Toward a brain-based componential semantic representation

Jeffrey R. Binder, Lisa L. Conant, Colin J. Humphries, Leonardo Fernandino, Stephen B. Simons, Mario Aguilar & Rutvik H. Desai

View supplementary material ☐

Published online: 16 Jun 2016.

Submit your article to this journal ☐

Article views: 2933

View related articles ☐

View Crossmark data ☐

Citing articles: 51 View citing articles ☐

Routledge
Taylor & Francis Group

Toward a brain-based componential semantic representation

Jeffrey R. Binder[a], Lisa L. Conant[a], Colin J. Humphries[a], Leonardo Fernandino[a], Stephen B. Simons[b], Mario Aguilar[b] and Rutvik H. Desai[c]

[a]Department of Neurology, Medical College of Wisconsin, Milwaukee, WI, USA; [b]Teledyne Scientific, LLC, Durham, NC, USA; [c]Department of Psychology, University of South Carolina, Columbia, SC, USA

ABSTRACT

Componential theories of lexical semantics assume that concepts can be represented by sets of features or attributes that are in some sense primitive or basic components of meaning. The binary features used in classical category and prototype theories are problematic in that these features are themselves complex concepts, leaving open the question of what constitutes a primitive feature. The present availability of brain imaging tools has enhanced interest in how concepts are represented in brains, and accumulating evidence supports the claim that these representations are at least partly "embodied" in the perception, action, and other modal neural systems through which concepts are experienced. In this study we explore the possibility of devising a componential model of semantic representation based entirely on such functional divisions in the human brain. We propose a basic set of approximately 65 experiential attributes based on neurobiological considerations, comprising sensory, motor, spatial, temporal, affective, social, and cognitive experiences. We provide normative data on the salience of each attribute for a large set of English nouns, verbs, and adjectives, and show how these attribute vectors distinguish a priori conceptual categories and capture semantic similarity. Robust quantitative differences between concrete object categories were observed across a large number of attribute dimensions. A within- versus between-category similarity metric showed much greater separation between categories than representations derived from distributional (latent semantic) analysis of text. Cluster analyses were used to explore the similarity structure in the data independent of a priori labels, revealing several novel category distinctions. We discuss how such a representation might deal with various longstanding problems in semantic theory, such as feature selection and weighting, representation of abstract concepts, effects of context on semantic retrieval, and conceptual combination. In contrast to componential models based on verbal features, the proposed representation systematically relates semantic content to large-scale brain networks and biologically plausible accounts of concept acquisition.

Introduction

Theories of semantic representation generally begin with analyses of conceptual content. This analytical or componential approach assumes that concepts can ultimately be represented by sets of features or attributes that are in some sense primitive or basic components of meaning. The present availability of brain imaging tools has enhanced interest in how concepts are represented in human neural systems, and accumulating evidence supports the claim that these representations are at least partly "embodied" in the perception, action, and other modal neural systems through which concepts are acquired (Binder & Desai, 2011; Kiefer & Pulvermüller, 2012; Meteyard, Rodriguez Cuadrado, Bahrami, & Vigliocco, 2012). In

this paper we explore the possibility of devising a sufficiently explanatory componential model of semantic representation based entirely on such functional divisions in the human brain.

An essential starting point is to define what types of entities are the components of such a representation and how these differ from those of more standard semantic theories. In classical category theory, concepts are defined by the presence or absence of binary features that are necessary and sufficient for category identification (Aristotle, 1995/350 BCE). The concept of *bird*, for example, is composed of such features as *wings, feathers, beak, flying, egg-laying*, and so on. An important extension of classical theory was the realization that most categories have somewhat fuzzy

boundaries due to varying degrees of prototypicality among members (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Thus a penguin could still be identified (though more slowly) as a bird, despite the fact that it does not fly, through a process of jointly weighing all of its features. Extensive work using feature generation tasks has documented what people think of as the features of entities, and how these features rank in importance for a given concept (Cree & McRae, 2003; Garrard, Lambon Ralph, Hodges, & Patterson, 2001; McRae, de Sa, & Seidenberg, 1997; Vinson, Vigliocco, Cappa, & Siri, 2003). Such data provide a powerful means of assessing degrees of similarity between concepts and of grouping concepts into hierarchical categories on the basis of this similarity structure.

A limitation that applies to these standard approaches, however, is that the features they employ to define conceptual content are typically also complex concepts. Physical features like wings, feathers, and beaks are components in the sense that they are parts of a larger entity, but they are no more primitive than the larger entities they define. Like the larger entities, these parts are concepts that must be learned through perceptual and verbal experience, and each has its own set of defining features, which in turn have their own defining features, and so on. The resulting combinatorial explosion presents a bracing problem for constraining feature-based theories. Across even the closed set of known physical objects, the possible set of physical features is extremely large, and the same is true for motion and action features such as flying, swimming, running, jumping, eating, singing, barking, howling, and so on. The inability of standard verbal feature-based theories to locate a set of atomic features from which all others are derived places hard limits on the explanatory value of such approaches in several areas.

One of these limitations, and the one of central concern here, is the lack of any known relationship between verbal features and neurobiological mechanisms. There are no neural systems specifically dedicated to representing feathers, wings, beaks, or flight, nor is it plausible that such dedicated neural systems exist for every conceivable feature. Even if it turns out that the brain representation of features and other complex concepts includes neurons or local neuron ensembles dedicated to a single concept (Barlow, 1972; Bowers, 2009), the mere existence of such "grandmother cells" in itself provides no account of the means by which external or internal stimuli lead to their activation. Without such a mechanistic account, the claim that each possible concept is represented in the brain by a cell or group of cells adds nothing to our understanding of concept representations beyond a mere listing of the concepts. What we seek instead is an understanding of why concepts are organized in the brain in the way that they actually are organized. What facts about the brain determine whether a concept is represented by one group of cells rather than another group? A closely related problem, for which feature-based conceptual representations are similarly limited in providing an answer, is the question of how features and other concepts come to be learned by the brain. The idea that concepts are represented in the brain by abstract "concept cells" also fails to address the deeper question of how concepts can be understood without a "grounding" in sensory–motor experience that enables reference to the external environment and the phenomenological qualia of conscious thought (Harnad, 1990).

The limitations of traditional feature-based semantic theories for understanding concept representation in the brain are, however, a direct reflection of what these theories aim to achieve, which is to describe the things that exist, their similarity structure, and their groupings into categories. The question of how such information is organized in the brain is outside the domain of these aims, so it should not be surprising that these theories have the limitations they have. In contrast, embodiment theories of knowledge representation provide a fairly straightforward analysis of conceptual content in terms of sensory, motor, affective, and other experiential phenomena and their corresponding modality-specific neural representations. In the theory outlined here, these neurobiologically defined "experiential attributes" provide a set of primitive features for the analysis of conceptual content, while simultaneously grounding concepts in experience and providing a mechanistic account of concept acquisition.

Defining conceptual features in terms of neural processes represents a radical departure from traditional verbal feature analysis. Prior work along these lines includes the Conceptual Semantics program of Jackendoff (Jackendoff, 1990), which conceives of

semantic content in terms of experiential primitives, emphasizing in particular the role of space, time, event, and causality phenomena in understanding propositional content. Other related early work includes the sensory–functional theory proposed by Warrington and colleagues to account for category-related object processing impairments in neurological patients (Warrington & Shallice, 1984). In this theory, the relevant semantic content of physical entities is essentially reduced to two experiential attributes processed in two distinct brain networks. Subsequent work has generally recognized the limitations of a simple sensory–functional dichotomy (Allport, 1985; Warrington & McCarthy, 1987), leading to investigation of an ever-expanding list of sensory, motor, and affective attributes of concepts and the neural correlates of these attributes.

Studies of modality-specific contributions to conceptual knowledge have usually focused on a single attribute or small set of attributes. Lynott and Connell collected normative ratings for a sample of 423 concrete adjectives that describe object properties (Lynott & Connell, 2009) and 400 nouns (Lynott & Connell, 2013) on strength of association with each of the five primary sensory modalities. Participants were asked to rate how strongly they experienced a particular concept by hearing, tasting, feeling through touch, smelling, and seeing. Gainotti et al. (Gainotti, Ciaraffa, Silveri, & Marra, 2009; Gainotti, Spinelli, Scaricamazza, & Marra, 2013) expanded on this set by dividing the visual domain into three dimensions (shape, colour, and motion) and adding manipulation experience as a motor dimension. Ratings were obtained on 49 animal, plant, and artefact concepts, revealing highly significant differences between categories in the perceived relevance of these sensory–motor dimensions to each concept. Using a similar set of dimensions, Hoffman and Lambon Ralph (Hoffman & Lambon Ralph, 2013) obtained strength of association ratings for a set of 160 object nouns. Participants were asked, "How much do you associate this item with a particular (colour, visual form, observed motion, sound, tactile sensation, taste, smell, action)?". The resulting attribute ratings predicted lexical–semantic processing speed more accurately than verbal feature-based representations (McRae, Cree, Seidenberg, & McNorgan, 2005), and they supported a novel conceptual distinction between mechanical devices (e.g., vehicles) and other non-living objects, in that the former had strong sound and motion characteristics that made them more similar to animals.

