

Creating Meaningful Word Vectors and Examining their use as Representations of Word Meaning

Jerry T. Ball, Stuart M. Rodgers*, Roger W. Schvaneveldt, Alan R. Ball
Institute for Defense Analyses

Abstract

We identify three shortcomings of word vectors as representations of the full meaning of words: 1) the dimensions of the vectors are implicit and difficult to interpret, 2) the vectors entangle all the meanings and uses of words, and 3) the vectors are unstructured. We propose solutions to each of these shortcomings and explore the implications. Our goal is to integrate word, phrase, and clause level vectors representing fine-grained, associative aspects of meaning into grammatical analysis, to support the resolution of structural ambiguities that cannot be grammatically resolved.

Keywords: word vector; semantic primitive, associative meaning, grammatical analysis; structural ambiguity; Preference Semantics; vector semantics; LLM

Introduction

A key component of Large Language Models (LLMs) (Wikipedia, 2024) are word embeddings that encode the meanings of words as vectors in a high dimensional space (Mikolov et al. 2013). The word vectors made available by Google (2024), and used in this study, have 300 dimensions. The utility of word vectors in transformer based LLMs has recently been demonstrated (Vaswani et al., 2017). Adequately trained LLMs produce a sequence of words that are grammatically and semantically coherent, starting with an initial prompt which is itself a sequence of words. However, as representations of full word meaning, word vectors suffer from at least three shortcomings: 1) the meaning of the vector dimensions is implicit and largely uninterpretable by humans, 2) the word vectors entangle all the different meanings and uses of ambiguous words, and 3) the dimensions of the word vectors are unstructured—i.e. each dimension is assumed to be orthogonal to every other dimension to support vector manipulation and parallel computation. The use of transformers in LLMs addresses the second shortcoming. Transformers are trained to adjust the values of the dimensions of word vectors based on the surrounding context. They allow the LLM to predict the next word given the contextually adjusted values, improving performance significantly. Unfortunately, it is difficult for humans to interpret the adjustments that transformers make because the dimensions of the values that get adjusted are only implicitly meaningful. It is also difficult to improve the word vectors outside of training them using machine learning with huge amounts of data.

Our goal is to use word vectors to represent the fine-grained, associative meaning of words to resolve structural ambiguities that cannot be grammatically resolved, as part of a near human scale computational cognitive model of grammatical analysis (Ball, 2023, in preparation). To

achieve this goal, our plans are 1) to integrate word vectors representing the vector semantics of associative word meaning (Jurafsky & Martin, 2023) into grammatical representations at the lexical level, 2) to disentangle the associative meanings encoded by the word vectors by adjusting the values of the appropriate dimensions based on the evolving grammatical context during grammatical analysis, and 3) to integrate word vectors together into phrase and clause level vectors that represent the vector semantics of phrases and clauses within structured grammatical representations. Basically, we want to combine vector semantic associative meaning representations with structured grammatical representations. Since associative meaning vectors will be integrated into grammatical representations, they can be used to resolve structural ambiguities that cannot be resolved grammatically, like the ambiguities discussed below. The contextually and dynamically disentangled word vectors can also be used as representations of the associative meaning of words that avoid the need for static word senses (Kilgarriff, 1997).

Another important goal is to be able to adjust the word vectors to correct shortcomings and eliminate noise. Achievement of this goal will be facilitated if the word vectors are interpretable by humans. Although word vectors represent associative aspects of word meaning, the learning mechanisms used to create the vectors—relying primarily on co-occurrence statistics and stochastic gradient descent—result in dense vectors that are often flawed and noisy. These word vectors need to be improved either manually or automatically prior to integration into structured grammatical representations.

Creating Meaningful Word Vectors

We begin by demonstrating a simple way to transform an implicit word vector into an explicitly meaningful vector—addressing the shortcoming that the dimensions of implicit word vectors are very difficult to interpret. Explicitly meaningful word vectors facilitate manual or automated adjustment of the vector dimensions. They can also be used to support the resolution of structural ambiguities that cannot be grammatically resolved—especially prepositional phrase attachment ambiguity, and resultative vs. adjunct phrase or clause ambiguity, as demonstrated by the following examples:

*The man **saw** the planet **with a telescope***
*The man saw **the planet with two moons***
The man** saw the planet **in his pajamas
*The man **hammered** the metal **flat***
*The man **hammered** the metal **fast***

The man hammered **the metal** (that was) **rusted**
The man hammered the metal **naked**

We highlight the phrase whose attachment is grammatically ambiguous in bold black and the preferred attachment in bold blue or green, based on our own judgments.

