



Language, Cognition and Neuroscience

ISSN: 2327-3798 (Print) 2327-3801 (Online) Journal homepage: http://www.tandfonline.com/loi/plcp21

From words-as-mappings to words-as-cues: the role of language in semantic knowledge

Gary Lupyan & Molly Lewis

To cite this article: Gary Lupyan & Molly Lewis (2017): From words-as-mappings to words-ascues: the role of language in semantic knowledge, Language, Cognition and Neuroscience, DOI: 10.1080/23273798.2017.1404114

To link to this article: <u>https://doi.org/10.1080/23273798.2017.1404114</u>

(1	(1

Published online: 05 Dec 2017.



🖉 Submit your article to this journal 🗹

Article views: 26



View related articles 🗹



View Crossmark data 🗹

Full Terms & Conditions of access and use can be found at http://www.tandfonline.com/action/journalInformation?journalCode=plcp21

REGULAR ARTICLE

Check for updates

Taylor & Francis Group

Routledge

From words-as-mappings to words-as-cues: the role of language in semantic knowledge

Gary Lupyan^a and Molly Lewis^{a,b}

^aDepartment of Psychology, University of Wisconsin-Madison, Madison, WI, USA; ^bComputation Institute, University of Chicago, Chicago, IL, USA

ABSTRACT

Semantic knowledge (or semantic memory) is knowledge we have about the world. For example, we know that knives are typically sharp, made of metal, and that they are tools used for cutting. To what kinds of experiences do we owe such knowledge? Most work has stressed the role of direct sensory and motor experiences. Another kind of experience, considerably less well understood, is our experience with language. We review two ways of thinking about the relationship between language and semantic knowledge: (i) language as mapping onto independently-acquired concepts, and (ii) language as a set of cues to meaning. We highlight some problems with the words-as-mappings view, and argue in favour of the words-as-cues alternative. We then review some surprising ways that language impacts semantic knowledge, and discuss how distributional semantics models can help us better understand its role. We argue that language has an abstracting effect on knowledge, helping to go beyond concrete experiences which are more characteristic of perception and action. We conclude by describing several promising directions for future research.

ARTICLE HISTORY

Received 21 July 2016 Accepted 7 November 2017

KEYWORDS

Concepts; language; words; semantics; knowledge

1. Introduction

A central goal of cognitive science and cognitive neuroscience is to understand how people represent and organise semantic knowledge (e.g. Mahon & Hickok, 2016; Yee, Jones, & McRae, in press). Semantic knowledge is a broad construct, including everything one knows about dogs, fruit, knives, things that are green, time machines, and Holden Caulfield from The Catcher in the Rye, etc. Researchers have attempted to understand the format in which semantic knowledge is represented (e.g. Barsalou, 2008; Edmiston & Lupyan, 2017; Laurence & Margolis, 1999; Mahon & Caramazza, 2008; Murphy, 2002) and how it develops (Carey, 2009; Clark, 1973; Saji et al., 2011; Wagner, Dobkins, & Barner, 2013). A guestion that has received relatively less attention is to what kinds of experiences do we owe such knowledge?

We can broadly distinguish between two kinds of experiences. The first includes all our interactions with the world through nonverbal perception and action. For example, we can learn that knives are used for cutting by observing their use and by using them ourselves; we can learn that limes and grasshoppers are similarly green by observing their colours. Although there continues to be disagreement about whether the representations that comprise our semantic knowledge have the same format as the representations used in perception and action, few deny the importance of these experiences for the development of semantic knowledge.¹

The second kind of experience is language.² It is through linguistic experiences that we learn names for things, e.g. that certain cutting tools are called "knives" and that "knives are made from steel." To what extent does our semantic knowledge depend on experiences with language?

One answer is: it depends entirely on the domain. Knowledge in some domains may depend heavily on language. For example, without reading or talking about Catcher in the Rye, we would know nothing about Holden Caulfield (if a movie version were ever released, the linguistically derived knowledge would be further supplemented by the visual experiences of watching the movie). That we know some things by virtue of learning about them through language is not especially controversial (e.g. Bloom, 2002; Painter, 2005), but sometimes a point of confusion. For example, after eviscerating the idea that the language one speaks can affect one's conceptual structure, Devitt and Sterelny concluded that "the only respect in which language clearly and obviously does influence thought turns out to be rather banal: languages provides us with most of our concepts" (1987, p. 178). Presumably

the concepts the authors had in mind were those acquired with the help of language, and perhaps could not be acquired in its absence. Like others who have been puzzled by this example of supposed banality (e.g. Gentner & Goldin-Meadow, 2003), we believe that the possibility that some domains of semantic knowledge owe themselves primarily or even entirely to language is an important observation deserving of greater study. In contrast, knowledge in some domains seems to have little or nothing to do with language. For example, our knowledge that knives are sharp or what a lemon tastes like appears to be independent of language and instead depend entirely on sensorimotor experiences. In this paper, we argue that, in fact, much more of our semantic knowledge may derive from language than is often assumed.

2. Two accounts of the relationship between verbal and nonverbal knowledge

In this section we review two perspectives concerning the relationship between verbal and nonverbal knowledge. The first perspective is that verbal knowledge *maps onto* nonverbal knowledge. This position is sometimes glossed in the literature as the "cognitive priority" hypothesis (Bowerman, 2000) or more commonly, the notion that words map onto concepts. For example, when Snedeker and Gleitman ask "Why is it hard to label our concepts?" (2004), they are assuming that there are nonlinguistic and independently acquired (or else innate non-acquired) concepts that comprise our semantic knowledge, and then we learn words as labels for those concepts (Figure 1(A)). On this perspective, while language enables us to effectively communicate what we know, it plays no significant role in acquiring that semantic knowledge:

The meanings to be communicated, and their systematic mapping onto linguistic expressions, arise independently of exposure to any language. (Gleitman & Fisher, 2005, p. 133)

The second perspective paints a very different picture (Figure 1(B)). Rather than simply mapping onto pre-existing conceptual representations, words help *construct* these representations. Rather than mapping onto meaning, words are *cues* to meaning (Elman, 2004, 2009; Rumelhart, 1979). On this perspective, our semantic knowledge reflects both perceptual and linguistic experiences, but, as we describe in more detail below, the two experiences differ in some key ways. On the words-as-cues view, language is afforded a much more central role in the formation of semantic knowledge, not only as a *source* of knowledge (of the sort you



Figure 1. (A) According to the words-as mapping perspective, words map onto pre-existing concepts. This view places limits on the potential for language to inform semantic knowledge. (B) On the alternative words-as-cues perspective, semantic knowledge derives from both verbal and nonverbal experiences. Both types of experiences contribute to a common representational space (schematiczed here as a manifold). On this view, language can distort semantic knowledge derived from perception/action or even be the sole source of knowledge for some domains.

might obtain by reading a book), but as a way of augmenting knowledge derived from direct motor and perceptual experience. The words-as-cues perspective is logically compatible with the view that humans are born with certain pre-linguistic conceptual primitives that "seed" lexical systems. But by viewing linguistic experience as playing a potentially causal role in creating semantic knowledge, this perspective opens up a variety of empirical questions concerning the ways in which language does (and does not) shape semantic knowledge.

2.1. Perspective 1: words-as-mapping

Perhaps the most widespread view regarding the relationship between verbal and nonverbal knowledge is that words map onto "concepts". It is this view that Li and Gleitman have in mind when they claim that "humans invent words that label their concepts" (Li & Gleitman, 2002, p. 266). The mapping view dominates the literature on word learning (for further discussion, see Barrett, 1986; Tomasello, 2001; Xu & Tenenbaum, 2007). For example, researchers routinely talk about the child "traditional language 'mapping problem' [wherein] children attach the forms of language to what they know about objects, events and relations in the world" (Bloom, 1995, pp. 21–22), and ask "[h]ow do infants begin to map words to concepts, and thus establish their meaning?" (Waxman & Leddon, 2010, p. 180).³ Most of this literature says little about where the concepts that words map onto come from, but rather assumes they are generated by some process independent of language. Siskind provides an illustrative example:

Suppose that you were a child. And suppose that you heard the utterance John walked to school. And suppose that when hearing this utterance, you saw John walk to school. And suppose, following Jackendoff (1983), that upon seeing John walk to school, your perceptual faculty could produce the expression GO(John, TO(**school**)) to represent that event. And further suppose that you would entertain this expression as the meaning of the utterance that you just heard. At birth, you could not have known the meanings of the words John, walked, to, and school, for such information is specific to English. Yet, in the process of learning English, you come to possess a mental lexicon that maps the words John, walked, to, and school to representations like John, GO(x, y), TO(x), and school, respectively. (Siskind, 1996)

For words to map onto concepts, the concepts must exist prior to the words (see Devitt & Sterelny, 1987; Pinker, 1994; cf. Bowerman, 2000; Levinson, 1997; Malt et al., in press) and so it makes little sense to ask about the effect of learning the word "school" or indeed that schools are the sort of thing that "John" can "go to" on our semantic knowledge of schools. If "these linguistic categories and structures are more-or-less straightforward mappings from a preexisting conceptual space, programmed into our biological nature" (Li & Gleitman, 2002, p. 266), then asking what effect the learning and use of language has on their development becomes self-defeating.⁴

In the two sections below, we will describe two problems with the words-as-mapping view: the contextdependence of word meaning and the challenge of cross-linguistic differences. We believe that when taken together these problems are insurmountable and that an alternative way of thinking about the relationship between language and "nonverbal" semantic knowledge is needed. We describe this alternative in Section 3.

2.1.1. The context dependence of word-meaning

The first problem with the idea of mapping words to meanings is that the meaning of a word depends on the context in which it is used. That words are polysemous is of course not news (see e.g. Murphy, 2010; Nerlich, 2003; Sandra & Rice, 1995; Tyler & Evans, 2001). The reason polysemy is such a problem for the words-as-mapping view is that it complicates the process of mapping arguably to the point of implausibility. Consider these sentences:

- (1) The man took the candy
- (2) The man took the car.
- (3) The man took the picture.
- (4) The man took a seat.
- (5) The man took a stand.
- (6) The man took the stand.

How many distinct concepts does the word "take" map onto in sentences 1–6? Does taking the car map onto the same concept of TAKE as taking the picture? To claim that all these instances map onto the same concept is problematic because it fails to account for the differences in meaning. Claiming that all the instances correspond to distinct concepts is also problematic because it fails to capture the similarity relationships that the various senses of "take" share (see also Lakoff, 1990).⁵

The problem of context dependence, however, runs deeper than words having multiple senses. Consider these sentences:

- (7) The fish attacked the swimmer.
- (8) The fish avoided the swimmer.