Although prior work in this area has focused overwhelmingly on concrete object and action concepts, more recent studies explore dimensions of experience that may be more important in the representation of abstract concepts. Several authors have emphasized the role of affective and social experiences in abstract concept acquisition and embodiment (Borghi, Flumini, Cimatti, Marocco, & Scorolli, 2011; Kousta, Vigliocco, Vinson, Andrews, & Del Campo, 2011; Vigliocco, Meteyard, Andrews, & Kousta, 2009; Wiemer-Hastings & Xu, 2005; Zdrazilova & Pexman, 2013). Crutch et al. (Crutch, Williams, Ridgway, & Borgenicht, 2012) obtained ratings for 200 abstract and 200 concrete nouns on the relatedness of each word to concepts of time, space, quantity, emotion, polarity (positive or negative), social interaction, morality, and thought, in addition to the more physical dimensions of sensation and action. These representations successfully predicted performance by a severely aphasic patient on an abstract noun comprehension task in which semantic similarity between target and foil responses was manipulated, whereas similarity metrics (latent semantic analysis cosine similarity) based on patterns of word usage (Landauer & Dumais, 1997) were not predictive (Crutch, Troche, Reilly, & Ridgway, 2013). Troche et al. (Troche, Crutch, & Reilly, 2014) provide further analyses of these representations, identifying three latent factors (labelled perceptual salience, affective association, and magnitude) that accounted for 81% of the variance. Abstract nouns were rated higher than concrete nouns on the dimensions of emotion, polarity, social interaction, morality, and space.

Building on this prior work, our aim in the present study was to develop a more comprehensive conceptual representation based on known modalities of neural information processing. This more comprehensive representation would ideally capture aspects of experience that are central to the acquisition of event concepts as well as object concepts, and abstract as well as concrete concepts. The resulting representation, described in the following section, contains entries corresponding to specialized sensory and motor processes; affective processes; systems for processing spatial, temporal, and causal information; social cognition processes; and abstract

cognitive operations. We hope this approach stimulates further empirical and theoretical efforts toward a comprehensive brain-based semantic theory.

Neural components of experience

The following section briefly describes the proposed components of a brain-based semantic representation, which are listed in Tables 1 and 2. Each of these components is defined by an extensive body of physiological evidence that can only be hinted at here. Examples are provided of how each component applies to a range of concepts. Two fundamental principles guided selection of the components. First, we assume that all aspects of mental experience can contribute to concept acquisition and therefore concept composition. These "aspects of mental experience" include not only sensory perceptions, but also

affective responses to entities and situations, experience with performing motor actions, perception of spatial and temporal phenomena, perception of causality, experience with social phenomena, and experience with internal cognitive phenomena and drive states. Second, we focus on experiential phenomena

Table 1. Sensory and motor components, organized by domain.

Domain	Component	Description (with reference to nouns)
Vision	Vision	something that you can easily see
Vision	Bright	visually light or bright
Vision	Dark	visually dark
Vision	Colour	having a characteristic or defining colour
Vision	Pattern	having a characteristic or defining visual texture or surface pattern
Vision	Large	large in size
Vision	Small	small in size
Vision	Motion	showing a lot of visually observable movement
Vision	Biomotion	showing movement like that of a living thing
Vision	Fast	showing visible movement that is fast
Vision	Slow	showing visible movement that is slow
Vision	Shape	having a characteristic or defining visual shape or form
Vision	Complexity	visually complex
Vision	Face	having a human or human-like face
Vision	Body	having human or human-like body parts
Somatic	Touch	something that you could easily recognize by touch
Somatic	Temperature	hot or cold to the touch
Somatic	Texture	having a smooth or rough texture to the touch
Somatic	Weight	light or heavy in weight
Somatic	Pain	associated with pain or physical discomfort
Audition	Audition	something that you can easily hear
Audition	Loud	making a loud sound
Audition	Low	having a low-pitched sound
Audition	High	having a high-pitched sound
Audition	Sound	having a characteristic or recognizable sound or sounds
Audition	Music	making a musical sound
Audition	Speech	someone or something that talks
Gustation	Taste	having a characteristic or defining taste
Olfaction	Smell	having a characteristic or defining smell or smells
Motor	Head	associated with actions using the face, mouth, or tongue
Motor	Upper limb	associated with actions using the arm, hand, or fingers
Motor	Lower limb	associated with actions using the leg or foot
Motor	Practice	a physical object YOU have personal experience using

Table 2. Spatial, temporal, causal, social, emotion, drive, and attention components.

Domain	Name	Description (with reference to nouns)
Spatial	Landmark	having a fixed location, as on a map
Spatial	Path	showing changes in location along a particular direction or path
Spatial	Scene	bringing to mind a particular setting or physical location
Spatial	Near	often physically near to you (within easy reach) in everyday life
Spatial	Toward	associated with movement toward or into you
Spatial	Away	associated with movement away from or out of you
Spatial	Number	associated with a specific number or amount
Temporal	Time	an event or occurrence that occurs at a typical or predictable time
Temporal	Duration	an event that has a predictable duration, whether short or long
Temporal	Long	an event that lasts for a long period of time
Temporal	Short	an event that lasts for a short period of time
Causal	Caused	caused by some clear preceding event, action, or situation
Causal	Consequential	likely to have consequences (cause other things to happen)
Social	Social	an activity or event that involves an interaction between people
Social	Human	having human or human-like intentions, plans, or goals
Social	Communication	a thing or action that people use to communicate
Social	Self	related to your own view of yourself, a part of YOUR self-image
Cognition	Cognition	a form of mental activity or a function of the mind
Emotion	Benefit	someone or something that could help or benefit you or others
Emotion	Harm	someone or something that could cause harm to you or others
Emotion	Pleasant	someone or something that you find pleasant
Emotion	Unpleasant	someone or something that you find unpleasant
Emotion	Happy	someone or something that makes you feel happy
Emotion	Sad	someone or something that makes you feel sad
Emotion	Angry	someone or something that makes you feel angry
Emotion	Disgusted	someone or something that makes you feel disgusted
Emotion	Fearful	someone or something that makes you feel afraid
Emotion	Surprised	someone or something that makes you feel surprised
Drive	Drive	someone or something that motivates you to do something
Drive	Needs	someone or something that would be hard for you to live without
Attention	Attention	someone or something that grabs your attention
Attention	Arousal	someone or something that makes you feel alert, activated, excited, or keyed up in either a positive or negative way

for which there are likely to be corresponding distinguishable neural processors, drawing on evidence from animal physiology, brain imaging, and neurological studies. In particular, we focus on macroscopic neural systems that can be distinguished with in vivo human imaging methods. We do not require or propose that these neural processors are self-contained "modules" or that they are highly localized in the brain, only that they process information that primarily represents a particular aspect of experience.

A detailed listing of the queries used to elicit association ratings for each component is available for download (see link prior to the References section).

Visual components

Visual and other sensory components are listed in Table 1. Components of visual experience for which there are likely to be corresponding distinguishable neural processors include luminance, size, colour, visual texture, visual shape, visual motion, and biological motion. Within the visual shape perception network are distinguishable networks that primarily process faces, human body parts, and three-dimensional spaces.

Luminance is a basic property of vision, but is also relevant to such concepts as *sun, light, gleam, shine, ink, night, dark, black*, and so on, that are defined by brightness or darkness. Luminance is an example of a scalar quantity—that is, one that represents a continuous one-dimensional variable. Linguistic representation of scalars tends to focus on ends of the continuum (e.g., bright/dark, hot/cold, heavy/light, loud/soft, etc.), which raises a general question of whether or to what extent opposite ends of a scalar (e.g., bright and dark) are represented in distinguishable neural processors. Are there non-identical neural networks that encode high and low luminance? The answer depends to some degree on the spatial scale in question. It is known that some neurons respond somewhat selectively to stimuli with high luminance while others respond more to stimuli with low luminance (i.e., ON and OFF cells in retina and V1), resulting in two networks that are distinguishable at a microscopic scale. If these and other luminance-tuned neurons intermingle within a single network, however, the high- and low-luminance networks might be indistinguishable, at least at the macroscopic scale of in vivo imaging.

In contrast, the scalar quantity visual size is linked to large-scale retinotopic organization throughout the visual cortical system and has been proposed as a major determinant of medial-to-lateral organization even at the highest levels of the ventral visual stream (Hasson, Levy, Behrmann, Hendler, & Malach, 2002). At these higher levels, medial-to-lateral organization appears not to depend on the actual retinal extent of a given object exemplar, which varies with distance from the observer, but rather on a "typical" or generalized perceptual representation. Images of buildings, for example, produce stronger activation in medial regions than do images of insects, and the opposite pattern holds for lateral regions, even when the images are the same physical size (Konkle & Oliva, 2012). By extension, the general theory that concepts are represented partly in perceptual systems would predict similar medial–lateral activation differences in the ventral visual stream for concepts representing objects that are reliably large or reliably small.

