

Event Knowledge vs. Verb Knowledge

Jon A. Willits (willits@wisc.edu)

Department of Psychology, 1202 W. Johnson Street
Madison, WI 53706 USA

Rachel Shirley Sussman (rssussman@wisc.edu)

Department of Psychology, 1202 W. Johnson Street
Madison, WI 53706 USA

Michael S. Amato (amato@wisc.edu)

Department of Psychology, 1202 W. Johnson Street
Madison, WI 53706 USA

Abstract

Two kinds of knowledge are often confounded: knowledge about events occurring in the physical world, and knowledge about words occurring in language. We present two studies designed to show that these are distinct, that this distinction is important, and that it is difficult to determine which kind of knowledge is responsible for certain behaviors. In Study 1, we show that priming from verbs to nouns is both capable of and perhaps best explained by knowledge from distributions of word use rather than distributions of experiential knowledge of events. In Study 2, we show that differences in the scope of generalizations one might make, based on distributional statistics in language, can allow us to make testable predictions about whether the source of the knowledge is language-based or experience-based. This often-overlooked difference between experiential knowledge and word knowledge is important for developing a theory about how all knowledge is structured and used.

Keywords: Language understanding; Semantics; Syntax

Event Knowledge vs. Verb Knowledge

Much of our knowledge of the world around us can be characterized as knowledge about events, which may in turn give rise to broad expectations about which entities are likely to act in particular ways, in particular locations, and with particular participants. Furthermore, the way we understand events in the world and the structures we attribute to them will influence how that knowledge is used during other tasks and processes. However, our experience of events and their structure comes not only through our own observation of events, but structured language that we use to represent these events and their components. In this way, any act of communication reflects at least two sources of information: 1) *experiential knowledge*, obtained by interacting with entities and events in the world, and 2) *lexical knowledge* about the words we use to refer to and represent those entities and events.

This paper has three aims. The first is to suggest that the distinction between these two kinds of knowledge is important, because the kind of knowledge each represents is distinct and complimentary in important ways. The second is to point out that that some behaviors in experimental

situations are often misattributed to experiential knowledge about real-world events, when in fact those behaviors may be better explained in terms of lexical knowledge about words. The final aim is to offer a method for determining if a given behavior is due to experiential knowledge, knowledge about words, or some interaction of the two.

Event Schemas & Distributional Statistics

The best-supported theory about how our knowledge about events is structured is the notion of *event schemas* (Ferretti, McRae, & Hatherall, 2001; McRae, Ferretti, & Amyote, 1997). These event schemas are organized such that the important components of an event (like its typical agents, patients, instruments, and locations) are a part of the schema, and become more accessible whenever the schema is invoked. For example, hearing or reading the verb *flying* will activate the event schema FLY and its typical components like PILOT, AIRPORT, and PLANE. A further important aspect of event schema theory is that the specific schemas that are invoked are computed in a manner that is dependant on the sentential and situational context. For example, the specific schema a verb activates is sensitive to the syntax in which the verb is embedded (e.g. “*she was arrested by the cop*” vs. “*she arrested the cop*”, Ferretti et al, 2001) and the verb’s aspectual morphology (e.g. “*the diver had snorkeled*” vs. “*the diver was snorkeling*”, Ferretti, Kutas, & McRae, 2007). The consequence is a theory of event knowledge based on distributed, prototype-like schemas that are very similar to those suggested by connectionist theories. As such, Ferretti and colleagues argue that event knowledge is very general, both in terms of the types of relations such schemas specify, the tokens (of objects) that can fill those relations, and the semantic features which characterize which objects are most characteristic of those roles.

Ferretti and colleagues have amassed considerable support for the event schema account. They have shown that verbs prime characteristic agents, patients, instruments, and locations (Ferretti et al, 2001) and that this effect is modulated by the syntactic or morphological context of the verb (Ferretti et al, 2007). For example, imperfect aspect

verbs (of the inflectional form “*was cooking*”) prime locations, but perfect aspect verbs (*had cooked*) do not. They have also shown that the time it takes to read a syntactically ambiguous sentence is sensitive to what they call the *thematic fit* of the verb and its noun relations (Ferretti et al, 2001; Hare, Tanenhaus, & McRae, 2007). They argue that the notion of thematic fit (the match in semantic features between a specific relation and its slot in an event schema) better accounts for data in sentence comprehension experiments than theories based on selectional restrictions like “agents must be animate” (Gropen et al, 1991) or the causal control of the sentence’s head noun (McKoon & Ratcliff, 2003).

