

Effects of Anaphoric Dependencies and Semantic Representations on Pronoun Interpretation

Elsi Kaiser

University of Southern California, Department of Linguistics,
Los Angeles, California 90089, USA
`emkaiser@usc.edu`

Abstract. We describe three experiments that use priming methodology to investigate the nature of the abstract mental representations activated during pronoun resolution, in order to contribute to our understanding of how structural representations and semantic coherence representations influence pronoun interpretation. The results of Experiment 1 suggest that there exists a level of abstract anaphoric dependency representations. Experiments 2 and 3 indicate that general coherence representations are activated during pronoun interpretation and thus provide direct evidence for the existence of shared abstract representations between (non-pronominal) coherence-related inferencing and pronoun resolution processes. Moreover, Experiment 3 has implications for our understanding of the connections between linguistic and non-linguistic cognitive processes.

Keywords: Discourse coherence, Pronoun interpretation, Psycholinguistics, Eye-tracking, Priming, Anaphoric dependencies.

1 Introduction

Pronouns are central for communication, but the processes and representations underlying pronoun resolution are not yet fully understood. Competing accounts disagree regarding the contribution of syntactic factors, semantic factors and factors influencing referent prominence more generally. For example, coherence-based accounts (e.g. [6,7,8,9,12]) attribute a more central role to intersentential semantic relations (e.g. *X is a result of Y*, *W is similar to Z*) than to syntactic representations, in contrast to more structurally-oriented approaches. We describe three experiments that use priming methodology to investigate the nature of the abstract mental representations activated during pronoun resolution, in order to contribute to our understanding of how structural representations and semantic coherence representations influence pronoun interpretation. More broadly, this research aims to further our understanding of how humans perceive the semantic relations between objects, events and situations.

Priming is a well-known phenomenon according to which prior exposure to a stimulus influences (often facilitates) subsequent processing of a similar stimulus (or the same stimulus). For example, existing work has shown robust effects of syntactic priming in production [2,11], indicating that producing a particular syntactic structure facilitates subsequent production of the same structure.

Recent work suggests that syntactic priming also exists in comprehension (e.g. [5,16]).

In the experiments discussed here, we build on the logic of existing research and use priming as a tool to test whether two processes make use of the same (or overlapping) underlying representations. Although existing work has found evidence for priming of abstract syntactic representations, the abstract representations activated during pronoun interpretation have not yet been directly investigated. We explored two related questions: (i) *Experiment 1*: Does pronoun interpretation result in the activation of abstract anaphoric dependency representations? (ii) *Experiments 2 and 3*: Does pronoun interpretation result in the activation of abstract, possibly domain-general, coherence representations?

2 Experiment 1: Abstract Referential Dependencies

This experiment tested for the possibility of priming on the level of anaphoric dependency relations. Does processing a certain coreferential configuration facilitate subsequent comprehension of the same configuration? In other words, if a comprehender has recently constructed a certain kind of anaphoric dependency (e.g. interpreted an object pronoun as referring to the subject of the preceding clause), is s/he biased to construct the same kind of dependency again when faced with an ambiguous pronoun? If we find evidence for such priming, this would provide evidence for the existence of a distinct level of anaphoric dependency representations.

To investigate these issues, we conducted a comprehension-priming experiment. Participants (n=24) read sentences on a computer screen. Nonsense verbs were used to avoid effects of verb semantics. The prime sentences (ex. (1)) used gender cues to force a subject interpretation (ex. (1a)) or an object interpretation (ex. (1b)) of the pronoun in object position (Note that the syntactic structures of (1a) and (1b) are the same, and thus, any differences between subject primes and object primes cannot be attributed to syntactic priming). Neutral primes ended in intransitives (ex. (1c)). All sentences (targets, primes, fillers) were followed by questions. There was no noun or verb overlap between primes and targets.