There is now strong evidence not only for a fairly specialized colour perception network in the human brain (Bartels & Zeki, 2000; Beauchamp, Haxby, Jennings, & DeYoe, 1999), but also for activation in or near this processor by concepts with colour content (Hsu, Frankland, & Thompson-Schill, 2012; Hsu, Kraemer, Oliver, Schlichting, & Thompson-Schill, 2011; Kellenbach, Brett, & Patterson, 2001; Martin, Haxby, Lalonde, Wiggs, & Ungerleider, 1995; Simmons et al., 2007). The question of whether there are distinguishable neural processors activated by particular colour concepts (e.g., *red* vs. *green* vs. *blue*) has not been studied, but this appears unlikely. Colour-sensitive neurons in the monkey inferior temporal cortex form a network of millimetre-sized "globs", within which there is evidence for spatial clustering of neurons by colour preference (Conway & Tsao, 2009). However, the spatial scale of these clusters is on the order of 50–100 μm, beyond the spatial resolution of current in vivo imaging methods. More importantly, it is not clear that the spatial organization of colour *concept* representations should necessarily reflect this perceptual organization. Division of the visible light frequency spectrum into colour concepts varies considerably across cultures (Berlin & Kay, 1969), whereas the spatial organization of colour-sensitive neurons presumably does not. As with the attribute of luminance, we have assumed for these reasons that a single neural processor encodes the colour

content of concepts, regardless of the particular colour in question. A concept like *lemon* that is reliably associated with a particular colour is expected to activate this neural processor to approximately the same degree as a concept like *coffee* that is reliably associated with a different colour.

Visual motion perception involves a relatively specialized neural processor known as MT or V5 (Albright, Desimone, & Gross, 1984; Zeki, 1974). Responses in MT vary with motion velocity and degree of motion coherence (Lingnau, Ashida, Wall, & Smith, 2009; Rees, Friston, & Koch, 2000; Tootell et al., 1995). Motion velocity is relevant for our purposes, as it is strongly represented at a conceptual level. For example, movements may be relatively fast or slow in velocity, and this scalar quantity is central to action concepts like *run, dash, zoom, creep, ooze, walk,* and so on, as well as entity concepts like *bullet, jet, hare, snail, tortoise, sloth*, and so on. Movements may also have a specific pattern or quality, as illustrated, for example, by the large number of verbs that express manner of motion (*turn, bounce, roll, float, spin, twist,* etc.). We are not aware of any evidence for large-scale spatial organization in the cortex according to type or manner of motion; therefore the proposed semantic representation is designed to capture strength of association with visually observable characteristic motion in general, regardless of specific type. The exception to this rule is the perception of biological motion, which has been linked with a neural processor that is clearly distinct from MT, located in the posterior superior temporal sulcus (STS; Grossman & Blake, 2002; Pelphrey, Morris, Michelich, Allison, & McCarthy, 2005). Biological motion contains higher order complexity that is characteristic of face and body part actions. Perception of biological motion plays a central role in understanding face and body gestures and thus the affective and intentional state of other agents (i.e., theory of mind; Allison, Puce, & McCarthy, 2000). Biological motion is thus one of several key attributes distinguishing intentional agents (*boy, girl, woman, man*, etc.) from inanimate entities.

Visual shape perception involves an extensive region of extrastriate cortex, including a large lateral extrastriate area known as the lateral occipital complex (Malach et al., 1995) and an adjacent area on the ventral occipital–temporal surface known as VOT (Grill-Spector & Malach, 2004). In human functional magnetic resonance imaging (fMRI) studies these areas respond more strongly to images of objects than to images in which shape contours of the objects have been disrupted ("scrambled" images), similar to response patterns observed in primate ventral temporal neurons (Tanaka, Saito, Fukada, & Moriya, 1991). Shape information is generally the most important attribute of concrete object concepts and distinguishes concrete concepts from a range of abstract (e.g., affective, social, and cognitive) entities and event concepts. Variation in the salience of visual shape also occurs within the domain of concrete entities. Shape is typically not a defining feature of substance nouns (*metal, plastic, water*), but it has relevance for some substance-like mass nouns (*butter, coffee, rice*). Concrete objects also vary in shape complexity. Animals, for example, tend to have more complex shapes than plants or tools (Snodgrass & Vanderwart, 1980). We hypothesize that both shape salience and shape complexity determine the degree of activation in shape processing regions during concept retrieval.

More focal areas within or adjacent to these shape processing regions show preferential responses to human faces (Kanwisher, McDermott, & Chun, 1997; Schwarzlose, Baker, & Kanwisher, 2005) and body parts (Downing, Jiang, Shuman, & Kanwisher, 2001). A specialized mechanism for face perception is consistent with the central role of face recognition in social interactions. Visual perception of body parts probably plays a crucial role in understanding observed actions and mapping these onto internal representations of the body schema during action learning. These processors would be differentially activated by concepts pertaining to faces (*nose, mouth, portrait*) and body parts (*hand, leg, finger*). As a first approximation, we propose that both of these neural processors are activated to some degree by all concepts referring to human beings, which by necessity include a face and other human body parts. Activation of the face processor may be minimal in the case of concepts that represent general human types and roles (*boy, man, nurse*, etc.), as these concepts do not include information about any specific face, whereas concepts referencing specific individuals (*brother, father, Einstein*) might include specific facial feature information and therefore activate the face processor to a greater degree.

Perception of three-dimensional spaces involves a specialized neural processor in the posterior

parahippocampal region (Epstein & Kanwisher, 1998; Epstein, Parker, & Feiler, 2007). Although this processor is typically studied using visual stimuli and is often considered a high-level visual area, it more likely integrates visual, proprioceptive, and perhaps auditory inputs, all of which provide correlated information about three-dimensional space. Further discussion of this component of experience is provided in the section below on Spatial Components (Table 2).

Somatosensory components

Somatosensory networks of the cortex represent body location (somatotopic representation), body position and joint force (proprioception), pain, surface characteristics of objects (texture and temperature), and object shape (Friebel, Eickhoff, & Lotze, 2011; Hendry, Hsiao, & Bushnell, 1999; Lamm, Decety, & Singer, 2011; Longo, Azanon, & Haggard, 2010; McGlone & Reilly, 2010; Medina & Coslett, 2010; Serino & Haggard, 2010). Somatotopy is an inherently multi-modal phenomenon, because the body location touched by a particular object is typically the same body location as that used during exploratory or manipulative motor actions involving the object. Conceptual and linguistic representations are more likely to encode the action as the salient characteristic of the object rather than the tactile experience (we say "I threw the ball" to describe the event, not "the ball touched my hand"), therefore somatotopic knowledge was encoded in the proposed representation using action-based rather than somatosensory attributes (see Motor Components).

Temperature, tactile texture, and pain were included as components of the representation, as was a component referring to the salience of object weight. Weight is perceived mainly via proprioceptive representations and is a salient attribute of objects with which we interact. Temperature and weight are two more examples of scalar entities that may or may not have a macroscopic topography in the brain (Hendry et al., 1999; Mazzola, Faillenot, Barral, Mauguiere, & Peyron, 2012). That is, although there is evidence that these types of somatosensory information are distinguishable from each other, no evidence has yet been presented for a scalar topography that separates, for example, hot from cold representations. Tactile texture can refer to a variety of physical properties; we operationalized this

concept as a continuum from smooth to rough, and thus the texture component is another scalar entity. For all three of these scalars, we elected to represent the entire continuum as a single component. That is, the temperature component is intended to capture the salience of temperature to a concept, regardless of whether the associated temperature is hot or cold.

The published literature on conceptual processing of somatosensory information has focused almost exclusively on impaired knowledge of body parts in neurological patients (Coslett, Saffran, & Schwoebel, 2002; Kemmerer & Tranel, 2008; Laiacona, Allamano, Lorenzi, & Capitani, 2006; Schwoebel & Coslett, 2005), though no definitive localization has emerged from these studies. A number of functional imaging studies suggest an overlap between networks involved in real pain perception and those involved in empathic pain perception (perception of pain in others) and perception of "social pain" (Eisenberger, 2012; Lamm et al., 2011), suggesting that retrieval of pain concepts involves a simulation process. To our knowledge, no studies have yet examined the conceptual representation of temperature, tactile texture, or weight.

Finally, object shape is an attribute learned partly through haptic experiences (Miquée et al., 2008), but the degree to which haptic shape is learned independently from visual shape is unclear. Some objects are seen but not felt (i.e., very large and very distant objects), but the converse is rarely true, at least in sighted individuals. It was therefore decided that the visual shape query would capture knowledge about shape in both modalities, and that a separate query regarding extent of personal manipulation experience (see Motor Components) would capture the importance of haptic shape relative to visual shape.

Auditory components

The auditory cortex includes low-level areas tuned to frequency (tonotopy), amplitude, and spatial location (Schreiner & Winer, 2007), as well as higher levels showing specialization for auditory "objects" such as environmental sounds (Leaver & Rauschecker, 2010) and speech sounds (Belin, Zatorre, Lafaille, Ahad, & Pike, 2000; Liebenthal, Binder, Spitzer, Possing, & Medler, 2005). Corresponding auditory attributes of the proposed conceptual representation capture the

degree to which a concept is associated with a low-pitched, high-pitched, loud (i.e., high amplitude), recognizable (i.e., having a characteristic auditory form), or speech-like sound. The attribute "low in amplitude" was not included because it was considered to be a salient attribute of relatively few concepts—that is, those that refer specifically to a low-amplitude variant (e.g., *whisper*). Concepts that focus on the absence of sound (e.g., *silence, quiet*) are probably more accurately encoded as a null value on the "loud" attribute, as fMRI studies have shown auditory regions where activity increases with sound intensity, but not the converse (Röhl & Uppenkamp, 2012).

