Title
On-line Reference Assignment for Anaphoric and Non-Anaphoric Nouns: A Unified, Memory-Based Model in ACT-R

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 29(29)

ISSN
1069-7977

Authors
Pyke, Aryn
West, Robert L.
LeFevre, Jo-Anne

Publication Date
2007

Peer reviewed

On-line Reference Assignment for Anaphoric and Non-Anaphoric Nouns: A Unified, Memory-Based Model in ACT-R

Aryn Pyke (aryn.pyke@gmail.com)
Institute of Cognitive Science, Carleton University, 1125 Colonel By Drive
Ottawa, ON K1S 5B6 Canada

Robert L. West (rlwest@connect.carleton.ca)
Institute of Cognitive Science, 1125 Colonel By Drive
Ottawa, ON K1S 5B6 Canada

Jo-Anne LeFevre (jlefevre@connect.carleton.ca)
Department of Psychology, 1125 Colonel By Drive
Ottawa, ON K1S 5B6 Canada

Abstract

The computational model in present paper confirms that memory-based accounts are sufficient to account for a high rate of success at first-pass referent retrieval for anaphoric (and non-anaphoric) nouns. Because even definite noun phrases can often be non-anaphoric (e.g., Poesio & Vieira, 1998), an adequate model must account for how a reader makes an explicit or implicit decision about the anaphoric status of a noun (herein: The Anaphoric Classification Problem). We explain why we are inclined to reject the conventional intuition that: the failure to find/retrieve a referent within the discourse then, serially, leads to treating a (possibly anaphoric) noun as a new referent. Instead, we extend the memory-based account to address this classification problem. We suggest that LTM contains both generic referent types and specific referent tokens, which simultaneously compete for retrieval via resonance. The nature of what is retrieved (token vs. type) determines whether the reader effectively treats a noun as anaphoric or not. Our model predicts whether an anaphor in a given text will be misinterpreted as a new referent during first-pass processing. The influence of anaphor word choice is explained, and encompasses metaphoric anaphors.

Keywords: noun anaphora; memory-based text processing; resonance; reference assignment; cognitive modeling; ACT-R

Introduction

An anaphoric noun is one that denotes a referent that was previously mentioned in the discourse, but possibly using a different antecedent term. For example, in (2) "fruit" (or "apple") is used anaphorically to denote the referent introduced by the antecedent "apple" in (1).

(1) John bought an apple.
(2) John enjoyed the fruit/apple.

Readers are often able to re-activate the intended referent[1] almost immediately after encountering an anaphoric noun (Dell, McKoon, & Ratcliff, 1983; Sanford & Garrod, 1989) that is, after the first-pass processing of the noun. To account for this on-line ability, some models suggest that when a reader encounters an anaphoric noun, he/she undertakes a strategic search for an antecedent through a representation of the discourse context (e.g., Clark & Sengul, 1979; Kintsch & van Dijk, 1978; O'Brien, Plewes, & Albrecht, 1990). In such strategic-search models, the discourse might be mentally represented as a proposition network, and the reader might systematically troll backwards through it in search of the antecedent that (according to some criterion) could constitute a match to the current anaphor term. After all, how else could we account for the fact that readers come across the right referent (i.e., the particular one mentioned earlier in the discourse)?

An answer to this "How else" question is furnished by the memory-based view of text processing (see Gerrig & O'Brien, 2005 for a review). According to the memory-based view, the successful retrieval of a referent need not require (or constitute evidence of) a strategic, proactive search, because passive general-purpose memory processes often can automatically bring the intended referent to mind. In particular, under the resonance model (e.g., Gernsbacher, 1989; Myers & O'Brien, 1998) current information in working memory (i.e., the anaphoric noun) serves as a cue that automatically boosts activation of other entities throughout long-term memory -- including, ideally, the intended referent -- in accord with their conceptual overlap with the cue. Thus, at the time the anaphor 'fruit' in (2) is encountered, the apple referent can be automatically re-activated via resonance in virtue of its conceptual overlap with the anaphor (a pre-existing conceptual association).