To create explicitly meaningful word vectors, we start with the semantic primitives originally proposed by Wierzbicka (2021). These primitives are attested to exist as words within a wide variety of languages. They include the following core primitives, using English words, quoted from Goddard (2010), and Goddard & Wierzbicka (2014):

Substantives: *I, you, someone, something/thing, people, body*

Relational Substantives: *kind, part*

Determiners: *this, the same, other/else*

Quantifiers: *one, two, some, all, much/many*

Evaluators: *good, bad*

Descriptors: *big, small*

Mental Predicates: *know, think, want, feel, see, hear*

Speech: *say, words, true*

Actions, Events, Movements, Contact: *do, happen, move, touch*

Location, Existence, Possession, Specification: *be (somewhere), there is, have, be (someone/something)*

Life and Death: *live, die*

Time: *when/time, now, before, after, a long time, a short time, for some time, moment*

Space: *where/place, here, above, below, far, near, side, inside*

Logical Concepts: *not, maybe, can, because, if*

Intensifier, Augmentor: *very, more*

Similarity: *like/way*

Goddard & Wierzbicka (2014) claim that every language includes words that capture the meaning of these core semantic primitives. Goddard & Wierzbicka (2014) supplement these core primitives with additional primitives that are to some extent language specific, and add a collection of useful words for minimal English.

From this broader collection of primitives, we extracted 270 single word primitives. After analysis, we added another 30+ primitive words with the goal of having at least one primitive word to function as a proxy for all grammatically salient dimensions of meaning. We use these 300+ primitives as the dimensions to create meaningful word vectors for each of the ~100,000 words and multi-word units in the mental lexicon of our grammatical analysis system. For each primitive dimension, we compute the cosine similarity between the implicit vector for the word whose meaningful word vector we are creating and the implicit vector for the primitive word, and set the value of this dimension to that similarity. We repeat this process for each word in the mental lexicon. The result is a new set of explicitly meaningful word vectors. For a complete listing of the primitives and a considerably more detailed

discussion of key issues, see the long version of this paper (Ball, Rodgers & Ball, 2024, in preparation).

We first examine the meaningful word vectors for several of the words in the preceding examples. Since it is difficult to examine 300+ dimensions of meaning for each word, and since there is a considerable amount of noise in the derived values—making dimensions with lower values less useful as meaning elements—we use a threshold to limit the number of dimensions to be examined in this paper. However, untargeted vector computations typically make use of all 300+ dimensions. Setting the threshold to 0.25, we show the values for primitive dimensions greater than or equal to 0.25, ordered from most to least similar:

man :	saw :	planet :
man: 1.0	see: 0.515	earth: 0.680
woman: 0.766	did: 0.441	moon: 0.502
person: 0.534	could: 0.360	creature: 0.364
someone: 0.496	when: 0.323	country: 0.346
soldier: 0.475	like: 0.298	desert: 0.336
him: 0.454	called: 0.296	sun: 0.320
father: 0.420	sharp: 0.277	environment: 0.302
creature: 0.357	there: 0.277	island: 0.290
people: 0.339	want: 0.275	stars: 0.286
wife: 0.329	front: 0.274	sky: 0.284
God: 0.320	back: 0.274	sea: 0.267
child: 0.316	think: 0.263	biological: 0.265
doctor: 0.314	the: 0.262	thing: 0.262
dead: 0.311	second: 0.258	God: 0.257
knife: 0.310	big: 0.257	ice: 0.254
dog: 0.308		yours: 0.253
cat: 0.299		skin: 0.250
who: 0.298		
snake: 0.281		
body: 0.278		
cow: 0.271		
night: 0.256		
nurse: 0.254		
teacher: 0.250		

These meaningful word vectors make explicit the essentially associative nature of the implicit word vectors provided by Google. The dimensions can also be ordered to highlight the most important dimensions of meaning. An examination of the most similar dimensions for each word reveals that most of the primitives are related either paradigmatically (e.g. *man~woman*, *planet~earth*) or syntagmatically (e.g. *man~dead*, *saw_{too}~sharp*) to the word, with some less obvious relations (e.g. *planet~skin*).