On any strictly linguistic analysis, "fish" would appear to mean the same thing in (7) and (8). But how would we know if "fish" actually maps onto the same concept of FISH in both cases? One way to find out is by using a cued recall task. In this procedure, participants read a list of sentences that includes either (7) or (8), and are then cued with various words, and asked to recall the sentence most related to the cued word. The word "fish" is similarly effective for cuing (7) and (8). However, recall of (7) can be doubled by cuing people with the word "shark". Using more specific cues such as the names of various typical fish either reduces or has no effect on the recall of (8) (Anderson et al., 1976; see also Anderson & Shifrin, 1980; Garnham, 1979).⁶ On the mapping view, the word "fish" would therefore seem to "map" onto SHARK in (7) but not in (8). Although these problems have long been appreciated as presenting difficulties for word learning, they also point to limitations of the words-as-mapping view as a framework for thinking about the relationship between words and meanings.

2.1.2. Cross-linguistic differences in patterns of naming

If words map to prelinguistic concepts, we might think that the vocabularies of all languages are largely the same, varying only to the extent that speakers of different languages are likely to have different artifacts, plants, and animals in their environment that need to be named. But this is *not* what we find (e.g. Wierzbicka, 1996). The notion of even a small "core vocabulary" common to all languages has been described as a "mystical concept" (Borin, 2012). This poses a fatal problem for the words-as-mapping view. If languages differ in their vocabulary, how can we ever know what concept a word of a language supposedly maps onto?

One solution is to posit a set of common prelinguistic concepts which all languages draw on, but do not necessarily lexicalise in the same way, i.e. that the mapping between languages and concepts is not-oneto-one and the differences in the vocabularies of different languages owe themselves to different mappings to an otherwise common conceptual space. Indeed, many point out that "not all words map onto concepts" (Carruthers & Boucher, 1998, p. 185) and that "words do not always map onto concepts in a one-to-one manner" (Hoff, 2013, p. 163). Hoff's need to clarify this point highlights the propensity for researchers to assume just such a one-to-one relationship, e.g. see Laurence and Margolis (1999, p. 4).

To see why this solution fails, we need to ask what concepts people who take the word-as-mapping view have in mind when they talk about words mapping onto concepts. Often, the answer leads right back to language. The reasoning seems to be: if there is a word "bird" in English, it is because the words of English pick out the pre-existing concepts; hence our concepts (or at least the basic ones) just correspond to the words of English. For example, Laurence and Margolis (1999) write that they "won't worry about the possibility that one language may use a phrase where another uses a word," and assume that simple expressions correspond to simple concepts, e.g. "BIRD rather than BIRDS THAT EAT REDDISH WORMS IN THE EARLY MORNING HOURS" (p. 4). The assumption here is that there is something inherently simple and basic about the category BIRD and it is because of this simplicity that it maps onto a single English word. The problem is that this reasoning is potentially circular. If language plays a causal role in conceptual development, then the apparent naturalness of BIRD may owe itself to experience with the word. Malt et al. (2015; see also Wierzbicka, 2013) paint a dire picture of the extent to which researchers have relied on words to make inferences about nonverbal semantic knowledge:

... the prevailing assumption seems to be that many important concepts can be easily identified because they are revealed by words – in fact, for many researchers, the words of English.... words such as "hat", "fish", "triangle", "table", and "robin". [Researchers often take] English nouns to reveal the stock of basic concepts that might be innate ... [and] work on conceptual combination has taken nouns such as "chocolate" and "bee" or "zebra" and "fish" to indicate what concepts are combined (Malt et al., 2015)

However basic the meaning of "bird" may be to English speakers, the *English* meaning which excludes bats and grasshoppers, but includes penguins and emus is, in fact, not universal. For a speaker of Nunggubuyu, the closest word denotes a category that includes both bats and grasshoppers. One may surmise that their intuition would be that the English meaning of "bird" is the more unnatural one (Wierzbicka, 1996, p. 151). Even words like "eat" and "drink" which seem to map onto actions in which all humans engage are *not* lexical universals (Wierzbicka, 2009).⁷

This non-universality of word meanings is a problem for the words-as-mapping perspective because it means that we cannot use words as proxies for prelinguistic concepts. For example, imagine that you are an English-speaking researcher interested in understanding the nonlinguistic semantic representations of body parts. Quite reasonably, you may assume that the concepts HAND and ARM are basic nonlinguistic meanings which map onto the nonlinguistic categories of "hand" and "arm". But now suppose you are a Russian-speaking researcher who has learned, as part of learning Russian, to use the term ("ro'ka") which refers to a region that in English corresponds to both the "hand" and the "arm". You might naturally assume that the word "ruka," neatly maps onto the prelinguistic and basic concept RUKA⁸ with the part of the RUKA picked out by "hand" an example of a conceptual elaboration. Who is right?

If this sounds like what Whorf was describing when he wrote of a "principle of relativity which holds that all observers are not led by the same physical evidence to the same picture of the universe, unless their linguistic backgrounds are similar" (Whorf, 1956, p. 214), it is the observers in this case are researchers and the "picture of the universe" concerns the conceptual inventory humans are thought to possess. There is some irony that many of the researchers who rely on the language they happen to speak to make inferences about universal concepts also dismiss Whorf's observations about how the linguistic constructs that we have internalised can skew our perspective of what is natural (see Leavitt, 2011 for a historical perspective of this point). Although some concepts are indeed more basic than others and conceptual complexity is likely to correlate strongly with linguistic complexity (Lewis & Frank, 2016; Shepard, Hovland, & Jenkins, 1961), making valid inferences about conceptual complexity from language requires exhaustive cross-linguistic comparison.

Taken together, we believe that contextual dependence of word meanings and their cross-linguistic variability pose two insurmountable problems for the wordsas-mapping perspective.

2.2. Perspective 2: words as cues

The alternative to thinking about words deriving meanings by mapping onto a separate conceptual landscape, is to think of words as helping to construct meaning, a framework we will gloss as words-as-cues (e.g. Elman, 2004, 2009; Lupyan, 2016; Lupyan & Bergen, 2015; Rumelhart, 1979). On this view, the meaning of a word is "revealed by the effects it has on [mental] states" (Elman, 2004, p. 301). What are the mental states on which words have their effects? They are the same mental states that are activated in response to nonverbal stimuli: sensory states, motor states, affective states, and all the combinations thereof. Just as seeing a raspberry is made meaningful by our prior interaction with it (we learn what they taste like, that red ones are ripe, etc.), so hearing or seeing the word "raspberry" is made meaningful by activating the same types of mental states. On its face, this account is just the embodied (or grounded) cognition account (Allport, 1985; Barsalou, 2008, 2016 for a defense against recent critiques). But instead of asking whether meaning is grounded in sensorimotor states or fully abstracted from them (see Borghi & Cimatti, 2009; Louwerse, 2011; Zwaan, 2014 for thoughtful articulations of various middle grounds), thinking of words as cues encourages us to ask a somewhat different line of inquiry: what do words – and experience with language more broadly – *do* to our mental states that comprise semantic knowledge in the course of learning and using language?

Consider our semantic knowledge of colours. Putting aside the question concerning its format, what kinds of experiences are responsible for colour knowledge? For example, how do we know that cherries and bricks have a roughly similar colour? One answer is that we know this through perceptual experiences. Language would seem to have nothing to do with it. However, many individuals who are congenitally blind can also know the characteristic colours of various objects as well as the relationship between colour, e.g. that orange is more similar to red than to green (Connolly, Gleitman, & Thompson-Schill, 2007; Marmor, 1978). This knowledge could not have been acquired from perception and owes itself solely to *linguistic* experiences. That someone with no direct perceptual experiences of colour can nevertheless know something (indeed, guite a bit!) about colour solely from language raises the question of the extent to which language can inform our semantic knowledge when combined with perceptual experiences. This is the topic of the next section.

3. The consequences of the word-as-cues perspective for understanding the structure of semantic knowledge

In the sections that follow we briefly consider three questions suggested by the words-as-cues approach: (1) How rich is the input from language? (2) Which aspects of semantic knowledge can we learn from language and which can we not? (3) What, if anything, is special about using words to construct meaning? What is the difference between thoughts about red raspberries as cued by as actual raspberries and thoughts cued by linguistic phrases like "red" and "raspberry"?

3.1. How rich is the input from language?

Imagine a learner whose sole source of knowledge is language. No perception, no action, and no interaction. How much, in principle, is it possible to learn from such input? The answer is surprising, at least to us. Consider: In 2011, IBM's Watson computer handily beat Ken Jennings and Brad Rutter, the two best *Jeopardy* players. It won by using knowledge gleaned *entirely* from ingesting language. This language was, of course, previously generated by real, behaving, embodied human being (and we imagine was heavily curated by engineers to maximise Watson's chances of winning). But even so, Watson was able to win with knowledge conveyed entirely through language. However contrived *Jeopardy* may be, this example begins to hint at just how much structure there exists in language that can be potentially exploited by a learner. The possibility of learning conceptual structure *from* linguistic experiences suggests that rather than existing prior to language – as implied by the words-as-mapping view – aspects of conceptual structure are, in principle, learnable from language itself.

It has been long known that linguistic regularities can be used to bootstrap learning. Hearing the phrase "here's some sib" allows even young children to assume that sib is a type of substance, and if one heard "here's a sib", one can infer that "sib" is a countable object (Brown, 1958; Fisher, Hall, Rakowitz, & Gleitman, 1994). When presented with the seemingly nonsensical phrase "The gostak distims the doshes", we may not know what a gostak is, but we know that it is doing something (distimming) and it is doing it to the doshes (Ingraham, 1903).⁹ In more morphologically rich languages, the power of grammar to support meaning is substantially enhanced. The Russian version of the gostak example is "Glokaya kuzdra shteko budlanula bokra i kurdyachit bokryonka" (Uspenskij, 1962) with the bold text highlighting the (meaningless) word stems; the rest are inflectional suffixes. This "nonsense" string communicates a wealth of information to the listener despite lacking "real" words.¹⁰

Knowing something about the grammar helps to construct meaning without needing to first associate the terms with lexical concepts (see also Goldberg, 2003 on constructing meaning from syntax). Now let us up the ante. Imagine the linguistic input is just lots of structured text, for example articles from Wikipedia or a large collection of news-stories. The text is presented as a (very) long string without any tagging or parsing or hand tweaking to a naïve learner that knows nothing about language or the world. How much is it possible to learn from such an input? Perhaps because psychologists have long lived in the shadow of the "poverty of the stimulus", the answer might be "not very much." For example, Gleitman rightly points out that seeing a door open is probably more likely to be accompanied by linguistic expressions like "Look who's home!" rather than anything to do with opening or doors (Gleitman, 1990, p. 21). And so, to the uninitiated, the empirical successes of using the signals present in language to derive meaning are nothing short of astounding.