Certainly higher-level and pragmatic processes may also play a role in comprehension. However, to account for readers' frequent success at referent reactivation during (first-pass) anaphor processing, we agree that there may be "no need to invoke any process other than general memory processing" (Gerrig & O'Brien, 2005, p. 230). The computational model in the present paper constitutes an existence proof that memory-based models are indeed sufficient, not only in principle, but in practice, to account for a high rate of success at first-pass referent retrieval.

The present paper and model also identify and address a fundamental, but we believe, previously neglected and under-estimated problem: The Anaphoric Classification Problem. In particular, how (and how accurately) can readers judge whether or not a noun is anaphoric during first-pass processing? Our model demonstrates how the memory-based view can be operationalized to address this classification problem, and in particular, to predict when a

[1] The term 'referent' is being used here in the cognitive sense (as in Gundel, Hedberg, & Zacharski, 2001) to mean the *mental representation* of the entity (person or object) in question.

particular anaphor will be initially misinterpreted as a new referent. To foreshadow, we suggest that it is *not* the failure to find/retrieve a referent that then, serially, leads to treating a noun (non-anaphoric or anaphoric) as a new referent. Rather, we suggest that retrieval, per se, typically always succeeds, but the nature of what is retrieved (particular referent token vs. referent type) determines whether the reader effectively treats a noun as anaphoric or not. In our model, generic type referents and specific discourse referents simultaneously compete for retrieval.

The Anaphoric Classification Problem

Anaphoric nouns are not a natural kind. Any noun in the lexicon can serve in an anaphoric or non-anaphoric capacity. So, how do readers determine whether or not a given occurrence of a noun is anaphoric? The definite article "the" is often assumed to indicate that the intended referent is already familiar to the reader (e.g., Clark & Sengul, 1979; Garnham, 1989; Just & Carpenter, 1987). However, corpus analyses have established that definite noun phrases (DNPs) are equally likely to introduce new referents into the discourse as to denote referents that have already been mentioned (Gundel et al., 2001; Poesio & Vieira, 1998). Thus, readers are not privy to the anaphoric status of a noun a priori, and may misclassify an anaphoric noun as non-anaphoric (and vice versa). In the literature, the reader's task is sometimes dubbed *anaphor resolution*, but this terminology seems to under-represent the full scope of the reader's (and modeler's) task.

Criteria for a Referent Retrieval Model

In view of the classification problem, a cognitive model of the (first-pass) referent retrieval process for nouns must meet the following criteria:

i) It must account for how the reader explicitly or implicitly makes an anaphoric/non-anaphoric classification of a noun

ii) In so doing it must serve to predict, for particular texts and anaphors, when the reader will be apt to: a) misclassify an anaphoric noun as a new referent; and b) misclassify a non-anaphoric noun as an anaphor.

The classification problem is addressed by some non-cognitive Natural Language Processing algorithms (e.g., Bean & Riloff, 1999; Vieira & Poesio, 2000), however, they involve multiple passes forward and backward through the text, and do not directly speak to the development of a *cognitive* model of *on-line* processing. We now describe two possible solutions to the classification problem, which, respectively, extend the strategic-search approach, and the memory-based approach to meet the above criteria. The second, memory-based approach is our positive proposal.

A (Non-ideal) Strategic-Search Approach

The statistical unreliability of the definite article to signal anaphoricity raises further doubts about the plausibility of automatic strategic-search accounts. Readers apparently have no sure-fire local cue to determine when they should expect -- and thus proactively search for -- an antecedent for a given noun phrase. One could argue that a reader (and the model) must make a default assumption of anaphoricity, at least for *definite* NPs. Thus the reader would always launch a strategic antecedent search upon encountering a DNP.