Ferretti and colleagues’ theory of event schemas is a good basic theory of how event knowledge might be structured. However, one critical way in which it is underspecified is in the difference between event knowledge (things we know by interacting with the real world) versus verb knowledge (things we know by hearing, reading, and talking about events that occur in the world). Ferretti and colleagues are very clear that both types of knowledge are important, and that the interaction between these two kinds of knowledge is critical for explaining behavior (as in the case of syntactic constraints upon what schemas are activated or computed). But the specific details about how this knowledge is divided up remains unspecified. For a particular behavior (like the priming results that show a relationship between imperfect aspect verbs with locations), is the priming behavior due to the structure of event knowledge, or due to the structure of verb knowledge, or to some interaction of the two?

Ferretti and colleagues argue that most of our important knowledge about events is obtained by interacting with the physical world; words serve merely as cues that compute the correct event schemas for a given linguistic context. Under this view, one can imagine a model of event knowledge like that shown in Figure 1.




Figure 1. *Was cooking* primes *kitchen* via links from an event schema, as in Ferretti et al (2007).

In this model, on reading the phrase “*the chef had cooked*”, the combination of the verb *cook*, and its morpho-syntactic frame “*had X-ed,*” combine to compute a COOK schema that is “completed” in terms of its aspect, and this concept is much more likely to have a high association strength with direct object concepts like PASTA and very unlikely to have high association strength with location concepts like KITCHEN. This differential activation of PASTA relative to KITCHEN in turn activates lexical items like *pasta* more

than the lexical items like *kitchen*. In this model, the relevant aspect/argument biases observed in the experiment are in the form of our knowledge about events as they happen in the world, not our knowledge about verbs as they are used in language.

But a second model, shown in Figure 2, is also consistent with the data. In this model, the lexical form *cook*, regardless of its aspectual form, activates a schema for COOK. But critical to this model, the activation of this event schema is completely irrelevant to the observed priming effects. In this case, the source of those effects is knowledge within the lexical system itself. It is possible that the word *cook*, or the verb phrase *had cooked*, co-occurs more often with (and is thus more highly associated with) object words like *pasta* than locations like *kitchen*, and vice versa for *cooking* and *was cooking*. In this model, associations between lexical items (and syntactic structures) are the source of the effect, without important differential input from our knowledge about experiences in the world. A model such as this imbues much more power to lexical knowledge, and suggests that at least some behavioral effects are due to verb knowledge and not event knowledge.




Figure 2. *Was cooking* primes *kitchen* directly due to knowledge of their co-occurrence in language

We should be clear that we are not arguing that a such a word-driven theory, without any support from experiential knowledge, is actually a likely explanation in this case or in general. More likely, both kinds of knowledge are operating and interacting to produce the observed behaviors. But previous research has shown that simple distributional statistics of words within language are capable of predicting a vast range of both grammatical and semantic knowledge (see Redington and Chater, 1998, for a review). For example, Willits, Duran, and D’Mello (2007) showed that verbs are more likely to co-occur with, and be distributionally similar to, related arguments compared to unrelated arguments. This is even true for very specific kinds of semantic relations, like whether or not an instrument like *sword* is semantically obligatory to its verb (*cut* vs. *behead*). Such evidence suggests that a model of word knowledge based on the distributional statistics of words within language would be quite enriched and capable of doing considerable semantic work; work that would typically be considered to be derived from interacting directly with the world. It is therefore conceivable that the kind of effects observed by Ferretti and colleagues could be attributed to our knowledge about how verbs used in

language, and not about or knowledge of events in the world. Such a hypothesis is worth testing, because understanding what behaviors are attributable to word knowledge will give us a better idea of what is not explainable by word knowledge, and therefore what must be explained by reference to experiential knowledge.