We will focus on distributional models of semantics which attempt to discover structure in language by tracking the contexts in which words are used. In accordance with the classic dictum "you shall know a word by the company it keeps" (Firth, 1957), words used in similar contexts tend to have similar meanings. By computing contextual similarities, it is possible to capture similarity relations between words that occur in similar contexts even if they never occurred in the same context (Hollis & Westbury, 2016; Lenci, 2008). Beyond capturing similarity relations between single words, it is possible to compose word vectors (Mitchell & Lapata, 2010) to capture similarity not only between sentences that were part of the training set, but entirely novel sentences, for example, recognising the similarity between "cookie dwarves hop under the crimson planet" and "gingerbread gnomes dance under the red moon".¹¹

Although distributional models have had a long history in psychology (Andrews, Vigliocco, & Vinson, 2009; Louwerse, 2011; see Hollis & Westbury, 2016; Lenci, 2008; Yee et al., in press for reviews), arguably the most exciting recent development has been a massive scaling up of these models using more efficient training methods and large corpora of text. The best known of these is Google's word2vec (Mikolov, Chen, Corrado, & Dean, 2013; Mikolov, Sutskever, Chen, Corrado, & Dean, 2013), a three-layer neural network that is trained by being presented one word at a time and attempting to predict the words that surround the input word.¹² Although differing in details, the logic of this predictive model is broadly similar to that of Elman's Simple Recurrent Network (Elman, 1990) which, when trained with a toy corpus was able to extract lexical classes (nouns, verbs) and broad semantic fields (animals, edible items, etc.) (Elman, 2004). The new models go much further by capturing a considerable amount of variance of human word-to-word similarity ratings (e.g. Gerz, Vulić, Hill, Reichart, & Korhonen, 2016; Hill, Reichart, & Korhonen, 2016; Levy & Goldberg, 2014). Here are some similarity relations word2vec captures by simply attempting to predict words from surrounding words:

Similarity (red, orange)	<pre>> Similarity (red, green) > Similarity (blanket.</pre>
bed)	couch)
Similarity (lemon, sour)	<pre>> Similarity (watermelon, sour)</pre>
Similarity (New York, Boston)	<pre>> Similarity (New York, Milwaukee)</pre>

That it is possible to derive such similarity relationships from strings of words is both interesting and relevant for understanding the kinds of knowledge that is implicitly conveyed by language, but would not surprise those familiar with Latent Semantic Analysis (LSA; Landauer & Dumais, 1997) and Hyperspace Analogue to Language (HAL; Lund & Burgess, 1996) models of 20 years ago. However, models like *word2vec* capture not only similarity relations between words (which, as mentioned above, can also be used to derive similarity relations of larger utterances), but learn a kind of semantic compositionality. By calculating a direction vector between one word (or group of words) and another and then translating it onto another word, the network's learned representations are able to perform a kind of analogical reasoning. Figure 2 shows some examples from a *word2vec* skip-gram model trained on Google News corpus (Mikolov, Sutskever, et al., 2013),¹³ and a version of the skip-gram model (fast-text) trained on Wikipedia, (Bojanowski, Grave, Joulin, & Mikolov, 2016) (and available in multiple languages¹⁴). The first example (Figure 2(A)) has by now been widely circulated: the directional vector projecting from "man" to "woman" acts as a kind of feminising operator such that when it is applied to "king", the closest resultant word is "queen." The model is not performing explicit analogy as such. Rather, it is attempting to find the word that minimises the dot product between it and king — man + woman (Levy & Goldberg, 2014).

Figure 2(B) shows additional subtlety captured by these semantic embeddings. The projection from the "United States" to "Washington", when applied to



Figure 2. Sample "analogies" derived from a skip-gram model trained on Google News or Wikipedia (see text for details). The visualisations are schematic projections of the 300-dimensional semantic embedding.

Australia, yields its capital, Canberra. The projection from the United States to New York when applied to Australia, yields Sydney, capturing something like city size or prominence. The projection from the "United States" to "Yosemite" appears to convey a park/wilderness dimension and when applied to Australia yields Tasmania. Figure 2(C) shows an example of some of the morphological structure that the models extracted, having learned that the difference between "walking" and "walked" is analogous to that between "swimming" and "swam" (given that the models are not trained on phonology or orthography, these relationships are derived entirely from patterns of usage). Figure 2(D) shows some interesting hints that the models also learn some deeper relationships. The projection from "breakfast" to "dinner" appears to capture something about relative time. When applied to "morning", it yields "afternoon" (with "evening" and "night" as competitors immediately); when applied to "today," it yields "tomorrow." Of course the difference between "breakfast" and "dinner" is not only one of time. We know, for example, that dinners tend to be larger than breakfasts. It appears that the network does as well. When the lunch:dinner projection is applied to "small", the result is "large".

By computing a vector from multiple words it is possible to capture even more abstract relationships. Projecting from "animal" (or "mammal") to several wild carnivores ("wolf" and "tiger") and then applying the projection to "fish" yields "shark" – a prototypically dangerous fish (Figure 2(E)). Projecting from "animal" (or "mammal") to several common pets ("dog" and "cat") and then projecting to "fish" yields "goldfish," a common pet fish (Figure 2(F)). This is, of course, the very sort of compositionality that was once thought impossible without the use of explicit compositional semantics (Fodor, 2001; Fodor & Lepore, 1996; Gleitman, Armstrong, & Connolly, 2012). It is interesting to consider to what extent such relational sensitivity is an implementation of conceptual theories (Murphy & Medin, 1985).

To be sure, these models are far from perfect and make errors no healthy person would make (see below). To perform well, they require exposure to very large amounts of text; the performance declines considerably when the models instead receive the kind of input that is closer to what children actually hear (Asr, Willits, & Jones, 2016).¹⁵ On the other hand, human learners benefit from many sources of information that are not available to the models. Recall that the sole input to the models are strings of text. They do not benefit from pragmatic inference and, as we mentioned earlier, lack all direct exposure to the world (though of course people who generate the language on which the models are trained, do have such direct access).

The above examples suggest that learning the relationships shown in Figure 2 could, in principle, be driven by language. The information is there. The extent to which people use it remains an open question.

3.2. What kinds of semantic knowledge can we learn from language and what can we not?

The examples discussed in section 3.1 show some of the kinds of knowledge it is possible to obtain from language. Imagine a hypothetical learner whose only input was naturally occurring language. What kinds of knowledge would be difficult or impossible to learn from this input? What kinds of knowledge would be relatively easy to learn? What kinds of knowledge might we be learning *only* by virtue of using language? We will not fully answer these questions here but we highlight below some potentially useful directions for making some progress.

One may suppose that the hypothetical learner whose input is purely linguistic would learn nothing about what things look like, feel like, or sound like. Nevertheless, language captures a surprising amount of perceptual knowledge (and this is why someone who is congenitally blind knows quite a bit about perceptual qualities like colour). However, it is not a coincidence that the examples most frequently used to highlight the semantic savvy of models like word2vec are of the man:woman :: king: gueen variety. In our informal analysis, the model's performance on analogies like ball:round :: banana:? is dismal. None of the top 30 of the model's responses even pertain to shape. Although the model "knows" that apple:red :: banana:yellow, the colours blue, purple, and pink are in close competitors to yellow while green is not. Investigating perceptual gualities like taste and feel, likewise reveals large gaps. Although Similarity(pillow, soft) > Similarity(pillow, hard), the top 30 semantic neighbours of "pillow" do not include "soft" which is one of the most frequent human-produced associations (Nelson, McEvoy, & Schreiber, 2004). Similarly, the models learn that "Firestone" is a kind of tire, but judge tires to be more similar to squares than circles.

A general hypothesis then is that language input is especially useful for generating abstract and relational knowledge and poorer at generating concrete perceptual knowledge (see Gentner & Boroditsky, 2001 for related discussion). To our knowledge, this hypothesis has not been comprehensively tested. Some supportive evidence comes from a series of analyses by Hill et al. (2016) comparing the similarity spaces generated by *word2vec* and related models to human similarity judgments. As mentioned above, the correlation between

the models and human judgments is impressively high (with correlations above .7 being common). However in much of the work the semantic embeddings learned by the models were compared to human judgments of relatedness between word pairs (e.g. Baroni, Dinu, & Kruszewski, 2014). On this measure, leashes are judged as being highly related to both dogs and ropes (confusingly this measure is often called "similarity"). When the learned semantic embeddings are instead correlated with human ratings gathered to prioritise perceptual similarity (e.g. by instructing human raters to rate pairs such as car-tire as low because cars and tires share few perceptual features), the correlations drop to around 0.4 (Hill et al., 2016). This suggests that the distributional semantic models are failing to capture a considerable amount of perceptual similarity. Hill et al. (2016) also reported that the correlations between model and human ratings were stronger for abstract words than for concrete words. However, these results were confounded by lexical class and word frequency. In our own analysis (to be reported elsewhere), we compared word2vec's similarity to human similarity ratings (Hill et al., 2016) for words varying in concreteness (Brysbaert, Warriner, & Kuperman, 2014) while controlling for lexical class, frequency, and number of senses using WordNet's synsets. The correlation between word2vec and human similarity was substantially worse for words with more senses (i.e. more polysemous words), but the effect of concreteness interacted in complex ways with polysemy and lexical class, suggesting that the story is more complex than a simple main effect of concreteness.