We have some reservations about this suggestion. First, the referent assignment process will not be unified in the sense that it will be different depending on whether the noun is proceeded a definite article or an indefinite article (e.g., Just & Carpenter, 1987; Kintsch, 1998). While primarily a cosmetic objection, if comparable performance can be achieved by a simpler, unified model (as we hope to demonstrate it can), then parsimony will argue for the unified model. A second reservation about a model which launches a strategic search for every DNP pertains to cognitive efficiency. According to corpus studies (e.g., Poesio & Vieira, 1998), such nouns will frequently be non-anaphoric, and any time and effort invested in an antecedent search will in vain. Under this account, non-anaphoric nouns (or anaphoric nouns misinterpreted as non-anaphoric) will presumably require two stages to process: 1) a strategic search for an antecedent that should fail, followed by 2) some method to initialize/associate a new referent to the noun. However, when readers sometimes treated an anaphoric noun as a new referent, Levine, Guzman and Klin (2000) found no accompanying delay in sentence reading time. Their results suggest that readers may not invest time actively seeking an antecedent in the discourse before resorting to a new referent interpretation.

Having outlined some limitations of the strategic-search solution, we now turn to our positive proposal.

A Unified Memory-Based Approach

Here we extend the memory-based view to provide a unified procedural treatment applicable to both anaphoric and non-anaphoric nouns. We suggest that a *preliminary* determination of anaphoric status is *not* based upon the article preceding the noun, but is made implicitly, based on the nature of what is activated by resonance. We, however, acknowledge that the article may impact *subsequent* verification, reanalysis, and pragmatic processes.

However, we propose that first-pass referent retrieval (and implicitly, the anaphoric classification task) is driven by general-purpose memory mechanisms. The referent initially retrieved for a noun will simply be the one that is the most accessible (i.e., most active), while that noun is being processed. As discussed above, the noun itself exerts a relatively immediate and important impact on the relative accessibilities of referents via resonance. The impact of any prior word or process (memory-based or otherwise) will be entirely mediated by its lingering effect on the accessibility of the referents in LTM. Thus, a referent's accessibility during first-pass noun processing owes to two components: i) a resonance boost from the noun term currently being processed; and ii) its 'context dependent' pre-noun activation level – the net effects of prior processing on referent activations. An example of a known context dependent influence is the recency of use of a referent (or the distance in words or sentences between the anaphor and antecedent, e.g., Duffy & Rayner, 1990)

Our treatment differs from conventional resonance accounts in that it provides a more comprehensive, quantitative and real-time operationalization of contextual influences on the LTM referents' pre-anaphor activation levels. Importantly, our treatment also encompasses the anaphoric classification problem. We suggest that a classification decision about a noun's anaphoric status may be determined by the outcome of a competition for retrieval between specific referents and generic referents (e.g., schemas/types) in LTM. Thus, we characterize the sources of referent activation with an eye to their impact on the *relative* accessibility of the specific intended referent, versus the accessibility of the applicable generic referent.

Our implementation was developed using an extension of the ACT-R cognitive architecture (Anderson et al., 2004), in its python implementation (Stewart & West, 2006). ACT-R was an appropriate development framework because its memory system that can simulate real-time activation fluctuation and decay for mental representations ("chunks").

Generic versus Specific Referents

In our model, LTM contains two types of referents: particular referents we've had prior experience (e.g., the apple you had for lunch), as well as referent templates (e.g. generic apple concept). To capture this type/token distinction in the syntax of the model and this paper, generic referents are given the identifier 0, whereas particular tokens of referents in memory are individuated with unique, non-zero identifiers. For example the generic apple referent will be designated [apple0], whereas the particular apple you had for breakfast might be [apple7]. The working assumption is that when we read about a new, unfamiliar apple -- as you are, this instant, reading about the apple I have in my bag -- we use the generic apple referent [apple0] to create a new referent token for the new apple, and assign it a unique ID, say [apple8]. We note that referents are conceptual, not lexical, representations.

