Article Title: Building semantic memory from embodied and distributional language experience

Article Type:

⊙ OVERVIEW

Authors:

| |
|---|
| **Charles P. Davis**
Department of Psychological Sciences, University of Connecticut
Connecticut Institute for the Brain and Cognitive Sciences, University of Connecticut
email: charles.davis@uconn.edu
ORCID: 0000-0002-7293-2769 |
| **Eiling Yee**
Department of Psychological Sciences, University of Connecticut
Connecticut Institute for the Brain and Cognitive Sciences, University of Connecticut
email: eiling.yee@uconn.edu
ORCID: 0000-0001-6614-9025 |

Abstract

Humans seamlessly make sense of a rapidly changing environment, using a seemingly limitless knowledgebase to recognize and adapt to most situations we encounter. This knowledgebase is called *semantic memory*. Embodied cognition theories suggest that we represent this knowledge through simulation: understanding the meaning of *coffee* entails re-instantiating the neural states involved in touching, smelling, seeing, and drinking coffee. Distributional semantic theories suggest that we are sensitive to statistical regularities in natural language, and that a cognitive mechanism picks up on these regularities and transforms them into usable semantic representations reflecting the contextual usage of language. These appear to present contrasting views on semantic memory, but do they? Recent years have seen a push toward combining these approaches under a common framework. These hybrid approaches augment our understanding of semantic memory in important ways, but current versions remain unsatisfactory in part because they treat sensory-perceptual and distributional-linguistic data as interacting but distinct types of data that must be combined. We synthesize several approaches which, taken together, suggest that linguistic and embodied experience should instead be considered as inseparably entangled: just as sensory and perceptual systems are reactivated to understand meaning, so are experience-based representations endemic to linguistic processing; further, sensory-perceptual experience is susceptible to the same distributional principles as language experience. This conclusion produces a characterization of semantic memory that accounts for the interdependencies between linguistic and embodied data

that arise across multiple timescales, giving rise to concept representations that reflect our shared and unique experiences.

Keywords: semantic memory, concepts, distributional semantics, embodied cognition

Graphical/Visual Abstract and Caption
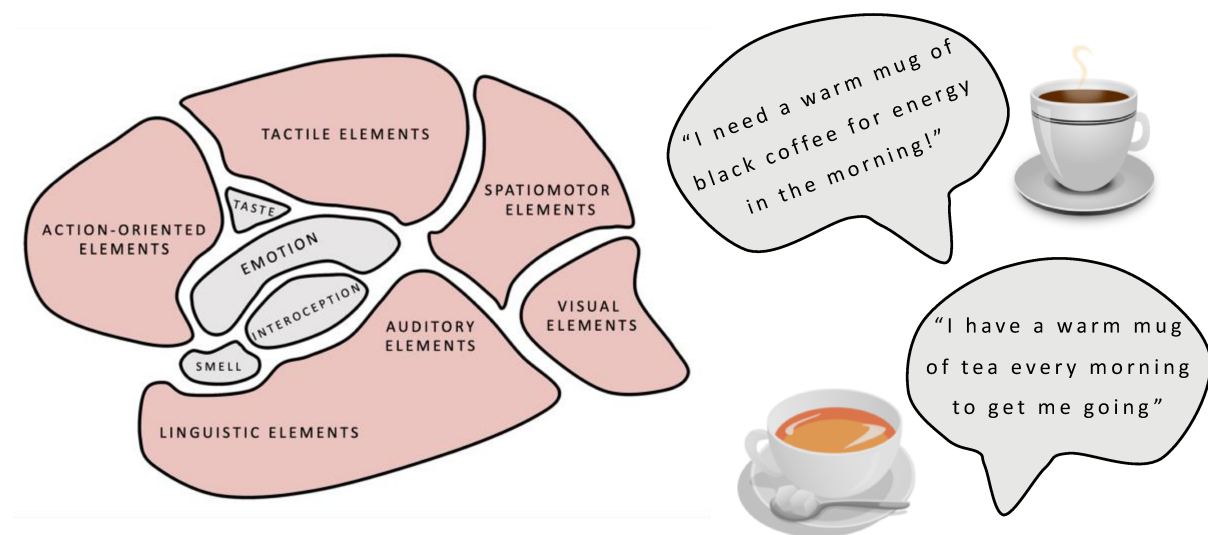
Semantic memory is the lens through which we comprehend the environments which we pass through. Two approaches—embodied cognition and distributional semantics—offer what are, at first glance, competing accounts of semantic memory. We describe interdisciplinary efforts to combine these approaches into one framework.

Introduction

Understanding how we represent *meaning*—the core problem of semantic memory research—is central to understanding how humans operate in the world: how we recognize things as the same even when encountered in different conditions, how we remember objects and events and deploy that knowledge in different situations, how we know (or think we know) what is similar to what, and how we communicate with each other despite experiencing the world differently. None of these would be possible without semantic memory.[1] This paper will explore and review complementarities between two influential accounts of meaning—*embodied cognition* and *distributional semantic models*. Embodied and distributional accounts have typically been presented as contrasting views of semantic memory, owing largely to a lack of cross-disciplinary communication: embodied theories typically derive from experimental psychology, while distributional semantic models emerge from

---

[1] While we refer to "semantic memory" as a singular construct in order to focus specifically on two accounts of meaning, we do not intend to suggest that semantic memory independently supports the representation of meaning. Rather, our contention is that semantic memory is part of an integrated memory system, influencing and influenced by, among other cognitive functions, episodic memory as well as more implicit forms of memory like procedural memory (for discussion, see Yee et al., 2018).

computational linguistics. In this paper, we synthesize several approaches to reconciliation, which taken together suggest that embodied experience is a type of distributional information similar to that captured in distributional semantic models[2] (e.g., Andrews et al., 2009; Hoffman et al., 2018; Johns & Jones, 2012; Steyvers, 2010), and language and its distributional characteristics reflect a type of embodied experience (e.g., Clark, 2006; Dove, 2018, 2019; Borghi et al., 2019). That is, we can consider distributional and embodied information as fundamentally the same type of data, entangled and mutually influencing each other across multiple timescales. The representations that emerge through this interaction reflect the characteristics of our unique and shared environments.

In this paper, we first provide a brief historical review of embodied and distributional semantic models, outlining major advantages and criticisms of each. Next, we review reconciliation efforts and outline remaining issues. We then synthesize these interdisciplinary efforts toward reconciliation, providing a rough sketch of how both embodied and distributional knowledge can be obtained via the same sensitivity to regularities present in experience-based input. This sensitivity to experience-based input engenders semantic representations reflective of our shared and unique experiences. We conclude by offering suggestions for future (interdisciplinary) work considering embodied and distributional data in a common framework.

HISTORICAL OVERVIEW

Modern semantic memory research has chiefly operated on two independent paths: distributional semantic models (e.g., Landauer & Dumais, 1997; Griffiths et al., 2007; Lund & Burgess, 1996) and embodied cognition (e.g., Allport, 1985; Barsalou, 1999; Damasio, 1989; Glenberg, 1997). Until the last decade or so, these fields have been largely independent due to apparent differences in the characterization of meaning.[3] Distributional models suggest that meaning can be inferred from the contexts (almost always operationalized as language contexts) in which words appear—this is well illustrated by Firth's (1957) supposition that "You shall know a word by the company it keeps." Meaning in (language-based) distributional models is derived from and represented in terms of statistical patterns of co-occurrence with other words in a language. Embodied approaches, on the other hand, suggest that meaning is grounded in our sensory, perceptual, motor, interoceptive, and introspective experiences with the world (e.g., Barsalou, 1999). Context is also important in embodied theories (for reviews see, e.g., Yee & Thompson-Schill, 2016; Yeh & Barsalou, 2006), but this context is *situated* or *grounded*: to access the meaning of a word, we simulate the bodily states associated with experiencing that concept 'in the wild' (Barsalou, 1999), and this simulation varies according to the current context as well as an individual's history.

While recent years have seen growing interest in reconciling distributional and embodied perspectives in a common model (e.g., Andrews et al., 2009), such approaches often consider the information (i.e., the input) that goes into distributional models and embodied information as two

---

[2] Throughout this paper, we use the word "distributional" in two different ways: distributional *models* refer to the cognitive mechanisms, implemented in formal models, that are sensitive to distributional statistics, and that are traditionally implemented using language as the input. Distributional *statistics* (or distributional *information*) refers to properties of the environment (whether language or other). While language input is distributional in nature, we do not intend to conflate language input with distributional statistics—other sorts of input (e.g., perceptual) are similarly characterized by distributional statistics.
[3] We characterize these differences as "apparent" because we view them primarily as a result of methodological differences due to the type of input that is the focus of study in different disciplines.

distinct types of data in need of combination. And so the question remains: what is the relationship between distributional and embodied information? Before speculating on a solution, we provide a historical review of embodied and distributional semantics as they have developed independently.