To summarise, distributional semantic models learn an impressive amount of abstract information (see section 3.1), but fail to capture some seemingly basic perceptual information such as tires being round. An exciting recent development are multimodal semantic models which combine in a common set of network weights perceptual information (typically representations learned by deep convolutional neural networks), mirroring the schematic of the words-as-cues schematic (Figure 1(B)). These weights ought to correlate with human knowledge better than the approximation given by linguistic patterns or by perceptual information alone, and it appears that they do (Anderson, Bruni, Bordignon, Poesio, & Baroni, 2013; Bruni, Boleda, Baroni, & Tran, 2012; Bruni, Tran, & Baroni, 2014; Frome et al., 2013; Silberer, Ferrari, & Lapata, 2016). This work, still in its early stages, is highly promising for providing further insights into what kinds of knowledge perception and language offer to the learner. It bears mention that although the performance of these models shows that it is - in principle - possible to learn these semantic embeddings from the input, it does not necessarily follow that people's semantic knowledge is learned in this fashion.

3.3. Knowledge through language versus knowledge through perception: language as a means of abstraction

So far we have described some of the surprising richness of language as a guide to semantic knowledge. Linguistic input is surprisingly informative about space, time, relational knowledge, and conveys a surprising amount of what people ordinarily think of as basic perceptual information. The ability to derive – from ungrounded strings of symbols alone – that goldfish are pet fish and that breakfasts come before dinners is, of course, only possible because language is produced by people with grounded experiences and there are limits to perceptual knowledge that language tends to encode (indeed, it may be the most evident perceptual facts may be missing from the language signal precisely because they are perceptually evident). Understanding these limits is an important future direction.

In this section, we ask a question that follows from sections 3.1 and 3.2: To the extent that certain kinds of knowledge can be conveyed by both language and perception, are they conveyed in the same way? For example, we can learn that bananas are yellow by seeing them, or through language (not only by hearing people describe them as yellow, but also by noting the similarity of the contexts in which the words occur). But is there a systematic difference in the kinds of semantic knowledge that may be formed from language versus from perceptual experience? The difficulty of answering this question is due in part to the difficulty in assessing the provenance of one's knowledge (but see Bellebaum et al., 2013; Cross et al., 2012 for examples of effectively manipulating it). However, some insight comes from studies in which linguistic factors are experimentally manipulated while people attempt to learn new categories or use existing knowledge to recognise or make inferences about familiar categories. Below, we briefly review evidence that, under the influence of language, semantic knowledge may become more categorical with consequences for behaviour ranging from basic perception to reasoning (see Lupyan, 2012, 2016; Lupyan & Bergen, 2015; Perry & Lupyan, 2014 for more extended discussion.). To foreshadow: Language appears to promote abstraction. We perceive specific objects and events, but we talk about them categorically. As a consequence, learning and using verbal labels appears to augment perceptual representations to make them more categorical - taking on a form that emphasises category diagnostic features, and thereby

becoming separated (i.e. abstracted) from the specific experience.

Studies of category learning provide one source of evidence that something interesting happens when learning is augmented by language. Infants and toddlers appear to learn categories more effectively when the categories are accompanied by labels (e.g. Fulkerson & Waxman, 2007: Waxman & Markow, 1995: cf. Robinson, Best, Deng, & Sloutsky, 2012). Studies of category learning in adults have likewise shown language to facilitate learning of new categories. When participants were tasked with learning which of two species of "aliens" should be approached versus avoided, they learned about twice as guickly when the to-be-learned categories were labelled (Lupyan & Casasanto, 2015; Lupyan, Rakison, & McClelland, 2007). These results importantly complement the studies with infant and older children (e.g. Casasola, 2005; Fulkerson & Waxman, 2007) in that the adult all knew that there were two categories to be learned, but were nevertheless able to learn them more easily when the categories were accompanied by labels.

Labels continue to aid categorisation even of previously learned very familiar items. Lupyan and Thompson-Schill (2012) conducted a series of cued recognition experiments. On each trial participants heard a cue - either a verbal cue such as "dog" or a non-verbal cue such as a dog-bark. Following the cue, participants saw a picture that either matched the cue at the basic level (e.g. a picture of a dog) or did not match (e.g. a picture of a car). Participants had to indicate whether the cue matched the image. All of the auditory cues were normed to be maximally unambiguous and participants heard each cue and saw each image dozens of times over the course of the experiment. The results showed that hearing words led to consistently faster categorisation of subsequently presented pictures. This label-advantage persisted for new categories of "alien musical instruments" for which participants learned to criterion either names or corresponding sounds - further evidence that the advantage did not arise from the non-verbal cues being less familiar or inherently more difficult to process. The effects of labels were not uniform, but affected the most typical category members more than the less typical instances, as expected if the labels were activating a representation that emphasised category-diagnostic features and abstracted over more idiosyncratic features. Subsequent studies (Edmiston & Lupyan, 2015) tested and confirmed the prediction that non-verbal cues such as dog barks activated states that were better matches to specific dogs whereas verbal labels appeared to activate states that were more abstracted - taking on a form that emphasised category-diagnostic features.

To further understand the mechanisms of this label advantage, Boutonnet and Lupyan (2015) measured electrophysiological responses (ERPs) elicited by images that matched or mismatched previously presented verbal and nonverbal cues, e.g. a picture of a dog or car following the word "dog" or a barking sound. Both verbal and nonverbal mismatching trials elicited identically strong N400 responses - often interpreted to mean that both cuepicture mismatches were equally "surprising" at a semantic level. However, only verbal cues elicited differential early visual (P100) responses to matching vs. mismatching pictures. These early visual responses predicted behavioural categorisation responses occurring half a second later. These results show that when requiring recognition of images at the level of basic categories, a label activates more categorical representations which allow for more efficient visual recognition of the image - a form of category-based attention (Lupyan, 2017). Further evidence that verbal labels elicit more categorical representations comes from the domain of colour. In a series of colour discrimination tasks, Forder and Lupyan (2017a, 2017b) found that colour names – e.g. hearing the word "green" - affected people's visual discrimination accuracy by facilitating discrimination of category members from nonmembers and distinguishing category typical from atypical colours. Ostensibly the same information conveyed via visual cues had no comparable effect on visual discrimination. Combined, these results suggest that language activates visual representations that are partly constitutive of visual knowledge (Edmiston & Lupyan, 2017), but in so doing, augments them into a more categorical form than when ostensibly the same representations are activated by nonlinguistic inputs.

Just as adding linguistic experience can enhance categorisation, interfering with language can impair it. In one study (Lupyan, 2009), participants completed an odd-one out task where they had to choose which picture or word did not belong based on colour, size, or thematic relationship while under conditions of verbal interference. Based on prior work suggesting that individuals with anomic aphasia, a subtype of aphasia characterised by poor naming combined with good language comprehension, had specific difficulty with categorisation tasks requiring focus on a specific perceptual dimension (Davidoff & Roberson, 2004), we predicted that verbal interference would specifically affect colour and size categorisation blocks while leaving thematic categorisation unaffected. This is precisely what was found. Unlike appreciating the similarity between a saucepan and a refrigerator (items which cohere together in a variety of ways), appreciating the similarity between a cherry and a brick requires projecting the representations onto a colour dimension, and

representing the overlap. In subsequent work, we have confirmed earlier reports that this type of "low-dimensional" categorisation is impaired in individuals with anomic aphasia (Lupyan & Mirman, 2013). In another test of the idea that verbal labels promote the formation of low-dimensional categories, Perry and Lupyan (2014) had participants learn simple visual categories of "minerals" (Gabor patches) defined by two dimensions: orientation (more vs. less steep) and spatial frequency (higher vs. lower contrast). The space of exemplars was set up so that participants could use a single dimension - either orientation or spatial frequency - or integrate both dimensions. The results showed that when implicit labelling was subtly interfered with using cathodal stimulation over Wernicke's area, participants were more likely to learn categories that incorporated both dimensions. In contrast, control participants were more likely to learn a more one-dimensional representation of the categories. It is this process of dimensional reduction, we think is aided by language. Although seeing colours certainly does not require or depend on language, categorising objects by their colour similarity is another matter, as Koemeda-Lutz et al., remarked, "red cherries and red bricks may be judged to be alike mainly via what is concentrated and coined in the verbal label 'red'" (Koemeda-Lutz, Cohen, & Meier, 1987).

Why do labels have the effect of making our representations more categorical? Consider colour as an example. All of our perceptual experiences with colour involve specific objects with specific size, texture, location, and all the other properties intrinsic to perceiving visual objects. We do not see red. Rather, we see specific instances of redness. Contrast this with the experience of hearing or reading something described as "red." The word abstracts over shades of red in a way that perceptual experiences of redness do not; an actual experience of redness cannot have an ambiguous hue. A linguistic experience of redness (i.e. talk about "red things") can (although, frequently more specific shades can often be recovered from context: cf. "red hair" vs. "red car"). We have previously referred to this aspect of language as motivation (Edmiston & Lupyan, 2015; Lupyan & Bergen, 2015). Perceptual cues are motivated. For example, a dog bark is *motivated* by such factors as the dog's size; the larger the dog the lower the pitch of the bark. Just about every perceptual input is motivated in this way. In contrast, verbal cues are (for the most part) unmotivated. Normally, we cannot tell from how one says "dog" what size dog they have in mind.¹⁶ The uncertainty inherent in language promotes the formation - in both developmental time and in the moment - of representations that represent category diagnostic information abstract over idiosyncratic information. The and

consequences of these more categorical representations are substantial, spanning basic perceptual tasks (Lupyan, 2008; Lupyan & Spivey, 2008, 2010a, 2010b; Lupyan & Ward, 2013) and higher level reasoning (Lupyan, 2015).

4. Outstanding questions and further directions

We have argued for a shift in thinking about the relationship between verbal and nonverbal knowledge and the influence of language on semantic knowledge: away from thinking of words as mapping onto pre-existing concepts, and toward focusing on words as cues, which along with perception and action help construct our semantic knowledge across developmental and inthe-moment timescales. In the sections above, we have described a number of what we believe are particularly exciting advancements such as the development of large-scale distributional semantics models and the exciting development of models that learn multimodal embeddings that combine information from perception and language. In this last section, we describe some additional research directions and experimental hypotheses suggested by thinking about semantic knowledge from a words-as-cues framework. These are organised around three themes: (1) expanding the study of context, (2) consequences of cross-linguistic differences, and (3) consequences of the learning source on semantic knowledge.