An important question is how the meaningful word vectors compare to implicit word vectors at representing the cosine similarity between words. We want the meaningful word vectors to be at least comparable in performance to the implicit word vectors. We also want the implicit and meaningful word vectors to match human expectations—using our own judgments as stand-ins.

As an example, consider the expression *the man hammered the metal naked*. In this expression, the preferred

interpretation of the adjective *naked* is that it is functionally related to the noun phrase *the man*—i.e. it is the man who is naked. For our initial analysis, we provide comparisons of the word vector similarities for the head nouns *man* and *metal*, and the verb *hammered*, to the adjective *naked*—ignoring the influence of the determiner *the*, for now. To support disambiguation of the grammatical function of the adjective *naked* in this expression, we want the vector for the word *naked* to be more similar to the vector for the word *man* than to the vectors for the words *hammered* or *metal*. Computing the cosine similarity using the implicit vectors and the meaningful vectors, we get the following:

	Implicit	Meaningful
<i>man~naked:</i>	0.254	0.766
<i>metal~naked:</i>	0.086	0.652
<i>hammered~naked:</i>	0.053	0.533

For both the implicit and meaningful word vectors, the cosine similarity between *naked* and *man* is the largest, matching our expectations. Examining a few additional nouns gives:

	Implicit	Meaningful
<i>woman~naked:</i>	0.303	0.759
<i>rock~naked</i>	0.101	0.758
<i>child~naked:</i>	0.096	0.735

The vectors for the words *woman* and *child* have higher similarities to the vector for the word *naked* than the vector for the word *metal*, reflecting their common use as human nouns. However, the vector for the word *rock* has a higher similarity to the vector for the word *naked* than the vector for the word *child* does, perhaps due to the considerable ambiguity of this word (e.g. *a naked rock star*).

For the expression *the man hammered the metal flat*, we get the following:

	Implicit	Meaningful
<i>metal~flat:</i>	0.196	0.658
<i>hammered~flat:</i>	0.187	0.688
<i>man~flat:</i>	0.130	0.616

Note that this expression is ambiguous between having an adjunct phrase headed by the adverb *flat* that modifies the verb *hammered* (i.e. *to hammer flat*), and having a bare resultative clause headed by the adjective *flat* that indicates the resulting state of the object *the metal* (i.e. *the metal is flat*). The implicit word vectors give the cosine similarity between *flat* and *metal* the highest similarity—indicating the preferred grammatical function of the adjective *flat* as functionally related to *metal*. The meaningful word vectors give the cosine similarity between *flat* and *hammered* the highest similarity—indicating the preferred grammatical function of the adverb *flat* as a modifier of *hammered*. Our expectations support either of these attachments.

As an example in which the noun *metal* should have the highest cosine similarity, consider the expression *the man hammered the metal (that was) rusted*.

	Implicit	Meaningful
<i>rusted~metal:</i>	0.344	0.803
<i>rusted~hammered:</i>	0.207	0.567
<i>rusted~man:</i>	0.148	0.661

In both the implicit and meaningful word vectors, *rusted* is more similar to *metal* than *hammered* or *man*, matching our expectations.

Although the results on these examples are promising, the results on other examples are mixed, and using raw cosine similarity between word vectors of head words to establish functional relationships needs to be improved. For example, if we replace *the man* with *John* as in *John hammered the metal naked*, we get:

	Implicit	Meaningful
<i>metal~naked:</i>	0.086	0.652
<i>hammered~naked:</i>	0.053	0.533
<i>John~naked:</i>	0.046	0.520

The implicit and meaningful word vectors for the proper noun *John* are less similar to the vector for *naked* than either the vectors for *metal* or *hammered*. This result strongly suggests that the vector representation for the proper noun *John* needs to be improved. Examining the vector similarities for additional proper nouns, we get:

	Implicit	Meaningful
<i>Susan~naked:</i>	0.097	0.605
<i>Mary~naked:</i>	0.026	0.555
<i>Joe~naked:</i>	-0.008	0.583
<i>Detroit~naked:</i>	0.019	0.549

Only the implicit vector for the female proper noun *Susan* is more similar to the vector for *naked* than the vectors for the noun *metal* (0.086) or the verb *hammered* (0.053). There is also a lot of noise in the implicit vector representations for proper nouns, with the cosine similarity between the vectors for the adjective *naked* and the proper noun *Joe* being negative. Since none of the meaningful word vectors for proper nouns are more similar to *naked* than either *metal* or *hammered*, the derivation of meaningful word vectors appears to have decreased the noise—in the sense that we get the same result for all four proper nouns. However, the meaningful word vectors do not give the expected result.

These initial results suggest that the vector representations for nouns like *man*, *woman*, and *child* match our expectations with respect to *naked* better than proper nouns like *John*, *Joe*, *Mary* and *Susan*. An examination of several pronouns shows that they align with nouns:

	Implicit	Meaningful
<i>she~naked:</i>	0.176	0.803
<i>it~naked:</i>	0.124	0.739
<i>he~naked:</i>	0.118	0.765

The implicit and meaningful word vectors for these pronouns are all more similar to the vector for *naked* than the vectors for *metal* or *hammered*. However, the implicit vector for the pronoun *it* is unexpectedly more similar to the

vector for *naked* than the implicit vector for the pronoun *he*. On the other hand, the meaningful vector for *he* is more similar to *naked* than the meaningful vector for *it*, as expected.

How might we improve the vector representations of proper nouns? We need some mechanism for strengthening the representation of the animacy = human and gender = male or female dimensions of meaning. Since pronouns are tightly constrained in meaning, we can use them as proxies for these weakly encoded dimensions of meaning of proper nouns. Adding the vector for the appropriate pronoun to the vectors for proper nouns to better represent animacy and gender, we get:

	Implicit	Meaningful
<i>Susan+she~naked:</i>	0.164	0.855
<i>Mary+she~naked:</i>	0.120	0.840
<i>John+he~naked:</i>	0.100	0.815
<i>Joe+he~naked</i>	0.057	0.825

The meaningful word vectors for the combination of the proper noun with the appropriate pronoun are all more similar to the vector for *naked*, than the vectors for *metal* (0.652) or *hammered* (0.533). The implicit word vectors pattern similarly, except for the implicit vector for *Joe* plus *he*, likely due to noise in the implicit vector for *Joe*. Since the representation of words in the mental lexicon includes representations of their grammatical features—including animacy and gender in the case of proper nouns (Ball, Chapter 4, 2023, in preparation)—this knowledge can be used to automatically adjust the vectors for proper nouns to strengthen the representation of these two dimensions of meaning, by adding the appropriate pronoun. More testing is needed to determine if untargeted addition of the vector for the appropriate pronoun is sufficient, but initial results are encouraging.

Transforming Word Vectors to Add Context

We next consider how to transform or adjust meaningful word vectors to incorporate context. It is our intention to use meaningful word vectors to represent the fine-grained, associative meaning of words as part of a near human scale computational cognitive model of grammatical analysis. An important element of grammatical analysis is determination of the part of speech of the words in the input. We recently demonstrated a capability to determine the part of speech of words in input expressions at a level of accuracy that is competitive with state-of-the-art machine learning and deep learning systems (Ball & Rodgers, 2023; Ball & Rodgers, 2024, in preparation). Once the part of speech of a word is determined, the high level referential and relational type, and grammatical features of the word become available, and the primitive dimensions that are proxies for these grammatically salient dimensions of meaning can be targeted for adjustment. For example, once it is determined that *saw* is functioning as a verb in *the man saw the planet with a telescope*, the meaningful word vector for *saw*—

which entangles the noun meaning of *saw* as a cutting tool—can be adjusted to reflect the relational type action, the referential type situation, the verb part of speech, and the past tense grammatical feature as encoded by the past tense verb *saw* in the mental lexicon—partially disentangling the associative meaning of *saw* based on the grammatical context. The adjusted word vector can then be used to support grammatical analysis—especially the resolution of structural ambiguities that cannot be grammatically resolved, like the examples considered in this paper.