In general, when a reader encounters the noun "apple", this will resonate with all conceptually related referents in LTM, be they generic or specific. Generic referents may thereby 'compete' with specific referents for retrieval. Consequently, if a generic referent is most accessible as the noun is processed, it will be retrieved, and the noun will effectively be interpreted as a new referent. If a specific referent from the discourse is most accessible, the noun will effectively be interpreted as anaphoric.

Ideally, generic referents are retrieved and used to initialize new discourse referents for non-anaphoric nouns, as [apple0] is used to token [appleANT] in (1). They can also, however, compete with specific referents for retrieval when processing anaphoric nouns, as [fruit0] will compete for retrieval with [appleANT] while processing "fruit" in (2). Thus it is not the absolute level of activation that the intended referent achieves that is germane, it is its level of activation relative to its generic competitor(s).

Memory-Based 'Context Dependence'

We now discuss how the discourse context, up to and including the referring noun, influences the present accessibility levels of the generic and specific referents in LTM. In our model, the referent retrieval process in play during the first-pass processing of the anaphoric noun is strictly memory-based (resonance), so the impact of prior discourse processing will be entirely mediated by its lingering effect on the accessibility of the various referents in LTM. Thus, to model the first-pass processing of a noun, we must also simulate the processing of the preceding discourse, but only so far as is necessary to approximate the accessibilities of the various referents in LTM at the time the noun of interest is encountered. Some pre-anaphor processes/effects are outlined below, along with supporting empirical evidence. To situate these effects, we first provide an overview of our full simulation.

Overview of Our Text-Processing Simulation

0. LTM is seeded with generic referents for various discourse concepts, including, importantly the antecedent concept and also (if different) the anaphor concept. For example, [apple0] and [fruit0] will be placed in memory.
1. Words of a story are then read and processed serially.
2. Each content word (e.g., noun, adjective, verb) spreads activation to both specific and generic referents according to the conceptual similarity between the word and the referent.
3. If the current word is a referring term (noun, name, pronoun), a referent is retrieved from memory (e.g., the most accessible one, be it specific or generic). This referent is thus boosted in activation, and so are its companion referents from previous sentences (see below). If the retrieved referent is generic, the generic referent is used to create a new specific referent.
4. At the end of each sentence, during wrap-up, participating referents are reactivated and associated.

In the simulation, the processing of each word takes about 150-200 ms/word, because only 3-4 basic operations (or ACT-R "productions") are required. Thus the time-course of processing is psychologically plausible to better capture the real-time decay of activation for the various referents.

Factors Affecting the Accessibility of Referents

1. Resonance with the Referring Noun. Each referent in LTM (specific and generic) will receive an activation boost in accord with its conceptual association with the current noun. In this model (see also Budiu & Anderson, 2004), the strengths of conceptual associations were estimated using Latent Semantic Analysis (LSA) values (lsa.colorado.edu). To produce LSA estimates, a large text corpus is analyzed and each lexical concept is represented as a vector in a multidimensional semantic space (Landauer, Foltz, &

Laham, 1998). Similarity between concepts is computed as the cosine of the angle between their vectors, so the maximum similarity, of a concept to itself, is 1. We discuss the impact of anaphor word choice more extensively later on. The remaining factors below relate to pre-anaphor 'context' effects on the accessibilities of referents.

2. Spread of Activation from Pre-Anaphor Words. Just as the anaphor resonates with, or spreads activation to, referents, our model assumes that such activation spread generally occurs as each content word in the discourse is encountered. The activation boosts received by referents may persist (as do lexical priming effects, e.g., Collins & Loftus, 1975) even when the reader progresses on to the next word. While such effects decay they may exert a cumulative effect on a referent's activation.