A few studies of sound-related knowledge (Goldberg, Perfetti, & Schneider, 2006; Kellenbach et al., 2001; Kiefer, Sim, Herrnberger, Grothe, & Hoenig, 2008) reported activation in or near the posterior superior temporal sulcus (STS), an auditory association region implicated in environmental sound recognition (Leaver & Rauschecker, 2010; J. W. Lewis et al., 2004). Trumpp et al. (Trumpp, Kliese, Hoenig, Haarmeier, & Kiefer, 2013) described a patient with a circumscribed lesion centred on the left posterior STS who displayed a specific deficit (slower response times and higher error rate) in visual recognition of sound-related words. All of these studies examined general environmental sound knowledge; nothing is yet known regarding the conceptual representation of specific types of auditory attributes like frequency or amplitude.

Localization of sounds in space is an important aspect of auditory perception, yet auditory perception of space has little or no representation in language independent of (multimodal) spatial concepts; there are no concepts with components like "makes sounds from the left", because spatial relationships between sound sources and listeners are almost never fixed or central to meaning. Like shape attributes of concepts, spatial attributes and spatial concepts are learned through strongly correlated, multimodal experiences involving visual, auditory, tactile, and motor processes. In the proposed semantic representation model, spatial attributes of concepts are represented by modality-independent spatial components (see Spatial Components).

A final salient component of human auditory experience is the domain of music. Music perception, which includes processing of melody, harmony, and rhythm elements in sound, appears to engage relatively specialized, right lateralized networks, some of which may also process nonverbal (prosodic) elements in speech (Zatorre, Belin, & Penhune, 2002). Many words refer explicitly to musical phenomena (*sing, song, melody, harmony, rhythm,* etc.) or to entities (instruments, performers, performances) that produce musical sounds. Therefore, the model proposed here includes a component to capture the experience of processing music.

Gustatory and olfactory components

Gustatory and olfactory experiences are each represented by a single attribute that captures the relevance of each type of experience to the concept. These sensory systems do not show clear large-scale organization along any dimension, and neurons within each system often respond to a broad spectrum of items. Both gustatory and olfactory concepts have been linked with specific early sensory and higher associative neural processors (Goldberg et al., 2006; González et al., 2006).

Motor components

The motor components of the proposed conceptual representation (Table 1) capture the degree to which a concept is associated with actions involving particular body parts. The neural representation of motor action verbs has been extensively studied, with many studies showing action-specific activation of a high-level network that includes the supramarginal gyrus and posterior middle temporal gyrus (Binder, Desai, Conant, & Graves, 2009; Desai, Binder, Conant, & Seidenberg, 2009). There is also evidence for somatotopic representation of action verb representation in primary sensorimotor areas (Hauk, Johnsrude, & Pulvermuller, 2004), though this claim is somewhat controversial (Postle, McMahon, Ashton, Meredith, & de Zubicaray, 2008). Object concepts associated with particular actions also activate the action network (Binder et al., 2009; Boronat et al., 2005; Martin et al., 1995), and there is evidence for somatotopically specific activation depending on which body part is typically used in the object interaction (Carota, Moseley, & Pulvermüller, 2012). Following precedent in the neuroimaging literature (Carota et al., 2012; Hauk et al., 2004; Tettamanti et al., 2005), a three-

way partition into head-related (face, mouth, or tongue), upper limb (arm, hand, or fingers), and lower limb (leg or foot) actions was used to assess somatotopic organization of action concepts. A reference to eye actions was felt to be confusing, because all visual experiences involve the eyes; in this sense, any visible object could be construed as "associated with action involving the eyes".

Finally, the extent to which action attributes are represented in action networks may be partly determined by personal experience (Calvo-Merino, Glaser, Grezes, Passingham, & Haggard, 2005; Jäncke, Koeneke, Hoppe, Rominger, & Hänggi, 2009). Most people would rate "piano" as strongly associated with upper limb actions, but most people also have no personal experience with these specific actions. Based on prior research, it is likely that the action representation of the concept "piano" is more extensive in the case of pianists. Therefore, a separate component (called "practice") was created to capture the level of personal experience the rater has with performing the actions associated with the concept.

Spatial components

Spatial and other more abstract components are listed in Table 2. Neural systems that encode aspects of spatial experience have been investigated at many levels (Andersen, Snyder, Bradley, & Xing, 1997; Burgess, 2008; Kemmerer, 2006; Kranjec, Cardillo, Schmidt, Lehet, & Chatterjee, 2012; Woods et al., 2014). As mentioned above, the acquisition of basic spatial concepts and knowledge of spatial attributes of concepts generally involves correlated, multimodal inputs. The most obvious spatial content in language is the set of relations expressed by prepositional phrases, which have been a major focus of study in linguistics and semantic theory (Landau & Jackendoff, 1993; Talmy, 1983). Lesion correlation and neuroimaging studies have implicated various parietal regions, particularly the left inferior parietal lobe, in the processing of spatial prepositions and depicted locative relations (Amorapanth, Widick, & Chatterjee, 2010; Noordzij, Neggers, Ramsey, & Postma, 2008; Tranel & Kemmerer, 2004; Wu, Waller, & Chatterjee, 2007). Prepositional phrases appear to express most, if not all, of the explicit relational spatial content that occurs in English; a semantic component targeting this type of content therefore appears unnecessary (i.e., the set of words with high values on this component would be identical to the set of spatial prepositions). The spatial components of the proposed representation model focus instead on other types of spatial information found in nouns, verbs, and adjectives.

Perception of three-dimensional spaces involves a specialized neural processor in the posterior parahippocampal region (Epstein & Kanwisher, 1998; Epstein et al., 2007). The perception of three-dimensional space is based largely on visual input, but also involves correlated proprioceptive information whenever three-dimensional space is experienced via movement of the body through space. Spatial localization in the auditory modality also probably contributes to the experience of three-dimensional space. Some concepts refer directly to interior spaces, such as *kitchen, foyer, bathroom, classroom*, and so on, and seem to include information about general three-dimensional size and layout, as do concepts about types of buildings (*library, school, hospital*, etc.). More generally, such concepts include the knowledge that they can be physically entered, navigated, and exited, which makes possible a range of conceptual combinations (*entered the kitchen, ran through the hall, went out of the library*, etc.) that are not possible for objects lacking large interior spaces (*entered the chair). The salient property of concepts that determines whether they activate a three-dimensional representation of space is the degree to which they evoke a particular "scene"—that is, an array of objects in a three-dimensional space, either by direct reference or by association. Space names (*kitchen*, etc.) and building names (*library*, etc.) refer directly to three-dimensional scenes. Many other object concepts evoke mental representations of scenes by virtue of thematic relations, such as the evocation of *kitchen* by *oven*. An object concept like *chair* would probably fail to evoke such a representation, as chairs are not linked thematically with any particular scene. Thus there is a graded degree of mental representation of three-dimensional space evoked by different concepts, which, we hypothesize, results in a corresponding variation in activation of the three-dimensional space neural processor.

Large-scale (topographical) spatial information is critical for navigation in the environment. Entities that are large and have a fixed location, such as geological formations and buildings, can serve as landmarks within a mental map of the environment. We

hypothesize that such concepts are partly encoded in environmental navigation systems located in the hippocampal, parahippocampal, and posterior cingulate gyri (Burgess, 2008; Maguire, Frith, Burgess, Donnett, & O'Keefe, 1998; Wallentin, Østergaarda, Lund, Østergaard, & Roepstorff, 2005), and we assess the salience of this "topographical" attribute by inquiring to what degree the concept in question could serve as a landmark on a map.

Spatial cognition also includes spatial aspects of motion, particularly direction or path of motion through space (location change). Some verbs that denote location change refer directly to a direction of motion (*ascend, climb, descend, rise, fall*) or type of path (*arc, bounce, circle, jump, launch, orbit, swerve*). Others imply a horizontal path of motion parallel to the ground (*run, walk, crawl, swim*), while others refer to paths toward or away from an observer (*arrive, depart, leave, return*). Some entities are associated with a particular type of path through space (*bird, butterfly, car, rocket, satellite, snake*). Visual motion area MT features a columnar spatial organization of neurons according to preferred axis of motion; however, this organization is at a relatively microscopic scale, such that the full 180° of axis of motion is represented in 500 μm of cortex (Albright et al., 1984). In human fMRI studies, perception of the spatial aspects of motion has been linked with a neural processor in the intraparietal sulcus (IPS) and surrounding dorsal visual stream (Culham et al., 1998; Wu, Morganti, & Chatterjee, 2008).

A final aspect of spatial cognition relates to peripersonal proximity. Intuitively, the coding of peripersonal proximity is a central aspect of action representation and performance, threat detection, and resource acquisition, all of which are neural processes critical for survival. Evidence from both animal and human studies suggests that the representation of peripersonal space involves somewhat specialized neural mechanisms (Graziano, Yap, & Gross, 1994; Makin, Holmes, & Zohary, 2007). We hypothesize that these systems also partly encode the meaning of concepts that relate to peripersonal space—that is, objects that are typically within easy reach and actions performed on these objects. Given the importance of peripersonal space in action and object use, two further components were included to explore the relevance of movement away from or out of versus toward or into the self. Some action verbs (*give, punch, tell, throw*) indicate movement or relocation of something away from the self, whereas others (*acquire, catch, receive, eat*) indicate movement or relocation toward the self. This distinction may also modulate the representation of objects that characteristically move toward or away from the self (Rueschemeyer, Pfeiffer, & Bekkering, 2010).