Embodied semantics

Philosophical and historical background

Embodied approaches can be traced back to British Associationism of the 17th and 18th centuries. According to empiricist philosopher John Locke (1689), conceptual knowledge is built through experiencing the environment, where the environment is experienced in terms of elementary perceptual attributes, and concept representations are built incrementally upon these attributes. Thus, *coffee* is represented in terms of its color, smell, taste, manipulability, and so on (more contemporarily, this idea was articulated by Allport, 1985; see Figure 1). For David Hume (1748), conceptual knowledge could also be traced to world experiences. Critically, Hume suggested that consistent *association* of the senses leads those sensory experiences to become connected in the mind, such that with more frequent conjoint experience, the appearance of one sensory experience can activate its associates.
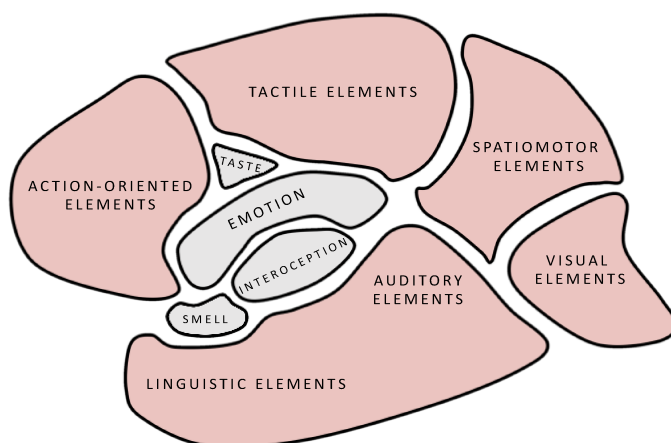

Figure 1. A cartoonized brain depicting how distributed brain regions might contribute to conceptual knowledge under an embodied cognition framework. Pink areas are roughly meant to correspond to cortical regions, and grey areas are roughly meant to correspond to subcortical areas. Figure adapted from Allport (1985) and Thompson-Schill et al. (2006). This figure is licensed under a CC-BY 4.0 International License.

In 1900, neurologist Carl Wernicke outlined a neurophysiological framework that incorporated many of the principles of British Associationism, bearing strong resemblance to contemporary theories of conceptual knowledge (e.g., Damasio, 1989). Wernicke (1900) suggested that recurring stimuli engender the same pattern of activation with each instantiation, and those stimuli retain their mutual association as a memory trace distributed across sensory-perceptual areas. For instance, when we drink coffee, we jointly experience its smell, taste, and color, and this association is not only reflected in joint activation of olfactory, gustatory, and visual brain systems, but their (neural) association is also retained in memory. These associative networks are established via neural connectivity (in more contemporary terms, between *convergence zones*, which are bundles of

neurons sensitive to co-occurring inputs, e.g., the convergence of information about sound and visual motion; Damasio, 1989). And once an associative network is established, Wernicke suggested (as Hume did centuries earlier), partial activation of the network can trigger activation of the entire network, where the network corresponds to all information associated with the concept (Gage & Hickok, 2005). This means that *any* experience associated with the concept can activate relevant conceptual knowledge—for example, because the word *coffee* tends to be experienced in the same settings as actually experiencing *coffee,* merely hearing or seeing the word will activate a network of sensory and perceptual brain areas involved in seeing, smelling, and tasting coffee (see also Pulvermüller, 2013, who describes these phenomena in terms of Hebbian learning, and Yee & Thompson-Schill, 2016, who describe conceptual knowledge as "the flow of activation […] through a network of connections that cumulatively reflect prior experience", p. 1022, reflecting Jeff Elman's (1990) approach to knowledge formation using simple recurrent networks; see Box 1).

Box 1. The simple recurrent network.

The simple recurrent network (SRN; Elman, 1990) is a classic example of how, through accumulated experience over time with words (or, in principle, any experience), concepts and categories can emerge without explicit feedback. The network retains a copy of its previous state (i.e., the context) and uses this to predict the next element in a sequence. Importantly, although the hidden units of the SRN are undifferentiated computationally, the representations that emerge after learning— which reflect accumulated knowledge about the contexts in which we experience things—are *not* undifferentiated functionally. Due in part to this computational plasticity, SRNs have been used to understand how abstract structure emerges in many contexts, including both embodied (e.g., Yee & Thompson-Schill, 2016, who based their account on conceptual principles borrowed from the SRN) and distributional (e.g., Hoffman et al., 2018—an implemented model) experience frameworks. In our view, then, knowledge of a concept is no more than the knowledge of the contexts (whether embodied or distributional) in which that concept (or the word(s) that refer to it) occurs (see also Elman, 2009; Yee & Thompson-Schill, 2016). Understanding the computational principles of the SRN is critical to understanding some of the concepts—as well as the implemented hybrid models (e.g., Hoffman et al., 2018)—discussed in this paper. We also discuss insights generated by the SRN with respect to abstraction and cognition more broadly elsewhere (Davis et al., 2020b).
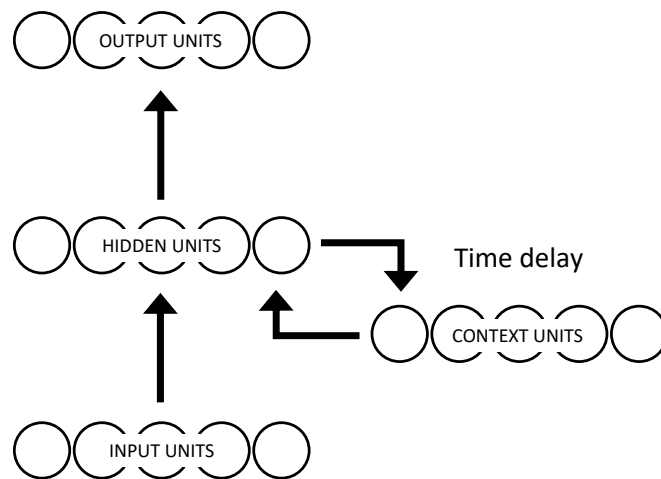
Figure 2. A simple recurrent network. Each layer consists of one or more units, and information (e.g., words, semantic features) flows first from input units, to hidden units, and then to output units. At every timepoint, the context units propagate to the hidden layer, giving the network access to its 'memory' of prior states.

Empirical evidence

Is it true that areas of the brain that, for example, are involved in perceiving the color of coffee or in guiding action when sipping from a mug also become active when thinking about the color associated with coffee or when reading the word *sip*? Sometimes. Early neuroimaging studies did not find *direct* overlap between brain regions involved in perception and action and the representation of sensorimotor information. Rather, they observed that, e.g., areas of the brain *adjacent* to those involved in perceiving color are activated when we say what an object's most typical color is (e.g., *brown* for *coffee*) as compared to saying its name (Chao & Martin, 1999; Martin et al., 1995), and that when passively reading action verbs implicating movement of the hands, legs, and mouth, areas *adjacent* to the corresponding motor region are activated for words like *sip* (Hauk et al., 2004). However, more recent studies *have* observed evidence of direct overlap. For instance, when color-perception areas were identified using a more demanding color perception task (i.e., making subtle judgments about differences in the hues of presented colors) direct overlap was observed between areas involved in color perception and color knowledge (Simmons et al., 2007). This implies that at least some part of the system supporting color perception also represents color knowledge (for further discussion, see Martin, 2016). In addition to the visual system, similar findings have been reported for motor (in premotor cortex; Willems et al., 2010), auditory (e.g., Kiefer et al., 2008), and emotional systems (Ziegler et al., 2018). These are only a few examples of evidence supporting embodied cognition—there is now abundant empirical work (described in greater detail elsewhere; e.g., Barsalou, 2016; Meteyard et al., 2012) suggesting that conceptual knowledge is (at least partially) sensorimotor-based.

Critiques

Arguments advanced against embodied cognition have pointed out that this overlap or adjacency of activation need not imply that sensory and/or motor regions are *functionally* involved in conceptual processing. Rather, it has been argued that activation in these areas may simply be a (downstream) *consequence* of conceptual processing that actually occurs without any functional involvement of sensory or motor regions (see, e.g., Mahon & Caramazza, 2008; Mahon, 2015). However, there is now evidence from neuropsychological, neurostimulation, and behavioral studies suggesting that not only are the same regions for perceiving objects *active* when thinking about those objects in their absence, but also that those regions are to some degree *necessary* for comprehension. For example, compared with age-matched controls, patients with Parkinson's disease have difficulty accessing the meaning of words and sentences referring to motor action (Fernandino et al., 2013a, 2013b; see also Buccino et al., 2018). This suggests that the motor system (which is compromised in Parkinson's disease) is necessary for understanding the meaning of manually experienced concepts (for related evidence in various sensorimotor domains and in both healthy and patient populations, see e.g., Davis et al., 2020c; Trumpp et al., 2013; Vukovic et al., 2017; Yee et al., 2013).

Another critique levied against embodied approaches is that there are many concepts, e.g., *idea* or *justice* (typically referred to as "abstract" concepts) for which it is not obvious that sensory or motor systems would be routinely involved when we experience them. We have only just begun to understand the representational substrates of such concepts, but there is emerging evidence that we understand concepts like *justice* at least in part by re-activating the emotion systems involved in *feeling* justice (e.g., Kousta et al., 2011; Vigliocco et al., 2013), the social systems involved in *understanding* justice (e.g., Rice et al., 2018), the memory systems involved in encoding environmental cues to justice (e.g., Davis et al., 2020a), the interoceptive systems that process internal bodily sensations associated with experiencing an instance of *justice* (e.g., perhaps a steadying heartrate and reduction in muscle tension; Connell et al., 2018), the magnitude systems involved in comprehending quantity (e.g., Wilson-Mendenhall et al., 2013), the temporal brain systems involved in processing time and duration (for discussion, see Binder et al., 2016; Davis et al., 2020a) and the linguistic systems involved in *communicating about* justice (e.g., Borghi & Zarcone, 2016).

Concepts that are supported by these systems more than by sensory or motor systems tend to be considered more "abstract." But even highly abstract concepts like *idea* involve some sensorimotor experience (see Lynott et al., 2020). Indeed, it is increasingly accepted that abstractness is a continuum—there is no real dichotomy between abstract and concrete concepts. Instead, where a given concept falls on the abstract-to-concrete continuum is determined by the relative contributions of sensorimotor vs. these other systems (for further discussion, see Vigliocco et al., 2009; for a more detailed discussion of "abstract" concepts and embodied frameworks, see Borghi et al., 2017, 2019).