3. Recency and Frequency of Use. These general factors affect any mental representation's accessibility. Evidence indicates that the further back that an antecedent is (in sentences, and consequently in time), the more challenging it is to process the anaphor (e.g., Clark & Sengul, 1979; Duffy & Rayner, 1990; Levine, Guzman, & Klin, 2000). A referent that was referred to many times in a text, and/or was referred to in the sentence preceding the critical anaphor should be more cognitively accessible, ceteris paribus, than a referent mentioned only once several sentences back. In the ACT-R cognitive architecture (and in our model), the effect of recency and frequency of use on the activation level of representation i is quantified by B_i, where t_j is the time since use j, and d is a constant whose default value is .5 (Anderson et al., 2004).

$$B_i = \ln \sum_{j=1}^{N} t_j^{-d}$$

4. Sentence Wrap-Up Effects. Just and Carpenter (1980) suggested that integrative processes occur at sentence end, which is why readers tend to spend relatively longer on the final word in each sentence. These wrap-up processes may result in sentence-end activation effects (Balogh, Zurif, Prather, Swinney, & Finkel, 1998). Probe studies suggest that a referent mentioned early in a sentence may also produce facilitation effects at sentence end (e.g., Dell et al., 1983; McKoon & Ratcliff, 1980). In our model, we assume that integrative processes at sentence end result in an activation boost (weight 1) of the specific referents mentioned in the sentence.

5. Discourse Dependent Associations. In addition to pre-existing associations like those we are modeling with LSA, discourse dependent associations may be formed in memory. Spread of activation through such associations may produce intermittent (yet cumulatively significant) activation contributions to an intended referent during pre-anaphor processing. For example, each sentence (and proposition) in a discourse may contain several referents. In Dell et al. (1983, see also McKoon & Ratcliff, 1980), two referents which have appeared in a common sentence are called *companions*. The comprehension process appears to forge a lasting association between companion referents in memory, possibly during sentence wrap-up, such that when a referent is subsequently encountered, its companion(s) from prior sentences also become re-activated right away, and to a comparable degree (Dell et al., 1983). In our model, whenever a referent is retrieved (assigned), its companions are also boosted in activation, thereby making them more accessible as referent candidates for upcoming nouns.

How Our Memory-Based Approach Resolves the Classification Problem

(How) does our model account for the fact that we often correctly interpret a noun as anaphoric or not, during first-pass processing? We'll use example (1)-(2) to detail how referent retrieval is influenced, and how anaphoric classification is arbitrated, by the accessibility factors described in the preceding section.

Interpreting Anaphoric Nouns

Will an anaphoric noun be correctly classified as such? Let's assume that when processing the antecedent term "apple", the reader created a specific referent [appleANT] in memory. Ideally, this will be the referent retrieved when processing the anaphor "fruit". However, we suggest that another possibility is that the word "fruit" might instead bring to mind a generic fruit referent [fruit0]. If the generic referent is retrieved, the reader will effectively have treated (misclassified) the noun as non-anaphoric. In first-pass processing, which referent is retrieved (generic or specific) is determined by which is more accessible. In our model, two activation components determine the relative accessibilities of referents: i) the activation boost each receives in virtue of resonance with the anaphor term "fruit", and ii) the influence of prior 'context' (e.g., recency of use) on their relative accessibilities. We'll see that the resonance component of activation favours the generic referent, whereas the context influence tends to favour the specific referent. Thus, if the intended referent [appleANT] is to be retrieved, the contextual influence on its activation must more than compensate for its resonance disadvantage.

Resonance Boost from Anaphor. First consider the relative resonance boosts provided by the anaphoric noun "fruit" to the two discourse referents [appleANT] and [John1]. Resonance boosts depend on (usually discourse-independent) conceptual associations. LSA estimates the fruit-apple association at 0.47 and the fruit-John association at 0.03. Thus, resonance arguably removes [John1] from the running. However, the [appleANT] is still not necessarily retrieved, because there are other potential competitors. In particuar, the noun ("fruit") also boosts activation of non-discourse referents in LTM such as the generic fruit referent [fruit0], which would otherwise come into play for non-anaphoric uses of the term "fruit" (e.g.