Mathematical cognition, particularly the representation of numerosity, is a somewhat specialized information processing domain with well-documented neural correlates (Dehaene, 1997; Harvey, Klein, Petridou, & Dumoulin, 2013; Piazza, Izard, Pinel, Bihan, & Dehaene, 2004). In addition to number names, which refer directly to numerical quantities, many words are associated with the concept of numerical quantity (*amount, number, quantity*) or with particular numbers (*dime, hand, egg*). Although not a spatial process per se, since it does not typically involve three dimensions, numerosity processing seems to depend on directional vector representations similar to those underlying spatial cognition. Thus, a component was included to capture association with numerical quantities and was considered loosely part of the spatial domain.

Temporal and causal components

Event concepts include several distinct types of information. Temporal information pertains to the duration and temporal order of events. Some types of event have a typical duration, in which case the duration may be a salient attribute of the event (*breakfast, lecture, movie, shower*). Some event-like concepts refer specifically to durations (*minute, hour, day, week*). Typical durations may be very short (*blink, flash, pop, sneeze*) or relatively long (*childhood, college, life, war*). Events of a particular type may also recur at particular times and thus occupy a characteristic location within the temporal sequence of events. Examples of this phenomenon include recurring daily events (*shower, breakfast, commute, lunch, dinner*), recurring weekly or monthly events (*Monday, weekend, payday*), recurring annual events (holidays and other cultural events, seasons and weather phenomena), and concepts associated with particular phases of life (*infancy, childhood, college, marriage, retirement*). The neural correlates of these types of information are not yet clear (Ferstl, Rinck, & von Cramon, 2005; Ferstl & von Cramon, 2007; Grondin,

2010; Kranjec et al., 2012). Because time seems to be pervasively conceived of in spatial terms (Evans, 2004; Kranjec & Chatterjee, 2010; Lakoff & Johnson, 1980), it is likely that temporal concepts share their neural representation, at least in part, with spatial concepts. Components of the proposed conceptual representation model are designed to capture the degree to which a concept is associated with a characteristic duration, the degree to which a concept is associated with a long or a short duration, and the degree to which a concept is associated with a particular or predictable point in time.

Causal information pertains to cause and effect relationships between concepts. Events and situations may have a salient precipitating cause (*infection, kill, laugh, spill*), or they may occur without an immediate apparent cause (e.g., some natural phenomena, random occurrences). Events may or may not have highly likely consequences. Imaging evidence suggests involvement of several inferior parietal and temporal regions in low-level perception of causality (Blos, Chatterjee, Kircher, & Straube, 2012; Fonlupt, 2003; Fugelsang, Roser, Corballis, Gazzaniga, & Dunbar, 2005; Woods et al., 2014), and a few studies of causal processing at the conceptual level implicate these and other areas (Kranjec et al., 2012; Satpute et al., 2005). The proposed conceptual representation model includes components designed to capture the degree to which a concept is associated with a clear preceding cause, and the degree to which an event or action has probable consequences.

Social components

Theory of mind (TOM) refers to the ability to infer or predict mental states and mental processes in other volitional beings (typically though not exclusively other people). We hypothesize that the brain systems underlying this ability are involved when encoding the meaning of words that refer to intentional entities or actions because learning these concepts is frequently associated with TOM processing. Essentially, this attribute captures the semantic feature of intentionality. In addition to the query pertaining to events with social interactions as well as the query regarding mental activities or events, we included an object-directed query assessing the degree to which a thing has human or human-like intentions, plans, or goals, and a corresponding verb

query assessing the degree to which an action reflects human intentions, plans, or goals. The underlying hypothesis is that such concepts, which mostly refer to people in particular roles and the actions they take, are partly understood by engaging a TOM process.

There is strong evidence for the involvement of a specific network of regions in TOM. These regions most consistently include the medial prefrontal cortex, posterior cingulate/precuneus, and the temporoparietal junction (Schilbach et al., 2012; Schurz, Radua, Aichhorn, Richlan, & Perner, 2014; Van Overwalle, 2009; van Veluw & Chance, 2014) with several studies also implicating the anterior temporal lobes (Ross & Olson, 2010; Schilbach et al., 2012; Schurz et al., 2014). The temporoparietal junction (TPJ) is an area of particular focus in the TOM literature, but it is not a consistently defined region and may encompass parts of the posterior superior temporal gyrus/sulcus, supramarginal gyrus, and angular gyrus. Collapsing across tasks, TOM or intentionality is frequently associated with significant activation in the angular gyri (Carter & Huettel, 2013; Schurz et al., 2014), but some variability in the localization of peak activation within the TPJ is seen across different TOM tasks (Schurz et al., 2014). In addition, while some studies have shown right lateralization of activation in the TPJ (Saxe & Wexler, 2005), several meta-analyses have suggested bilateral involvement (Schurz et al., 2014; Spreng, Mar, & Kim, 2009; Van Overwalle, 2009; van Veluw & Chance, 2014).

Another aspect of social cognition likely to have a neurobiological correlate is the representation of the self. There is evidence to suggest that information that pertains to the self may be processed differently from information that is not considered to be self-related. Some behavioural studies have suggested that people show better recall (Rogers, Kuiper, & Kirker, 1977) and faster response times (Kuiper & Rogers, 1979) when information is relevant to the self, a phenomenon known as the self-reference effect. Neuroimaging studies have typically implicated cortical midline structures in the processing of self-referential information, particularly the medial prefrontal and parietal cortices (Araujo, Kaplan, & Damasio, 2013; Northoff et al., 2006). While both self- and other-judgments activate these midline regions, greater activation of medial prefrontal cortex for self-trait judgments than for other-trait judgments has

been reported, with the reverse pattern seen in medial parietal regions (Araujo et al., 2013). In addition, there is evidence for a ventral-to-dorsal spatial gradient within medial prefrontal cortex for self- relative to other-judgments (Denny, Kober, Wager, & Ochsner, 2012). Preferential activation of left ventrolateral prefrontal cortex and left insula for self-judgments and of the bilateral temporoparietal junction and cuneus for other-judgments has also been seen (Denny et al., 2012). The relevant query for this attribute assesses the degree to which a target noun is "related to your own view of yourself" or the degree to which a target verb is "an action or activity that is typical of you, something you commonly do and identify with".

A third aspect of social cognition for which a specific neurobiological correlate is likely is the domain of communication. Communication, whether by language or nonverbal means, is the basis for all social interaction, and many noun (e.g., *speech, book, meeting*) and verb (e.g., *speak, read, meet*) concepts are closely associated with the act of communicating. We hypothesize that comprehension of such items is associated with relative activation of language networks (particularly general lexical retrieval, syntactic, phonological, and speech articulation components) and gestural communication systems compared to concepts that do not reference communication or communicative acts.

Cognition component

A central premise of the current approach is that *all* aspects of experience can contribute to concept acquisition and therefore concept composition. For self-aware human beings, purely cognitive events and states certainly comprise one aspect of experience. When we "have a thought", "come up with an idea", "solve a problem", or "consider a fact", we experience these phenomena as events that are no less real than sensory or motor events. Many so-called abstract concepts refer directly to cognitive events and states (*decide, judge, recall, think*) or to the mental "products" of cognition (*idea, memory, opinion, thought*). We propose that such concepts are learned in large part by generalization across these cognitive experiences, in exactly the same way as concrete concepts are learned through generalization across perceptual and motor experiences.

Domain-general cognitive operations have been the subject of a very large number of neuroimaging studies, many of which have attempted to fractionate these operations into categories such as "working memory", "decision", "selection", "reasoning", "conflict suppression", "set maintenance", and so on. No clear anatomical divisions have arisen from this work, however, and the consensus view is that there is a largely shared bilateral network for what are variously called "working memory", "cognitive control", or "executive" processes, involving portions of the inferior and middle frontal gyri ("dorsolateral prefrontal cortex"), mid-anterior cingulate gyrus, precentral sulcus, anterior insula, and perhaps dorsal regions of the parietal lobe (Badre & D'Esposito, 2007; Derrfuss, Brass, Neumann, & von Cramon, 2005; Duncan & Owen, 2000; Nyberg et al., 2003). We hypothesize that retrieval of concepts that refer to cognitive phenomena involves a partial simulation of these phenomena within this cognitive control network, analogous to the partial activation of sensory and motor systems by object and action concepts.