We turn to an exploration of how to transform word vectors further to incorporate context beyond the grammatically determined referential and relational type, part of speech, and grammatical features of a word. One of the early stages of processing in transformer based LLMs involves a mechanism for encoding positional information. Because the transformer architecture does not have access to explicit symbolic or structural knowledge about the input at this early stage in analysis—which could be used to support positional encoding—it performs that encoding based on mechanisms that do not require such knowledge. A commonly used mechanism is to rotate the dimensions of the implicit word vectors based on the position of the word, with the vectors for words occurring earlier in the input receiving more rotation than words occurring later in the input. Since our adjustments to meaningful word vectors occur in the context of incremental and interactive grammatical analysis (cf. Altmann & Mirkovic, 2009), we are able to use the results of that analysis to reflect positional encoding. We are also able to go beyond positional encoding, since the grammatical analysis mechanism determines the grammatical function of the phrases in the input. In particular, grammatical analysis identifies the noun phrases and prepositional phrases that function as the subject, object, indirect object, and locative argument, among others, within clausal expressions. The phrases performing these functions are available in grammatical function specific working memory buffers. Grammatical analysis also identifies the predicate head of a clause. In addition, the first and most recent phrase of a given type are also identified and stored in working memory buffers to model primacy and recency effects. Given this grammatical knowledge, we can adjust the cosine similarities to reflect not just positional encoding, but grammatical function. For example, in the expression *the man hammered the metal flat*, the noun phrase *the man*, with head *man*, functions as the subject, the verb *hammered* functions as the predicate head, the noun phrase *the metal*, with head *metal*, functions as the object, and the adjective or adverb *flat* either functions as the head of a resultative clause or as the head of a predicate modifier. As an initial exploration of positional encoding, we have assigned the value 1.05 to the head of the subject, 1.10 to the predicate head, 1.15 to the head of the object, 1.15 to the head of the most recent phrase after the object, and 1.0 otherwise. When the cosine similarity between the heads is computed, these values are adjusted to reflect positional encoding and

grammatical function. We show the results below with and without positional encoding (PE):

	Meaningful	Meaningful+PE
<i>hammered~flat:</i>	0.688 (1.0)	0.756 (1.10)
<i>metal~flat:</i>	0.658 (1.0)	0.756 (1.15)
<i>man~flat:</i>	0.616 (1.0)	0.647 (1.05)

Note that positional encoding changes the preferred functional relationship from that of using the hammer to flatten the metal to either that of the metal being flat or that of using the hammer to flatten the metal. Both relationships are semantically and grammatically acceptable.

For the expression *the man hammered the metal naked*, with positional encoding, we get:

	Meaningful	Meaningful+PE
<i>man~naked:</i>	0.766 (1.0)	0.804 (1.05)
<i>metal~naked:</i>	0.652 (1.0)	0.749 (1.15)
<i>hammered~naked:</i>	0.533 (1.0)	0.586 (1.10)

Although the word *man* occurs early in the expression, because it functions as the head of the noun phrase functioning as the subject, it is strengthened (1.05) based on its grammatical function, and it still has the highest similarity to the adjective *naked*, even though the word *metal* which functions as the head of the noun phrase functioning as the object, receives more strengthening (1.15) based on its grammatical position and function.

We have also begun to explore mechanisms for adjusting the word vectors to reflect the associative meaning of surrounding words. Vector addition is not sufficient since it would result in all the adjusted word vectors for the input words being the same. However, vector addition is still feasible if the vector that is being adjusted is given more weight than the vectors for surrounding words that are being added—e.g. for *the man*, give the adjusted vector for *man* 3 times the weight of the vector for *the*, and vice versa. Then normalize. Initial explorations in this direction are discussed in Ball, Rodgers & Ball (2024, in preparation).