"Sally gave Mark a piece of fruit"). In LSA, the fruit-fruit association is 1, so the anaphor term, "fruit", will actually provide a larger boost of activation to [fruit0] than to the specific apple referent [appleANT]. In principle, a noun anaphor will always afford a relatively higher boost of activation to its generic referent than to the specific referent (initialized using a different antecedent concept). However, the reader is not doomed to retrieve the generic referent, because a referent's net activation is also influenced by the pre-anaphor processing (context).

Our example illustrates that resonance may play a big role in giving a selective advantage to the intended referent [appleANT] *relative to other specific discourse referents*. However, contextual influences must in turn play the role of outweighing the relative advantage that resonance gives to the generic referent [fruit0] relative to the specific one. (Even) when the anaphor and antecedent are the same ("apple"), the resonating anaphor will spread comparable activation to the generic apple referent [apple0] as to the discourse specific one, so 'contextual' influences are still in play to determine whether the generic [apple0] referent or the specific intended referent [appleANT] is retrieved.

Pre-anaphor Contextual Influence on Activation. Some sources of activation furnish a relative advantage to specific discourse referents, to enable them to win out over a generic referent. The frequency and recency of usage of specific referents in the discourse will selectively bolster their activation level (vs. their generic counterparts). Furthermore, in our model, sentence wrap-up and discourse-specific associations (companion spreading) augment the activation of specific referents but not generic ones.

Figure 1 presents the activation profiles produced by our model for [appleANT], [John1], and the generic referents: [fruit0] and [apple0]. When considering the classification problem for the noun "fruit" in the second sentence, the crucial contrast is between the solid line at the top designating [appleANT], and the dotted line at the bottom of the graph designating [fruit0]. Letters (a)-(d) in the figure flag pre-anaphor events of interest. When processing "apple" at (a), our referent [appleANT] is mentioned (used, boosted) for the first time. At (b) [appleANT] receives a boost during the wrap-up of the first sentence, because it was a referent in that sentence. At (c), referent John is encountered (retrieved), however [appleANT] will also receive a comparable boost in activation because it was a companion of the referent John in the first sentence. In all, the cumulative effects of (a) recent usage, (b) first sentence wrap-up, (c) discourse-dependent companion association and (d) resonance, render the intended referent [appleANT] more accessible overall than the generic fruit referent (or the generic apple referent). Thus our model predicts that the correct referent will be retrieved (and thus that the correct classification will be made) during the first-pass processing of the anaphor "fruit" in this discourse. However, point (d) illustrates that resonance alone produced a greater boost in activation for the generic fruit referent than for the specific apple referent.

Interpreting Non-Anaphoric Nouns

Following the same memory-based, standard operating procedure (i.e., to retrieve the most active referent after resonance), our model also aptly handles the processing of the non-anaphoric noun "apple" in the first sentence. Prior to the noun "apple", the referent for [John1] was most active, due to its recent mention. The term "apple" resonates strongly (LSA=1) with the generic apple referent [apple0] in the LTM knowledge base, but has no strong pre-existing association with the John referent (apple-John LSA is 0.09). So, at point (a) in Figure 1, the generic referent [apple0] becomes the most active referent, and is retrieved. Consequently the noun "apple" is correctly interpreted as non-anaphoric, and a specific referent [appleANT] is tokened using the retrieved generic one. Thus, the same on-line process (i.e., retrieving the most active referent after resonance) can account for referent retrieval (and correct classification) for anaphoric and non-anaphoric nouns.

