Elsi Kaiser

Looking beyond personal pronouns and beyond English: Typological and computational complexity in reference resolution

Elsi Kaiser: Department of Linguistics, University of Southern California, 3601 Watt Way, GFS 301, Los Angeles, CA 90089-1693, U.S.A. E-mail: emkaiser@usc.edu

1 Introduction

In their paper, Kehler and Rohde aim to reconcile two divergent approaches to pronoun interpretation, what they refer to as 'coherence-driven' and 'Centering-driven', by means of a probabilistic Bayesian approach that incorporates insights from both traditions. On the one hand, according to coherence-based approaches based on Hobbs (1979, 1990), pronoun interpretation is a side-effect that falls out from general discourse-level processing: Comprehenders draw inferences about the coherence relations between sentences and interpret pronouns accordingly. On the other hand, instantiations of Centering Theory (e.g., Grosz, Joshi & Weinstein, 1995) tend to focus not on the semantic relations between sentences but rather operate on a ranked set of 'centers' (entities mentioned in the discourse) and use information-structural factors – in particular notions linked to topicality – to guide pronoun interpretation. Centering Theory allows for the centers to be ranked in different ways (to capture cross-linguistic variation), though researchers have often assumed that the ranking depends on grammatical role, such that subjects are ranked above direct objects, which are ranked above obliques, and so on. As Kehler and Rohde note, under coherence-based approaches to pronoun interpretation, semantics, world knowledge and inferences are crucial, but under Centering-based approaches, information-structural notions such as topicality (related to grammatical roles) are of central concern.

In the present paper, Kehler and Rohde present data from a series of experiments, showing that neither coherence-driven nor Centering-driven approaches alone are able to fully capture the data. Instead, they build on ideas introduced in Kehler 2002 and Kehler et al. (2008) and argue for a Bayesian approach in which

pronoun interpretation is influenced both by (i) comprehenders' expectations about coherence relations and by (ii) production-based constraints that modulate speakers' decisions about whether to produce a pronoun, a process which is sensitive to Centering-style topicality effects.

Right at the start, I would like to say that I am sympathetic to this approach and believe that reconciling coherence-based approaches and Centering-Theory type approaches is valuable step, both empirically and theoretically. My own research on both English (Kaiser, 2010, Kaiser, Li and Holsinger 2011) and German (Kaiser, 2011b) also indicates that the use and production of pronouns is guided both by coherence relations and by interpretation biases associated with pronouns:

In Kaiser (2010), I report the results of a series of sentence-completion experiments, some of which provided participants with a pronoun prompt and others in which participants could choose what referring expression to generate (cf. the pronoun/no-pronoun experiments described by Kehler and Rohde). The aim was to investigate how focus in *it*-clefts influences discourse-level processing. In essence, I show that distinguishing P(referent), P(referent|pronoun), and P(pronoun|referent) allows us to straightforwardly capture a pattern of results that might otherwise seem contradictory, and shed new light on how contrastive focus influences the use and interpretation of subsequent referential forms.

In Kaiser, Li and Holsinger (2011), we report two experiments similar to the experiments with passive voice that Kehler and Rohde report in the current paper. However, instead of using implicit causality (IC) verbs like Kehler and Rohde, we used agent-patient verbs such as *kick, tickle* in active and passive voice. In our study, participants heard sentence fragments with or without pronoun prompts and provided spoken continuations (*Mary slapped Lisa at the zoo. As a result (she) . . . / Lisa was slapped by Mary at the zoo. As a result (she) . . .*). We used 'as a result' to keep the coherence relation constant. Interestingly, the bigger-picture implications of our results fit with the probabilistic Bayesian approach, but the details differ somewhat from what Kehler and Rohde found with implicit-causality verbs like *amaze*: In Exp 1, when a pronoun prompt was given, we found a clear effect of thematic role on pronoun interpretation: People tend to interpret a pronoun as referring to the preceding *patient* (the object in actives and the subject in passives). In Exp2, with no pronoun prompt, when participants chose to produce a pronoun, there were again indications of a patient pronominalization preference. However, overall, the patient was not consistently the entity with the highest likelihood-of-mention: In active voice conditions, most continuations started with the preceding object/patient. However, in the passive conditions, continuations were evenly split between the subject (the promoted patient) and the agent that has been demoted to the by-phrase. Thus,