Emotion and evaluation components

There is strong evidence for the involvement of a specific network of regions in affective processing in general. These regions principally include the orbitofrontal cortex, dorsomedial prefrontal cortex, anterior cingulate gyri (subgenual, pregenual, and dorsal), inferior frontal gyri, anterior temporal lobes, amygdala, and anterior insula (Etkin, Egner, & Kalisch, 2011; Hamann, 2012; Kober et al., 2008; Lindquist, Wager, Kober, Bliss-Moreau, & Barrett, 2012; Phan, Wager, Taylor, & Liberzon, 2002; Vytal & Hamann, 2010). Activation in these regions can be modulated by the emotional content of words and text (Chow et al., 2013; Cunningham, Raye, & Johnson, 2004; Ferstl et al., 2005; Ferstl & von Cramon, 2007; Kuchinke et al., 2005; Nakic, Smith, Busis, Vythilingam, & Blair, 2006). However, the nature of individual affective states has long been the subject of debate. One of the primary families of theories regarding emotion are the dimensional accounts, which propose that affective states arise from combinations of a small number of more fundamental dimensions, most commonly valence (hedonic tone) and arousal (Russell, 1980). In current psychological construction accounts, this input is then interpreted as a particular emotion

using one or more additional cognitive processes, such as appraisal of the stimulus or the retrieval of memories of previous experiences or semantic knowledge (Barrett, 2006; Lindquist, Siegel, Quigley, & Barrett, 2013; Posner, Russell, & Peterson, 2005). Another highly influential family of theories characterizes affective states as discrete emotions, each of which have consistent and discriminable patterns of physiological responses, facial expressions, and neural correlates. Basic emotions are a subset of discrete emotions that are considered to be biologically predetermined and culturally universal (Hamann, 2012; Lench, Flores, & Bench, 2011; Vytal & Hamann, 2010), although they may still be somewhat influenced by experience (Tracy & Randles, 2011). The most frequently proposed set of basic emotions are those of Ekman (Ekman, 1992), which include anger, fear, disgust, sadness, happiness, and surprise.

There is substantial neuroimaging support for dissociable dimensions of valence and arousal within individual studies; however, the specific regions associated with the two dimensions or their endpoints have been inconsistent and sometimes contradictory across studies (Anders, Eippert, Weiskopf, & Veit, 2008; Anders, Lotze, Erb, Grodd, & Birbaumer, 2004; Colibazzi et al., 2010; Gerber et al., 2008; P. A. Lewis, Critchley, Rotshtein, & Dolan, 2007; Nielen et al., 2009; Posner et al., 2009; Wilson-Mendenhall, Barrett, & Barsalou, 2013). With regard to the discrete emotion model, meta-analyses have suggested differential activation patterns in individual regions associated with basic emotions (Lindquist et al., 2012; Phan et al., 2002; Vytal & Hamann, 2010), but there is a lack of specificity. Individual regions are generally associated with more than one emotion, and emotions are typically associated with more than one region (Lindquist et al., 2012; Vytal & Hamann, 2010). Overall, current evidence is not consistent with a one-to-one mapping between individual brain regions and either specific emotions or dimensions; however, the possibility of spatially overlapping but discriminable networks remains open. Importantly, studies of emotion perception or experience using multivariate pattern classifiers are providing emerging evidence for distinct patterns of neural activity associated with both discrete emotions (Kassam, Markey, Cherkassky, Loewenstein, & Just, 2013; Kragel & LaBar, 2014; Peelen, Atkinson, & Vuilleumier, 2010; Said, Moore, Engell, & Haxby, 2010) and specific combinations of valence and arousal (Baucom, Wedell, Wang, Blitzer, & Shinkareva, 2012).

The semantic representation model proposed here includes separate components reflecting the degree of association of a target concept with each of Ekman's basic emotions. The representation also includes components corresponding to the two dimensions of the circumplex model (valence and arousal). There is evidence that each end of the valence continuum may have a distinctive neural instantiation (Anders et al., 2008; Anders et al., 2004; Colibazzi et al., 2010; Posner et al., 2009), and therefore separate components were included for pleasantness and unpleasantness.

Drive components

Drives are central associations of many concepts and are conceptualized here, following Mazlow (Mazlow, 1943), as needs that must be filled to reach homeostasis or enable growth. They include physiological needs (food, rest/sleep, sex, emotional expression), security, social contact, approval, and self-development. Drive experiences are closely related to emotions, which are produced whenever needs are fulfilled or frustrated, and to the construct of reward, conceptualized as the brain response to a resolved or satisfied drive. Here we use the term "drive" synonymously with terms such as "reward prediction" and "reward anticipation". Extensive evidence links the ventral striatum, orbital frontal cortex, and ventromedial prefrontal cortex to processing of drive and reward states (Diekhof, Kaps, Falkai, & Gruber, 2012; Levy & Glimcher, 2012; Peters & Büchel, 2010; Rolls & Grabenhorst, 2008). The proposed semantic components represent the degree of general motivation associated with the target concept and the degree to which the concept itself refers to a basic need.

Attention components

Attention is a basic process central to information processing, but is not usually thought of as a type of information. It is possible, however, that some concepts (e.g., "scream") are so strongly associated with attentional orienting that the attention-capturing event becomes part of the conceptual representation. A component was included to assess the degree to which a concept is "someone or something that grabs your attention".

Attribute salience ratings for 535 English words

Ratings of the relevance of each semantic component to word meanings were collected for 535 English words using the internet crowdsourcing survey tool Amazon Mechanical Turk (https://www.mturk.com/mturk/). The aim of this work was to ascertain the degree to which a relatively unselected sample of the population associates the meaning of a particular word with particular kinds of experience. We refer to the continuous ratings obtained for each word as the *attributes* comprising the word's meaning. The results reflect subjective judgments rather than objective truth and are likely to vary somewhat with personal experience and background, presence of idiosyncratic associations, participant motivation, and comprehension of the instructions. With these caveats in mind, it was hoped that mean ratings obtained from a sufficiently large sample would accurately capture central tendencies in the population, providing representative measures of real-world comprehension.

As discussed in the previous section, the attributes selected for study reflect known neural systems. We hypothesize that all concept learning depends on this set of neural systems, and therefore the same attributes were rated regardless of grammatical class or ontological category.

The word set

The concepts included 434 nouns, 62 verbs, and 39 adjectives. All of the verbs and adjectives, and 141 of the nouns, were pre-selected as experimental materials for the Knowledge Representation in Neural Systems (KRNS) project (Glasgow et al., 2016), an ongoing multi-centre research program sponsored by the US Intelligence Advanced Research Projects Activity (IARPA). Because the KRNS set was limited to relatively concrete concepts and a restricted set of noun categories, 293 additional nouns were added to include more abstract concepts and to enable richer comparisons between category types. Table 3 summarizes some general characteristics of the word set. Table 4 describes the content of the set in terms of major category types. A complete list of the words and associated lexical data are available for download (see link prior to the References section).

Table 3. Characteristics of the 535 rated words.

Characteristic	Mean	SD	Min	Max	Median	IQR
Length in letters	5.95	1.95	3	11	6	3
Log(10) orthographic frequency	1.08	0.80	−1.61	3.05	1.17	1.12
Orthographic neighbours	4.19	5.33	0	22	2	6
Imageability (1–7 scale)	5.17	1.16	1.40	6.90	5.58	1.96

Note: Orthographic frequency values are from the Westbury Lab UseNet Corpus (Shaoul & Westbury, 2013). Orthographic neighbour counts are from CELEX (Baayen, Piepenbrock, & Gulikers, 1995). Imageability ratings are from a composite of published norms (Bird, Franklin, & Howard, 2001; Clark & Paivio, 2004; Cortese & Fugett, 2004; Wilson, 1988). IQR = interquartile range.

Query design

Queries used to elicit ratings were designed to be as structurally uniform as possible across attributes and grammatical classes. Some flexibility was allowed, however, to create natural and easily understood queries. All queries took the general form of {head}{relation}{content}, where {head} refers to a lead-in phrase that was uniform within each word class, {content} to the specific content being assessed, and {relation} to the type of relation linking the target word and the content. The {head} for all queries regarding nouns was: "To what degree do you think of this thing as . . . ", whereas the {head} for all verb queries was: "To what degree do you think of this as . . . ", and the {head} for all adjective queries was: "To what degree do you think of this property as . . . ". Table 5 lists several examples for each grammatical class.

For the three scalar somatosensory attributes Temperature, Texture, and Weight, it was felt necessary for the sake of clarity to pose separate queries for each end of the scalar continuum. That is, rather than a single query such as "To what degree do you think of this thing as being hot or cold to the touch?", two separate queries were posed about hot and cold attributes. The Temperature rating was then computed by taking the maximum rating for the word on either the Hot or the Cold query. Similarly, the Texture rating was computed as the maximum rating on Rough and Smooth queries, and the Weight rating was computed as the maximum rating on Heavy and Light queries.

A few attributes did not appear to have a meaningful sense when applied to certain grammatical classes. The attribute Complexity addresses the degree to which an entity appears visually complex and has no obvious sensible correlate in

Table 4. Cell counts for grammatical classes and categories included in the ratings study.

