referent was the previous subject, and likewise towards a name when the referent was a non-subject. This result may at first seem contradictory: If participants have a clear preference to use pronouns to refer to the previous subject and names to refer to non-subjects, why would the pronoun interpretation bias ever be toward a non-subject, as was the case for prompts like (3a)?

Kehler et al. (2008) offered an explanation for the apparent contradiction by modeling the relationship between production and interpretation in terms of Bayes' Rule, as shown in equation (4).

(4) $P(referent \mid pronoun) =$
$$\frac{P(pronoun \mid referent)\, P(referent)}{\sum\limits_{referent \in referents} P(pronoun \mid referent)\, P(referent)}$$

The term $P(referent \mid pronoun)$ represents the interpretation bias: the probability, given that a pronoun has occurred, of it being used by the speaker to refer to a particular referent. On the other hand, the term $P(pronoun \mid referent)$ represents the production bias: the probability, assuming that a particular entity is being referred to, that the speaker would have used a pronoun to refer to it. Bayes' Rule says that these biases are not mirror images of each other, but instead are related by the prior $P(referent)$, which represents the NEXT-MENTION bias: the probability that a particular referent will get mentioned next regardless of the referring expression used.[1] Equation (4) thus explains why there is nothing contradictory about having both a strong production bias toward pronominalizing the previous subject (and not pronominalizing non-subjects) and yet a lack of a subject bias in interpretation, as long as the prior $P(referent)$ points strongly enough away from the subject referent, as it does for object-biased IC verbs.

According to this model, a comprehender's interpretation bias thus relies jointly on his estimates of the likelihood that a particular referent will be mentioned next (regardless of form of reference) and of the likelihood that the speaker would have chosen a pronoun (instead of another form of reference) to refer to that referent. The predictions of the analysis can be tested using passage completion experiments: The values for the terms on the right-hand side of equation (4) can be estimated from the data collected in free prompt conditions (3b), which will yield a prediction for $P(referent \mid pronoun)$. If the Bayesian characterization is correct, this predicted bias should be highly correlated with the actual interpretation biases estimated directly in pronoun-prompt conditions (3a) in otherwise identical contexts.

Kehler et al.'s proposal went further than positing a Bayesian characterization in an important respect, specifically by claiming that the two terms in the numerator of equation (4) are conditioned by different sets of contextual factors. On the one hand, the data they surveyed suggested that the factors that condition the next-mention bias $P(referent)$ are primarily semantic (e.g., verb type) and pragmatic (e.g., coherence relations). On the other hand, the factors that condition the production bias $P(pronoun \mid referent)$ appear to be grammatical and/or information structural (e.g., based on grammatical role obliqueness or topichood, both of which

---

[1] The denominator of equation (4) is simply the probability that a pronoun is the form of reference chosen by the speaker ($P(pronoun)$), which can be computed by summing the numerator over all referents that are compatible with the pronoun. This term has the effect of normalizing the probabilities so that they sum to 1.

amount to a preference for sentential subjects). Considering this in light of the asymmetry between production and interpretation captured by equation (4), this picture makes a striking prediction: That the speaker's decision about whether or not to pronominalize a reference will be insensitive to the semantically- and pragmatically-driven contextual factors that in part determine the comprehender's interpretation biases. This hypothesis is surprising because it violates the intuition, represented in much previous work, that speakers will pronominalize mentions of referents in just those cases in which their comprehenders would be expected to interpret the pronouns to those same referents.

As unintuitive as this may seem, the results of several recent passage completion studies have provided preliminary support for this prediction (Rohde, 2008, Exps V–VII; Fukumura and van Gompel, 2010). We briefly describe one reported on by Rohde and Kehler (2014; see also Rohde, 2008, Experiment VI) which, unlike the other studies, examined pronoun usage in referentially-ambiguous contexts. Rohde and Kehler used the 3-way context manipulation in (2), adapted so that the possible referents were of the same gender, and a 2-way prompt manipulation (pronoun or free). As expected, the pronoun interpretation bias varied with the IC bias across the context types, with subject mentions being most frequent for (2a), least in (2b), and (2c) in between. However, this context difference did not affect rate of pronominalization in the free prompt condition. Instead, only grammatical role mattered, with participants pronominalizing subject references far more often than non-subject ones.