active sentences with causal connectives show a bias for patient-initial continuations. However, in passive sentences placing the semantically-prominent agent in the *by*-phrase boosted the likelihood of agent-initial continuations. We suggest that this asymmetry is a markedness effect that stems from the syntax-semantics mismatch that arises when semantically-prominent arguments – agents – are placed in a syntactically low-prominence position, the *by*-phrase. Although the specifics of these results differ from what Kehler and Rohde found – presumably due to differences in verb class and coherence relations – both sets of results highlight the importance of separating out P(referent), P(referent|pronoun) and P(pronoun|referent).

On the whole, in my opinion Kehler and Rohde present convincing evidence that their conciliatory, probabilistic Bayesian approach captures the data better than coherence-driven or centering-driven accounts on their own. In my commentary, I focus primarily on areas for further work – in other words, issues that would benefit from more research before one can claim that the Kehler-Rohde approach offers a broadly-applicable, crosslinguistically-tested account of pronoun interpretation. In other words, although I agree that their approach does better than purely coherence-based or purely Centering-based approaches, at the same time I think that some important questions still remain open.

The structure of my commentary is as follows. I start by discussing questions regarding the broader aims of the theory: Should we regard this as (i) a theory of how people use and interpret singular personal pronouns in English in particular, as (ii) a theory of how people use and interpret the most reduced referring expressions in language, or, can it be conceptualized more broadly, as (iii) a theory of how anaphoric expressions more generally are produced and interpreted in language? Building on these issues, I then turn to questions regarding crosslinguistic variability in anaphoric paradigms. I summarize existing work in this area and identify questions for future work. In the final section of my paper, I explore questions regarding the computational tractability of the Kehler-Rohde approach, in particular (i) the question of what inferences are (and aren't) generated during real-time processing (and related questions regarding the potential processing cost of drawing inferences), and (ii) questions regarding the relevant set of coherence relations – for example, are we dealing with a finite set generated from a small number of primitives, or a finite but unstructured set, or perhaps a more open-ended set of possible coherence relations?

Importantly, as Kehler and Rohde point out, "our analysis should not be construed as claiming that no factors other than those we have discussed influence pronominal reference" (p. 30). My aim in this commentary is to explore some of

the other factors that may play a role, and more generally to identify questions that I regard as potentially fruitful avenues for future work.

2 Looking beyond English, beyond personal pronouns

Kehler and Rohde's paper focuses on English and on personal pronouns. In fact, they make an explicit connection between pronominalization and topichood, building on the claims of Centering Theory. Since pronouns are the most reduced referring expressions in English, Kehler and Rohde's claims regarding personal pronouns can be described as claims about how language users produce and comprehend the *most reduced referring expressions* in their language. Thus, one might ask whether their Bayesian probabilistic approach is intended to also apply to other kinds of referring expressions (demonstrative pronouns, null vs. overt pronouns etc), or whether it is a theory specifically about how we use the most reduced (and topicality-associated?) forms available in a given language. The question of how broadly the theory should be construed is important, because the phenomenon of pronoun production and interpretation is widely regarded as part of a larger system of reference resolution, which in many languages includes other forms such as demonstrative pronouns, phonologically reduced pronouns and null pronouns (often referred to as pro-drop or topic-drop). This typological richness brings up the question of whether Rohde and Kehler's claims for (personal) pronouns extend to other referential forms. In this section, I discuss some of the existing reference resolution research on forms other than personal pronouns and on languages other than English.

When we look **crosslinguistically**, it soon becomes clear that even the seemingly straightforward case of referring to a human antecedent – accomplished with personal pronouns 'he' and 'she' in English – can be done with more than one referential form in other languages. For example, in Finnish, the proximal demonstrative *tämä* 'this' can be used for human antecedents, in addition to the gender-neutral personal pronoun *hän* 's/he' (ex.1). In German, as in ex.(2), humans can be referred to with personal pronouns (*er, sie, es* 'he, she, it') or with so-called d-pronouns/demonstrative pronouns (*der, die, das*). The longer forms of the proximal demonstrative (*dieser, diese, dieses* 'this') can also be used in German. Dutch, another Germanic language, uses full and reduced forms of personal pronouns as well as the distal demonstrative *die* 'that', as in ex.(3). (For work on Finnish, see e.g., Kaiser & Trueswell 2008; for German, see Bosch & Umbach, 2007; for Dutch, see Comrie 2000, Kaiser 2011a. For further typological discussion of other languages, see Comrie 2000, Diessel 2012).