4.1. Expanding the study of context

If a language acts "directly on mental states" (Elman, 2004) - the same mental states that constitute all our general knowledge about the world, then, in principle, anything a speaker knows can be brought to bear on understanding a given linguistic utterance (Casasanto & Lupyan, 2014 for discussion). That is, the effect that a linguistic input has on the hearer's mind depends on the hearer's current state of mind. The role of context in language processing has been studied extensively, but much of this work has been focused on how expectations set up by language influence subsequent linguistic processing (e.g. Tabossi & Johnson-Laird, 1980). We advocate for a much broader view of context. For example, the meaning of "puzzle" - its effect on your mental states - ought to be different if said by a child (which may bring to mind a big colourful jigsaw puzzle) than by a grandma (in which case it might bring to mind cross-word or Sudoku puzzles). Thus although in both cases the word may be activating a more categorical state than one induced by seeing a particular puzzle, the details of this abstraction may differ considerably depending on context.

There is some existing evidence that is consistent with this prediction. For example, van Berkum et al. (Van Berkum, van den Brink, Tesink, Kos, & Hagoort, 2008) found a larger N400 ERP response (interpreted as indexing a semantic anomaly) when participants heard sentences like "Every evening I drink some wine before I go to sleep" spoken by a child speaker, or "I have a large tattoo on my back" spoken in an upper-class accent. The difference in the N400 component was found at the same latency as those for purely semantic anomalies ("Dutch trains are sour and blue"), about 200 ms after the acoustic onset of the unpredicted information (e.g. "tattoo", "sour"), showing that putatively pragmatic information is operating on the timescale as classically "semantic" information, blurring the boundary between these domains. An important direction for future research will be to better understand such effects (Hagoort & Van Berkum, 2007; Van Berkum, 2008; see also Willits, Amato, & MacDonald, 2015).

4.2. Consequences of cross-linguistic differences for semantic knowledge

Different languages vary in the ease of expressing the "same" idea. Skeptics should attempt translation (Eco, 2008). One problem, discussed in section 2.1.2, is that seemingly equivalent words generalise differently in different languages. For example, in English, one can "open" an envelope, a drawer, and a bag. Korean uses different verbs for these actions, but conflates opening an envelope, unwrapping a package, and removing wallpaper (Bowerman & Choi, 2003). At issue is not whether an English speaker can conceptualise what opening an envelope, unwrapping a package, and removing wallpaper have in common or whether a Korean speaker can tell them apart. At issue is what are the consequences of having linguistic experiences that make these actions more or less similar by virtue of the language statistics. The input from English highlights doors and envelopes as having something in common (the ability of being opened) while the input from Korean does not.

It may appear that such differences in linguistic patterns are trivial in the face of perceptual evidence. We think otherwise. Consider again the case of colour categories. One could argue that the fact that Russian distinguishes between light and dark blue using distinct lexical items is trivial given that ostensibly the same meanings can be expressed in English simply by adding a modifier: instead of "siniy", "dark blue;" instead of "goluboy", "light blue". However, even this seemingly superficial difference has potentially profound

consequences. Although one can certainly describe colours as "light blue" in English, one cannot describe a colour as simply "blue" in Russian. Second, although in English "dark blue" means a darker blue than just "blue", what counters as dark? In our analysis of a large database of English colour naming (Munroe, 2010), we found that the only colour names showing substantial agreement (>50%) were those that were named by single lexical items (blue, green, red, purple, etc.). That is, English speakers agree far more on what colours are "blue" than what colours are "light blue". We would expect Russian speakers to show considerably more agreement on what colours are "siniye," a prediction that has not been tested, to our knowledge. If confirmed, it would speak to the power of a common language to help align semantic knowledge across individuals (Lupyan & Bergen, 2015).

Cross-linguistic differences in meanings can also be observed even when words seem to have one-to-one translations. One widely used method for investigating word meanings (and semantic knowledge more generally) is asking people to generate associates given a word. The assumption is that two people whose representation of, say, "snow" is similar, will tend to generate similar associates when cued with "snow." Do words that seem to have one-to-one translations mean the same thing in different languages? Consider the words "snow", "idea", "cheese", and "jealousy" – words with straightforward translations into Dutch.

Figure 3 shows patterns of word associations for four words in English and their closest Dutch translations (De Deyne, Navarro, & Storms, 2013). Association patterns for "snow" and "idea," have very similar associates. If we take word associations to reveal something about semantic knowledge, we can conclude that the semantic representations of "snow" and "idea" in English and Dutch speakers are quite similar. In other cases, however, speakers of English and Dutch show substantial divergence. For example, given the cue "cheese," English speakers are more likely than Dutch speakers to produce the word "cheddar" and "mouse" while Dutch speakers are more likely to produce the words "yellow" and "holes." Some of these differences probably derive from differences in direct experience. Cheddar cheese is far more popular in the United States than in the Netherlands, and (we surmise) the opposite may be true for cheese varieties with "holes". Other differences however, likely have owe themselves to differences in the *linguistic* input. For example, the association between "cheese" and "mouse" - stronger for English than Dutch speakers probably has more to do with differences in patterns of language use than differences in direct experience. Observed differences in word associations for abstract



Figure 3. Word associates from English and Dutch speakers for two relatively concrete ("cheese" and "snow") cue words and two more abstract cue words ("idea" and "jealousy"). For each cue word, we show the top 10 associates in English (left) and their Dutch translations (right). "Snow" and "idea" show similar patterns of associates; "cheese" and "jealousy" differ more substantially.

words like "jealousy" reinforce this point. Dutch speakers are more likely to associate "jealousy" with "man" and "woman" while English speakers are more likely to associate it with "anger" and "rage," a difference we also think is more likely due to language than by differences in direct experiences of jealousy. Preliminary evidence for hypothesis comes from analyses of comparing word-embeddings generated from Dutch and English input (Wikipedia-trained fast-text vectors). Correlating the semantic structure of the English semantic space surrounding the four words in Figure 3 revealed a .85 correlation between the Dutch and English semantic neighbourhood of "idea", but only a .3 correlation between the neighbourhoods of "jealousy" (the two concrete nouns, "cheese" and "snow" have correlations of .75 and .55, respectively, going against the pattern shown by the word associations). Further investigation of this issue is clearly necessary.

4.3. Does the source of knowledge matter? Investigating individual differences

If our semantic knowledge is derived from a combination of perceptual and linguistic experiences and linguistic experiences help to make perceptually derived knowledge more categorical and abstracted, then people whose knowledge in a certain domain is more linguistic in origin may have a more categorical and abstracted representation than someone whose knowledge was more perceptually grounded. One consequence is that a person who learned, e.g. that alligators are green from books may be less likely to display interference of this knowledge by visual interference (Edmiston & Lupyan, 2017). To the extent that verbally learned knowledge also leads to convergence in the face of varying perceptual inputs, relying on language may also help people converge onto common semantic representations.

As a preliminary test of the latter prediction, in a recent study (Lupyan & Lewis, in prep.) we asked people to complete a word-association task and to answer several questions about the extent to which they experience thinking as an inner monologue (a measure on which people substantially differ, though for reasons still not well understood). People who responded as engaging more in inner monologue produced word associates that were less idiosyncratic and more similar to one another - as expected if words promote category abstraction. A guestion for future research is to explore the extent to which different experiences with language (e.g. reading more versus less, regularly talking to more vs. fewer different people, being monolingual vs. multilingual) affect the degree of alignment, or "success," in the context of communication tasks (e.g. Clark & Wilkes-Gibbs, 1986).

5. Summary

To what kinds of experiences do we owe our knowledge about the world? One kind, extensively studied, is engaging with the world through perception and action. Another kind, of which we know preciously little, is all the information we may be learning from language. Traditionally, many researchers have assumed that the relationship between language and semantic knowledge is one of mapping such that words derive meaning from being mapped onto conceptual representations that are formed independently of language. We criticise the words-as-mapping view as untenable, and argue for an alternative wherein our semantic knowledge is structured by both direct perceptual and action experiences as well as linguistic experiences. On this view, words, like other perceptual inputs, are cues to meaning and help to construct our conceptual repertoire (Elman, 2004, 2009; Lupyan, 2016; Lupyan & Thompson-Schill, 2012). This view places language alongside perception and action in its ability to structure semantic knowledge. One source of supporting evidence for the potential of language to structure knowledge comes from distributional semantics models that demonstrate the impressive amount of information that language conveys about space, time, relations, and even some basic perceptual facts. Recent work showing that distributional semantics mirror common social biases (Caliskan,

Bryson, & Narayanan, 2017) further raise the stakes for understanding the causal connection between language and semantic knowledge.

Another source of evidence for the influence of language on knowledge comes from empirical studies showing that language augments category learning and the dynamics of activation of semantic/perceptual knowledge. Language appears to make semantic/perceptual representations more categorical.

Despite some tantalising hints at the power of language to structure semantic knowledge, many key issues remain unanswered. These include developing a better understanding of the role of language on a developmental timescale and the impact that different patterns of lexicalisation in different languages may have on the structure of semantic knowledge.

Notes

- According to the Language of Thought hypothesis (Fodor, 1975, 2010), although many semantic facts are learned (for example that switchblade knives are illegal in some places), the (lexical) concept KNIFE – and all other lexical concepts – are innate. This assertion is largely ignored by practicing cognitive scientists and cognitive neuroscientists, but generative linguistics (or at least the minimalist programme version of it) seems to depend on the *a priori* existence of lexical concepts because otherwise the *merge* operation would have nothing to merge (Chomsky, 2010; see Bickerton, 2014 for discussion).
- 2. Language use of course involves perception and action, but for our purposes it is useful to distinguish between perceiving nonverbal stimuli (e.g. a real dog or a picture of a dog) and perceiving linguistic/symbolic stimuli (the word "dog"), and, likewise, distinguishing actions involved in making a sandwich and the action of using language to describe making a sandwich.
- A somewhat different sense of "the mapping problem" involves figuring out the *local* reference of a word, i.e. that an utterance of "apple" refers to the *particular* apple sitting on the table (Lewis & Frank, 2013; McMurray, Horst, & Samuelson, 2012).
- 4. Not everyone who relies on the words-as-mapping view denies the role of words on conceptual development. For example, Waxman and colleagues have long argued that words facilitate infants' and older children's category learning (Balaban & Waxman, 1997; Fulkerson & Waxman, 2007; Waxman & Markow, 1995). Findings that words facilitate category learning are unexpected on the view that the categories existed prior to word learning.
- 5. One especially well-studied case of ambiguity is the word "some" which can either mean "some but not all" or "all". One proposed solution to such ambiguities is to posit that there is a core meaning of "some" which is then modified by pragmatics making the word only look polysemous (Grice, 1957; Huang & Snedeker, 2009; see also Jackendoff, 1990). While the role of pragmatics in constructing

meaning is indisputable, we are unsure of what independent evidence supports the existence of "core meanings" in the minds of the speakers – meanings that are simultaneously abstract and precise enough to give rise to all the attested senses.