In order to test the question of whether or not a behavior is attributable to word knowledge, one needs to first show that the distributional patterns in language are *sufficient* for predicting the observed behavioral effects. Next, if knowledge derived solely from word statistics is sufficient a sufficient explanation, one would like to show whether or not they are a *necessary* (in addition to experiential knowledge) for explaining the behavioral effects. Well-developed hypotheses about the patterns that might exist in language can help one do this, as we will show. Study 1 is designed to test both the sufficiency and necessity of knowledge about word use for predicting the relationship observed by Ferretti et al (2007).

Finally, the strongest (and most difficult) test is whether knowledge about word use, in addition to being both necessary and sufficient, is in actuality the only kind of knowledge being used. Study 2 is a demonstration about the level of generalized knowledge about verb, aspect, and noun relationships one could make based on the language statistics. Study 2, in combination with the results from Study 1, helps us make strong predictions about the behaviors to expect if the knowledge is strongly language based, and follow-up behavioral experiments one could conduct based on these predictions is described in the Discussion.

Study 1

Ferretti et al (2007) found priming from verbs to related locations, but only when the verbs were in the imperfect aspectual form (e.g. faster reaction times for *bathroom* when the prime was “*was showering*” compared to when the prime was “*was drawing*”, but not when the prime was “*had showered*” compared to “*had drawn*.” Such a result is interesting in that it shows us how knowledge is constrained, namely that aspect impacts the kinds of relations that are primed. However, Ferretti et al’s result does not tell us whether that knowledge is derived from real-word experience or from linguistic experience. In this study, we will investigate the extent to which distributional statistics concerning verb-noun co-occurrences account for the pattern of priming that Ferretti observed. More specifically, we will compare the probability of co-occurrence of a verb’s aspectual forms and its related and unrelated locations (*cooking/kitchen*, *cooked/kitchen*, *worshipping/kitchen*, and *worshipped/kitchen*). We will also compare the probability of co-occurrence of the entire verb phrase (*was cooking/kitchen*, *had cooked/kitchen*, *was cooking/kitchen*, and *had worshipped/kitchen*). Comparing co-occurrences of the verb and verb phrase with location will allow us to test both the sufficiency of lexical information for predicting the observed behavioral effects,

and the necessity using lexical knowledge (like the syntactic frames) compared to using experiential knowledge alone. This is illustrated by comparing three possible outcomes:

1.) Verb-location co-occurrence probability predicts verb-location relatedness, regardless of the verb’s aspectual form (*-ing* vs. *-ed*) or syntactic frame (*was* vs. *had*). If this is the case, it is evidence of the ability of lexical information to predict semantic relationships. But in this case it would also show that this lexical information is too powerful, as it would predict that priming would occur under any conditions of verb-location relatedness, and Ferretti et al (2001 and 2007) strongly suggests this is not the case.

2.) Another possibility is that both the verb and verb phrases co-occur more with related locations, but only when they are of the correct aspectual form. In this case *cooking* and *was cooking* would co-occur more often with *kitchen* than unrelated verbs or verb phrases like *worshipping* or *was worshipping*. In addition, *cooking* and *was cooking* also co-occur more often with *kitchen* than do related pairs of the opposite aspectual form, like *cooked* or *had cooked*. This would be evidence that the lexical statistics are sufficient for predicting the behavioral results found by Ferretti et al (2007), and that model based solely on word knowledge could not be eliminated as a possible cause of the effects. But this would also not eliminate the possibility that experiential knowledge is actually the cause, and that the potential word knowledge observed in the language statistics is just redundant or epiphenomenal.

3.) A third possibility is that only full verb phrases, in the imperfect form, are more likely to co-occur with locations. This would suggest not only that knowledge based on word co-occurrence is sufficient, but also that at a minimum some word knowledge (in the form of the distributional structure of the syntactic frames *has* and *was*) is necessary to produce knowledge structure that predicts the behavioral effects. It may even suggest that the patterns of word co-occurrence are the main source of the effect, and that conceptual event knowledge is not the source of the behavior. This would be true because it would show that the verbs alone (and perhaps their event referents), are too unconstrained to predict the actual behavior effects.