(1) Finnish:

Pekka halusi pelata tennistä Matin kanssa, mutta {hän/tämä} oli sairas.

Pekka-NOM wanted play-INF tennis-PART Matti-GEN with, but {s/he-NOM/ this-NOM} was sick

'Pekka wanted to play tennis with Matti, but {he/demonstrative} was sick.'

(2) German:

Peter wollte mit Paul Tennis spielen. Doch {er/der} war krank. (from Bosch & Umbach 2007)

Peter wanted with Paul Tennis play-INF. But {he/dem} was sick.

'Peter wanted to play tennis with Paul, but {he/demonstrative} was sick.'

(3) Dutch:

Peter wilde met Paul gaan tennissen. Maar {hij/die} was ziek. (from Bosch & Umbach 2007)

Peter wanted with Paul go tennis-play-INF. But {he/dem} was sick.

'Peter wanted to play tennis with Paul, but {he/demonstrative} was sick.'

In these kinds of sentences with subject-object word order,[1] the personal pronouns tend to be interpreted as referring to the preceding subject and the demonstratives as referring to the preceding object. In fact, the demonstratives are often characterized as referring to the non-topical entity or shifting focus to a new topic. Interestingly, it seems that the referential biases of the demonstratives are more strict, in that they are often felt to be less ambiguous than pronouns (see e.g. Comrie 2000, Bosch & Umbach 2007, Kaiser 2011b for discussion). For example, in ex.(2) from German, the demonstrative *der* is felt to have a strong object preference, whereas the pronoun *er* is potentially more ambiguous (i.e., its preference for the subject is weaker).

In addition to personal pronouns and demonstrative pronouns, other forms are also used for referring to human antecedents, including null pronouns (e.g., pro-drop in Romance and topic-drop in East Asian languages), phonologically reduced pronouns and so-called bound pronouns (see e.g. Kibrik 2001 for cross-linguistic discussion). In sum, even if we only restrict ourselves to considering reference to human antecedents, we see that although personal pronouns do the brunt of the heavy lifting in English, many other languages have a broader array

---

1 The referential biases of personal pronouns and demonstratives can be interestingly different for sentences with noncanonical word order, such OVS in Finnish. See Kaiser & Trueswell 2008 and Kaiser & Vihman 2009 for discussion of results showing that, at least in Finnish and Estonian, the division of labor between personal pronouns and demonstratives is not symmetrical.

of referential forms. In light of the typological complexity in this area, we are faced with the question of whether the principles that Rohde and Kehler propose for the interpretation of personal pronouns in English also apply in the broader crosslinguistic context, and in referential systems where speakers have an option not just between personal pronouns and names/full nouns, but between personal pronouns, demonstrative pronouns and in some cases null pronouns as well (e.g. consider Spanish, where humans can be referred to with null pronouns, overt pronouns (*él, ella*) as well as demonstratives (*éste, ésta*)).

Existing experimental research in this domain partly supports the general claims of the probabilistic Bayesian account that Kehler and Rohde propose for personal pronouns, but at the same time makes it very clear that not all referring expressions work the same way. Let us briefly review some of the relevant studies:

In Kaiser (2011b), I conducted a passage completion study, using a method very much like Kehler and Rohde, to explore the interpretation of German personal pronouns and demonstratives. Rather than using transfer-of-possession verbs or implicit causality verbs, I used simple agent-patient verbs (e.g. *tickle, push*), followed by the ambiguous connective *dann* 'then' and a personal pronoun or demonstrative pronoun in subject position, as shown below:

(4) Sample stimulus from Kaiser (2011b)
Die Schauspielerin hat die Schneiderin gekitzelt und dann hat {sie/die}
The actress has the seamstress tickled and then has {she/dem} . . .
'The actress tickled the seamstress and then {she/demonstrative} . . .'