Study

In this study we aim to evaluate the model within a single experiment that brings all of the foregoing predictions together. The design employs a novel manipulation that utilizes a potential source of pragmatic enrichment, particularly the fact that relative clauses (RCs) attached to direct objects can be inferred to provide explanations of the matrix event (Rohde, Levy, & Kehler, 2011). Example stimuli are shown in (5):

(5) a. The boss fired the employee who was embezzling money. He _____
[ExplanationRC, PronounPrompt]

　 b. The boss fired the employee who was hired in 2002. He _____
[NoExplanationRC, PronounPrompt]

　 c. The boss fired the employee who was embezzling money. _____
[ExplanationRC, FreePrompt]

　 d. The boss fired the employee who was hired in 2002.
_____
[NoExplanationRC, FreePrompt]

On a typical interpretation, the relative clause in (5a) does more than constrain the reference of the noun phrase to which it attaches. That is, while not entailed, (5a) invites the pragmatic enrichment that the employee was fired because of the

embezzling. Crucially, this inference is not necessary to make the sentence felicitous; (5b) is fine without inferring an analogous causal link between the firing and the hiring.

Accounts of pronoun interpretation that appeal primarily to surface-level characteristics of the context (first-mention, subject assignment, grammatical role parallelism, and so forth) find little to distinguish (5a-b). The Bayesian analysis does predict a difference, however, based on an interconnected sequence of referential and coherence-driven interdependencies, as illustrated in Figure 1. First, we manipulate RC type because we expect that participants will write fewer Explanation continuations in (5a) and (5c) than (5b) and (5d) respectively, since the RCs in (5a) and (5c) already provide a cause (Simner & Pickering, 2005; Kehler et al., 2008; Bott & Solstad, 2012). Second, this difference is in turn predicted to yield a difference in next-mention biases ($P(referent)$): Since object-biased IC verbs impute causality to the object, a greater number of explanation continuations for (5b) and (5d) should lead to a greater number of next-mentions of the object as well. Third, the analysis predicts that rates of pronoun production ($P(referent|pronoun)$) as measured in the free prompt condition (5c-d) should only be affected by grammatical role (favoring pronominalizations of the subject), but display no interaction with RC type (with participants no less likely to pronominalize subject mentions, and likewise no more likely to pronominalize object mentions, in 5d than 5c). Finally, RC type is expected to affect interpretation biases as measured in the PronounPrompt condition, with a greater number of object interpretations in (5b) than (5a), since $P(referent|pronoun)$ is determined in part by next-mention expectations ($P(referent)$). As such, we expect the RC manipulation to have an effect on pronoun interpretation, but not production. Since interpretation biases are also determined in part by production biases, an effect of grammatical role favoring subjects is also predicted in (5a-b) compared to their free prompt counterparts in (5c-d).
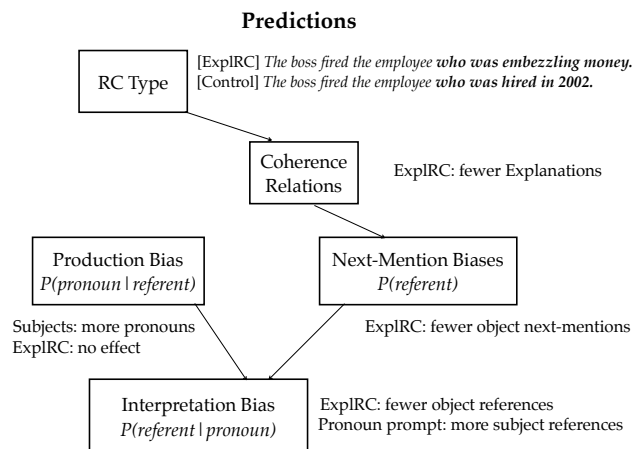
Predictions

Figure 1: Predictions as a Graphical Model

Design

A 2x2 design was employed using stimuli like (5a-d), which varied RC type (ExplanationRC or NoExplanationRC) and prompt type (PronounPrompt or FreePrompt). The experiment utilized 24 target items and 36 fillers. Context sentences for the stimuli always used object-biased IC verbs in the matrix clause. Two pieces of clip art were also displayed that indicated the gender of each event participant (always the same for both event participants for stimuli, so that reference in the PronounPrompt condition would be ambiguous). Participants (n=40) completed passages on Mechanical Turk, seeing each stimulus in a single condition in a fully balanced design. Two judges who were blind to the hypothesis annotated the data for coherence relations (Explanation or Other), first-mentioned referent (Subject, Object, or Other), and form of reference in the FreePrompt condition (Pronoun or Other). Outcomes were modeled using mixed-effects logistic regression with maximal random effects structure when supported by the data.