Further, implicit in experience-based, embodied theories of semantic memory is the idea that conceptual representations are *individualized*. That is, because we have different experiences, my representation of *coffee* is different from that of my local barista. But if we all have different semantic representations (e.g., if *coffee* means something different to me than it does to you), how can we communicate with each other? This poses a difficult—though not insurmountable—problem

for embodied theories (see e.g., Yee & Thompson-Schill, 2016). Later in this paper, we speculate on how uniting distributional and embodied data under a common framework provides a potential solution to both the problem of abstract concepts and the question of how shared meaning is achieved.

Overall, it is becoming increasingly evident that in order to comprehend the meaning of something, it helps to (at least partially) reengage the neural systems that are involved in actually experiencing that thing. This suggests that information in these neural systems constitutes part of a concept's meaning. However, most of the effects seen in embodied cognition research are relatively small: the ability to identify a *hammer*, for example, is not completely lost when a patient suffers damage to motor areas of the brain. One might imagine that this lack of catastrophic interference is a problem for embodied accounts. But there are two reasons for this lack of catastrophic interference. First, because concepts are distributed over multiple sensorimotor modalities (e.g., motor, visual, auditory; see Figure 1), when one modality is interfered with, much of the representation may still be available. Second, concepts are also supported by knowledge that is not directly sensorimotor. This includes information that does not have obvious correlates in any individual sensory or motor modality (and may be, e.g., emotional, social, or interoceptive, or stored in higher levels of the semantic system), and it also includes language. Indeed, although much of semantic knowledge comes from direct experience with objects and actions, much also comes from spoken (and written) language—we have knowledge of places that we have never been, and of people we have never met. We turn now to semantic knowledge derived from language input, before considering how sensorimotor and language knowledge may mutually reinforce one another.

Distributional semantics

Philosophical and historical background

Distributional semantic models have been developed based on the distributional hypothesis, which suggests that a word derives meaning as a function of the 'company it keeps'—that is, the words and linguistic contexts with which it tends to occur (e.g., Firth, 1957; Harris, 1954). While some have suggested that the learning mechanisms in distributional models are general mechanisms that could in principle handle any type of data[4] (e.g., events, experiences; Landauer & Dumais, 1997; see also Günther et al., 2019), the vast majority of implemented models use linguistic corpus data as input.). Thus, in line with the field more broadly (see e.g., Lenci, 2018), in this paper, when we refer to distributional models we are generally referring to models based on language corpus data.

Despite the fact that meaning in distributional semantic models is almost invariably derived from language input only, they have been remarkably successful at capturing, in broad strokes, important aspects of the organization of semantic knowledge. For instance, these models can make human-like judgments about category membership, and about overall semantic similarity (for review, see Lenci, 2018). In what follows, we present a brief overview of the main properties of the primary families of distributional semantic models. This will set the stage for later discussion of hybrid models that incorporate sensory-perceptual data.

---

[4] An exciting new field of work has begun to implement distributional principles for visual scenes (Sadeghi et al., 2015), often using computer vision techniques (e.g., Bruni et al., 2014).

The philosophical history of distributional theories can be traced to Wittgenstein (1953), who suggested that word meaning is characterized by a word's use in language. Distributional frameworks, like embodied semantics, share the notion that our knowledge is derived from *association.* But in contrast to embodied semantics, which emphasize sensory, motor, and action associations, in distributional models, concept representations are computed from word-based co-occurrence vectors, on which we can measure the similarity of the contexts in which words appear (e.g., our representation of *coffee* includes lexical associations with the words *mug*, *sip*, and *brown*, and is related to *tea* because both tend to occur in similar contexts). That is, words have meaning by virtue of the frequency with which they co-occur with other words, or the extent to which they tend to occur in similar contexts. How exactly those statistical patterns of co-occurrence are extracted and analyzed differs among distributional approaches.

Types of models

Below, we provide a rough characterization of how different types of distributional semantic models build semantic knowledge. For a more complete picture of current and future directions, we refer the reader to Boleda (2020) and Lenci (2018), as well as Baroni et al. (2014), Mandera et al. (2017), and Wingfield and Connell (2019).

Early psychological research on distributional semantic models was dominated by *count* models, which count how many times a word appears in particular contexts, or how many times other words co-occur with it (for discussion, see also Baroni et al., 2014; Mandera et al., 2017). Of the count models, *latent abstraction models*—the most well-known of which is latent semantic analysis (LSA; Landauer & Dumais, 1997)—have had perhaps the most lasting impact on the field. These models, in line with the distributional hypothesis, compute co-occurrence frequencies across large corpora of linguistic contexts. They then apply a dimensionality reduction technique to the data to derive a matrix that is meant to reflect higher-order semantic relationships. This dimensionality reduction also results in patterns of similarity that extend beyond co-occurrence: that is, if words tend to occur in similar (but not necessarily the same) contexts, they come to be related. These features allow the models to perform well on semantic similarity judgments and English-language tests for nonnative speakers (Landauer & Dumais, 1997).

Passive co-occurrence models are similar in their counting of word co-occurrences across contexts, but they do not perform dimensionality reduction as in latent abstraction models. The most well-known of these models is the hyperspace analog to language (HAL; Lund & Burgess, 1996). Passive co-occurrence models slide a moving window (window size is typically on the order of several words or a sentence) over text corpora, allowing for incremental learning of semantic representations. These models thus generate representations based on a plausible yet simplified mechanism of human learning, that is, Hebbian learning: they accumulate co-occurrence information over time. Compared to LSA, the representations that emerge from HAL are sparse, but still predict semantic judgments. More modern instantiations (e.g., COALS; Rohde et al., 2009) perform even better on such tasks by changing the method of calculating co-occurrence and applying a decomposition technique (Riordan & Jones, 2011).

In *Bayesian models*, instead of simply counting co-occurrence statistics (e.g., relying on principles of Hebbian learning), the problem of human semantic representation is formulated as one of rational statistical inference. The most influential of these has been the topic model (Griffiths et al., 2007).

Such models think of semantic organization probabilistically: any given document is a probability distribution of topics, and each topic is a probability distribution of words, where the goal of the model is to estimate the distribution of topics in a given text (as opposed to representing a word in high-dimensional space). Importantly, these models are generative in that they can predict the composition of future documents given a particular mixture of topics. The topic model also allows words to have different meanings depending on the context: since topics are probability distributions over words, a given word differs in its likelihood of appearing in any number of topics, and thus, that word may have different meanings across topics. Such models are successful in accounting for synonym judgments, semantic priming in ambiguous words, and so on (see Griffiths et al., 2007).

Predict models on the other hand share with co-occurrence models the reliance on a context window to understand word meaning, but instead of simply counting those co-occurrences and representing them in vector form, predict models use neural networks to derive error-driven predictions about word characteristics. The most prominent family of predict models is Google's word2vec (Mikolov et al., 2013). These models learn to predict either the current word given the context (usually a window of a researcher-determined width surrounding the target word; the continuous bag of words model), or the context words given the current word (skip-gram model). One major advantage of predict models is that the cognitive mechanism—prediction—is well supported as an actual mechanism of human learning (see Mandera et al., 2017, for discussion). And the reader may notice that the contingencies between words and their contexts that are encoded by such models are similar to those learned in models such as the simple recurrent network (SRN; see Box 1), in which recurrence through time, coupled with a task to explicitly predict what will come next (Elman, 1990), leads to emergent representations that reflect the encoding of such contingencies. Whereas the basic SRN does not scale up to large vocabularies or long sequences of text (because information about long-distance dependencies is, in effect, swamped by more local information), contemporary predict models such as word2vec and its successors, as well as recurrent neural networks (RNNs) with long short-term memory units (LSTMs), which embody the basic computational principles of the SRN while avoiding problems with long-distance dependencies, do scale[5]. But the *nature* of the contingencies they are capable of encoding is relatively similar in all cases.

Critiques

Although the semantic "knowledge" generated by distributional semantic models can approximate human responses in many semantic tasks, from a psychological perspective, these models are of interest only to the extent that the principles that shape their operation lend insight into human semantic processing. While some specify basic mechanisms that might correspond to the way humans encode associations (e.g., passive co-occurrence models are loosely based on Hebbian learning), other mechanisms (e.g., LSA's dimensionality reduction) incorporate no method for learning across time (as is present in generative models like the topic model) or for prediction (as is present in predict models), making them seem implausible (i.e., children presumably do not acquire millions of words only to reduce them into semantic vectors at a later date). For this reason, some distributional models have been criticized as being mere methodological tools, not theories of

---

[5] To vocabularies of tens of thousands of words.

semantic memory (Perfetti, 1998). The most common critique, however, is one that—regardless of the psychological plausibility of the mechanisms by which they build semantic knowledge—applies to all distributional semantic models that use only linguistic corpus data as input: unlike embodied theories, typical distributional models provide no mechanism by which the symbols they process (i.e., words) are linked to the concepts to which they refer—words are understood through their relations to other words, but how do any of those words latch onto meaning out in the world? That is, how are they "grounded" in the real world?

The problem of symbol grounding is illustrated by Searle's (1980) *Chinese room problem* (see also Harnad, 1990). A variant of the problem goes like this: You are a monolingual speaker of English and isolated in a room with nothing but a huge book. You have been told how to use this book to look up (based on appearance) any sequence of Chinese characters to find a second "response" sequence. An interlocutor is outside the room, and you must communicate with her using only slips of paper slid under the door. She slides a piece of paper to prompt a response from you, and you search the book for an appropriate response. Ultimately, you find the unfamiliar squiggles that match her squiggles and submit your response. She is under the impression that you understand Chinese, but do you?