- 6. This phenomenon was termed "instantiation": a process wherein a more general (e.g. "fish") is interpreted as a more specific word (e.g. "shark"). Debate ensued as to whether instantiation effects were better understood as "refocusing" or "restructuring" (see Roth & Shoben, 1983 for discussion). For present purposes, we ignore the difference between these accounts, but note the parallel between such instantiation effects and those described by Zwaan, Stanfield, and Yaxley (2002) wherein people e.g. recognise a picture of an eagle with an outstretched wings faster after reading a sentence about an eagle in the sky.
- 7. We are not claiming that there are no universal dimensions to linguistically expressed meanings and are sympathetic to proposals such as Wierzbicka and Goddard's Natural Semantic Metalanguage (Goddard & Wierzbicka, 2002; Wierzbicka, 1996), but the semantic primes of this proposed metalanguage share little with the kinds of concepts (CHAIR, DOG, SCHOOL) to which words are often thought to map onto.
- 8. Russian speakers refer to the hand by using the phrase "kist' ruki", but this phrase refers to a part of the arm rather than to a separate body part. To the extent that English speakers endorse the claim that the hand is attached to the arm rather than being a part of the arm, the meanings of "kist' ruki" is not a direct translation of "hand". Of potential interest, the typical meaning of kist' is a [paint]brush, and is historically derived from the root that denoted a "bunch" or "bundle" (e.g. of twigs). The additional sense to refer to the hand is a later derivation (Fasmer, 2009). To speculate, there may be an analogy drawn between the fingers of the hand and the bristles of a brush, though such connections are unlikely to be psychologically real. For additional discussion of the semantic organisation of knowledge related to the body, see Majid (2015).
- See http://iplayif.com/?story=http://parchment.toolness. com/if-archive/games/zcode/gostak.z5.js to get a firsthand sense of semantics conveyed purely by English syntax and morphology.
- See also https://research.googleblog.com/2017/03/anupgrade-to-syntaxnet-new-models-and.html for an example of the latest version of Google's state-of-theart parser (Parsey McParseface) applied to such "meaningless" sentences.
- 11. The example comes from an online lecture by Baroni http://docplayer.net/31565915-Distributional-semantics. html.
- 12. We focus on this skip-gram instantiation of the word2vec model. An alternative instantiation is continuous bag of words (CBOW) which involves presenting the context as input and learning to predict the target word. Besides neural network based models which are trained gradually using sliding word or context windows, there are now large-scale models utilising global word co-occurrence counts. The most successful of these is GLoVE (Pennington, Socher, & Manning,

2014) which performs a bit better than word2vec when trained on very large corpora and somewhat worse when trained on smaller corpora.

- 13. https://code.google.com/archive/p/word2vec/.
- https://github.com/facebookresearch/fastText/blob/ master/pretrained-vectors.md.
- 15. The networks' performance can also be fragile. For example, the shark/goldfish analogies (Figure 2(E–F)) work for a model trained on Wikipedia, but not for one trained on the Google News corpus. The Wikipedia-trained model correctly relates *cats:leopards* to *dogs: wolves*, but fails to relate *cats:lions* to *dogs:wolves*, instead outputting *bulls, eagles, donkeys* in place of "wolves". It also fails relating *cat:leopard* to *dog:wolf*, outputting in place of "dog" *lion, boar*, and *leopard*. Note, however, that the model still shows sensitivity to the count status of the nouns.
- 16. The idea of motivation is related to Grice's distinction between natural and non-natural meaning (Grice, 1957), with natural meaning mapping being motivated and non-natural (i.e. conventional) being unmotivated. There are some differences though. For example, using applause to signal approval might be seen as nonnatural in that it is conventional and non-indexical, but it is nevertheless motivated in that the length of applause tends to correlate with the amount of approval. One cannot applaud without committing to some length and the length conveys meaning.

Acknowledgements

We thank Bill Thompson for critical help with the cross-linguistic analyses reported here.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was partially supported by National Science Foundation [grant number PAC 1331293] to G.L and prepared while the first author was in residence at the Max Planck Institute for Psycholinguistics.

References

- Allport, D. A. (1985). Distributed memory, modular subsystems and dysphasia. In S. K. Newman & R. Epstein (Eds.), *Current perspectives in dysphasia* (pp. 32–60). New York: Churchill Livingstone.
- Anderson, A. J., Bruni, E., Bordignon, U., Poesio, M., & Baroni, M. (2013). Of words, eyes and brains: Correlating image-based distributional semantic models with neural representations of concepts. In *Proceedings of the conference on empirical methods on natural language processing* (pp. 1960–1970). Seattle, WA. Retrieved from http://www.aclweb.org/ anthology/D13-1202
- Anderson, R. C., Pichert, J. W., Goetz, E. T., Schallert, D. L., Stevens, K. V., & Trollip, S. R. (1976). Instantiation of general

terms. Journal of Verbal Learning and Verbal Behavior, 15(6), 667–679.

- Anderson, R. C., & Shifrin, Z. (1980). The meaning of words in context. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 331–348). Hillsdale, NJ: Erlbaum.
- Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, *116*(3), 463–498. doi:10. 1037/a0016261
- Asr, F., Willits, J., & Jones, M. (2016). Comparing predictive and co-occurrence based models of lexical semantics trained on child-directed speech. In *Proceedings of the 37th meeting of the cognitive science society* (pp. 1092–1097). Retrieved from https://mindmodeling.org/cogsci2016/ papers/0197/paper0197.pdf
- Balaban, M. T., & Waxman, S. R. (1997). Do words facilitate object categorization in 9-month-old infants? *Journal of Experimental Child Psychology*, 64(1), 3–26.
- Baroni, M., Dinu, G., & Kruszewski, G. (2014). Don't count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of the* 52nd annual meeting of the association for computational linguistics (pp. 238–247). Baltimore, MD. Retrieved from http:// anthology.aclweb.org/P/P14/P14-1023.pdf
- Barrett, M. D. (1986). Early semantic representations and early word-usage. In S. A. I. Kuczaj & M. D. Barrett (Eds.), *The devel*opment of word meaning (pp. 39–67). New York: Springer. doi:10.1007/978-1-4612-4844-6_2
- Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, *59*(1), 617–645. doi:10.1146/annurev.psych.59. 103006.093639
- Barsalou, L. W. (2016). On staying grounded and avoiding quixotic dead ends. *Psychonomic Bulletin & Review*, 1–21. doi:10.3758/s13423-016-1028-3
- Bellebaum, C., Tettamanti, M., Marchetta, E., Della Rosa, P., Rizzo, G., Daum, I., & Cappa, S. F. (2013). Neural representations of unfamiliar objects are modulated by sensorimotor experience. *Cortex; A Journal Devoted to the Study of the Nervous System and Behavior*, 49(4), 1110–1125. doi:10.1016/j.cortex. 2012.03.023
- Bickerton, D. (2014). Some problems for biolinguistics. *Biolinguistics*, 8, 73–96.
- Bloom, L. (1995). *The transition from infancy to language: Acquiring the power of expression* (Revised ed.). Cambridge: Cambridge University Press.
- Bloom, P. (2002). *How children learn the meanings of words*. Cambridge, MA: A Bradford Book.
- Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2016). Enriching word vectors with subword information. arXiv:1607.04606 [Cs]. Retrieved from http://arxiv.org/abs/ 1607.04606
- Borghi, A. M., & Cimatti, F. (2009). Words as tools and the problem of abstract words meanings. In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st annual conference of the cognitive science society*. Retrieved from http:// csjarchive.cogsci.rpi.edu/proceedings/2009/papers/536/ index.html
- Borin, L. (2012). Core vocabulary: A useful but mystical concept in some kinds of linguistics. In D. Santos, K. Lindén, & W. Ng'ang'a (Eds.), *Shall we play the festschrift game?* (pp. 53– 65). Berlin: Springer. doi:10.1007/978-3-642-30773-7_6

- Boutonnet, B., & Lupyan, G. (2015). Words jump-start vision: A label advantage in object recognition. *The Journal of Neuroscience*, 35(25), 9329–9335. doi:10.1523/JNEUROSCI. 5111-14.2015
- Bowerman, M. (2000). Where do children's word meanings come from? Rethinking the role of cognition in early semantic development. In L. Nucci & E. Saxe (Eds.), *Culture, thought, and development* (pp. 199–230). Mahwah, NJ: Lawrence Erlbaum. Retrieved from https://books.google.com/books? hl=en&lr=&id=TsZ5AgAAQBAJ&oi=fnd&pg=PA173&ots=wQ71 Nw1nVD&sig=uG37FVCKBMuJP3-Z2VamRRthWbs
- Bowerman, M., & Choi, S. (2003). Space under construction: Language-specific spatial categorization in first language acquisition. In D. Gentner & S. Goldin-Meadow (Eds.), Language in mind: Advances in the study of language and thought (pp. 387–427). Cambride, MA: MIT Press.
- Brown, R. (1958). Words and things. New York: Free Press.
- Bruni, E., Boleda, G., Baroni, M., & Tran, N.-K. (2012). Distributional semantics in technicolor. In *Proceedings of* the 50th annual meeting of the association for computational linguistics: Long papers-volume 1 (pp. 136–145). Association for Computational Linguistics. Retrieved from https://dl. acm.org/citation.cfm?id=2390544
- Bruni, E., Tran, N. K., & Baroni, M. (2014). Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49(1), 1–47.
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3), 904–911. doi:10.3758/s13428-013-0403-5
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, *356*(6334), 183–186. doi:10. 1126/science.aal4230
- Carey, S. (2009). The origin of concepts. New York, NY: Oxford University Press.
- Carruthers, P., & Boucher, J. (1998). *Language and thought: Interdisciplinary themes.* New York, NY: Cambridge University Press.
- Casasanto, D., & Lupyan, G. (2014). All concepts are ad hoc concepts. In E. Margolis & S. Laurence (Eds.), *Concepts: New directions* (pp. 543–566). Cambridge, MA: MIT Press.
- Casasola, M. (2005). Can language do the driving? The effect of linguistic input on infants' categorization of support spatial relations. *Developmental Psychology*, 41(1), 183–192. doi:10. 1037/0012-1649.41.1.188
- Chomsky, N. (2010). Some simple evo devo theses: How true might they be for language? In R. K. Larson, V. Déprez, & H. Yamakido (Eds.), *The evolution of human language: Biolinguistic perspectives* (pp. 45–62). Cambridge University Press. doi:10.1017/CBO9780511817755.003
- Clark, E. V. (1973). What's in a word? On the child's acquisition of semantics in his first language. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 65–110). New York: Academic Press.
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. Cognition, 22(1), 1–39.
- Connolly, A. C., Gleitman, L. R., & Thompson-Schill, S. L. (2007). Effect of congenital blindness on the semantic representation of some everyday concepts. *Proceedings of the National Academy of Sciences*, 104(20), 8241–8246. doi:10. 1073/pnas.0702812104