Results

All of the hypotheses were confirmed. Recall that the first prediction is that we would see a greater percentage of Explanation relations in the NoExplanationRC condition than in the ExplanationRC condition. The results, shown in Figure 2, confirm the hypothesis (β=2.06; p<.001).
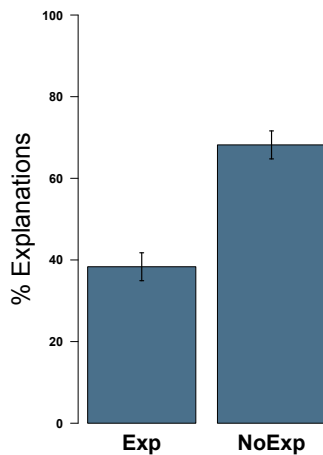

Figure 2: Percentage of Explanation Continuations in each RC condition (Standard Errors over Participant Means)

This leads to the second hypothesis, which is that for the FreePrompt condition, a greater percentage of next mentions of the object referent is predicted in the NoExplanationRC condition than in the ExplanationRC condition. The results, shown in Figure 3, confirmed this as well (β=.720; p<.05).

Our third hypothesis is that the rate of pronominalization during production in the FreePrompt condition will *not* be similarly affected by RC condition; instead, it should only be affected by grammatical role, favoring pronominalization of
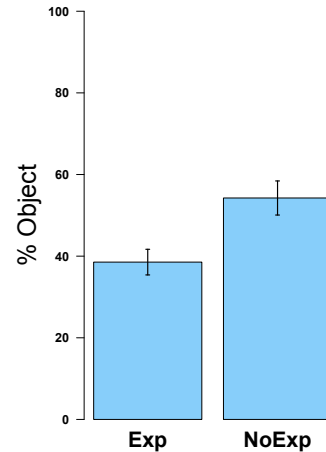

Figure 3: Percentage of Next-Mentions of Direct Object in each RC condition in FreePrompt Condition

the subject referent. The results, shown in Figure 4, confirm the effect of grammatical role (β=4.11; p<.001) and lack of interaction with RC condition (β=0.12; p=.92).[2]
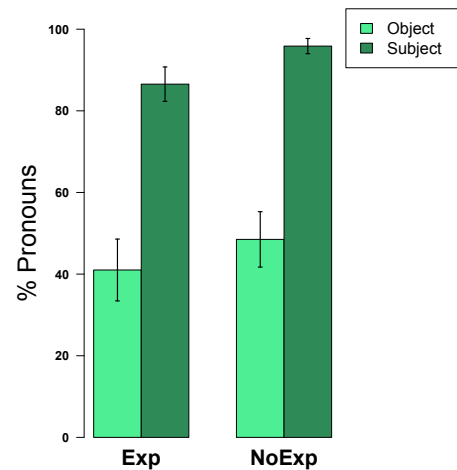

Figure 4: Rate of Pronominalization of Next-Mentioned Entity in FreePrompt Condition, by Grammatical Role and RC Condition

Lastly, we have two predictions for pronoun interpretation biases $P(referent \mid pronoun)$ as measured by the Pronoun-Prompt condition. The results are shown in Figure 5. (Note that the bars for the FreePrompt condition, which represent the prior $P(referent)$, are those shown in Figure 3.) First, we

---

[2]There was also a marginal main effect of RC condition (β=0.94; p=.078), whereby there were numerically more pronominalizations in the NoExplanationRC condition than the ExplanationRC condition. Whereas the reason for this is not clear, what is important for the hypothesis is the lack of interaction between grammatical role and RC type, i.e. that there was no pattern by which participants pronominalized subject referents less and object referents more in the NoExplanationRC condition.

predict a greater percentage of object mentions in the NoExplanationRC condition than in the ExplanationRC condition. This was confirmed (β=1.17; p<.005). Second, we predict a greater percentage of object mentions in the FreePrompt condition than the PronounPrompt condition. This was also borne out (β=-1.27; p<.001).[3] The results therefore confirm that a subtle pragmatic manipulation of the context – whether or not an RC invites the inference of a cause of the matrix event – influences pronoun interpretation biases, but not production biases. The effect on interpretation is a surprise for theories based on surface-level characteristics of the context, which find little to distinguish contexts like (5a-b).
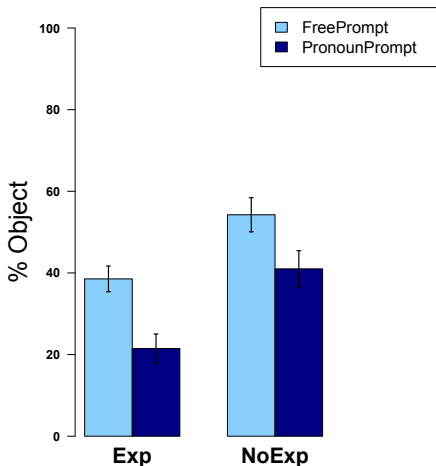