- (1a) William swooked Betty and Kevin brucked him. [Subject Prime]
- (1b) William swooked Betty and Kevin brucked her. [Object Prime]
- (1c) William swooked Betty and Kevin brucked. [Neutral Prime]
- (2) Target: Stephen tolvered Peter and Diane churbited him.
[Question: Diane churbited —. Stephen Peter]

The critical target sentences contained ambiguous object-pronouns (ex. (2)), whose interpretation (preceding subject or object?) was probed by the subsequent question, to test whether the preceding prime influences pronoun interpretation. Targets were preceded by subject primes, object primes or neutral primes, as we wanted to analyze whether the anaphoric dependency in the prime influences participants' interpretation of the ambiguous pronoun in the target. We measured and analyzed participants' responses to the questions and the speed of these responses.

2.1 Predictions

We expected targets to show a baseline object preference, given that we were testing object pronouns and that existing research has repeatedly found effects of structural parallelism (e.g. [13,3]), i.e., pronouns prefer antecedents in parallel structural positions. However, if it is the case that activation/construction of a particular kind of anaphoric dependency makes participants more likely to activate/construct that same kind of dependency again, then we predict that the object preference should be weakened by subject primes (e.g. 1a). In other words, subject responses should be more likely to occur or easier to process after subject primes than after neutral or object primes.

2.2 Responses

Overall, there were more object responses (89%) than subject responses (10%) as expected given structural parallelism. Crucially, participants' responses were modulated by preceding primes: There were roughly twice as many subject-responses after subject primes as after object primes or neutral primes ($p < .05$).

The effect of the primes can be seen visually in Fig. 1 below, which shows the *object advantage score* for each of the three conditions, derived by subtracting the proportion of subject responses from the proportion of object responses. As Fig. 1 shows, although all three conditions have an overall object preference (as shown by the relatively high object advantage score in all three conditions), the object advantage score is lower after subject primes than after object primes and neutral primes.

2.3 Speed of Responses

Object-responses were approximately twice as fast as subject-responses. There was a numerical effect of prime on response speed: Subject-responses were faster after subject primes (<1500ms) than object primes (>1600ms) or neutral primes (approximately 2000ms).

2.4 Discussion

The results of Experiment 1 suggest that if a comprehender has recently processed/constructed a particular kind of anaphoric dependency, s/he is more likely to construct the same kind of dependency again when faced with an ambiguous pronoun.

In other words, it seems that anaphoric dependencies can be primed, even in the absence of noun/verb overlap between targets and fillers. This finding fits well with the idea that there exists a distinct level of anaphoric dependency representations, whose activation can linger and influence the interpretation of subsequent pronouns.

However, we should also keep in mind the question of how anaphoric dependency representation relate to coherence representations, because, as discussed

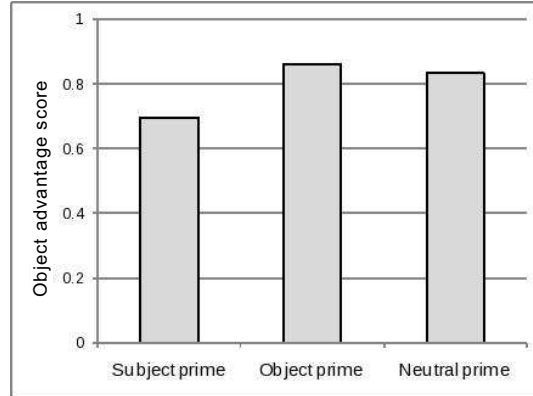


Fig. 1. The object advantage score for target sentences, shown as a function of prime type (The object advantage score is a measure of how strongly the object was preferred over the subject, and it was computed by subtracting the proportion of subject responses from the proportion of object responses)

in Section 3 below, certain anaphoric dependency patterns tend to be correlated with certain coherence relations. We will consider the implications of Experiment 1 again in Section 4, after presenting Experiments 2 and 3.