With these kinds of sentences in English (with personal pronouns in subject position), 'result' relations tend to be associated with object reference (see Rohde 2008) and 'occasion/narrative' relations tend to be associated with subject reference (suggested by Kehler 2002). In other words, the idea is that if the coherence relation is occasion, P(referent|pronoun) is higher for subject referents, and if the coherence relation is result, P(referent|pronoun) is higher for object referents. My study had two main aims: First, I wanted to test whether German personal pronouns will also show this kind of coherence sensitivity. I was especially interested in testing whether German personal pronouns can be pushed towards the object by the presence of a result relation, even though German has a specific form for referring to objects (the demonstrative pronoun). My second main aim was to explore whether demonstratives influence comprehenders' expectations about coherence, even if they are consistently strongly object-biased (i.e., presumably cannot be pushed towards subject reference). For example, if a connective is ambiguous between a result relation and a narrative relation, can comprehenders'

assumptions about which relation to activate be influenced by the referential biases of the anaphor?

With regard to coherence, I found that interpretation of personal pronouns is influenced by coherence relations – even in a language like German where more specific forms for object reference are available. In addition, regarding the interpretation of demonstrative pronouns, I found that although coherence does not modulate the antecedent choice of anaphoric demonstratives to the same extent that it influences the interpretation of personal pronouns (demonstrative pronouns consistently prefer objects), demonstrative pronouns nevertheless interact with coherence-related processing by guiding comprehenders' expectations of coherence relations. In particular, we find that object-biased expressions (demonstrative pronouns) trigger an expectation of a result relation (see also Rohde 2008, Rohde & Kehler 2008 on English pronouns).

As a whole, these results confirm that coherence-based factors influence the interpretation of personal pronouns (P(referent|pronoun)) in languages with richer anaphoric paradigms as well, but also show that at least some anaphoric forms are less susceptible to the effects of coherence relations. In German, P(referent|demonstrative) is not modulated by coherence relations as much as P(referent|pronoun). Thus, we should not assume that coherence effects will occur equally in all languages or with all anaphoric forms, though my results do show that they are not restricted just to the most-reduced forms that a language has. These findings also connect to the form-specific multiple-constraints approach proposed by Kaiser & Trueswell (2008), which claims that the interpretation of different anaphoric forms can differ in how sensitive they are to different kinds of information.

Kaiser & Trueswell's claim that different referential forms are not all equally sensitive to different types of information also receives support from work by Ueno & Kehler (2010, 2011), who investigated the interpretation of Japanese null and overt pronouns. Ueno & Kehler found that Japanese overt pronouns patterned like English overt pronouns in sentences with transfer-of-possession verbs, but null pronouns appear to exhibit a more grammaticalized subject preference. These findings are largely confirmed in their second study (2011) with implicit causality verbs. Interestingly, the finding that null pronouns – the most reduced, default forms in Japanese – tend to have a subject bias and do not show a clear sensitivity to coherence-related factors suggests that it is not correct to say that it's the most reduced form of a language that shows the strongest sensitivity to coherence relations.

Generally speaking, an important question for future work has to do with the broader crosslinguistic applicability of Kehler and Rohde's account. On the one hand, when it comes to making inferences regarding coherence relations,

languages are not expected to differ, and we expect the fundamental patterns to be universal. On the other hand, when it comes to the use and interpretation of specific referring expressions, languages differ considerably in their anaphoric paradigms and existing research already suggests that the interpretation of different kinds of pronouns – overt personal pronouns, null personal pronouns, overt demonstrative pronouns – is not equally sensitive to coherence-related factors. This highlights the need for further crosslinguistic work in this area.

So far, we have focused on data from other languages. However, even before we look at other languages or other forms, interesting questions come up regarding the pronoun *it* in English (as well as demonstrative pronouns like *this* and *that*, cf. Webber 1991 on discourse deixis). Consider the example below from Hutchins & Somers (1992), where all versions have the same basic coherence relation (explanation):

(5) a. The monkey ate the banana because it was hungry.
    b. The monkey ate the banana because it was ripe.
    c. The monkey ate the banana because it was tea-time.