Methods

Stimuli The stimuli from Ferretti et al’s (2007) priming experiment were used, consisting of 24 related verb-location items in imperfect form (*was cooking, kitchen*), 24 related items in perfect form (*had cooked, kitchen*), 24 unrelated items in imperfect form (*was worshipping, kitchen*), and 24 unrelated items in perfect form (*had worshipped, kitchen*). The items were counterbalanced such each verb, in each aspect form, occurred twice; paired each time with a related location and with an unrelated location.

Corpus & Procedure We calculated the probability of co-occurrence between the verbs (V) and verb phrases (VP) and locations in a corpus of text. The corpus was obtained by downloading the online encyclopedia Wikipedia (download date 07/17/2006). The corpus contained

approximately 532 million word tokens, and was cleaned and tokenized by the same procedure as in Willits et al (2007). For each V or VP, we counted the number of times each location occurred within ten words preceding or following the V or VP. These co-occurrence totals for each V-location and VP-location pair were then divided by the frequency of the V or VP (to control for differences in word frequency between the two verbs or verb phrases). This resulted in an estimate of the likelihood, given the V or VP, that the location occurred within ten words before or after the V or VP. These V-location (and VP-location) co-occurrence probabilities were then analyzed in a 2 (related vs. unrelated) x 2 (aspect: perfect vs. imperfect) x 2 (item type: verb vs. verb phrase) ANOVA, testing for the effect of these variables on the likelihood of co-occurrence.

Results

Study 1 means are shown in Figure 3, for which there was significant three-way interaction ($F_{(1,23)} = 4.52, p < .05$). Planned follow-up tests revealed that the verbs alone were more likely to co-occur with related verbs than unrelated verbs in perfect ($F_{(1,23)} = 7.674, p = 0.011$) and imperfect form ($F_{(1,23)} = 17.241, p < 0.001$). Related verbs in imperfect form were *not* significantly more likely than related verbs in perfect form to co-occur with their related locations ($F_{(1,23)} = 3.101, p = 0.092$). In contrast, for verb phrases, the verb phrases in imperfect form were significantly more likely to co-occur with related locations than unrelated locations ($F_{(1,23)} = 7.353, p < 0.012$), and more likely to co-occur than locations and verb phrases in perfect form ($F_{(1,23)} = 7.128, p < 0.014$). Related verb phrases in perfect form were not significantly different from unrelated verb phrases in perfect form ($F_{(1,23)} = 0.619, p = 0.439$).

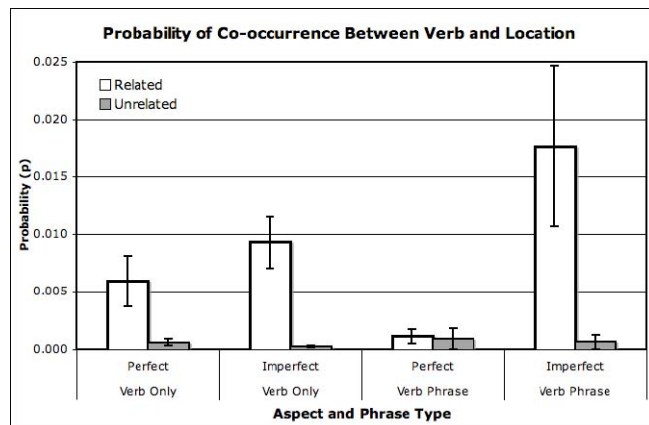


Figure 3. Co-occurrence probability of verbs and locations.

Discussion

Verbs alone were more likely to co-occur with their related locations, and did not vary as a function of their aspectual form. In contrast, verb phrase co-occurrence showed the exact pattern that predicts the pattern of priming found by Ferretti et al (2007), where only the related verb phrase in imperfect form is likely to co-occur with its related location.

This provides further evidence that some of the knowledge being used in these tasks is knowledge about the distribution of words, and not purely experiential knowledge about the distribution of events. This is because the verb alone does not predict the behavioral data; only the verb instantiated in the correct verb phrase does so.