Of course not (at least, not in any intuitive sense of understanding). It has been argued that, like you in the Chinese room, distributional models do not truly understand situations because, like the symbols in the Chinese room problem, the symbols in distributional models are not tied to real experience (Glenberg & Robertson, 2000). To illustrate the contrast, imagine sitting in your apartment, writing a paper while finishing a coffee, when suddenly the ceiling springs a leak. You gulp down your coffee and position the mug under the leak while you search for a larger vessel. How did you know to use the mug in this novel way? According to embodied theories, you perceive the mug's *affordances* (i.e., possibilities for action; Gibson, 1979; see also Glenberg, 1997), that is, it can hold liquid. You also have *experience* pouring liquid into mugs.

But what about distributional models? At least some fail at this task. Glenberg and Robertson (2000) generated several settings (e.g., "Zoey's roof sprung a leak while she was writing") and a set of sentences including an afforded (e.g., "In place of a bucket, she used her mug to catch the water"), a non-afforded (e.g., "In place of a bucket, she used her computer to catch the water"), and a related sentence (e.g., "In place of a bucket, she used a pot to catch the water"). Human participants had no difficulty distinguishing the afforded sentence from the non-afforded sentence: the afforded sentence was rated as a sensible solution. However, LSA did not make the same distinction—that is, cosine values were the same between afforded, non-afforded, and related sentences, and they did not predict human sensibility judgments.

While distributional semantic models that take only linguistic data as input may be unable to make use of affordances, what if distributional models are fed with embodied data (e.g., Johns & Jones, discussed below)? Are they then able to use that sensory-perceptual data—as humans are—to recognize the affordances of the mug? Can they then recognize that a mug can be used to temporarily contain a leaky ceiling?

RECONCILIATION: A REVIEW OF HYBRID APPROACHES

In addition to differences in the type of data considered important for representing meaning, a divide across disciplinary boundaries has exaggerated the difference between distributional and embodied approaches. For example, while embodied theorists typically rely on methods from experimental psychology and cognitive neuroscience, researchers working on distributional semantics are more likely to use computational methodologies. In the last decade, however, several attempts have been made to unite distributional and embodied approaches under a single framework. Here, we discuss the progress made by two lines of research. First, in computational cognitive science, researchers have implemented "hybrid" computational models that combine proxies for embodied data (typically feature-based representations, e.g., McRae et al., 2005) with distributional language data to assess whether combining the two types of data produces more human-like semantic representations (e.g., Andrews et al., 2009). Second, experimental cognitive science has examined the relative contributions of embodied and distributional language information to human lexical-semantic processing (e.g., Louwerse & Jeuniaux, 2010).

Hybrid (distributional + embodied) models

The first attempt to combine distributional linguistic data and a proxy for embodied data in a single model (Andrews et al., 2009; see also Steyvers, 2010) used a probabilistic Bayesian model based on the topic model (Griffiths et al., 2007) to create a joint distribution of distributional linguistic data and perceptual feature-based data. The semantic representations that emerged from this joint distribution matched human behavior better than if the model was fed either (a) each distribution individually or (b) both distributions independently. This suggests that the emergent representations are not simply the sum-total of feature-based and distributional linguistic representations: it is the interaction between experiential and linguistic data that allows for more human-like semantic knowledge to emerge. (Andrews et al., 2009). An important feature of this model is that it can essentially perform inference, providing a potential solution to the grounding problem for words experienced only via language. For example, let's say we have considerable sensorimotor experience with coffee, but we have never drunk tea before. Via these sensorimotor experiences with coffee, we have a grounded representation of coffee, where coffee is typically hot, has a dark color, is drunk for its stimulating properties, is served in a mug, and so on. The words *coffee* and *tea* happen to occur in similar contexts. Thus, even if the model has never directly experienced tea, it ascribes qualities to tea that are typical of (already grounded) words seen in similar contexts—that is, coffee.

More recent efforts have made this inference process more explicit by training a model to infer the sensory-perceptual properties of a concept based on the lexical associates of that concept (Johns & Jones, 2012; Hoffman et al., 2018). For example, Johns and Jones (2012) used a global memory model, inspired by Barsalou's (1999) *perceptual symbol systems,* that integrated distributional data (word-by-context co-occurrence vectors from Wikipedia; Willits et al., 2007) and multiple proxies for sensory-perceptual data (feature norms, McRae et al, 2005; Vinson & Vigliocco, 2008; and modality exclusivity norms, which indicate the extent to which a given concept is experienced across five modalities, Lynott & Connell, 2009). Not every word in the linguistic corpus has a sensory-perceptual representation, and so the model iteratively generates inferred perceptual representations for those words based on their similarity to all of the words that do have some sensory-perceptual

representation. But can these inferred perceptual representations use affordances to differentiate situations?

To test this, they used action words (e.g., *hang*) to stand in for sentences like "Hang the coat on the ____"[6] and computed their average cosine with object words—the object words consisted of realistic words, afforded words, and non-afforded words (e.g., *rack*, *vacuum*, and *cup*, respectively), as in Glenberg and Robertson (2000). The average cosine was highest for realistic words (*rack*), then for afforded words (*vacuum*), and lowest for non-afforded words (*cup*) when the inferred perceptual representations were used. Thus, although distributional semantic models may not be able to recognize affordances when provided with only language input (Glenberg & Robertson, 2000), when the model had access to perceptual data, it was able to 'recognize' affordances, just as humans do. The model was also sensitive to sensory-perceptual-based priming effects, where for example *typewriter* primes *piano* due to overlap in how the objects are manipulated, despite not sharing an associative or taxonomic semantic relationship (Myung, Blumstein, & Sedivy, 2006). These findings show that equipping a distributional-style model with sensory-perceptual property data may effectively simulate embodied phenomena.

Most hybrid computational models have, like Johns & Jones (2012), approached the problem of combining distributional and sensory-perceptual data by inputting distributional linguistic data, in the form of semantic vectors that have been derived from large corpora of text, alongside a proxy for embodied data into a distributional model. However, even without using previously derived semantic vectors as the language input, it is possible for a system to learn semantic relationships from co-occurrence. This was demonstrated by Hoffman et al. (2018), who combined an SRN (Elman, 1990; see Box 1) with a *hub-and-spoke* architecture, an influential model of semantic memory which suggests that conceptual knowledge consists of spatially distributed modality-specific information that converges in a central hub (Rogers et al., 2004; for reviews, see Lambon Ralph et al., 2017; Patterson et al., 2007). The model derives semantic representations from event-like sequences of verbal inputs and sensorimotor units, and predicts verbal and sensorimotor output (see Hoffman et al., 2018, Figure 1). Within this model, the hub functionally corresponds to the hidden layer illustrated in Box 1, and as such receives, together with the current input, a record of its prior state (essentially an encoding of its successively prior states). This prior state encodes contextual dependencies accumulated over recent experience.

The success of Hoffman et al.'s (2018) model shows that semantic representations can be derived from a continuous sequence of events, as we might imagine the process unfolding in humans. And like the other hybrid models described in this section, their model can also ascribe sensorimotor properties to concepts that do not (in the model) originally have them, as a function of their co-occurrence with concepts that do.

Thus, a critical upshot of the models described in this section is that they shed light on potential mechanisms by which concepts that have not been directly experienced can acquire an embodied representation: as long as we have experienced some concepts (e.g., *coffee*), we can use their sensory-perceptual characteristics to build a representation of concepts that appear in similar

---

[6] The change from sentences to words was made because LSA (and the model of Johns & Jones, 2012) is not a model of sentence comprehension.

contexts but which we have not directly experienced (e.g., *tea*).[7] Through this kind of "acquired embodiment," these models also suggest a mechanism by which more abstract concepts can acquire sensory-perceptual associations (e.g., *death* occurs in similar contexts to *funeral*, which is associated with sensory-perceptual properties like *black*).

This is not to say that these more abstract concepts would be devoid of meaning if they did not "acquire" embodiment from language. Rather, sensorimotor experience may form our earliest representations of all concepts, even highly abstract ones. However, because more abstract concepts tend to occur in variable, spatiotemporally extended contexts (consider that *game* can refer to a game of chess or a game of hockey, or that understanding *justice* entails apprehension of events spread over space and time; see Davis et al., 2020a, for discussion), they are abstracted further away from those experiences (see Pulvermüller, 2013). In addition, as discussed earlier, systems for affect, social cognition, magnitude, temporal properties, interception, and so on may all contribute to the embodied experience of more abstract concepts, and thus, their representation (see also Barsalou, 1999; Vigliocco et al., 2009).

Thus, the "acquired embodiment" mechanism proposed by Hoffman et al. (2018) might allow us to infer, from similar concepts, experiential properties for concepts which have relatively fewer or less stable sensorimotor associations, but it need not be the only path to embodiment of highly abstract concepts, nor is it exclusive to paradigmatically abstract concepts. The following section further probes the interdependencies between distributional linguistic and embodied data and reviews experimental evidence that questions whether this acquired embodiment mechanism is plausible for humans when learning new concepts.

Effects of embodied and distributional linguistic information on semantic processing

The previous section reviewed possible architectures within which distributional linguistic and embodied data might be integrated. The experimental literature, however, complicates this picture: embodied and distributional linguistic data may be so entangled at multiple timescales—from learning and acquisition to real-time processing—that trying to treat them as separate and postulating a mechanism by which they are then combined may create a false dichotomy, and may be implausible as a mechanism by which humans incorporate distributional language and embodied data in building semantic representations.

To what extent does linguistic information contribute to effects that have typically been considered as emerging due to embodiment alone? Embodied cognition investigations of lexical-semantic processing typically do not assess the extent to which other factors (e.g., distributional variables like word cooccurrence frequencies) could explain effects attributed to embodiment. For example, evidence of perceptual simulation in language processing comes from studies showing that words are processed faster when placed on a screen in their iconically canonical location—*attic* is processed faster when presented at the top of a screen, while *basement* is processed faster when presented at the bottom of a screen (Zwaan & Yaxley, 2003). This has been taken to suggest that we simulate situations when processing language. The *symbol interdependency hypothesis* (Louwerse,

---

[7]Note that Barsalou's (1999) *perceptual symbol systems* can also, in principle, perform this function: as long as requisite component simulators have been activated (e.g., *cup*, *dark*, *hot*, *stimulating*), the system can exhibit productivity by combining known perceptual symbols into a novel concept.