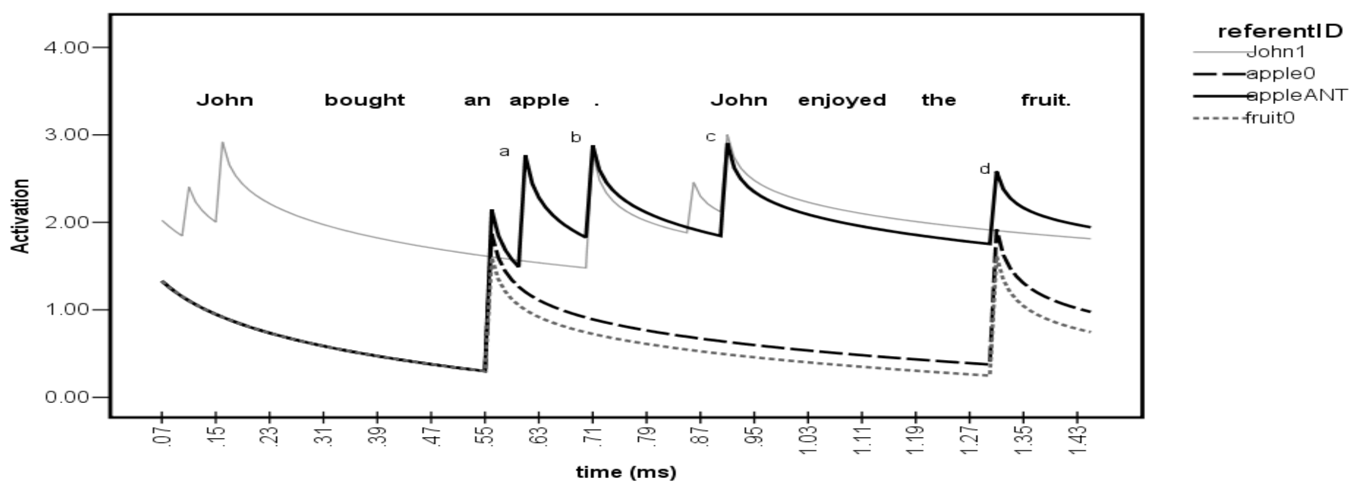

Figure 1: Model Output - Activation Profiles for Specific and Generic Candidate Referents in Discourse (1)-(2)

Misclassification and Anaphor Word Choice

The anaphoric status of "fruit" was correctly determined above by our model. However, evidence suggests that anaphors are sometimes misclassified (e.g., Levine et al., 2000). Our model accounts for potential first-pass misclassifications of an anaphor (i.e., retrieving a generic referent). Misclassification occurs if contextual influences fail to outweigh the relative advantage the anaphoric noun's generic referent gets from resonance. The anaphor term determines the most competitive generic referent in play. Consider the anaphors in (5), for the terrier [terrier1] in (3):

(3) John took his terrier_1 to the park.
(4) (The sun shone).
(5a) The dog…
(5b) The monster…

In (5a) "dog" maximally boosts [dog0] and moderately boosts [terrier1] (LSA=.72). In (5b) "monster" maximally boosts [monster0] and modestly boosts [terrier1] (LSA=.14), so context is relied upon more heavily in (5b) to compensate for the resonance discrepancy. Our process operates as normal for metaphoric anaphors, which was has precedent (Budiu & Anderson, 2004; Lemaire & Bianco, 2003) but the Classification Problem was not addressed. Here, if line (4) is omitted, our model correctly retrieves [terrier1] in (5a) and (5b). Including (4) makes [terrier1] less accessible (less recently active) and its activation falls below that of the generic referent [monster0] in (5b), so "monster" will be initially treated as a new referent. In general, the greater the activation 'head-start' the specific intended referent has due to pre-anaphor contextual influences, the greater the latitude in anaphor choice (i.e. the weaker the LSA anaphor-antecedent association can be, before resulting in misclassification).

Concluding Remarks

If LTM contains both generic (type) and specific (token) referents, readers can retrieve appropriate referents for both anaphoric and non-anaphoric nouns rapidly via resonance. Resonance is memory-based. Any influence of prior context is mediated by its impact on the activations levels of LTM referents. Anaphoric nouns may be misinterpreted as new referents (Anaphoric Classification Problem), if they prompt the retrieval of a generic rather than a specific referent. So, referent retrieval may be relatively immediate (Sanford & Garrod, 1989, Just & Carpenter, 1980) but not necessarily correct. The reader may subsequently revise this misinterpretation while reading the remainder of the sentence (beyond the current scope), so post-anaphor processing delays may reflect such revision processes. Our model was motivated by and is consistent with existing empirical evidence, and systematic verification with additional human data is currently being undertaken.