- Cross, E. S., Cohen, N. R., de Hamilton, A. F. C., Ramsey, R., Wolford, G., & Grafton, S. T. (2012). Physical experience leads to enhanced object perception in parietal cortex: Insights from knot tying. *Neuropsychologia*, 50(14), 3207– 3217. doi:10.1016/j.neuropsychologia.2012.09.028
- Davidoff, J., & Roberson, D. (2004). Preserved thematic and impaired taxonomic categorisation: A case study. *Language and Cognitive Processes*, *19*(1), 137–174.
- De Deyne, S., Navarro, D. J., & Storms, G. (2013). Better explanations of lexical and semantic cognition using networks derived from continued rather than single-word associations. *Behavior Research Methods*, 45(2), 480–498. doi:10. 3758/s13428-012-0260-7
- Devitt, M., & Sterelny, K. (1987). Language and reality: An introduction to the philosophy of language. Cambridge, MA: MIT Press.
- Eco, U. (2008). Experiences in translation (1st ed.). Toronto: University of Toronto Press, Scholarly Publishing Division. A. McEwen, Trans.
- Edmiston, P., & Lupyan, G. (2015). What makes words special? Words as unmotivated cues. *Cognition*, *143*, 93–100. doi:10. 1016/j.cognition.2015.06.008
- Edmiston, P., & Lupyan, G. (2017). Visual interference disrupts visual knowledge. *Journal of Memory and Language*, *92*, 281–292. doi:10.1016/j.jml.2016.07.002
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*(2), 179–211.
- Elman, J. L. (2004). An alternative view of the mental lexicon. *Trends in Cognitive Sciences*, 8(7), 301–306. doi:10.1016/j. tics.2004.05.003
- Elman, J. L. (2009). On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science*, 33(4), 547–582. doi:10.1111/j.1551-6709.2009. 01023.x
- Fasmer, M. (2009). Etimologicheskiy slovar russkogo yazyka. V 4 tomah. Tom 4. T-Yaschur. Moskva: AST.
- Firth, J. R. (1957). A synopsis of linguistic theory 1930–1955. In Studies in linguistic analysis (pp. 1–32). Oxford: Blackwell.
- Fisher, C., Hall, D. G., Rakowitz, S., & Gleitman, L. (1994). When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. *Lingua*, *92*, 333–375. doi:10.1016/0024-3841(94)90346-8
- Fodor, J. A. (1975). *The language of thought*. Cambride, MA: Harvard University Press.
- Fodor, J. A. (2001). Language, thought and compositionality. *Mind and Language*, *16*(1), 1–15. doi:10.1111/1468-0017. 00153
- Fodor, J. A. (2010). LOT 2: The language of thought revisited. Oxford[u.a.]: Oxford University Press.
- Fodor, J. A., & Lepore, E. (1996). The red herring and the pet fish: Why concepts still can't be prototypes. *Cognition*, *58*(2), 253–270.
- Forder, L., & Lupyan, G. (2017a). Hearing words changes color perception: Facilitation of color discrimination by verbal vs. visual cues. Retrieved from https://psyarxiv.com/f83au
- Forder, L., & Lupyan, G. (2017b). Facilitation of color discrimination by verbal and visual cues. *PsyArXiv Preprints*. doi:10. 17605/OSF.IO/F83AU
- Frome, A., Corrado, G. S., Shlens, J., Bengio, S., Dean, J., Ranzato, M. A., & Mikolov, T. (2013). DeViSE: A deep visual-semantic embedding model. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), Advances in neural

information processing systems (pp. 2121–2129). Curran Associates. Retrieved from http://papers.nips.cc/paper/5204-devise-a-deep-visual-semantic-embedding-model.pdf

- Fulkerson, A. L., & Waxman, S. R. (2007). Words (but not tones) facilitate object categorization: Evidence from 6- and 12month-olds. *Cognition*, 105(1), 218–228.
- Garnham, A. (1979). Instantiation of verbs. *Quarterly Journal of Experimental Psychology*, 31(2), 207–214. doi:10.1080/14640747908400720
- Gentner, D., & Boroditsky, L. (2001). Individuation, relational relativity and early word learning. In M. Bowerman, & S. C. Levinson (Eds.), *Language acquisition and conceptual devel*opment (pp. 215–256). Cambridge: Cambridge University Press.
- Gentner, D., & Goldin-Meadow, S. (2003). *Language in mind: Advances in the study of language and thought*. Cambridge, MA: MIT Press.
- Gerz, D., Vulić, I., Hill, F., Reichart, R., & Korhonen, A. (2016). Simverb-3500: A large-scale evaluation set of verb similarity. arXiv Preprint arXiv:1608.00869. Retrieved from https://arxiv. org/abs/1608.00869
- Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition*, 1(1), 3–55.
- Gleitman, L., Armstrong, S. L., & Connolly, A. C. (2012). Can prototype representations support composition and decomposition? In M. Werning, W. Hinzen, & E. Machery (Eds.), *Oxford handbook of compositionality* (pp. 123–142). New York: Oxford University Press.
- Gleitman, L., & Fisher, C. (2005). Universal aspects of word learning. In J. McGilvray (Ed.), *The Cambridge companion to Chomsky* (pp. 123–144). New York, NY: Cambridge University Press. Retrieved from https://books.google.com/books?hl= en&lr=&id=I6CZ6wpNKeEC&oi=fnd&pg=PA123&ots=ujaHwty K8O&sig=N-te0BeTbB0lyV7aOmuwBAgwtCk
- Goddard, C., & Wierzbicka, A. (2002). *Meaning and universal grammar: Theory and empirical findings*. Amsterdam: John Benjamins.
- Goldberg, A. E. (2003). Constructions: A new theoretical approach to language. *Trends in Cognitive Sciences*, 7(5), 219–224.
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66, 377–388.
- Hagoort, P., & Van Berkum, J. (2007). Beyond the sentence given. Philosophical Transactions of the Royal Society B: Biological Sciences, 362(1481), 801–811. doi:10.1098/rstb. 2007.2089
- Hill, F., Reichart, R., & Korhonen, A. (2016). Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*. Retrieved from http:// www.mitpressjournals.org/doi/abs/10.1162/COLI_a_00237
- Hoff, E. (2013). *Language development*. Belmont, CA: Cengage Learning.
- Hollis, G., & Westbury, C. (2016). The principals of meaning: Extracting semantic dimensions from co-occurrence models of semantics. *Psychonomic Bulletin & Review*, 23(6), 1744–1756. doi:10.3758/s13423-016-1053-2
- Huang, Y. T., & Snedeker, J. (2009). Online interpretation of scalar quantifiers: Insight into the semantics–pragmatics interface. *Cognitive Psychology*, 58(3), 376–415.
- Ingraham, A. (1903). Swain school lectures. Open Court. Retrieved from https://archive.org/details/swainschoollect00ingrgoog
- Jackendoff, R. (1983). Semantics and cognition. Cambridge, MA: MIT Press.

- Jackendoff, R. S. (1990). Semantic structures. Cambride, MA: MIT Press.
- Koemeda-Lutz, M., Cohen, R., & Meier, E. (1987). Organization of and access to semantic memory in aphasia. *Brain and Language*, 30(2), 321–337.
- Lakoff, G. (1990). *Women, fire, and dangerous things*. Chicogo, IL: University of Chicago Press.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*(2), 211–240.
- Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings* (pp. 3–81). Cambridge, MA: A Bradford Book.
- Leavitt, J. (2011). Linguistic relativities: Language diversity and modern thought. New York, NY: Cambridge University Press.
- Lenci, A. (2008). Distributional semantics in linguistic and cognitive research. *Rivista Di Linguistica*, 20(1), 1–31.
- Levinson, S. C. (1997). From outer to inner space: Linguistic categories and non-linguistic thinking. In J. Nuyts & E. Pederson (Eds.), *Language and conceptualization* (pp. 13–45). Cambridge: Cambridge University Press.
- Levy, O., & Goldberg, Y. (2014). Neural word embedding as implicit matrix factorization. In *Advances in neural information processing systems* (pp. 2177–2185). Retrieved from http://papers.nips.cc/paper/5477-scalable-non-linearlearning-with-adaptive-polynomial-expansions.pdf
- Lewis, M. L., & Frank, M. C. (2013). An integrated model of concept learning and word-concept mapping. In *Proceedings* of the 35th annual meeting of the cognitive science society (pp. 882–887). Austin, TX. Retrieved from http://citeseerx.ist. psu.edu/viewdoc/summary?doi=10.1.1.691.8652
- Lewis, M. L., & Frank, M. C. (2016). The length of words reflects their conceptual complexity. *Cognition*, 153, 182–195. doi:10. 1016/j.cognition.2016.04.003
- Li, P., & Gleitman, L. (2002). Turning the tables: Language and spatial reasoning. *Cognition*, 83(3), 265–294.
- Louwerse, M. M. (2011). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, 3(2), 273–302. doi:10.1111/j.1756-8765.2010.01106.x
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers, 28*(2), 203– 208. doi:10.3758/BF03204766
- Lupyan, G. (2008). The conceptual grouping effect: Categories matter (and named categories matter more). *Cognition*, *108* (2), 566–577.
- Lupyan, G. (2009). Extracommunicative functions of language: Verbal interference causes selective categorization impairments. *Psychonomic Bulletin & Review*, 16(4), 711–718. doi:10.3758/PBR.16.4.711
- Lupyan, G. (2012). What do words do? Towards a theory of language-augmented thought. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 57, pp. 255–297). Waltham, MA: Academic Press. Retrieved from http://www. sciencedirect.com/science/article/pii/B9780123942937000078
- Lupyan, G. (2015). The paradox of the universal triangle: Concepts, language, and prototypes. *Quarterly Journal of Experimental Psychology*. doi:10.1080/17470218.2015.1130730
- Lupyan, G. (2016). The centrality of language in human cognition. Language Learning, 66(3), 516–553. doi:10.1111/lang. 12155