As these sentences illustrate, *it* is ambiguous in English. It can refer to animals (5a) or inanimate objects (5b). It can also be non-referential (5c), and can even refer to events or discourse segments (e.g. '. . . it made the children happy', '. . . but it was a bad idea'). A basic question is whether the interpretation of *it* patterns in the same way as *he/she*, and whether this depends on the referential status of the antecedent. In general, it seems that both Centering Theory and coherence-based approaches have focused on personal pronouns. However, if our aim is to provide a model of how speakers produce and interpret pronouns, we also need to include *it* in that set. The null hypothesis is that *it* would pattern the same way as personal pronouns (*he/she*), but the fact that this form is ambiguous between different types of reference (animates, inanimates, events), brings up interesting questions regarding the terms used in Bayes' rule. For example, if we conceptualize P(referent|pronoun) in terms of subjecthood (i.e., what is the probability that a pronoun refers to the subject), then we'd like to know whether – holding other things constant – P(referent|pronoun) differs systematically for *he/she* and *it*. We can ask more specific questions about whether, from the perspective of the speaker, P(pronoun|referent) is conditioned by the animacy of the intended referent, and whether, from the perspective of the comprehender, P(referent|pronoun) is sensitive to the animacy of potential referents.

Other directions for future work include the production and comprehension of plural pronouns (*they*), as well as the referential properties of demonstrative pronouns when used for inanimate entities. Existing work by Brown-Schmidt,

Byron and Tanenhaus (2004, 2005) shows that the interpretation of *it* and *that* is sensitive to extra-linguistic information, such as how easily two objects can be viewed as a composite. In commands like 'Put the cup on the saucer. Now put {it/that} over by the shovel', the demonstrative pronoun 'that' resulted in more composite 'cup-and-saucer' interpretations and the pronoun 'it' resulted in more 'cup' (theme) interpretations. Importantly, as Kehler and Rohde explicitly state in their paper (p. 30), they do *not* claim that no other factors beyond the ones they discuss are relevant. Thus, the general question of how to best incorporate other kinds of information into the Bayesian model is an interesting avenue for future work. A strength of the Bayesian probabilistic approach is that it allows us to formulate open questions precisely and to explore the inclusion of new factors in explicit ways, using tools from probability theory.

3 Computational simplicity or complexity?

In addition to considering the typological richness of human languages, as mentioned in the preceding section, I would also like to explore some questions regarding the computational tractability of the Kehler-Rohde approach. Kehler and Rohde criticize the SMASH family of approaches because "under this model (SMASH), encountering a pronoun triggers a fairly complex process in the mind of the comprehender" (p. 4), and suggest that the complexity of the pronoun-interpretation data should come not from a complex theory but rather from "the operation of discourse-level inferential processes that we already independently know to exist" (p. 5). I largely agree with this logic. At the same time, it is important to not forget that the complexity is still present in the system, even if it's now been 'moved' mostly into the domain of discourse-level inferencing rather than pronoun use and interpretation. In this section, I will explore two points related to these themes: (i) the question of what inferences are (and aren't) generated during real-time processing (and related questions regarding the potential processing cost of drawing inferences), and (ii) questions regarding the relevant set of coherence relations.

When discussing existing research on pronoun resolution, Kehler and Rohde suggest that a mechanism that involves searching for potential antecedents and filtering out irrelevant ones would be very unwieldy from a processing perspective and thus seems to conflict with our impression that pronoun resolution is an 'easy' process that occurs extremely frequently during communication. They present this as evidence against a 'SMASH' type view. However, thinking about issues related to processing or computational complexity load brings up interesting questions regarding the processing of inferences. Coherence-based

accounts focus on the idea that successful language comprehension requires us to draw inferences about the relations between sentences, and as Rohde and Kehler point out, these coherence relations influence our expectations about what will be mentioned next in the discourse. Drawing inferences, however, also requires some amount of processing resources. There is a large body of research in cognitive psychology and psycholinguistics exploring different kinds of inferences and whether they are drawn on-line during language processing.