2007), however, suggests that sensory-perceptual information is reflected in our language, and because of this, effects of embodied variables can also emerge via frequency-based characteristics of language usage. For instance, given that English is read not only from left to right but (like most languages) from top to bottom, this account predicts that in written English, when a sentence contains both *attic* and *basement*, the relative location of those physical spaces will be reflected in the word order—i.e., the word *attic* will come first more often than *basement*. This is indeed the case. Moreover, word-order frequency explains the location iconicity effect *better* than a measure of location iconicity itself does, suggesting that language experience not only reflects but may also modulate the effects of perceptual experience on language processing (Louwerse, 2008).

This is not to say that embodied factors do not play an important role in lexical-semantic processing (we saw in our historical review of embodied semantics that they do). Rather, in this case, the distributional characteristics of a language *roughly encode* embodied characteristics of perceptual experience, and these 'less precise' representations can stand in for full perceptual simulation during rapid conceptual processing, whereas resources necessary for more detailed perceptual simulation are deployed during slower language processing or when processing images (Louwerse & Connell, 2011; Louwerse & Jeuniaux, 2010; see also Barsalou et al., 2008; Connell, 2019; Connell & Lynott, 2013; Louwerse et al., 2015; Santos et al., 2011). As another example, there is evidence that a word's meaning is influenced by the *embodied properties* of the contexts in which it appears. Specifically, when the emotional valence, arousal, and concreteness (properties that arguably reflect embodied experience) of a word's average context[8] are analyzed, each of these properties explains significant variance in lexical decision times and recognition memory for the word above and beyond that explained by these same properties of the word itself (Snefjella & Kuperman, 2016). Thus, experiential properties of the (average) contexts a word appears in become a part of that word's meaning (see also Elman, 2009; Yee & Thompson-Schill, 2016), and these properties are reflected in lexical-semantic processing. Data like these suggest that distributional linguistic and embodied information are intimately linked: the way we use language—and the distributional characteristics that reflect this usage—is reflective of our embodied experience.

This link is complicated further by work showing that, in addition to language taking on properties of our embodied experience, our sensory-perceptual experiences can be shaped by language. That is, labelling an object can carve boundaries into our experience, changing the way we activate knowledge about object concepts (e.g., Edmiston & Lupyan, 2015) at the earliest stages of visual processing (Boutonnet & Lupyan, 2015), and even determining whether we see something or not (Lupyan & Spivey, 2010; Lupyan & Ward, 2013). For further discussion of how language affects perception, see Lupyan et al. (2020). All of this work is consistent with the take-home message of this review: what have traditionally been considered "embodied" and "distributional" language-based data are so interconnected that a meaningful divide cannot be made.

Another important question is how embodied and distributional linguistic information interact in learning: how much can we learn from distributional linguistic information, and how do words come to take on the embodied properties of the contexts in which they occur? As described above, it has been proposed that we can learn embodied meanings of words through *acquired embodiment* (e.g., Hoffman et al., 2018), whereby sensory-perceptual properties can be attributed to, for example,

---

[8] In this case, the five words preceding and five words following each instance of a word in a large corpus.

more abstract concepts by virtue of them sharing linguistic contexts with concrete concepts, (e.g., a relatively abstract concept like *death* might become associated with *black* via associations with *funeral*). And indeed, novel words *can* acquire embodied-like representations from purely linguistic experience (Günther et al., 2020). Specifically, after novel words were learned in contexts implying upwards or downwards movements, action-congruency effects were found if participants had to access word meaning. This suggests that embodied representations can be acquired via contextual association (see also Snefjella et al., 2020; Snefjella & Kuperman, 2016). However, when people learned similar concepts but were tested for action-congruency effects using a task that did not require accessing the words' meanings (Günther et al., 2018) no action-congruency effects were observed. Thus, short-term experience with language may not be enough to produce effects typically interpreted as reflecting *obligatory* engagement of embodied conceptual knowledge, but such embodied properties of the linguistic context in which a novel word is learned may be recruited when we *explicitly* consider the meaning of that novel word.

Learning novel words in "distributional" language contexts also seems to facilitate category learning—so long as the novel words are presented with known words that have coherent semantic organization (Ouyang et al., 2017). For instance, even if you don't know anything about *yerba mate*, finding out that it is associated with *tea*, *cups*, and *cafes*, can facilitate recognition of sentences containing the novel word, assimilate novel words into categories (either animals or vehicles), and inductively associate novel words with the appropriate referents (Ouyang et al., 2017). In contrast, if you only know that *yerba mate* is associated with other unknown words like a *bombilla* and a *guampa*[9] is unlikely to help you understand its meaning. Knowing about the associates of *yerba mate* informs our semantic knowledge of *yerba mate* only if the associates are meaningful to us. The reader may notice the relevance of the symbol-grounding problem here—a new word defined in terms of other, ungrounded words cannot acquire deep meaning.

The findings reviewed above are consistent with the suggestion that distributional models that take only language data as input may be adequate for broadly capturing semantic similarity, and emerging evidence suggests that novel words may, to some degree, acquire embodied representations by virtue of the embodied properties of the contexts in which they appear (Günther et al., 2020; Snefjella et al., 2020; see also Snefjella & Kuperman, 2016). An interesting open question is whether embodied experience alone (e.g., Öttl et al., 2017) also facilitates category learning, or whether language is necessary for carving categorical boundaries into our experience (for review, see Lupyan, 2012). While it seems that some categories—dense ones, with highly overlapping sensory-perceptual features across exemplars (e.g., dark, hot beverages)—might form without language (though their formation is no doubt aided by language), others with more abstract rules for category membership might rely on language and its co-occurrence properties (e.g., mammals; Markman & Hutchinson, 1984; Sloutsky, 2010; see Davis & Yee, 2019, for discussion; see also Lupyan, 2009).

---

[9] *Yerba mate* is a kind of tea popular in South America. It is typically drunk out of a hollowed gourd called a *guampa* and sipped from a flattened metal straw with a filter to strain the infusion. The straw is called a *bombilla*.

Summary of hybrid approaches

The experimental work described in this section suggests that distributional linguistic information is parasitic on perceptual data—language structure comes to reflect our embodied experiences. On the converse, embodiment can, to some degree, emerge solely as a result of distributional associations (Günther et al., 2020; but see Günther et al., 2018). It is clear that there are non-trivial interdependencies between distributional linguistic information and embodied-perceptual information (see also Andrews et al., 2014). The way that language is structured reflects the nature of our shared embodied experiences (e.g., Louwerse, 2008), and the way we perceive our embodied experiences is shaped by language (see Lupyan, 2012; Lupyan et al., 2020 for reviews). Thus, some form of hybrid approach that accommodates these interdependencies is necessary for building an adequate account of semantic memory. But what is the nature of the interdependency between distributional language and embodied data, and how does it emerge? In the sections below, we discuss possible solutions for these issues—solutions that emerge from the accounts reviewed above—and suggest some concrete steps for future cross-disciplinary work.

IMPLICATIONS FOR THEORIES OF SEMANTIC MEMORY

As described earlier, all distributional models propose that semantic knowledge arises as a function of some cognitive mechanism which attends to, processes, and stores the statistical regularities and associations in natural language. Although some studies have emphasized differences between embodied and distributional accounts of semantic memory (e.g., Glenberg & Robertson, 2000), this mechanism is not so different from the mechanism by which sensorimotor information comes to comprise concept representations under embodied cognition theories, where the simultaneous firing of neuron assemblies sensitive to, for example, touch, taste, sight, and speech lead to the association of those experiences over time (e.g., Barsalou, 1999; Damasio, 1989; Pulvermüller, 2013). And indeed, in most hybrid computational models, distributional linguistic and sensory-perceptual data are processed by the same mechanism (e.g., Andrews et al., 2009; Hoffman et al., 2018). Just as reading or hearing a word entails activation of its (linguistic) contextual associates for distributional language models, reading or hearing a word entails activation of its sensory, motoric, and perceptual associates for embodied accounts. We know words by the "linguistic and perceptual company they keep" (Louwerse, 2018).

In line with the conclusions offered by hybrid accounts, it is perfectly compatible with embodied theories for linguistic labels to develop in concert with other perceptual attributes of a concept, with the difference simply being that the event giving rise to a linguistic label is the perception of an auditory or visual word instead of a non-linguistic sensation, perception, or action. The words are simply integrated into associated simulators in memory. Indeed, what is sometimes cited as a strong embodied view (Barsalou, 1999; for discussion, see Meteyard et al., 2012) actually captures both types of information, emphasizing statistical linguistic processing and embodiment. In this way, distributional and embodied information are necessarily linked from the earliest stages: there is no meaningful separation between them, because they are never separate. An accompanying idea is that linguistic labels represent "just" another feature of a concept, albeit one which can activate conceptual knowledge in important (or *privileged*) ways (for discussion, see e.g., Connell, 2019; Lupyan, 2012; Yee, 2019). Just as there are differences in the time course of activation for different sensorimotor features (e.g., function and shape; Yee et al., 2011), labels may be activated faster

than detailed sensorimotor information. These differences need not reflect qualitative differences between "types" of feature information, but rather differences in the level of abstraction at which each feature operates (or the contextual appropriateness of a given feature; for discussion, see Yee & Thompson-Schill, 2016). Moreover, the label is *invariant* in the sense that whereas some features may be more or less active, or entirely absent on different instantiations of a concept, the label is a feature that, being generally applicable, can act as an anchor that binds more variable features. Perhaps because of their invariance, labels are effective as a computationally inexpensive way to access conceptual knowledge.