3 Experiments 2 and 3: Coherence Relation Priming

Experiments 2 and 3 tested whether the process of pronoun interpretation also activates abstract coherence representations. According to coherence-based approaches, pronoun resolution is largely a by-product of general inferencing about inter-clausal relations, and a growing body of research agrees that a successful account of pronoun interpretation needs to take into account the coherence relation between the pronoun-containing clause and the preceding clause. For example, in ex. (3), a subject interpretation of ‘him’ is more likely with a Cause-effect relation (3a) than with a Resemblance relation (3b) (e.g. [7,8,9,17]):

- (3) Phil tickled Stan, and Liz poked him.
 (a) Phil tickled Stan, and [as a result] Liz poked him_{Phil}
 (b) Phil tickled Stan, and [similarly] Liz poked him_{Stan}

However, the nature of these coherence relations is not yet well understood. Experiments 2 and 3 aim to contribute to our understanding of the linguistic and cognitive properties of these representations. We address two main questions:

First, *Experiment 2* tested whether these representations are specific to reference resolution. According to [6], coherence relations are not restricted to the domain of pronoun resolution – Hobbs claims that inferences about coherence relations exist independently of pronoun interpretation. However, so far we have no

direct evidence as to whether the representations activated during pronoun interpretation are the same as the representations activated during (non-pronominal) coherence establishment. Experiment 2 tackles this question.

Second, *Experiment 3* investigated whether these coherence representations are specific to the linguistic level of representation. In other words, how domain-general are these representations? Are they restricted to the linguistic domain, or potentially shared between different cognitive domains? Given that relations such as cause-effect and similarity also exist in other domains, e.g. vision, one might well expect coherence relations to be domain-general.

3.1 Design and Methods

Using priming and visual-world eye-tracking, we tested whether processing a particular coherence relation influences interpretation of a subsequent ambiguous pronoun. We used *linguistic primes* (Experiment 2) and *visuo-spatial primes* (Experiment 3) of three types: (i) Cause-effect, (ii) Resemblance, and (iii) Neutral/Baseline. The Neutral primes are best regarded as a baseline, because they were designed to evoke other kinds of coherence relations that, crucially, are neither Cause-effect nor Resemblance.

The **linguistic primes** were visually-presented two-clause sentences (see ex. (4) for an example of a cause-effect prime), whereas the **non-linguistic/visuo-spatial primes** were silent video clips of moving geometric shapes of various colors (see ex. (5), Fig. 2 for an example of a cause-effect video clip).¹ The linguistic primes contained no subject or object pronouns, in order to prevent anaphoric dependency priming from occurring. Furthermore, both linguistic and visuo-spatial primes were normed beforehand to ensure that the intended coherence relation was clear.

(4) *Sample linguistic prime*

[Cause-effect condition]: The patient pressed the red emergency button near the bed and a nurse quickly ran into the room.

(5) *Description of sample video prime*

[Cause-effect condition]: A triangle knocks into a circle which falls off a ledge.²

Two tasks were used to ensure that participants attended to the primes. In Experiment 2, with linguistic primes, participants were shown a prime sentence on the computer screen and instructed to read it aloud and indicate whether they had seen it earlier during the experiment. In Experiment 3, with visuo-spatial primes, participants were instructed to watch the video and afterwards to use the mouse to trace the paths of the objects.

¹ In Experiment 3, the video prime component was played twice on critical trials (as well as some filler trials), to increase detectability of potential priming effects.

² No linguistic information accompanied the actual videos. Furthermore, the arrows in the example image are there simply for purposes of illustration; they were not present in the actual videos.




Fig. 2. Sample cause-effect prime (The arrows were not present in the actual videos)




Fig. 3. Sample target visual display

The **critical targets** (which were the same in Experiments 2 and 3) were pictures of three same-sex characters (see Fig. 3), accompanied by an auditory sentence with an ambiguous object pronoun (ex. (6)). As in Experiment 1, we used nonsense verbs to eliminate any potential effects of verb semantics, and participants had been familiarized with the characters' names before the main experiment began. Each experiment contained 15 critical trials (prime+target sequences).

In addition, both experiments contained a large number of filler trials of two types: Some fillers resembled the targets whereas others resembled the primes. Fillers and targets were interspersed such that although critical trials consisted of a prime+target sequence, this patterning was not apparent to participants.

(6) Mary linded Lisa and Kate hepped her.