As Graesser et al. (1997) note, it is generally agreed that not all possible inferences are generated during real-time language processing, due to limitations on working memory and the 'computational explosion' this would entail. The question of which inferences are generated during real-time processing is still unresolved, and seems to depend on the nature of text as well as readers' aims (see Graesser et al., 2002 for discussion). Generally, Graesser et al. (1997, see also Graesser et al. 2002) suggest that inferences involving explanation relations ('why') guide comprehension much more than inferences about what happens next, or when or where events are occurring (e.g. Graesser & Clark, 1985). If this extends to the kinds of sentences that Kehler and Rohde are focusing on, we might predict that comprehenders will be more sensitive to 'because'/ explanation relations than elaboration relations or occasion relations. However, it is worth keeping in mind that the term 'inferences' is used in different ways by different researchers and not all findings from inferencing in cognitive psychology translate directly to inferencing as it is construed in coherence-based approaches. Nevertheless, the broader point that these results bring up concerns the question of how much inferencing people are actually doing in the course of language comprehension.

This question connects to claims regarding 'good enough' processing (Ferreira et al., 2002; Ferreira & Patson, 2007). A number of experiments – mostly focusing on syntactic representations – suggest that comprehenders do not always construct a fully specified representation of the input. Instead, they may parse sentences only on a shallow level – resulting in potentially underspecified or even incorrect representations – using computationally 'cheaper' heuristics such as treating a noun-verb-noun in sequence as an agent-verb-patient sequence (cf. Bever 1970, Townsend & Bever 2001), which results in the misinterpretation of passive sentences. Although these 'good enough' results come largely from the domain of syntactic representations, one can ask whether something similar is occurring on the discourse level.

The idea of 'good enough' processing, combined with the research showing that not all possible inferences are generated during real-time processing, suggests that perhaps comprehenders are also defaulting to more shallow processing

on the level of discourse representations. If processing is shallow and compre-
henders do not construct the relevant inferences, then the basic assumptions of
coherence-based approaches may be problematic. For example, are there situa-
tions where P(referent) – the probability of a given referent being mentioned – is
largely computed from, say, heuristics making reference to grammatical role,
rather than from inferences regarding the likely coherence relations between sen-
tences? These are still open questions, and they offer intriguing directions for fu-
ture research. In my opinion, these questions do not invalidate the key points of
Kehler and Rohde's approach, but suggest that further research is needed to see
whether the kind of inferencing that is at the heart of coherence-based approach-
es consistently takes place during language processing, especially in situations
where comprehenders may be distracted and thus may be using only part of their
processing resources for language processing.

   A further important consideration for Rohde and Kehler's model – since it
relies on Hobbsian coherence relations – is what the relevant set of coherence
relations is. Existing literature on this topic has led to widely divergent conclu-
sions (see Knott & Sanders 1998, Hovy & Maier 1995 for discussion). On the one
hand, there is what Hovy and Maier (1995) call the 'profligate' position, consisting
of researchers who propose a fixed number of relations, often somewhere ranging
from five to twenty (or more) (e.g. Hobbs 1979, Mann & Thompson 1988, Asher &
Lascarides 2003). On the other hand, according to the 'parsimonious' position,
coherence relations should be regarded as an open-ended set characterized by
just two primitives (e.g. Grosz & Sidner, 1986). Researchers also differ in terms of
the role they attribute to the goals and intentions of the conversational partici-
pants (e.g. Grosz & Sidner, 1986 vs. Mann & Thompson 1988).

   In the end, many questions remain open: Can discourse relations be success-
fully characterized by a closed set of coherence relations, or is it impossible to
identify a finite set? If there is a finite set, can the relations be boiled down to
a smaller set of primitives, from which all relations can be 'computed/derived/
assembled'? If not, how is the set structured? If we are to construct a detailed
model of pronoun interpretation that builds on coherence relations, these ques-
tions will need further study. Hovy & Maier (1995) note that Grosz and Sidner's
parsimonius two-relation view, which has a focus on the intentional level, cannot
be straightforwardly reconciled with more profligate approaches. If we try to
apply these different kinds of approaches to the Bayesian, probabilistic approach
advocated by Kehler and Rohde, it seems that we may end up with different pat-
terns of predicted interpretation and production biases. This kind of concern also
arises if we compare approaches that posit a very large number of coherence rela-
tions (e.g. more than 30) to those with a much smaller set – a smaller, more gen-
eral set may obscure differences that could be detected with a more fine-grained

approach. In other words, the theory of coherence relations that one uses as the foundation for the Kehler-Rohde approach does matter. Thus, if our ultimate aim is to create a maximally fleshed-out model of pronoun production and interpretation, a clearer understanding of coherence relations will be necessary. A possible direction is suggested by Kaiser (2012), which uses a cross-modal priming approach to explore how fine-grained the set of coherence relations actually is, with a focus on causal relations in particular.