But with this data alone, one cannot distinguish which of the two kinds of knowledge is actually being used in this task. The pattern of data suggests that perhaps the entire effect is due to verb knowledge. Since the verb alone co-occurs with related locations regardless of its aspect, this suggests that perhaps verbs themselves (and their event referents by themselves) are not the source of the interesting knowledge on display in this task (that locations tend to co-occur with ongoing events). But this rationale is only true if the word statistics and experiential statistics are identical, and there are good reasons to think this is not the case (see the General Discussion for these reasons).

A more definitive way to tell the difference would be to use these word statistics to make predictions in behavioral experiments designed specifically to test the difference (such studies are underway by the authors). One such test would be a priming experiment using verb-location pairs for which the verb in its *-ed* form tends to co-occur more often with locations, in contrast to the general trend. For such items, would you still get priming to locations from verbs in an imperfect syntactic frame? If so, this would suggest that you cannot explain the data with language statistics alone, and that real world knowledge is actually the source of the effect. If instead the behavior closely mirrored the word co-occurrence statistics, this would be evidence that word knowledge was driving the effect.

Study 2

Another way to attempt to distinguish whether the behavioral effects are due to word knowledge or to experiential knowledge is to find ways in which these types of knowledge are structured differently, and therefore make different predictions about the kinds of generalizations people would make. One such example is the degree to which one is able to generalize the affect-noun relationship. Using language statistics alone, would someone be able to learn the general principle that verbs in “was X-ing” form co-occur more often with locations than verbs in “had X-ed” form? If one could not learn it generally, could one at least learn it about individual verbs? Or is the only thing one can learn from the language statistics that specific verbs and specific nouns are likely to co-occur? With this information, a follow-up behavioral experiment would be able to test whether or not people seem to generalize in a way consistent with what is available from the language statistics. If not, language statistics are probably not the source of the knowledge.

Methods

Stimuli The stimuli are the same 24 verbs, in both perfect and imperfect form, used in Study 1.

Procedure We calculated the number of times each verb, in both of its aspectual forms, co-occurred with the 50,000 most frequent words in the corpus, within a ten word window. Next, for each of the 50,000 co-occurring words, we computed aspect bias scores by subtracting the word’s co-occurrence probability with the verb’s imperfect form, from the co-occurrence probability of the word with the verb’s perfect form. For example, *restaurant* occurred within a 10-word window of *cooking* 0.054% of the time, but occurred within a 10-word window of *cooked* only 0.029% of the time, giving it a difference score of -0.025%. For each of the 24 verbs, we sorted the 50,000 words by their difference scores, ranking the words in terms of their aspect bias.

Next, we inspected the words and selected the 50 most “was X-ing” biased nouns and 50 most “had X-ed” biased nouns. This produced a list of 100 verb-noun pairs for each verb, half of which were strongly perfectly biased, and half strongly imperfectly biased. Next we randomized the order of the word pairs and then the three authors rated the nouns on a 0-3 scale, in terms of how likely the noun was to be a location for that verb (0 being not at all or never, 1 being possibly, 2 being at least some of the time, and 3 being highly likely). For example, a pair like *cooked-kitchen*, *kitchen* got a rating of 3 because it is highly likely that *kitchen* is a location argument for *cooked*; for a pair like *cooked-competition*, *competition* got a rating of 2 (where at least some of the time it was likely to be used in a way where it was not a location argument for the verb); for a pair like *cooked-TV*, *TV* got a rating of 1 (where the noun was possibly but not likely to be a location argument of the verb); and for pairs like *cooked-potatoes*; *potatoes* got a rating of 0. Finally, we computed average ratings for each verb and tested them to see if one aspect form or the other was more likely to co-occur with locations.

Results

The mean location rating across all items for “was X-ing”-biased words: was 0.702 (SE = 0.136). The mean location rating for “had X-ed”-biased words was 0.610 (SE = 0.125). This difference was significant, $t(23) = 2.13$, $p = 0.044$. Across all items, the correlation between a verb’s aspect and the noun’s location rating was $r = 0.072$.