Type	N	Examples
Nouns (434)		
Concrete objects (275)		
Living things (126)		
Animals	30	crow, horse, moose, snake, turtle, whale
Body parts	14	finger, hand, mouth, muscle, nose, shoulder
Humans	42	artist, child, doctor, mayor, parent, soldier
Human groups	10	army, company, council, family, jury, team
Plants	30	carrot, eggplant, elm, rose, tomato, tulip
Other natural objects (19)		
Natural scenes	12	beach, forest, lake, prairie, river, volcano
Miscellaneous	7	cloud, egg, feather, stone, sun
Artifacts (130)		
Furniture	7	bed, cabinet, chair, crib, desk, shelves, table
Hand tools	27	axe, comb, glass, keyboard, pencil, scissors
Manufactured foods	18	beer, cheese, honey, pie, spaghetti, tea
Musical instruments	20	banjo, drum, flute, mandolin, piano, tuba
Places/buildings	28	bridge, church, highway, prison, school, zoo
Vehicles	20	bicycle, car, elevator, plane, subway, truck
Miscellaneous	10	book, computer, door, fence, ticket, window
Concrete events (60)		
Social events	35	barbecue, debate, funeral, party, rally, trial
Nonverbal sound events	11	belch, clang, explosion, gasp, scream, whine
Weather events	10	cyclone, flood, landslide, storm, tornado
Miscellaneous	4	accident, fireworks, ricochet, stampede
Abstract entities (99)		
Abstract constructs	40	analogy, fate, law, majority, problem, theory
Cognitive entities	10	belief, hope, knowledge, motive, optimism, wit
Emotions	20	animosity, dread, envy, grief, love, shame
Social constructs	19	advice, deceit, etiquette, mercy, rumour, truce
Time periods	10	day, era, evening, morning, summer, year
Verbs (62)		
Concrete actions (52)		
Body actions	10	ate, held, kicked, laughed, slept, threw
Locative change actions	14	approached, crossed, flew, left, ran, put, went
Social actions	16	bought, helped, met, spoke to, stole, visited
Miscellaneous	12	broke, built, damaged, fixed, found, watched
Abstract actions	5	ended, lost, planned, read, used
States	5	feared, liked, lived, saw, wanted
Adjectives (39)		
Abstract properties	13	angry, clever, famous, friendly, lonely, tired
Physical properties	26	blue, dark, empty, heavy, loud, shiny

Table 5. Example queries for six attributes.

Attribute	Query	{relation}	{content}
Colour			
	Noun	. . . having	a characteristic or defining color?
	Verb	. . . being associated with	color or change in color?
	Adjective	. . . describing	a quality or type of color?
Taste			
	Noun	. . . having	a characteristic or defining taste?
	Verb	. . . being associated with	tasting something?
	Adjective	. . . describing	how something tastes?
Lower limb			
	Noun	. . . being associated with	actions using the leg or foot?
	Verb	. . . being	an action or activity in which you use the leg or foot?
	Adjective	. . . being related to	actions of the leg or foot?
Landmark			
	Noun	. . . having	a fixed location, as on a map?
	Verb	. . . being	an action or activity in which you use a mental map of your environment?
	Adjective	. . . describing	the location of something, as on a map?
Human			
	Noun	. . . having	human or human-like intentions, plans, or goals?
	Verb	. . . being	an action or activity that involves human intentions, plans, or goals?
	Adjective	. . . being related to	something with human or human-like intentions, plans, or goals?
Fearful			
	Noun	. . . being	someone or something that makes you feel afraid?
	Verb	. . . being associated with	feeling afraid?
	Adjective	. . . being related to	feeling afraid?

the verb class. The attribute Practice addresses the degree to which the rater has personal experience manipulating an entity or performing an action and has no obvious sensible correlate in the adjective class. Finally, the attribute Caused addresses the degree to which an entity is the result of some prior event or an action will cause a change to occur and has no obvious correlate in the adjective class. Ratings on these three attributes were not obtained for the grammatical classes to which they did not meaningfully apply.

Survey methods

Surveys were posted on Amazon Mechanical Turk (AMT), an internet site that serves as an interface between "requestors" who post tasks and "workers" who have accounts on the site and sign up to complete tasks and receive payment. Requestors have the right to accept or reject worker responses based on quality criteria. Participants in the present study were required to have completed at least 10,000 previous jobs with at least a 99% acceptance rate, and to have account addresses in the United States; these criteria were applied automatically by AMT. Prospective participants were asked not to participate unless they were native English speakers; however, compliance with this request cannot be verified. After reading a page of instructions, participants provided their age, gender, years of education, and current occupation. No identifying information was collected from participants or requested from AMT.

For each session, a participant was assigned a single word and provided ratings on all of the semantic components as they relate to the assigned word. A sentence containing the word was provided to specify the word class and sense targeted. Two such sentences, conveying the most frequent sense of the word, were constructed for each word, and one of these was selected randomly and used throughout the session. Participants were paid a small stipend for completing all queries for the target word. They could return for as many sessions as they wanted. An automated script assigned target words randomly, with the constraint that the same participant could not be assigned the same word more than once.

For each attribute, a screen was presented with the target word, the sentence context, the query for that attribute, an example of a word that "would receive a high rating" on the attribute, an example of a word that "might receive a medium rating" on the attribute, and the rating options. The options were presented as a horizontal row of clickable radio buttons with numerical labels ranging from 0 to 6 (left to right), and verbal labels (Not At All, Somewhat, Very Much) at appropriate locations above the buttons. An example trial is shown in Figure S1 of the Supplemental Material. An additional button labelled "Not Applicable" was positioned to the left of the "0" button. Participants were forewarned that some of the queries "will be almost nonsensical, because they refer to aspects of word meaning that are not even applicable to your word. For example, you might be asked, 'To what degree do you think of *shoe* as an event that has a predictable duration, whether short or long?', with *movie* given as an example of a high-rated concept, and *concert* given as a medium-rated concept. Since *shoe* refers to a static thing, not to an event of any kind, you should indicate either 0 or "Not Applicable" for this question". All "Not Applicable" responses were later converted to 0 ratings.

General results

A total of 16,373 sessions were collected from 1743 unique participants (average 9.4 sessions/participant). Of the participants, 43% were female, 55% were male, and 2% did not provide gender information. The average age was 34.0 years ($SD = 10.4$), and average years of education was 14.8 ($SD = 2.1$).

A limitation of unsupervised crowdsourcing is that some participants may not comply with task instructions. Informal inspection of the data revealed examples in which participants appeared to be responding randomly or with the same response on every trial. A simple metric of individual deviation from the group mean response was computed by correlating the vector of ratings obtained from an individual for a particular word with the group mean vector for that word. For example, if 30 participants were assigned word X, this yielded 30 vectors, each with 65 elements. The average of these vectors is the group average for word X. The quality metric for each participant who rated word X was the correlation between that participant's vector and the group average vector for word X. Supplemental Figure S2 shows the distribution of these values across the

sample after Fisher *z* transformation. A minimum threshold for acceptance was set at $r = .5$, which corresponds to the edge of a prominent lower tail in the distribution. This criterion resulted in rejection of 1097 or 6.69% of the sessions. After removing these sessions, the final data set consisted of 15,268 rating vectors with a mean intra-word individual-to-group correlation of .785 (median .803), providing an average of 28.6 rating vectors (i.e., sessions) per word (*SD* 2.0, range 17–33). Mean ratings for each of the 535 words, computed using only the sessions that passed the acceptance criterion, are available for download (see link prior to the References section).

Exploring the representations

Here we explore the dataset by addressing three general questions. First, how much independent information is provided by each of the attributes, and how are they related to each other? Although the component attributes were selected to represent distinct types of experiential information, there is likely to be conceptual overlap between attributes, real-world covariance between attributes, and covariance due to associations linking one attribute with another in the minds of the raters. We therefore sought to characterize the underlying covariance structure of the attributes. Second, do the attributes capture important distinctions between a priori ontological types and categories? If the structure of conceptual knowledge really is based on underlying neural processing systems, then the covariance structure of attributes representing these systems should reflect psychologically salient conceptual distinctions. Finally, what conceptual groupings are present in the covariance structure itself, and how do these differ from a priori category assignments?

Attribute mutual information

To investigate redundancy in the information encoded in the representational space, we computed the mutual information (MI) for all pairs of attributes using the individual participant ratings after removal of outliers. MI is a general measure of how mutually dependent two variables are, determined by the similarity between their joint distribution p(X,Y) and factored marginal distribution p(X)p(Y). MI can range from 0, indicating complete independence, to 1 in the case of identical variables.

MI was generally very low (mean = .049), suggesting a low overall level of redundancy (see Supplemental Figure S3). Particular pairs and groups, however, showed higher redundancy. The largest MI value was for the pair Pleasant–Happy (.643). It is likely that these two constructs evoked very similar concepts in the minds of the raters. Other isolated pairs with relatively high MI values included Small–Weight (.344), Caused–Consequential (.312), Practice–Near (.269), Taste–Smell (.264), and Social–Communication (.242).

Groups of attributes with relatively high pairwise MI values included a human-related group (Face, Speech, Human, Body, Biomotion; mean pairwise MI = .320), an attention-arousal group (Attention, Arousal, Surprised; mean pairwise MI = .304), an auditory group (Audition, Loud, Low, High, Sound, Music; mean pairwise MI = .296), a visual–somatosensory group (Vision, Colour, Pattern, Shape, Texture, Weight, Touch; mean pairwise MI = .279), a negative affect group (Angry, Sad, Disgusted, Unpleasant, Fearful, Harm, Pain; mean pairwise MI = .263), a motion-related group (Motion, Fast, Slow, Biomotion; mean pairwise MI = .240), and a time-related group (Time, Duration, Short, Long; mean pairwise MI = .193).