In closing, it should be noted that the questions concerning the set of coherence relations are not an issue for Kehler and Rohde's theory specifically, but rather for all theories that make reference to the relations between discourse segments. Thus, while a better understanding of the set of coherence relations is an important goal for future work, the lack of current consensus in this domain does not invalidate Kehler and Rohde's point that coherence relations play a crucial role in driving the production and interpretation of pronouns.

Conclusions

My impression is that reconciling coherence-based approaches and Centering-Theory type approaches is a valuable step, both empirically and theoretically. Kehler and Rohde present convincing evidence that the probabilistic Bayesian approach captures the data better than coherence-driven or centering-driven accounts on their own. They argue that pronoun interpretation is subject to two key components, namely (i) comprehenders' expectations about coherence relations and (ii) production-based constraints that modulate speakers' decisions about whether to produce a pronoun. The explicit nature of their model allows us to formulate questions for future work in explicit ways, which is advantageous. In my commentary, I focused on exploring some of these open questions, including the intended breadth of the theory and its crosslinguistic applicability, in light of the fact that many languages have richer anaphoric paradigms than English. I also explored questions regarding the computational tractability of their approach, in particular questions regarding the on-line generation of inferences and the definition of the relevant set of coherence relations. Broadly speaking, the probabilistic Bayesian approach allows us to take new steps towards reconciling long-standing debates in the field of pronoun interpretation, and has the potential to facilitate future work both across languages and on the different referential forms within languages.

References

Bever, Thomas G. 1970. The cognitive basis for linguistic structures. In *Cognition and Language Development,* ed. by R. Hayes, 279–362. New York: Wiley & Sons, Inc.

Bosch, Peter & Carla Umbach. 2007. *Reference Determination for Demonstrative Pronouns*. In D. Bittner & N. Gargarina (eds.), *Intersentential Pronominal Reference in Child and Adult Language*. ZAS Papers in Linguistics No 48. pp. 39–51.

Brown-Schmidt, Sarah, Byron, Donna & Michael Tanenhaus. 2004. That's not it and 'it's' not 'that': The role of conceptual composites in in-line reference resolution. In M. Carreiras & C. Clifton, Jr. (Eds.), *On-line sentence processing: ERPS, eye movements and beyond*. Hove, UK: Psychology Press.

Brown-Schmidt, Sarah, Byron, Donna & Michael Tanenhaus. 2005. Beyond salience: Interpretation of personal and demonstrative pronouns. *Journal of Memory and Language*, 53, 292–313.

Comrie, Bernard. 2000. Pragmatic Binding: Demonstratives as anaphors in Dutch. In *Proceedings of the Twenty-Third Annual Meeting of the Berkeley Linguistics Society*. Berkeley: Berkeley Linguistics Society, pp. 50–61.

Diessel, Holger. 1999. *Demonstratives: Form, Function, and Grammaticalization*. Amsterdam: John Benjamins.

Ferreira, Fernanda, Karl Bailey & Vittoria Ferraro. 2002. Good-enough representations in language comprehension. *Current Directions in Psychological Science* 11. 11–15.

Ferreira, Fernanda & Nikole Patson. 2007. The good enough approach to language comprehension. *Language and Linguistics Compass* 1. 71–83.

Graesser, Arthur, Keith Millis & Rolf Zwaan. 1997. Discourse comprehension. *Annual review of psychology*, *48,* 163–189.

Graesser, Arthur, Brent Olde & Bianca Klettke. 2002. How does the mind construct and represent stories? In M. C. Green, J. J. Strange, & T. C. Brock (Eds.), *Narrative impact: Social and cognitive foundations* (pp. 231–263). Mahwah NJ: Erlbaum.

Graesser, Arthur, & Leslie Clark. 1985. *Structures and procedures of implicit knowledge*. Norwood, NJ: Ablex.

Grosz, Barbara J., Aravind Joshi & Scott Weinstein. 1995. Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21(2):202–225.