In hybrid (embodied + distributional) computational models that incorporate sensory-perceptual information, a mechanism is built-in by which words that refer to things that are not experienced through the senses (i.e., words for highly abstract concepts) can "acquire" embodiment. But the evidence reviewed here suggests that even the most abstract concepts involve *some* sensorimotor experience, and that linguistic labels develop in concert with perceptual symbols. Accordingly, it is not *necessary* for words to 'acquire' embodiment through contextual association. Rather, even highly abstract concepts can start out with embodied associations, and then be abstracted further away from embodied experience if they occur in more variable contexts (as concepts that are relatively more abstract do; Hoffman et al., 2013). In such cases, language may be more important for representation and processing because of its ability to help organize regularities in the environment, especially when those regularities are more variable or are organized via abstract rules not consistent across sensory-perceptual experiences (for developmental review, see Markman, 1990). For example, the linguistic label in conjunction with a prefrontal cortex-based selection mechanism may help to group sparsely distributed features of a category (e.g., in highly abstract concepts like *equivalence*; Sloutsky, 2010) and/or language may augment the environment by mediating recall and directing attention toward the features of categories and concepts (Clark, 2006; see also Lupyan, 2012; Althaus & Mareschal, 2014; for a broad overview of these issues, see Yee, 2019).

In addition to their ability to help group perceptual features into categories, words can also acquire category structure through language context alone in an experimental setting (Ouyang et al., 2017), suggesting that learning from language context alone may also be sufficient to learn the types of categories important to a shared understanding of how to categorize the world. Further pursuit of these issues might yield insight into a common critique of embodied theories: if our concept representations are built of individual experiences, how is it that we can communicate at all? How do we know what each other is talking about? One rebuttal to this concern is that we largely experience the same world—while many personal experiences are different, you and your neighbors have roughly the same experiences with cars, coffee, and carpets, even if my favorite barista's coffee representation engenders a more detailed simulation than mine does. This may be good enough— perfect representational overlap is not required to achieve successful communication, as long as there is sufficient overlap given the current communicative context (for discussion see, e.g., Casasanto & Lupyan, 2015; Connell & Lynott, 2014; Taylor & Zwaan 2009; Yee & Thompson-Schill, 2016). But if we can acquire category structure through linguistic context alone, this suggests that language usage is also a powerful mechanism through which we gain access to—and assimilate new information into—categories of knowledge that are largely agreed upon within human societies. Given the impressive successes of distributional models at predicting major semantic phenomena

like similarity judgments and priming, it should be no surprise that learning through language usage facilitates learning categories and relationships thought to be shared across individuals.

The ideas synthesized here lead to a position similar in many ways to that of Clark (2006) and later, Dove (2018, 2019), who suggest that language provides another domain of perceivable objects (e.g., words) which are used—in the same way as the sensory-perceptual properties of the external environment—for building coherent semantic representations through our fundamental ability to learn through statistical and associative regularities. That is, in a sense, their argument is also that distributional linguistic and embodied information should not be considered distinct.[10] Having words for things helps us to abstract and organize across regularities in the environment, and the distributional properties of language may also help us to abstract and organize across regularities in words so as to reflect category structure shared across individuals. These features of language allow our conceptual knowledge to be largely shared across members of a community, even if finer-grained details differ with individual experience.

NEXT STEPS

Cross-disciplinary approaches must be a continuing point of emphasis for further progress on these issues. Experimental approaches must investigate how distributional and embodied information interact in learning and representation, and what type of learning is necessary for embodied-like representations to emerge (e.g., Öttl et al., 2017; Günther et al., 2020). Neural investigations of this may be particularly important: given that embodied theories hold that sensory-perceptual experience gives rise to concept representation in modality-specific neural systems, one might ask: is there a mechanism by which linguistic context-derived representations can come to be encoded in sensorimotor cortex? From a different angle, how do particular properties of our environments influence (or become influenced by) distributional properties of a language? And to what extent do these effects of our environments on distributional statistics of language interact with individual differences in embodied experiences (e.g., discounting vs. reinforcing experiences)?

Developmental science also has a crucial role to play towards understanding the roles of embodied and distributional language experience in building semantic memory. At the heart of both approaches is the idea that *experience* is central to building conceptual representations, and that the resultant structure of semantic memory reflects those idiosyncratic experiences. Work on embodied cognition suggests that sensory, perceptual, and motor *experience* contributes to concept representations—that is, the more you experience something in a particular modality, the more its corresponding concept is represented in that modality (e.g., Davis et al., 2020c, Yee et al., 2013). However, these investigations typically measure experience at a single timepoint, when conceptual knowledge is assessed. Understanding the process by which experiences are laid down in sensorimotor systems over time is a ripe topic for future work. Even more poorly understood is how the interaction between embodied and distributional language information might contribute to developmental milestones. Using computational models (e.g., Andrews et al., 2009), it may be

---

[10] If one takes the perspective that sensorimotor experience is continuous, whereas language is discrete, it may seem that they are fundamentally distinct types of data and thus ought to be processed differently. But in fact, linguistic information is *not* discrete—the speech signal, for example, like sensorimotor experience, is continuous in nature, but can be perceived categorically (Liberman et al., 1967; see Harnad, 2003, 2017 for review). The act of labelling may be particularly important for categorical perception (e.g., Lupyan, 2017).

possible to derive empirical predictions about how children might represent the meaning of concepts that they have not yet directly experienced, by virtue of distributional associations with other, related concepts. Addressing these sorts of questions will engender a better understanding of how linguistic and embodied inputs interact from acquisition.

Computational approaches should, on the other hand, acknowledge that distributional language and embodied information cannot be considered separately, and that even so-called abstract concepts, which seem at first to be dependent on the distributional statistics of language, are not amodal—that is, even what have traditionally been considered quite abstract concepts (e.g., *truth*) have embodied components (see Borghi et al., 2017, for review; see also Lynott et al., 2020). The controlled semantic cognition approach (Hoffman et al., 2018) is a step in this direction because it implements distributional properties (i.e., language co-occurrence) in real time using an SRN. However, this approach might benefit from allowing the sensorimotor nodes to reflect a broader array of properties (e.g., affective qualities) and from allowing all concepts (even highly abstract) to have *some* embodied experience from the outset. In addition to accounting for the context-dependent sensorimotor representation of concepts, such models should strive to account for the sorts of semantic phenomena (e.g., semantic similarity) that distributional models are so well-suited to explaining. If concept representations emerge from episodes of joint linguistic and embodied experience, another interesting avenue might be to integrate principles from exemplar-based models (e.g., Hintzman, 1986). Such models can incorporate perceptual and linguistic representations within a single memory trace and learn semantic representations from these exemplar memory traces (Johns & Jones, 2015).

Emerging work also suggests that there may be no "one-size-fits-all" distributional language model. That is, one model may capture behaviour better for a certain type of semantic task, while another model may be better at capturing behaviour for a different one (Wingfeld & Connell, 2019). Model characteristics (e.g., model type, corpus size and quality) influence how well the model captures behavior, and these effects differ as a function of the *level* at which the task operates—for instance, some distributional models optimally capture synonym judgment (a relatively low-level task), while others better explain semantic decisions (a relatively high-level task). Likewise, different semantic relationships are best captured by different distributional language models (Brown et al., 2020). Humans may flexibly use distributional semantic knowledge as a function of situational demands (Wingfield & Connell, 2019), much in the same way that embodied representations may be recruited to differing degrees depending on the context (see e.g., Connell, 2019; Yee & Thompson-Schill, 2016). This work calls for a better understanding of the interplay between (different types of) distributional language knowledge and (different types of) sensory-perceptual knowledge as conceptual processing unfolds. Models that integrate the properties described in this review could potentially explain (a) the rich, detailed, embodied properties of semantic representations (e.g., reliance of particular concepts on the *manual* modality, as in *hammer*), (b) broad semantic relations that are shared across individuals (e.g., similarity across concepts, where *hammer* is related to *wrench*), and (c) how distributional and embodied information interact to build concept representations based on experiential association.

CONCLUSIONS

Embodied and distributional perspectives, at first glance, appear to be distinct approaches to answering the same question: how do humans understand and represent the meaning of things? Because early distributional models handled only linguistic data, it was suggested that the representations they contained made no contact with the world. That is, they were ungrounded. But there is no reason for "distributional" to mean "linguistic," and it is increasingly recognized that distributional and embodied approaches are not mutually exclusive, and are even complementary (for further discussion, see Günther et al., 2019). While hybrid computational approaches have treated sensory-perceptual and distributional information as distinct but interacting data types, experimental approaches and embodied theories suggest that this divide has little traction in reality. Distributional and embodied information are entangled through experiential association from the earliest stages of conceptual development. The implications of this for the emergence of embodied concept representations and implementation of these principles in formal architectures remain ripe topics for cross-disciplinary endeavors.

References

Allport, D. A. (1985). Distributed memory, modular subsystems and dysphasia. In S. K. Newman & R. Epstein (Eds.), *Current perspectives in dysphasia* (pp. 207–244). Edinburgh, UK: Churchill Livingstone.

Althaus, N., & Mareschal, D. (2014). Labels direct infants' attention to commonalities during novel category learning. *PLOS ONE*, *9*(7), e99670.

Andrews, M., Frank, S., & Vigliocco, G. (2014). Reconciling embodied and distributional accounts of meaning in language. *Topics in Cognitive Science*, *6*(3), 359–370.

Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating Experiential and Distributional Data to Learn Semantic Representations. *Psychological Review*, *116*(3), 463–498.