Figure 5: Percentage of Next-Mentions of the Direct Object by Prompt Condition and RC Condition

Finally, we can use the data collected here to test our Bayesian Hypothesis, i.e., that equation (4) captures the relationship between pronoun production and interpretation biases. We compare the predictions of the Bayesian model against two competing models. For each model, the predicted interpretation bias is estimated using biases measured in the FreePrompt condition. These predicted biases are then compared against the actual interpretation biases measured in the PronounPrompt condition.

The first competing model is what we call the MIRROR MODEL, according to which the interpretation bias toward a referent is proportional to the likelihood that a speaker would produce a pronoun to refer to that referent. Capturing the intuition that speakers will choose a pronoun in those cases in which hearers will be biased toward the intended referent, the predicted interpretation bias for this model is estimated using the pronominalization rate $P(pronoun \mid referent)$ measured in the FreePrompt condition, normalized by the sum of the pronominalization rate for all referents. Interestingly, this equation is equation (4) without the prior terms:

---

[3]There was also a marginal interaction (β=0.85; p=.078), whereby there was a larger numerical difference between the prompt conditions in the ExplanationRC condition than in the NoExplanationRC condition.

$$(6) \quad P(referent \mid pronoun) = \frac{P(pronoun \mid referent)}{\sum\limits_{referent \in referents} P(pronoun \mid referent)}$$

The second competing model is what we call the EXPECTANCY MODEL, according to which the interpretation bias toward a referent equals the probability that the referent gets re-mentioned (Arnold, 2001). The predicted interpretation bias for this model is thus estimated to be the next-mention bias $P(referent)$ measured in the FreePrompt condition. Interestingly, this equation is equation (4) without the likelihood terms:

$$(7) \quad P(referent \mid pronoun) = \frac{P(referent)}{\sum\limits_{referent \in referents} P(referent)}$$

Finally, per equation (4), the predicted interpretation bias for the Bayesian analysis results from combining the probabilities utilized by each of these other models. We then compare the predicted interpretation values for all three models against the observed pronoun interpretation biases $P(referent \mid pronoun)$, as measured by the data collected in the PronounPrompt conditions.

	Actual	Bayesian	Mirror	Expectancy
ExplRC	.215	.229	.321	.385
NoExplRC	.410	.373	.334	.542
R^2		.48/.49	.34/.42	.14/.12

Table 1: Actual and Predicted Rates of Pronominal Reference to Object

Table 1 shows the observed and model-predicted rates at which a pronoun prompt was interpreted to refer to the object across the two RC conditions. As can be seen, the observed values are most closely matched by the Bayes-derived values. We also test the degree of correlation between the values observed in the pronoun-prompt condition and those generated by the three different analyses using linear models. The correlation is performed over participant and item means; each participant (or item) contributes a value for the four pronoun interpretation estimates in each of the verb bias $\times$ referent combinations.[4] We expect that the predictions of all models will reveal some degree of correlation with the observed data: The Mirror model should capture the differences in biases between subject and non-subject referents, whereas the Expectancy model should capture differences across context type. Crucially, however, in combining the biases captured by both models, we expect the Bayesian model to be more highly correlated than either of the other models alone. And this is in fact the case, as can be seen in the last row of Table 1.

---

[4]We excluded data from participants (or items) for which the Mirror-model value and the Bayes-derived value could not be estimated—specifically if a participant's FreePrompt responses for a particular verb type contained no mentions of a particular referent or no pronouns for either referent; in both cases computing the predicted probabilities of the Mirror and Bayesian models would involve division by zero.

Conclusion

To sum, the experiment reported here demonstrates that biases towards referents of pronouns are sensitive to a subtle manipulation of context – whether or not an RC invites the inference of a cause – whereas production biases are not. Whereas this mismatch may seem unintuitive, it reflects precisely the dissociation between production and interpretation predicted by the Bayesian analysis.