Grosz, Barbara J. & Candace L. Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.

Hobbs, Jerry R. 1979, 'Coherence and coreference'. *Cognitive Science* 3:67–90.

Hobbs, Jerry R. 1990, *Literature and Cognition*. CSLI Lecture Notes 21. Stanford, CA.

Hovy, Eduard & Elisabeth Maier. 1995. *Parsimonious or profligate: How many and which discourse relations?* Technical report, University of Southern California.

Hutchins, W. John & Harold Somers. 1992. *An Introduction to Machine Translation*. Academic Press Limited, London.

Kaiser, Elsi. 2010. Investigating the Consequences of Focus on the Production and Comprehension of Referring Expressions. *International Review of Pragmatics*, 2(2), 266–297.

Kaiser, Elsi. 2011a. Salience and contrast effects in reference resolution: The interpretation of Dutch pronouns and demonstratives, *Language and Cognitive Processes,* 26, 1587–1624.

Kaiser, Elsi. 2011b. On the relation between coherence relations and anaphoric demonstratives in German. In Reich, Ingo et al. (eds.), *Proceedings of Sinn & Bedeutung 15,* pp. 337–351. Saarland University Press: Saarbrücken, Germany. http://universaar.uni-saarland.de/monographien/volltexte/2011/30/pdf/SinnUndBedeutung15.pdf

Kaiser, Elsi. 2012. Taking action: a cross-modal investigation of discourse-level representations. *Frontiers in Psychology* 3:156.

Kaiser Elsi., David Li & Edward Holsinger. 2011. Exploring the lexical and acoustic consequences of referential predictability. In I. Hendrickx, A. Branco, S. Lalitha Devi, & R. Mitkov (Eds.), *Anaphora Processing and Applications*, Lecture Notes in Artificial Intelligence, Vol. 7099. Heidelberg: Springer, 171–183.

Kaiser, Elsi & John Trueswell. 2008. Interpreting pronouns and demonstratives in Finnish: Evidence for a form-specific approach to reference resolution. *Language and Cognitive Processes*, 23(5): 709–748.

Kaiser, Elsi & Virve-Anneli Vihman. 2009. On the referential properties of Estonian pronouns and demonstratives. In H. Götzsche (ed.), *Memory, Mind and Language*, pp. 193–205. Cambridge Scholars Publishing.

Kehler, Andrew. 2002. *Coherence, Reference, and the Theory of Grammar*. CSLI Publications. Stanford, CA.

Kehler, Andrew, Laura Kertz, Hannah Rohde & Jeffrey Elman. 2008. Coherence and Coreference Revisited, *Journal of Semantics* (Special Issue on Processing Meaning), 25:1, pp. 1–44.

Kibrik, Andrej. 2001. Reference-maintenance in discourse. In: M. Haspelmath, E. Koenig et al. (eds.), *Language typology and language universals. An international handbook*. Vol. 2. Berlin: Mouton de Gruyter, 1123–1141.

Knott, Alistair & Ted Sanders. 1998. The classification of coherence relations and their linguistic markers: An exploration of two languages. *Journal of Pragmatics* 30, 135–175.

Mann, William C. & Sandra A. Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text*, 8(3):243–281.

Asher, Nicholas & Alexis Lascarides. 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.

Rohde, Hannah. 2008. *Coherence-Driven Effects in Sentence and Discourse Processing*. PhD dissertation, University of California, San Diego.

Townsend, David & Thomas G. Bever. 2001. *Sentence comprehension: The integration of habits and rules*. Cambridge, MA: MIT Press.

Ueno, Mieko & Andrews Kehler. 2010. The Interpretation of Null and Overt Pronouns in Japanese: Grammatical and Pragmatic Factors, in *the Proceedings of the 32nd Annual Meeting of the Cognitive Science Society*, Portland, August 2010. pp. 2057–2062.

Ueno, Mieko & Andrews Kehler. 2011. *Implicit Causality Biases in Japanese Pronoun Interpretation*. Poster presented at the CUNY Conference on Human Sentence Processing, Stanford.

Webber, Bonnie. 1991. Structure and ostension in the interpretation of discourse deixis. *Language and Cognitive Processes*, 6(2): 107–135.