Creating Phrasal Vectors

In the previous sections, we explored the use of primitive words to function as proxies for all dimensions of grammatically salient meaning within meaningful word vectors. In this section, we begin to explore the integration of meaningful word vectors to support the representation of the associative meaning of multi-word phrases. We first consider the simple phrasal expression *the man*, consisting of the determiner *the* followed by the noun *man*. How might the meaningful word vectors for the words *the* and *man* be combined into a meaningful vector for the associative meaning of the phrase *the man*? We first explore a targeted approach, and then consider an untargeted approach using vector addition. We begin by showing the meaningful word vector for the word *the*, at a threshold of 0.30:

***the* :**
 the: 1.0 ~ **entity + def desc + determiner + definite**
 this: 0.593
 that: 0.526
 one: 0.466
 on: 0.431
 same: 0.391
 all: 0.378
 time: 0.375
 an: 0.364
 what: 0.365 second: 0.364 those: 0.356 with: 0.354
 not: 0.351 when: 0.350 have: 0.341 where: 0.338
 could: 0.332 some: 0.329 him: 0.325 kind: 0.322
 here: 0.321 place: 0.317 part: 0.317 there: 0.314
 back: 0.306 around: 0.305 two: 0.303 front: 0.303

The primitive word *the* functions as a proxy for the relational type entity, the referential type definite description (def desc), the part of speech determiner, and the grammatical feature definiteness = definite. The vector for the word *man* is shown on page 2. The primitive word *man* functions as a proxy for the relational type person, the part of speech noun, and the grammatical features number = singular, animacy = human, and gender = male.

Since the grammatical category noun phrase is not a part of speech, we must determine how the phrasal category is determined. We use the grammatical concept of a head to do this. Since the noun *man* is the head of the noun phrase *the man*—at least in our approach, and in traditional grammar—we assume that the noun *man* determines the phrase to be a noun phrase. In addition, we assume that the meaningful word vector for the head noun is the base vector for creation of the noun phrase vector. How do we integrate the base vector with the satellite determiner *the*? A simple targeted solution is to set the value of the primitive dimension for *the* in the base vector to 1.0. If we do this, we get the following, at a threshold of 0.3:

***the man* :**
 man: 1.0 ~ **person + NP + sing + human + male**
 the: **1.0** ~ **entity + def desc + determiner + definite**
 woman: 0.766
 person: 0.534
 someone: 0.496
 soldier: 0.475
 him: 0.454
 father: 0.420
 creature: 0.357 people: 0.339 wife: 0.329 God: 0.320
 doctor: 0.314 dead: 0.311 knife: 0.310 dog: 0.308

The resulting phrasal vector represents the associative meaning of the noun phrase (NP) *the man* as an NP (from *man*) that is a definite description (from *the*) of a person (from *man*) with grammatical features singular (sing), human, and male (from *man*), and definite (from *the*). For the targeted integration of satellite words which are not primitives, we propose to use the highest value dimension of the satellite word vector as a proxy, and set the value in

the base vector to the value in the satellite vector. Additional dimensions may also be targeted for adjustment.

Alternatively, we can adopt an untargeted approach of adding the vectors for the words *the* and *man* together, followed by setting the value of these dimensions to 1.0:

the man :

man: 0.918→**1.0** ~ pers + NP + sing + human + male
 the: **0.284**→**1.0** ~ entity + def desc + det + definite
 woman: 0.695
 person: 0.537
 him: 0.522
 the: 0.520
 someone: 0.510
 soldier: 0.434
 father: 0.374
 one: 0.373 people: 0.362 who: 0.352 this: 0.335
 creature: 0.331 when: 0.330 God: 0.314 that: 0.309
 I: 0.308 side: 0.307 child: 0.301 night: 0.300

We compare the performance of the targeted (1) vs. non-targeted approach (2) against the performance of the meaningful word vector for the head noun *man* alone:

	Meaningful Vectors
<i>the man~naked</i> (2):	0.813
<i>man~naked</i> :	0.766
<i>the man~naked</i> (1):	0.742
<i>the man~flat</i> (2):	0.718
<i>man~flat</i> :	0.616
<i>the man~flat</i> (1):	0.604
<i>the man~fast</i> (2):	0.712
<i>man~fast</i> :	0.578
<i>the man~fast</i> (1):	0.567

Untargeted vector addition (2) results in a higher similarity of *the man* to *naked*, *flat*, and *fast*. It has the undesirable side effect of strengthening the encoding of inanimacy, leading to a degradation in overall performance.