Baroni, M., Dinu, G., & Kruszewski, G. (2014, June). 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 (Volume 1: Long Papers*; pp. 238–247).

Barsalou, L. W. (1999). Perceptions of perceptual symbols. *Behavioral and Brain Sciences, 22*(4), 637–660.

Barsalou, L. W. (2016). On staying grounded and avoiding quixotic dead ends. *Psychonomic Bulletin & Review*, *23*(4), 1122–1142.

Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In *Symbols, embodiment, and meaning* (pp. 245–283).

Binder, J. R., Conant, L. L., Humphries, C. J., Fernandino, L., Simons, S. B., Aguilar, M., & Desai, R. H. (2016). Toward a brain-based componential semantic representation. *Cognitive Neuropsychology*, *33*, 130–174.

Boleda, G. (2020). Distributional semantics and linguistic theory. *Annual Review of Linguistics*.

Borghi, A. M., Barca, L., Binkofski, F., Castelfranchi, C., Pezzulo, G., & Tummolini, L. (2019). Words as social tools: Language, sociality and inner grounding in abstract concepts. *Physics of Life Reviews*, *29*, 120–153.

Borghi, A. M., Binkofski, F., Castelfranchi, C., Cimatti, F., Scorolli, C., & Tummolini, L. (2017). The challenge of abstract concepts. *Psychological Bulletin, 143*(3), 263–292.

Borghi, A. M., & Zarcone, E. (2016). Grounding abstractness: abstract concepts and the activation of the mouth. *Frontiers in Psychology*, *7*, 1498.

Boutonnet, B., & Lupyan, G. (2015). Words jump-start vision: A label advantage in object recognition. *Journal of Neuroscience*, *35*(25), 9329–9335.

Brown, K.S., Yee, E., Saltzman, E., Magnuson, J.S. & McRae, K. (2020). What do computers know about semantics anyway? Testing distributional semantics models against a broad range of relatedness ratings. Poster presented at the 42nd Annual Conference of the Cognitive Science Society. Virtual Conference.

Bruni, E., Tran, N. K., & Baroni, M. (2014). Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, *49*, 1–47.

Buccino, G., Dalla Volta, R., Arabia, G., Morelli, M., Chiriaco, C., Lupo, A., … Quattrone, A. (2018). Processing graspable object images and their nouns is impaired in Parkinson's disease patients. *Cortex*, *100*, 32–39. doi:10.1016/j.cortex.2017.03.009

Casasanto, D., & Lupyan, G. (2015). All concepts are ad hoc concepts. In E. Margolis & S. Laurence (Eds.), *The conceptual mind: New directions in the study of concepts* (pp. 543–566). Cambridge, MA: MIT Press.

Chao, L. L., & Martin, A. (1999). Cortical regions associated with perceiving, naming, and knowing about colors. *Journal of Cognitive Neuroscience*, *11*(1), 25–35.

Clark, A. (2006). Language, embodiment, and the cognitive niche. *Trends in Cognitive Sciences*, *10*(8), 370–374.

Connell, L. (2019). What have labels ever done for us? The linguistic shortcut in conceptual processing. *Language, Cognition and Neuroscience*, *34*(10), 1308–1318.

Connell, L., & Lynott, D. (2013). Flexible and fast: Linguistic shortcut affects both shallow and deep conceptual processing. *Psychonomic Bulletin and Review*, *20*(3), 542–550.

Connell, L., & Lynott, D. (2014). Principles of representation: Why you can't represent the same concept twice. *Topics in Cognitive Science*, *6*(3), 390–406.

Connell, L., Lynott, D., & Banks, B. (2018). Interoception: the forgotten modality in perceptual grounding of abstract and concrete concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1752), 20170143.

Damasio, A. R. (1989). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition*, *33*, 25–62.

Davis, C. P., & Yee, E. (2019). Features, labels, space, and time: factors supporting taxonomic relationships in the anterior temporal lobe and thematic relationships in the angular gyrus. *Language, Cognition and Neuroscience*, *34*(10), 1347–1357.

Davis, C. P., Altmann, G. T. M., & Yee, E. (2020a) Situational systematicity: A role for schema in understanding the differences between abstract and concrete concepts. *Cognitive Neuropsychology*.

Davis, C. P., Altmann, G. T. M., & Yee, E. (2020b). Language as a mental travel guide. Commentary on *Above and beyond the concrete: The diverse representational substrates of the predictive brain*. *Behavioral and Brain Sciences.*

Davis, C. P., Joergensen, G. H., Boddy, P., Dowling, C., & Yee, E. (2020c). Making it harder to "see" meaning: The more you see something, the more its conceptual representation is susceptible to visual interference. Psychological Science, 31(5), 505–517. doi:10.1177/0956797620910748

Dove, G. (2018). Language as a disruptive technology: abstract concepts, embodiment and the flexible mind. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1752), 20170135.

Dove, G. (2019). More than a scaffold: Language is a neuroenhancement. *Cognitive Neuropsychology*, 1–24.

Edmiston, P., & Lupyan, G. (2015). What makes words special? Words as unmotivated cues. *Cognition*, *143*, 93ˆ100.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*(2), 179–211.

Elman, J. L. (2009). On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science*, *33*(4), 547–582.

Fernandino, L., Conant, L. L., Binder, J. R., Blindauer, K., Hiner, B., Spangler, K., & Desai, R. H. (2013a). Where is the action? Action sentence processing in Parkinson's disease. *Neuropsychologia*, *51*(8), 1510–1517.

Fernandino, L., Conant, L. L., Binder, J. R., Blindauer, K., Hiner, B., Spangler, K., & Desai, R. H. (2013b). Parkinson's disease disrupts both automatic and controlled processing of action verbs. *Brain and Language*, *127*(1), 65–74.

Firth, J. R. (1957). A synopsis of linguistic theory 1930–1955. In *Studies in linguistic analysis* (pp. 1–32). Oxford, England: Blackwell Publishers.

Gage, N. Y., & Hickok, G. (2005). Multiregional cell assemblies, temporal binding and the representation of conceptual knowledge in cortex: a modern theory by a "classical" neurologist, Carl Wernicke. *Cortex*, *41*(6), 823–832.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston, MA: Houghton Mifflin.

Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences*, *20*, 1–55.

Glenberg, A. M., & Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language*, *43*(3), 379–401.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, *114*(2), 211–244.

Günther, F., Dudschig, C., & Kaup, B. (2018). Symbol grounding without direct experience: Do words inherit sensorimotor activation from purely linguistic context? *Cognitive Science*, *42*, 336–374.

Günther, F., Nguyen, T., Chen, L., Dudschig, C., Kaup, B., & Glenberg, A. (2020). Immediate sensorimotor grounding of novel concepts learned from language alone. *Journal of Memory and Language.*

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, *42*(1-3), 335–346.

Harnad, S. (2003). Categorical perception. In *Encyclopedia of Cognitive Science* (Vol. 67, No. 4). MacMillan: Nature Publishing Group.

Harnad, S. (2017). To cognize is to categorize: Cognition is categorization. In Lefebvre C., & H. Cohen (Eds.), *Handbook of categorization in cognitive science* (pp. 21–54). Elsevier.

Harris, Z. (1954). Distributional structure. *Word*, *10*, 146–162.

Hauk, O., Johnsrude, I., & Pulvermüller, F. (2004). Somatotopic representations of action words in human motor and premotor cortex. *Neuron*, *41*, 301–307.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*(4), 411–428.

Hoffman, P., Lambon Ralph, M. A., & Rogers, T. T. (2013). Semantic diversity: A measure of semantic ambiguity based on variability in the contextual usage of words. *Behavior Research Methods*, *45*(3), 718–730.

Hoffman, P., McClelland, J. L., & Lambon Ralph, M. A. (2018). Concepts, control and context: A connectionist account of normal and disordered semantic cognition. *Psychological Review*, *125*(3), 293–328.

Hume, D. (1748/1977). *An enquiry concerning human understanding.* Hackett Publishing: Cambridge, UK.

Johns, B. T., & Jones, M. N. (2012). Perceptual inference through global lexical similarity. *Topics in Cognitive Science*, *4*(1), 103–120.

Johns, B. T., & Jones, M. N. (2015). Generating structure from experience: A retrieval-based model of language processing. *Canadian Journal of Experimental Psychology*, *69*(3), 233–251.

Jones, M. N., Johns, B. T., & Recchia, G. (2012). The role of semantic diversity in lexical organization. *Canadian Journal of Experimental Psychology*, *66*(12), 1–39.

Kiefer, M., Sim, E. J., Herrnberger, B., Grothe, J., & Hoenig, K. (2008). The sound of concepts: four markers for a link between auditory and conceptual brain systems. *Journal of Neuroscience*, *28*(47), 12224–12230.

Kousta, S. T., Vigliocco, G., Vinson, D. P., Andrews, M., & Del Campo, E. (2011). The representation of abstract words: why emotion matters. *Journal of Experimental Psychology: General*, *140*(1), 14–34.

Lambon Ralph, M. A., Jefferies, E., Patterson, K., & Rogers, T. T. (2017). The neural and computational bases of semantic cognition. *Nature Reviews Neuroscience*, *18*(1), 42.

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*, *1*(2), 1–54.

Lenci, A. (2018). Distributional models of word meaning. *Annual Review of Linguistics*, *4*(1), 1065975714.

Locke, J. (1689/1975). *An essay concerning human understanding*. Oxford, England: Clarendon Press.

Louwerse, M. M. (2007). Symbolic or embodied representations: A case for symbol interdependency. In *Handbook of latent semantic analysis* (pp. 107–120).

Louwerse, M. M. (2008). Embodied relations are encoded in language. *Psychonomic Bulletin and Review*, *15*(4), 838–844.