Eighteen native English speakers participated in each of the two experiments. Participants saw a scene and heard a sentence while their eye-movements were recorded. On target-type trials, the participants' task was to use the mouse to click on the last-mentioned person or thing.

Eye-movements in the visual-world paradigm are well-suited for investigating reference resolution, because existing research has shown that eye-movements to objects or pictures in a display are closely time-locked to the potential referents that a listener considers as language unfolds over time ([4,14], for a review see [15]). Thus, we can use looks to the different characters to gain insights into what participants are considering as potential referents for the ambiguous pronouns.

3.2 Predictions

If the abstract representations involved in pronoun resolution are connected to those in (pronoun-independent) coherence-related inferencing, we predict that lingering activation from primes can bias interpretation of ambiguous pronouns, resulting in more consideration of the subject after Cause-effect primes than after Resemblance or Neutral primes. (Because the pronouns are in object position, we expect an overall/eventual object preference, e.g. [13]) Crucially, using both linguistic primes (Experiment 2) and visuo-spatial primes (Experiment 3) allowed us to test whether the effects of coherence relation activation are specific to linguistic input, or whether they are more domain-general.

3.3 Results

Mouse click results: Participants' mouse clicks showed a strong object preference in all conditions.

Eye-movement results: After both linguistic and visuo-spatial Resemblance primes, eye-movements showed an early, persistent object preference ($p < .05$) that emerged within 200 ms of pronoun onset. Neutral primes also resulted in an object preference, although it reached significance later than with Resemblance primes. In contrast, Cause-effect primes resulted in initial competition between subject and object. In the Cause-effect condition of Experiment 2 (linguistic primes), the object preference did not reach significance until the 600-800 ms time slice after pronoun-onset ($p < .01$). In the Cause-effect condition of Experiment 3 (video primes), the emergence of the object preference was also delayed relative to Resemblance primes and Neutral primes: With Cause-effect primes, the object preference did not reach significance ($p < .01$) until the 400-600 ms time slice. Thus, participants' eye-movement patterns show that Cause-effect primes resulted in relatively more consideration of the subject early on, in both Experiment 2 and Experiment 3. In sum, we found priming effects both with linguistic primes and with visuo-spatial primes.

4 Conclusions

As a whole, the outcomes of Experiment 1 (anaphoric dependency priming) and Experiments 2 and 3 (coherence representation priming) contribute to our understanding of what representations are activated during the pronoun interpretation process. The results of Experiment 1 indicate that priming exists in the domain of reference resolution: Processing a subject interpretation appears to make a subsequent subject interpretation more likely to occur and/or easier to process. These findings fit well with the idea that there exists a level of abstract coreference representations or procedures, whose activation can linger and thereby bias the interpretation of a subsequent ambiguous pronoun.

Furthermore, given existing research showing that certain anaphoric dependencies are associated with certain kinds of coherence relations (recall the subject

vs. object effects associated with Cause-effect relations and Resemblance relations in ex. (3)), we are faced with the interesting question of how (and whether) the anaphoric dependency priming of Experiment 1 is related to the priming of coherence relations. In other words, could the effects observed in Experiment 1 actually be due to activation originating from representations of coherence relations, which may in turn be activating particular anaphoric dependencies? Or perhaps a combination of anaphoric dependency priming and coherence relation priming? Although the design of Experiment 1 does not allow us to offer a definite answer to this question at this stage, the results of Experiments 2 and 3 certainly fit well with these possibilities.

The results of Experiments 2 and 3 show that pronoun interpretation can be primed by coherence relations in preceding linguistic input as well as preceding visual input, even when primes and targets are connected only on the level of abstract coherence relations, and even when they are not in the same modality. This shows that general coherence representations are activated during pronoun interpretation, and thus provides direct evidence for the existence of shared abstract representations between (non-pronominal) coherence-related inferencing and pronoun resolution processes.