We next explore the creation of phrasal vectors for prepositional phrases, using the example *the man saw the planet with a telescope*. In this example, our expectation is that the noun *telescope* prefers to be functionally related to the verb *saw*. However, if we ignore the influence of the preposition *with*, the cosine similarities of the head words do not match this expectation.

Vector Similarity:	Meaningful Vectors
<i>planet~telescope</i> :	0.763
<i>man~telescope</i> :	0.612
<i>saw~telescope</i> :	0.556

The vector for the noun *telescope* is more similar to the vectors for the nouns *planet* and *man* than the vector for the verb *saw*. However, the comitative preposition *with* is a good indicator of the instrument of an action. Given this, we expect the vector for the preposition *with* to be more similar to the vector for the verb *saw* than the vectors for

the nouns *planet* and *man*. Computing the vector similarity of *with* to *saw*, *planet*, and *man* gives the following:

Vector Similarity:	Meaningful Vectors
<i>saw~with</i> :	0.783
<i>planet~with</i> :	0.547
<i>man~with</i> :	0.540

As expected, the vector for the preposition *with* is more similar to the vector for the verb *saw* than the vectors for the nouns *planet* or *man*. The meaningful word vector for *with* appears to adequately represent the comitative or instrument meaning of *with*. We next use vector addition to create a meaningful word vector for the combined words *with* and *telescope*—ignoring the determiner *a*. Comparing the meaningful word vector for *with* plus *telescope* to *saw*, *planet*, and *man* gives the following:

Vector Similarity:	Meaningful Vectors
<i>saw~with+telescope</i> :	0.895
<i>planet~with+telescope</i> :	0.873
<i>man~with+telescope</i> :	0.745

Due to the influence of the preposition *with*, the combined vector for *with* and *telescope* is more similar to the vector for *saw* than the vectors for *planet* and *man*, matching our expectations, and demonstrating that prepositions may make important contributions to associative meaning.

Concluding Remarks

Our current computational cognitive modeling research originates in earlier research in Preference Semantics (Wilks, 1975; Wilks, Huang & Fass, 1983), semantic priming (Meyer & Schvaneveldt, 1971; Schvaneveldt, 2004), and visual word recognition (Paap et al., 1982). However, we now use the ACT-R cognitive architecture (Anderson, 2007; Anderson et. al, 2004; Salvucci, 2018) in place of the Prolog programming language used in that earlier research (Ball, 1992). ACT-R improves on Prolog by elegantly integrating incremental symbolic processing with parallel probabilistic mechanisms for choosing between competing alternatives, and by organizing declarative memory into a multiple inheritance hierarchy (Ball, 2013). Within our near human scale computational cognitive model of grammatical analysis, the integration of meaningful word vectors as representations of associative meaning provides a vector semantic mechanism for determining semantic preferences and for resolving structural ambiguities that cannot be grammatically resolved. Recent advances in the machine learning of word vectors, combined with the development of transformer based LLMs, have paved the way for the integration of a vector semantic capability into grammatical analysis, whether using our approach or other symbolic approaches. More testing is needed to determine when targeted adjustment of individual vector dimensions—made feasible by the existence of meaningful word vectors—is preferred over non-targeted vector manipulation.

Acknowledgments

We wish to acknowledge the important early influence of Yorick Wilks and his theory of Preference Semantics, on the research reported in this paper. Unfortunately, due to his recent passing, we were unable to get his insightful analysis and feedback on this research.