Louwerse, M. M. (2018). Knowing the meaning of a word by the linguistic and perceptual company it keeps. *Topics in Cognitive Science*.

Louwerse, M., & Connell, L. (2011). A taste of words: linguistic context and perceptual simulation predict the modality of words. *Cognitive Science*, *35*(2), 381–398.

Louwerse, M. M., & Jeuniaux, P. (2010). The linguistic and embodied nature of conceptual processing. *Cognition*, *114*(1), 96–104.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, and Computers*, *28*(2), 203–208.

Lupyan, G. (2009). Extracommunicative functions of language: Verbal interference causes selective categorization impairments. *Psychonomic Bulletin & Review* 16, 711–718. doi:10.3758/PBR.16.4.711

Lupyan, G. (2012). Linguistically modulated perception and cognition: The label-feedback hypothesis. *Frontiers in Psychology, 3.*

Lupyan, G. (2017). The paradox of the universal triangle: Concepts, language, and prototypes. *Quarterly Journal of Experimental Psychology*, *70*(3), 389–412. doi:10.1080/17470218.2015.1130730

Lupyan, G., Abdel Rahman, R., Boroditsky, L., & Clark, A. (2020, preprint). Effects of language on visual perception. *PsyArXiv*.

Lupyan, G., & Spivey, M. J. (2010). Redundant spoken labels facilitate perception of multiple items. *Attention, Perception, & Psychophysics*, *72*(8), 2236–2253.

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.

Lynott, D., & Connell, L. (2009). Modality exclusivity norms for 423 object properties. *Behavior Research Methods*, *41*(2), 558–564.

Lynott, D., Connell, L., Brysbaert, M., Brand, J., & Carney, J. (2020). The Lancaster Sensorimotor Norms: multidimensional measures of perceptual and action strength for 40,000 English words. *Behavior Research Methods, 1–21.*

Mahon, B. Z. (2015). What is embodied about cognition? *Language, Cognition and Neuroscience*, *30*(4), 420–429.

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.

Mandera, P., Keuleers, E., & Brysbaert, M. (2017). Explaining human performance in psycholinguistic tasks with models of semantic similarity based on prediction and counting: A review and empirical validation. *Journal of Memory and Language*, *92*, 57–78.

Markman, E. M., & Hutchinson, J. E. (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. Cognitive Psychology, 16(1), 1–27.

Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science*, *14*(1), 57–77.

Martin, A. (2016). GRAPES—Grounding representations in action, perception, and emotion systems: How object properties and categories are represented in the human brain. *Psychonomic Bulletin & Review*.

Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L., & Ungerleider, L. G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science*, *270*, 102–105.

McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, *37*(4), 547–559.

Meteyard, L., Cuadrado, S. R., Bahrami, B., & Vigliocco, G. (2012). Coming of age: A review of embodiment and the neuroscience of semantics. *Cortex*, *48*(7), 788–804.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems* (pp. 3111–3119).

Myung, J. Y., Blumstein, S. E., & Sedivy, J. C. (2006). Playing on the typewriter, typing on the piano: Manipulation knowledge of objects. *Cognition*, *98*(3), 223–243.

Öttl, B., Dudschig, C., & Kaup, B. (2017). Forming associations between language and sensorimotor traces during novel word learning. *Language and Cognition*, *9*(1), 156–171.

Ouyang, L., Boroditsky, L., & Frank, M. C. (2017). Semantic coherence facilitates distributional learning. *Cognitive Science*, *41*, 855–884.

Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, *8*(12), 976–987.

Perfetti, C. (1998). The limits of co-occurrence: Tools and theories in language research. *Discourse Processes*, *25*, 363–377.

Pulvermüller, F. (2013). How neurons make meaning: Brain mechanisms for embodied and abstract-symbolic semantics. *Trends in Cognitive Sciences*, *17*(9), 458–470.

Rice, G. E., Hoffman, P., Binney, R. J., & Lambon Ralph, M. A. (2018). Concrete versus abstract forms of social concept: an fMRI comparison of knowledge about people versus social terms. *Philosophical Transactions of the Royal Society B*, *373*(1752), 20170136.

Riordan, B., & Jones, M. N. (2011). Redundancy in perceptual and linguistic experience: Comparing feature-based and distributional models of semantic representation. *Topics in Cognitive Science*, *3*(2), 303–345.

Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., & Patterson, K. (2004). Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review*, *111*(1), 205–235.

Rohde, D. L. T., Gonnerman, L., & Plaut, D. C. (2009). An improved model of semantic similarity based on lexical co-occurrence. *Cognitive Science*.

Sadeghi, Z., McClelland, J. L., & Hoffman, P. (2015). You shall know an object by the company it keeps: An investigation of semantic representations derived from object co-occurrence in visual scenes. *Neuropsychologia*, *76*, 52–61.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, *3*(3), 417–457.

Simmons, W. K., Ramjee, V., Beauchamp, M. S., McRae, K., Martin, A., & Barsalou, L. W. (2007). A common neural substrate for perceiving and knowing about color. *Neuropsychologia*, *45*(12), 2802–2810.

Sloutsky, V. M. (2010). From perceptual categories to concepts: What develops? *Cognitive Science*, *34*(7), 1244–1286.

Snefjella, B., Lana, N., & Kuperman, V. (2020). How emotion is learned: Semantic learning of novel words in emotional contexts. *Journal of Memory and Language*, *115*, 104171. doi:10.1016/j.jml.2020.104171

Snefjella, B., & Kuperman, V. (2016). It's all in the delivery: Effects of context valence, arousal, and concreteness on visual word processing. *Cognition*, *156*, 135–146.

Steyvers, M. (2010). Combining feature norms and text data with topic models. *Acta Psychologica*, *133*(3), 234–243.

Taylor, L. J., & Zwaan, R. A. (2009). Action in cognition: The case of language. *Language and Cognition*, *1*(1), 45–58. doi:10.1515/LANGCOG.2009.003

Thompson-Schill, S. L., Kan, I. P., & Oliver, R. T. (2006). Functional neuroimaging of semantic memory. In R. Cabeza & A. Kingstone, *Handbook of functional neuroimaging of cognition* (vol. 2, pp. 149–190). Cambridge, MA: The MIT Press.

Trumpp, N. M., Kliese, D., Hoenig, K., Haarmeier, T., & Kiefer, M. (2013). Losing the sound of concepts: Damage to auditory association cortex impairs the processing of sound-related concepts. *Cortex*, *49*(2), 474-486.

Vigliocco, G., Kousta, S. T., Della Rosa, P. A., Vinson, D. P., Tettamanti, M., Devlin, J. T., & Cappa, S. F. (2013). The neural representation of abstract words: The role of emotion. *Cerebral Cortex*, *24*(7), 1767–1777.

Vigliocco, G., Meteyard, L., Andrews, M., & Kousta, S. (2009). Toward a theory of semantic representation. *Language and Cognition*, *1*(2), 219–247. doi:10.1515/LANGCOG.2009.011

Vinson, D. P., & Vigliocco, G. (2008). Semantic feature production norms for a large set of objects and events. *Behavior Research Methods*, *40*(1), 183–190.

Vukovic, N., Feurra, M., Shpektor, A., Myachykov, A., & Shtyrov, Y. (2017). Primary motor cortex functionally contributes to language comprehension: An online rTMS study. *Neuropsychologia*, *96*, 222–229.

Willems, R. M., Toni, I., Hagoort, P., & Casasanto, D. (2010). Neural dissociations between action verb understanding and motor imagery. *Journal of Cognitive Neuroscience*, *22*(10), 2387–2400.

Willits, J. A., D'Mello, S. K., Duran, N. D., & Olney, A. (2007, January). Distributional statistics and thematic role relationships. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 29, No. 29).

Wilson-Mendenhall, C. D., Simmons, W. K., Martin, A., & Barsalou, L. W. (2013). Contextual processing of abstract concepts reveals neural representations of nonlinguistic semantic content. *Journal of Cognitive Neuroscience*, *25*(6), 920–935.

Wingfield, C., & Connell, L. (2019, preprint). Understanding the role of linguistic distributional knowledge in cognition. *PsyArXiv*. doi:10.31234/osf.io/hpm4z

Wittgenstein, L. (1953/2010). *Philosophical investigations*. Oxford, UK: John Wiley & Sons.

Yee, E. (2019). Abstraction and concepts: When, how, where, what and why? *Language, Cognition and Neuroscience*, *34*(10).

Yee, E., Chrysikou, E. G., Hoffman, E., & Thompson-Schill, S. L. (2013). Manual experience shapes object representations. *Psychological Science*, *24*(6), 909-919.

Yee, E., Huffstetler, S., & Thompson-Schill, S. L. (2011). Function follows form: Activation of shape and function features during object identification. *Journal of Experimental Psychology: General*, *140*(3), 348–363. doi:10.1037/a0022840

Yee, E., Jones, M. N., & McRae, K. (2018). Semantic memory. *The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience* (vol. 3).

Yee, E., & Thompson-Schill, S. L. (2016). Putting concepts into context. *Psychonomic Bulletin & Review*, *23*(4), 1015–1027.

Yeh, W., & Barsalou, L. W. (2006). The situated nature of concepts. *The American Journal of Psychology*, *119*(3), 349–384.

Ziegler, J. C., Montant, M., Briesemeister, B. B., Brink, T. T., Wicker, B., Ponz, A., ... & Braun, M. (2018). Do words stink? Neural reuse as a principle for understanding emotions in reading. *Journal of Cognitive Neuroscience*, *30*(7), 1023–1032.

Zwaan, R. A., & Yaxley, R. H. (2003). Spatial iconicity affects semantic relatedness judgments. *Psychonomic Bulletin & Review*, *10*(4), 954–958.