Acknowledgments

We thank Terry Stewart for his valuable advice. This research was funded by NSERC.

References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. L. (2004). An integrated theory of the mind. *Psychological Review, 111,* 1036-1060.

Balogh, J., Zurif, E., Prather, P., Swinney, D., & Finkel, L. (1998). Gap-filling and end-of-sentence effects in real-time language processing. *Brain and Language, 61,* 169-182.

Bean, D. L. & Riloff, E. (1999). Corpus-based identification of non-anaphoric noun phrases. *In Proceedings of the 37th Annual Meeting of the Assoc. for Computational Linguistics,* 373-380.

Budiu, R. & Anderson, J. R. (2004). Interpretation-based processing: a unified theory of semantic sentence comprehension. *Cognitive Science, 28,* 1-44.

Clark, H. H. & Sengul, C. J. (1979). In search of referents for nouns and pronouns. *Memory & Cognition, 7,* 35-41.

Dell, G., McKoon, G., & Ratcliff, R. (1983). The activation of antecedent information during the processing of anaphoric reference. *J. of Verb. Learning and Verb. Behav., 22,* 121-132.

Duffy, S. A. & Rayner, K. (1990). Eye-movements and anaphor resolution - Effects of antecedent typicality and distance. *Language and Speech, 33,* 103-119.

Gerrig, R., & O'Brien, E. (2005). The scope of memory-based text processing. *Discourse Processes,* 39(2&3), 225-242.

Gernsbacher, M. A. (1989). Mechanisms that improve referential access. *Cognition, 32,* 99-156.

Gundel, J., Hedberg, N., & Zacharski, R. (2001). Definite descriptions and cognitive status in English. *Journal of English Language and Linguistics, 5,* 273-295.

Just, M. A. & Carpenter, P. A. (1987). *The psychology of reading and language comprehension*. Boston: Allyn and Bacon, Inc.

Kintsch, W. (1989). *Comprehension: A Paradigm for Cognition*. New York: Cambridge University Press.

Kintsch, W. & van Dijk, T. A. (1978). Toward a model of text comprehension. *Psychological Review, 85,* 363-394.

Landauer, T.,Foltz, P.& Laham,D.(1998).An introduction to latent semantic analysis. *Discourse Processes,25,* 259-284

Lemaire, B. & Bianco, M. (2003). Context effects on metaphor comprehension: Experiment and simulation. In *Proceedings of 5th Int.Conf. on Cognitive Modelling (ICCM)*, 153-158.

Levine, W. H., Guzman, A. E., & Klin, C. M. (2000). When anaphor resolution fails. *J. of Mem. and Language, 43,* 594-617.

McKoon, G. & Ratcliff, R. (1980). The Comprehension Processes and Memory Structures Involved in Anaphoric Reference. *Bulletin of the Psychonomic Society, 16,* 150.

Myers, J. L. & O'Brien, E. J. (1998). Accessing the discourse representation during reading. *Discourse Processes, 26,*131-157.

O'Brien, E. J., Raney, G. E., Albrecht, J. E., & Rayner, K. (1997). Processes involved in the resolution of explicit anaphors. *Discourse Processes, 23,* 1-24.

O'Brien, E. J., Plewes, P. S., & Albrecht, J. E. (1990). Antecedent retrieval-processes. *JEP: LMC, 16,* 241-249.

Poesio, M. & Vieira, R. (1998). A corpus-based investigation of definite description use. *Computational Linguistics, 24,*183-216.

Sanford, A. J. & Garrod, S. C. (1989). What, when, and how?: Questions of immediacy in anaphoric reference resolution. *Language and Cognitive Processes, 4,* 235-262.

Stewart, T.C. & West, R. L. (2006) Deconstructing ACT-R. In *Proceedings of 7th Int. Conf. on Cognitive Modelling* (ICCM).