- Lupyan, G. (2017). Changing what you see by changing what you know: The role of attention. *Frontiers in Psychology*, *8*, 1055. doi:10.3389/fpsyg.2017.00553
- Lupyan, G., & Bergen, B. (2015). How language programs the mind. *Topics in Cognitive Science*. doi:10.1111/tops.12155
- Lupyan, G., & Casasanto, D. (2015). Meaningless words promote meaningful categorization. *Language and Cognition*, 7(2), 167–193. doi:10.1017/langcog.2014.21
- Lupyan, G., & Mirman, D. (2013). Linking language and categorization: Evidence from aphasia. *Cortex*, 49(5), 1187–1194. doi:10.1016/j.cortex.2012.06.006
- Lupyan, G., Rakison, D. H., & McClelland, J. L. (2007). Language is not just for talking: Labels facilitate learning of novel categories. *Psychological Science*, 18(12), 1077–1083.
- Lupyan, G., & Spivey, M. J. (2008). Perceptual processing is facilitated by ascribing meaning to novel stimuli. *Current Biology*, 18(10), R410–R412.
- Lupyan, G., & Spivey, M. J. (2010a). Making the invisible visible: Auditory cues facilitate visual object detection. *PLoS ONE*, *5* (7), e11452. doi:10.1371/journal.pone.0011452
- Lupyan, G., & Spivey, M. J. (2010b). Redundant spoken labels facilitate perception of multiple items. Attention, Perception, & Psychophysics, 72(8), 2236–2253. doi:10.3758/APP.72.8.2236
- Lupyan, G., & Thompson-Schill, S. L. (2012). The evocative power of words: Activation of concepts by verbal and nonverbal means. *Journal of Experimental Psychology: General*, 141(1), 170–186. doi:10.1037/a0024904
- Lupyan, G., & Ward, E. J. (2013). Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences*, *110*(35), 14196–14201. doi:10. 1073/pnas.1303312110
- Mahon, B. Z., & Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology-Paris*, *102*(1), 59–70
- Mahon, B. Z., & Hickok, G. (2016). Arguments about the nature of concepts: Symbols, embodiment, and beyond. *Psychonomic Bulletin & Review*, 1–18. doi:10.3758/s13423-016-1045-2
- Majid, A. (2015). Comparing lexicons cross-linguistically. In John R. Taylor (Ed.), *The Oxford handbook of the word* (pp. 364–379). New York, NY: Oxford University Press.
- Malt, B. C., Gennari, S., Imai, M., Ameel, E., Saji, N., & Majid, A. (2015). Where are the concepts? What words can and can't reveal. In E. Margolis & S. Laurence (Eds.), *Concepts: New directions*. Cambridge, MA: MIT Press.
- Malt, B. C., Gennari, S. P., Imai, M., Ameel, E., Saji, N., & Majid, A. (in press). Where are the concepts? What words can and can't reveal. In E. Margolis & S. Laurence (Eds.), *Concepts: New directions*. Camrbridge, MA: MIT Press.
- Marmor, G. S. (1978). Age at onset of blindness and the development of the semantics of color names. *Journal of Experimental Child Psychology*, 25(2), 267–278. doi:10.1016/ 0022-0965(78)90082-6
- McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review*, 119(4), 831–877. doi:10.1037/a0029872
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. arXiv Preprint arXiv:1301.3781. Retrieved from https://arxiv.org/ abs/1301.3781
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed Representations of Words and Phrases and their

Compositionality. arXiv:1310.4546 [Cs, Stat]. Retrieved from http://arxiv.org/abs/1310.4546

- Mitchell, J., & Lapata, M. (2010). Composition in distributional models of semantics. *Cognitive Science*, 34(8), 1388–1429. doi:10.1111/j.1551-6709.2010.01106.x
- Munroe, R. P. (2010, May 4). Color survey results. Retrieved from http://blog.xkcd.com/2010/05/03/color-survey-results/
- Murphy, G. L. (2002). *The big book of concepts*. Cambride, MA: The MIT Press.
- Murphy, M. L. (2010). Meaning variation: Polysemy, homonymy, and vagueness. In *Lexical meaning*. Cambridge University Press. doi:10.1017/CBO9780511780684.009
- Murphy, G. L., & Medin, D. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*(3), 289–316. doi:10. 1037/0033-295X.92.3.289
- Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*, 36(3), 402–407. doi:10.3758/BF03195588
- Nerlich, B. (2003). Polysemy: Flexible patterns of meaning in mind and language. New York, NY: Walter de Gruyter.
- Painter, C. (2005). Learning through language in early childhood. New York, NY: A&C Black.
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *EMNLP* (Vol. 14, pp. 1532– 1543). Retrieved from https://pdfs.semanticscholar.org/ b397/ed9a08ca46566aa8c35be51e6b466643e5fb.pdf
- Perry, L. K., & Lupyan, G. (2014). The role of language in multidimensional categorization: Evidence from transcranial direct current stimulation and exposure to verbal labels. *Brain and Language*, 135, 66–72. doi:10.1016/j.bandl.2014.05.005

Pinker, S. (1994). The language instinct. New York: Harper Collins.

- Robinson, C. W., Best, C. A., Deng, W. (Sophia), & Sloutsky, V. M. (2012). The role of words in cognitive tasks: What, when, and how? *Frontiers in Psychology*, 3. doi:10.3389/fpsyg.2012.00095
- Roth, E. M., & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, 15(3), 346–378. doi:10.1016/0010-0285(83)90012-9
- Rumelhart, D. E. (1979). Some problems with the notion that words have literal meanings. In A. Ortony (Ed.), *Metaphor* and thought (pp. 71–82). New York, NY: Cambridge University Press.
- Saji, N., Imai, M., Saalbach, H., Zhang, Y., Shu, H., & Okada, H. (2011). Word learning does not end at fast-mapping: Evolution of verb meanings through reorganization of an entire semantic domain. *Cognition*, 118(1), 45–61.
- Sandra, D., & Rice, S. (1995). Network analyses of prepositional meaning: Mirroring whose mind—the linguist's or the language user's? *Cognitive Linguistics*, 6(1), 89–130. doi:10. 1515/cogl.1995.6.1.89
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied*, 75(13), 1–42. doi:10.1037/h0093825
- Silberer, C., Ferrari, V., & Lapata, M. (2016). Visually grounded meaning representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. doi:10.1109/TPAMI.2016. 2635138
- Siskind, J. M. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, *61*(1), 39–91.
- Snedeker, J., & Gleitman, L. (2004). Why is it hard to label our concepts? In D. G. Hall & S. R. Waxman (Eds.), Weaving a

lexicon (illustrated ed., pp. 257–294). Cambridge, MA: The MIT Press.

- Tabossi, P., & Johnson-Laird, P. N. (1980). Linguistic context and the priming of semantic information. *The Quarterly Journal of Experimental Psychology*, 32(4), 595–603. doi:10.1080/ 14640748008401848
- Tomasello, M. (2001). Could we please lose the mapping metaphor, please? *Behavioral and Brain Sciences*, 24(6), 1119–1120.
- Tyler, A., & Evans, V. (2001). Reconsidering prepositional polysemy networks: The case of over. *Language*, 77(4), 724–765.
- Uspenskij, L. (1962). Слово о словах (Slovo o slovah). Лениздат (Lenizdat). Retrieved from http://www.speakrus.ru/uspens/
- Van Berkum, J. J. (2008). Understanding sentences in context: What brain waves can tell us. *Current Directions in Psychological Science*, 17(6), 376–380.
- Van Berkum, J. J., van den Brink, D., Tesink, C. M. J. Y., Kos, M., & Hagoort, P. (2008). The neural integration of speaker and message. *Journal of Cognitive Neuroscience*, 20(4), 580–591. doi:10.1162/jocn.2008.20054
- Wagner, K., Dobkins, K., & Barner, D. (2013). Slow mapping: Color word learning as a gradual inductive process. *Cognition*, 127 (3), 307–317. doi:10.1016/j.cognition.2013.01.010
- Waxman, S. R., & Leddon, E. M. (2010). Early word-learning and conceptual development. In U. Goswami (Ed.), *The Wiley-Blackwell handbook of childhood cognitive development* (pp. 180–208). Wiley-Blackwell. doi:10.1002/9781444325485.ch7
- Waxman, S. R., & Markow, D. B. (1995). Words as invitations to form categories: Evidence from 12- to 13-month-old infants. Cognitive Psychology, 29(3), 257–302.
- Whorf, B. L. (1956). *Language, thought, and reality*. (J. B. Carroll, Ed.). Cambridge, MA: MIT Press.
- Wierzbicka, A. (1996). Semantics: Primes and universals. Oxford University Press. Retrieved from http://books.google.com/ books?hl=en&lr=&id=ZN029Pmbnu4C&oi=fnd&pg=PA3&dq= info:QazMfnERxREJ:scholar.google.com&ots=u-f4yd6owJ&sig =JYdRS-9zbrFun-1S4B_h0GOHFdY
- Wierzbicka, A. (2009). All people eat and drink. Does this mean that "eat" and "drink" are universal human concepts? In J. Newman (Ed.), *Typological studies in language* (Vol. 84, pp. 65–89). Amsterdam: John Benjamins. doi:10.1075/tsl.84. 05wie
- Wierzbicka, A. (2013). Imprisoned in English: The Hazards of English as a default language (1st ed.). Oxford: Oxford University Press.
- Willits, J. A., Amato, M. S., & MacDonald, M. C. (2015). Language knowledge and event knowledge in language use. *Cognitive Psychology*, 78, 1–27.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*(2), 245–272. doi:10. 1037/0033-295X.114.2.245
- Yee, E., Jones, K., & McRae. (in press). Semantic memory. In J. W. & S. Thompson-Schill (Eds.), The Stevens' handbook of experimental psychology and neuroscience, fourth edition. Wiley. Retrieved from http://onlinelibrary.wiley.com/book/10.1002/ 9781119170174.
- Zwaan, R. A. (2014). Embodiment and language comprehension: Reframing the discussion. *Trends in Cognitive Sciences*, 18(5), 229–234. doi:10.1016/j.tics.2014.02.008
- Zwaan, R. A., Stanfield, R. A., & Yaxley, R. H. (2002). Language comprehenders mentally represent the shapes of objects. *Psychological Science*, 13(2), 168–171. doi:10.1111/1467-9280.00430