A corollary to the analysis is that there is no unified notion of entity salience that mediates pronoun behavior. Truth be told, the notion of salience has always been somewhat problematic in the pronoun literature. Whereas it is often claimed that pronouns refer to salient entities, the factors that contribute to salience are usually determined by examining the properties of pronominal referents in data: A circular enterprise. Perhaps the most appropriate measure of entity salience is the next-mention bias as estimated in free prompt passage completions, as this metric is independent of the particular linguistic forms that speakers choose (Miltsakaki, 2007). As we have seen, however, these biases are not the same as those we witness for pronoun interpretation (pace Arnold (2001)).

Finally, the experimental results demonstrate the hierarchical causal structure among the factors that contribute to pronoun interpretation. Whereas canonical 'bag-of-cues' analyses typically model the different factors that affect interpretation as an undifferentiated set, the factors on the right-hand side of Figure 1 are modeled independently of pronouns. The analysis and experiment presented here therefore capture the idea that interpretation is not merely a bottom-up process that occurs when a pronoun is encountered, but is instead what happens when top-down expectations about entity mention come into contact with the bottom-up linguistic evidence that a pronoun provides (Kehler & Rohde, 2013).

Acknowledgments

References

Arnold, J. (2001). The effect of thematic roles on pronoun use and frequency of reference continuation. Discourse Processes, 21(2), 137-162.

Bott, O., & Solstad, T. (2012). The mechanics of causal interpretation: Explaining the implicit causality bias. New York, March. (Poster presented at the 25th Annual CUNY Conference on Human Sentence Processing)

Brown, R., & Fish, D. (1983). The psychological causality implicit in language. Cognition, 14, 237-273.

Caramazza, A., Grober, E., Garvey, C., & Yates, J. (1977). Comprehension of anaphoric pronouns. Journal of Verbal Learning and Verbal Behaviour, 16, 601-609.

Crawley, R. A., Stevenson, R. J., & Kleinman, D. (1990). The use of heuristic strategies in the interpretation of pronouns. Journal of Psycholinguistic Research, 19, 245-264.

Fukumura, K., & van Gompel, R. P. G. (2010). Choosing anaphoric expressions: Do people take into account likelihood of reference? Journal of Memory and Language, 62, 52-66.

Garvey, C., & Caramazza, A. (1974). Implicit causality in verbs. Linguistic Inquiry, 5, 459-464.

Grosz, B. J., Joshi, A. K., & Weinstein, S. (1995). Centering: A framework for modelling the local coherence of discourse. Computational Linguistics, 21(2), 203–225.

Hobbs, J. R. (1979). Coherence and coreference. Cognitive Science, 3, 67-90.

Kehler, A., Kertz, L., Rohde, H., & Elman, J. L. (2008). Coherence and coreference revisited. Journal of Semantics, 25, 1-44.

Kehler, A., & Rohde, H. (2013). A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. Theoretical Linguistics, 39(1-2), 1–37.

Koornneef, A. W., & van Berkum, J. J. A. (2006). On the use of verb-based implicit causality in sentence comprehension: Evidence from self-paced reading and eye-tracking. Journal of Memory and Language, 54, 445-465.

Terry Kit-fong Au. (1986). A verb is worth a thousand words: The causes and consequences of interpersonal events implicit in language. Journal of Memory and Language, 25, 104-122.

McKoon, G., Greene, S. B., & Ratcliff, R. (1993). Discourse models, pronoun resolution, and the implicit causality of verbs. Journal of Experimental Psychology, 19, 1040-1052.

Miltsakaki, E. (2007). A rethink of the relationship between salience and anaphora resolution. In A. Branco (Ed.), Proceedings of the 6th discourse anaphora and anaphor resolution colloquium (p. 91-96). Lago, Portugal.

Rohde, H. (2008). Coherence-driven effects in sentence and discourse processing. Unpublished doctoral dissertation, UC San Diego.

Rohde, H., & Kehler, A. (2014). Grammatical and information-structural influences on pronoun production. Language, Cognition and Neuroscience, 29(8), 912-927.

Rohde, H., Levy, R., & Kehler, A. (2011). Anticipating explanations in relative clause processing. Cognition, 118, 339–358.

Simner, J., & Pickering, M. J. (2005). Planning causes and consequences in discourse. Journal of Memory and Language, 52, 226-239.

Smyth, R. J. (1994). Grammatical determinants of ambiguous pronoun resolution. Journal of Psycholinguistic Research, 23, 197-229.

Stevenson, R., Crawley, R., & Kleinman, D. (1994). Thematic roles, focusing and the representation of events. Language and Cognitive Processes, 9, 519-548.