On a verb-by-verb basis, 15 of the verbs had higher location ratings for words that co-occurred with the verb in “was X-ing” form, and 9 of the verbs had higher location ratings for words that co-occurred with the verb in “had X-ed” form. After a Bonferroni adjustment to control for family-wise error, 5 of the 24 differences were significant; four in a direction such that locations were biased towards “was X-ing”, and one in a direction such that locations were biased towards “had X-ed. The means, variances, p-value for the significance test comparing the two conditions, and the correlation for each verb are shown in Table 1.

Discussion

This data presents an interesting picture of how knowledge might be constrained by structure in language. While the

overall bias towards locations for verbs in their imperfect form was significant, the effect size was small and there is

Table 1. Location scores for verbs in -ing and -ed form

Verb	-ing mean	-ed mean	p	r
act	0.96	0.65	0.012	-0.252
applaud	0.54	0.88	0.008	0.265
browse	0.83	0.86	0.817	0.023
bury	0.73	1.03	0.031	0.216
confess	0.32	0.45	0.183	0.134
cook*	0.83	0.47	< 0.001	-0.355
dance	0.76	0.55	0.068	-0.183
draw	0.67	0.72	0.682	0.041
drive	0.73	0.60	0.232	-0.121
eat	0.69	0.45	0.046	-0.200
exercise	0.39	0.46	0.441	0.078
fish*	0.87	0.47	0.002	-0.312
gamble*	0.76	0.28	0.015	-0.338
hunt*	0.83	0.35	< 0.001	-0.391
mourn	0.75	0.62	0.261	-0.113
pray	0.72	0.56	0.177	-0.136
preach	0.76	0.71	0.704	0.038
shower	0.67	0.40	0.012	-0.250
skate	0.49	0.53	0.723	0.036
sleep	0.86	0.58	0.132	-0.152
stroll	0.63	0.59	0.760	-0.031
study*	0.42	0.89	< 0.001	0.359
swim	0.99	0.77	0.173	-0.137
worship	0.64	0.76	0.374	0.090

*Significant after Bonferroni correction

considerable verb-by-verb variance. Few of the individual verbs are themselves strongly biased (only 5 significantly so), and one of them (*study*) is strongly biased such that “*had studied*” is more likely to co-occur with locations than “*was studying*.” Because of this (and because 37.5% of the verbs showed some location bias towards the “*had X-ed*” form) it is likely that learning that “verbs in imperfect form as a class tend to co-occur with locations” may not actually be the optimal knowledge to learn. However, because some verbs *are* strongly biased, it may make sense to generalize this knowledge at the word level, learning as an example that “*was gambling* tends to co-occur more often with locations than *had gambled*”, and that “*had studied* co-occurs with locations more than *was studying*.” There seems to be sufficient information to abstract beyond individual verb-noun associations and learn word-general knowledge, but to learn class-general knowledge.

This data makes strong predictions about behavior we would expect in an experiment designed to test the level of abstract knowledge people are making, and the degree to which it is derived from experience with words or experience with the world. For example, if it is optimal to form word general knowledge regarding the pairing of aspect and location, one could test this with an experiment involving unrelated locations. If people rely on word knowledge (as opposed to experiential knowledge) for these kinds of judgments, we would expect to find priming or a high degree of association between verbs that are “*was X-*

ing” biased and non-co-occurring but possible locations (like *was gambling* and *hospital*). We would expect the reverse for verbs that are “had X-ed” biased (more priming for pairs like *had studied* and *hospital*).

General Discussion

In this paper, we set out to show that the difference between experiential knowledge and word knowledge is not always clear, and sometimes behavior attributable to experiential knowledge is actually due to word knowledge. Our results strongly make the case that it is possible that this is the case for some kinds of knowledge, like the relationship between aspect and relation types like probability of a location being thematically important.