References

- Altmann, G. & Mirkovic, J. (2009). Incrementality and Prediction in Human Sentence Processing. *Cognitive Science*, 222, 583-609.
- Anderson, J. (2007). *How Can the Human Mind Occur in the Physical Universe?* NY: Oxford University Press.
- Anderson, J., Bothell, D., Byrne, M., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111, 1036-1060.
- Ball, J. (2023, in preparation). Double R Grammar. https://www.researchgate.net/publication/353369825_Double_R_Grammar_Book_Chapter_1_Introduction
- Ball, J. (2023, Chapter 4, in preparation). Written Word Recognition. https://www.researchgate.net/publication/353560016_Double_R_Grammar_Book_Chapter_4_Written_Word_Recognition
- Ball, J. (2013). The Advantages of ACT-R over Prolog for Natural Language Analysis. *Proceedings of the 21st Annual Conference on Behavior Representation in Modeling and Simulation*.
- Ball, J. (1992). PM, Propositional Model, a Computational Psycholinguistic Model of Language Comprehension Based on a Relational Analysis of Written English. Ann Arbor, MI: UMI Dissertation Information Service.
- Ball, J. & Rodgers, S. (2023). Cognitively and Linguistically Motivated Part of Speech Tagging: Quantitative Assessment of a Near Human-Scale Computational Cognitive Model. Paper presented at *Virtual MathPsych/ICCM 2023*. Via <https://mathpsych.org/presentation/1267>.
- Ball, J. & Rodgers, S. (2024, in preparation). Cognitively and Linguistically Motivated Part of Speech Tagging: Quantitative Assessment of a Near Human-Scale Computational Cognitive Model. Long paper. https://www.researchgate.net/publication/377463503_Cognitively_Linguistically_Motivated_Part_of_Speech_Tagging_Quantitative_Assessment_of_a_Near_Human_Scale_Computational_Cognitive_Model
- Ball, J., Rodgers, S. & Ball, A. (2024, in preparation). Creating Meaningful Word and Phrase Vectors for use as Representations of Associative Meaning supporting Grammatical Analysis. Long paper. https://www.researchgate.net/publication/378593200_Creating_Meaningful_Word_Vectors_and_Examining_their_use_as_Representations_of_Word_Meaning
- Goddard, C. (2010). The Natural Semantic Metalanguage approach. B. Heine & H. Narrog (eds.). *The Oxford Handbook of Analysis*. Oxford: Oxford University Press, 459-484.
- Goddard, C. & Wierzbicka, A. (2014). Global English, Minimal English: Towards better intercultural communication. https://intranet.secure.griffith.edu.au/_data/assets/pdf_file/0023/346028/Global_English_Minimal_English_Goddard_Wierzbicka_2014.pdf
- Google (downloaded, 2024). Pretrained word vectors. <https://code.google.com/archive/p/word2vec/>
- Jurafsky, D. & Martin, J. (2023). Speech and Language Processing, Chapter 6. Draft of February 3, 2024. <https://web.stanford.edu/~jurafsky/slp3/6.pdf>
- Kilgarriff, A. (1997). I don't believe in word senses. *Computers and the Humanities* 31 (2): 91-113.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient Estimation of Word Representations in Vector Space. <https://arxiv.org/pdf/1301.3781.pdf>
- Meyer, D., & Schvaneveldt, R. (1971). Facilitation in recognizing pairs of words: evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90, 227-234.
- Paap, K., Newsome, S., McDonald, J., & Schvaneveldt, R. (1982). An Activation-Verification Model of Letter and Word Recognition: The Word-Superiority Effect. *Psychological Review*, 89, 573-594.
- Salvucci, D. (2018). *Java ACT-R, a Java Simulation & Development Environment for the ACT-R Cognitive Architecture*. <http://cog.cs.drexel.edu/act-r/>.
- Schvaneveldt, R. W. (2004). Finding meaning in psychology. In A. F. Healy (Ed.), *Experimental cognitive psychology and its applications: Festschrift in honor of Lyle E. Bourne, Jr., Walter Kintsch, and Thomas Landauer*. Washington, DC: American Psychological Association.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, L. & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 5998-6008.
- Wierzbicka, A. (2021). "Semantic Primitives", fifty years later. *Russian Journal of Linguistics* 25 (2), 317-342. DOI: <https://doi.org/10.22363/2687-0088-2021-25-2-317-342>
- Wikipedia (2024). LLMs. https://en.wikipedia.org/wiki/Large_language_model
- Wilks, Y. (1975). Preference Semantics. In *Formal Semantics*. Edited by E. Keenan. NY: Cambridge University Press.
- Wilks, Y., Huang, X. & Fass, D. (1985). Syntax, Preference and Right Attachment. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence (IJCAI-85)*. Los Angeles, CA, pp. 779-784.