The finding that visuo-spatial primes have an effect on pronoun interpretation suggests that the abstract coherence representations may in fact be *domain-general*, i.e., shared between linguistic and non-linguistic domains. This is an important question that deserves further study. In particular, we are faced with the question of whether the participants in Experiment 3 were perhaps re-coding the visuo-spatial information into linguistic information. To investigate this issue, an articulatory suppression experiment is currently underway. When participants are asked to do an articulatory suppression task (e.g., say the syllable ‘the’ repeatedly), the phonological component of working memory is engaged. This prevents subvocal rehearsal and as a result, participants’ ability to re-code the visuo-spatial information into linguistic form is expected to be significantly impaired (see e.g. [10,1]). Thus, the outcomes of this experiment will help us to find out whether the abstract representations underlying coherence relations are domain general or specific to language.

In sum, the results of the three experiments discussed here shed light on the cognitive processes and representations activated during the process of pronoun resolution – in particular, they provide evidence for the existence of a distinct, non-pronoun-specific level of coherence representations – and also have implications for our understanding of the relationship between linguistic and non-linguistic representations.

Acknowledgements. Special thanks to Edward Holsinger for help with programming, stimulus creation and data collection. I would also like to thank three anonymous reviewers for their useful comments and insightful feedback. This research was partially supported by the funding from the Advancing Scholarship in the Humanities and Social Sciences Fund at the University of Southern California.

References

1. Baddeley, A.D.: Working memory. Oxford University Press, New York (1986)
2. Bock, K.: Syntactic persistence in language production. *Cognitive Psychology* 18, 355–387 (1986)
3. Chambers, G.C., Smyth, R.: Structural parallelism and discourse coherence: A test of Centering Theory. *Journal of Memory and Language* 39, 593–608 (1998)
4. Cooper, R.M.: The control of eye fixation by the meaning of spoken language: A new methodology for the real-time investigation of speech perception, memory and language processing. *Cognitive Psychology* 6, 84–107 (1974)
5. Fedorenko, E., Gibson, E.: Syntactic priming in comprehension. Poster presented at the 18th CUNY Conference on Human Sentence Processing, Tucson, Arizona (2005)
6. Hobbs, J.R.: Coherence and coreference. *Cognitive Science* 3, 67–90 (1979)
7. Kehler, A.: Coherence, Reference, and the Theory of Grammar. CSLI Publications, Stanford (2002)
8. Kehler, A., Kertz, L., Rohde, H., Elman, J.: Coherence and Coreference Revisited. *Journal of Semantics* (Special Issue on Processing Meaning) 2, 1–44 (2008)
9. Kertz, L., Kehler, A., Elman, J.: Grammatical and coherence-based factors in pronoun interpretation. In: Proceedings of the 28th Annual Conference of the Cognitive Science Society, pp. 1605–1610 (2006)
10. Murray, D.J.: Articulation and acoustic confusability in short-term memory. *Journal of Experimental Psychology* 78, 679–684 (1965)
11. Pickering, M.J., Branigan, H.P.: The representation of verbs: Evidence from syntactic persistence in written language production. *Journal of Memory & Language* 39, 633–651 (1998)
12. Rohde, H., Kehler, A., Elman, J.: Pronoun Interpretation as a Side Effect of Discourse Coherence. In: Proceedings of the 29th Annual Conference of the Cognitive Science Society, pp. 617–622 (2007)
13. Smyth, R.: Grammatical determinants of ambiguous pronoun resolution. *Journal of Psycholinguistic Research* 23, 197–229 (1994)
14. Tanenhaus, M.K., Spivey-Knowlton, M.K., Eberhard, K.M., Sedivy, J.E.: Integration of visual and linguistic information in spoken language comprehension. *Science* 268, 632–634 (1995)
15. Tanenhaus, M., Trueswell, J.: Eye movements and spoken language comprehension. In: Traxler, M., Gernsbacher, M.A. (eds.) *Handbook of Psycholinguistics*, pp. 863–900. Academic Press, Amsterdam (2006)
16. Traxler, M., Pickering, M.: Syntactic priming in comprehension: Evidence from eye-movements. Talk presented at the 18th CUNY Conference on Human Sentence Processing, Tucson, Arizona (2005)
17. Wolf, F., Gibson, E., Desmet, T.: Discourse coherence and pronoun resolution. *Language and Cognitive Processes* 19, 665–675 (2004)