In Study 1, we showed that the patterns of verb phrase co-occurrences precisely fit behavioral data demonstrated by Ferretti et al (2007). In doing so, we make the case both for the sufficiency and necessity of word knowledge for explaining this particular behavioral phenomena, and are able to make several strong predictions (that can be tested by follow-up experiments) about whether it is, in actuality, word knowledge or experiential knowledge that people are using in these tasks. These predictions are that: 1) if behavior is being driven by word knowledge, then the differential priming of the type observed by Ferretti et al should be dependent on the verbs’ instantiation in the correct syntactic frame, 2) follow-up priming experiments can test if the behavior is truly based on word knowledge rather than experiential knowledge by picking items that violate the “*was X-ing / location*” relationship.

In Study 2, we investigated the predictions we make about the abstractness of knowledge we would expect people to have if they are relying on word knowledge. We found that word statistics predict that this knowledge might generalize for each verb, but likely not across the class of verbs as a whole. As such, one could do a follow-up priming experiments testing to see if behavior actually follows this pattern of generalization.

We think this work highlights the importance of the difference between word knowledge and experiential knowledge, and begins to show how we can tell the two apart and how they are used differently. We really do not think this is an empty or meaningless distinction. It is often argued that distributional statistics in language have the structure they do because they are a “reflection of the structure that exists in the experiential world.” But this is just wrong, as it ignores that the purpose of language is communicating information that is not already shared, not cataloging our experiences with the real world.

Consider a simple example, the word *carrot*. Our experiential concept of a carrot very likely is comprised of a set of characteristic properties of carrots: it is a vegetable, it is orange, pointed, grows in the ground, has a bushy green top, tastes slightly sweet, and so on. But by and large, these are not things we say, because they are already common knowledge. If one looks at the adjectives that are most likely to co-occur with carrot, you get words like *fried*, *organic*,

raw, and *purple*. What is interesting about these words is that they are definitely not characteristic features of *carrot*, but rather possible but not probable states or properties of particular *carrots* in particular situations. What is also interesting is that these co-occurrence properties are just as good at producing taxonomic-like category structures for words, as words like *organic*, *raw*, and *fried* are all more likely descriptors for foods and vegetables than they are for clothing or human-made artifacts (Willits, Schellin, & Burgess, 2004). Both experiential knowledge and word knowledge are likely to be structured in useful ways. But this structure is not redundant; it is complimentary in interesting ways. Understanding this complimentary structure, and it interacts, will allow us to construct a better model of conceptual structure.

Acknowledgements

We would like to thank Mark S. Seidenberg and Maryellen C. MacDonald for helpful feedback on the studies and paper. This work was supported by NIH grant T32 HD049899-01.

References

- Hare, M., Tanenhaus, M. K., & McRae, K. (2007). Understanding and producing the reduced relative construction: Evidence from ratings, editing, and corpora. *Journal of Memory and Language*, *56*, 410-435.
- Ferretti, T. R., McRae, K., & Hatherell, A. (2001). Integrating verbs, situation schemas, and thematic role concepts. *Journal of Memory and Language*, *44*, 516-547.
- Ferretti, T. R., Kutas, M., & McRae, K. (2007). Verb aspect and the activation of event knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *33*, 182-196.
- Gropen, J., Pinker, S., Hollander, M., & Goldberg, R. (1991). Affectedness and direct objects: The role of lexical semantics in the acquisition of verb argument structure. *Cognition*, *41*, 153-191.
- McRae, K., Ferretti, T. R., & Amyote, L. (1997). Thematic roles as verb-specific concepts. *Language and Cognitive Processes*, *12*, 137-176.
- McRae, K., Hare, M., Elman, J., & Ferretti, T. (2005). A basis for generating expectancies for verbs from nouns. *Memory & Cognition*, *33*, 1174-1184.
- Redington, M., & Chater, N. (1998). Connectionist and statistical approaches to language acquisition: A distributional perspective. *Language and Cognitive Processes*, *13*, 129-191.
- Willits, J. A., Schellin, K., & Burgess, C. (2004). The induction of hierarchical structure from global co-occurrence information. Poster presented at the 44th Annual Meeting of the Psychonomics Society, Vancouver, CA.
- Willits, J. A., Duran, N. D., & D’Mello, S. K. (2007). The distributional structure of verbs. Paper presented at the 21st Annual meeting of the Cognitive Science Society.