Attribute covariance structure

Factor analysis was conducted to investigate potential latent dimensions underlying the attributes. Principal axis factoring (PAF) was used with subsequent promax rotation given that the factors were assumed to be correlated. Mean attribute vectors for the 496 nouns and verbs were used as input; adjective data were excluded so that the Practice and Caused attributes could be included in the analysis. To determine the number of factors to retain, a parallel analysis (eigenvalue Monte Carlo simulation) was performed using PAF and 1000 random permutations of the raw data from the current study (Horn, 1965; O'Connor, 2000). Factors with eigenvalues exceeding the 95th percentile of the distribution of eigenvalues derived from the random data were retained and examined for interpretability.

The Kaiser–Meyer–Olkin Measure of Sampling Accuracy was .874, and Bartlett's Test of Sphericity was significant, $\chi^2(2016) = 41{,}129.17$, $p < .001$,

indicating adequate factorability. Sixteen factors had eigenvalues that were significantly above random chance. These factors accounted for 81.0% of the variance. Based on interpretability, all 16 factors were retained. Table 6 lists these factors and the attributes with the highest loadings on each.

As can be seen from the loadings, the latent dimensions revealed by PAF mostly replicate the groupings apparent in the mutual information analysis. The most prominent of these include factors representing visual and tactile attributes, negative emotion associations, social communication, auditory attributes, food ingestion, self-related attributes, motion in space, human physical attributes, surprise, characteristics of large places, upper limb actions, reward experiences, positive associations, and time-related attributes.

Together, the mutual information and factor analyses reveal a modest degree of covariance between the attributes, suggesting that a more compact representation could be achieved by reducing redundancy. A reasonable first step would be to remove either Pleasant or Happy from the representation, as these two attributes showed a high level of redundancy. For some analyses, use of the 16 significant factors identified in the PAF rather than the complete set of attributes may be appropriate and useful. On the other hand, these factors left ~20% of the variance unexplained, which we propose is a reflection of the underlying complexity of conceptual knowledge. This complexity places limits on the extent to which

the representation can be reduced without losing explanatory power. In the following analyses we have included all 65 attributes in order to investigate further their potential contributions to understanding conceptual structure.

Attribute composition of some major semantic categories

Mean attribute vectors representing major semantic categories in the word set were computed by averaging over items within several of the a priori categories outlined in Table 4. Figures 1 and 2 show mean vectors for 12 of the 20 noun categories in Table 4, presented as circular plots.

Vectors for three verb categories are shown in Figure 3. With the exception of the Path and Social attributes, which marked the Locative Action and Social Action categories, respectively, the three verb categories evoked a similar set of attributes but in different relative proportions. Ratings for physical adjectives reflected the sensory domain (visual, somatosensory, auditory, etc.) to which the adjective referred.

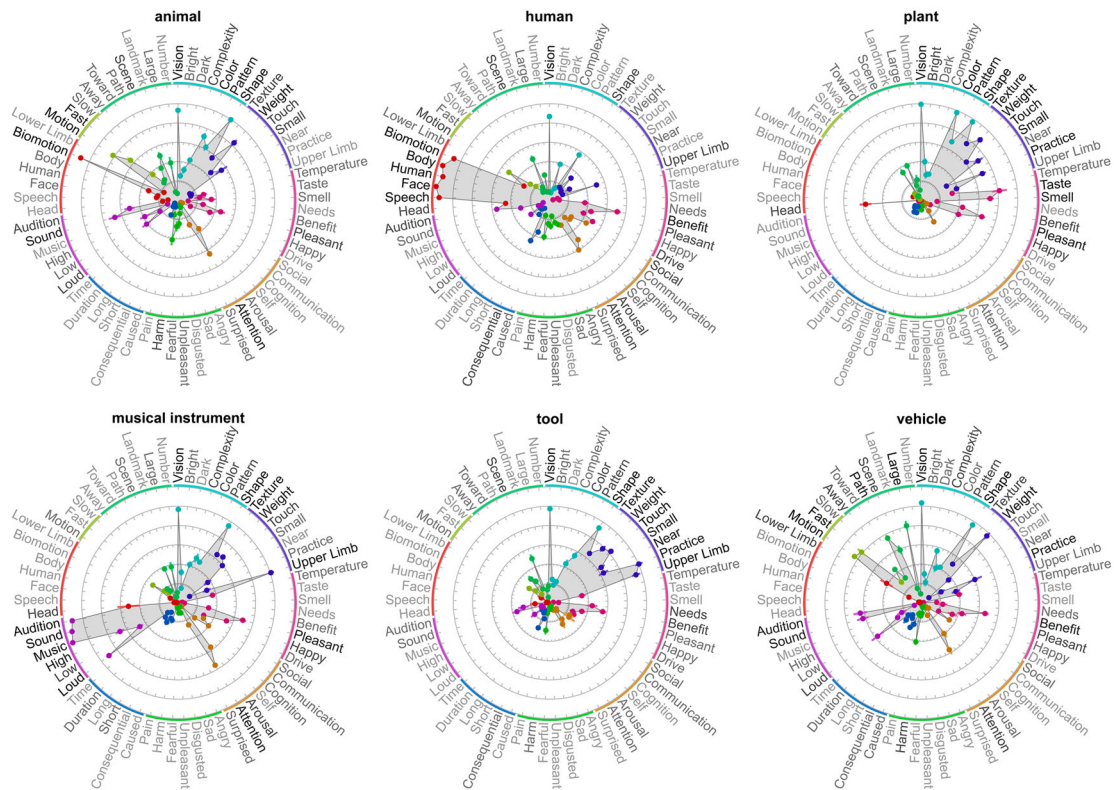## Quantitative differences between a priori categories

The a priori category labels shown in Table 4 were used to identify attribute ratings that differ significantly between semantic categories and therefore may contribute to a mental representation of these category distinctions. For each comparison, an unpaired *t*-test was conducted for each of the 65 attributes, comparing the mean ratings on words in one category with those in the other. It should be kept in mind that the results are specific to the sets of words that happened to be selected for the study and may not generalize to other samples from these categories. For that reason, and because of the large number of contrasts performed, only differences significant at $p < .0001$ will be described.

Initial comparisons were conducted between the superordinate categories Artefacts ($n = 132$), Living Things ($N = 128$), and Abstract Entities ($N = 99$). As shown in Figure 4, numerous attributes distinguished these categories. Not surprisingly, Abstract Entities had significantly lower ratings than the two concrete categories on nearly all of the attributes related to sensory experience. The main exceptions to this were
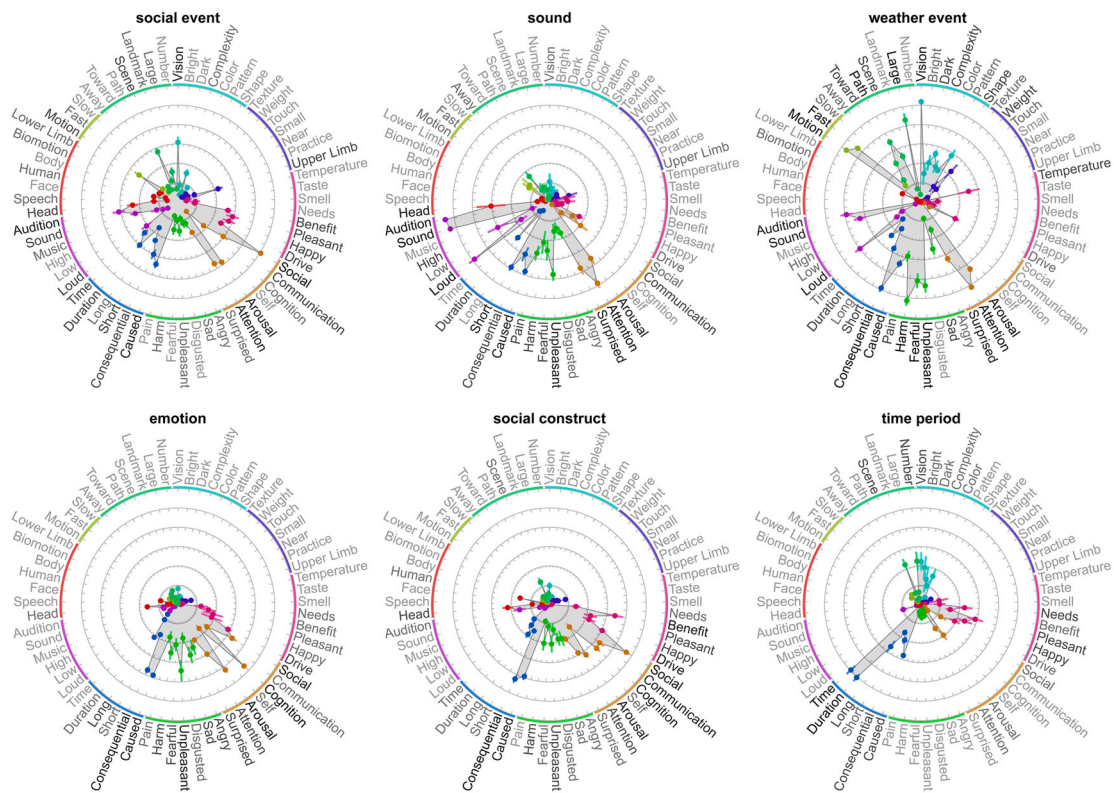
Table 6. Summary of the attribute factor analysis.

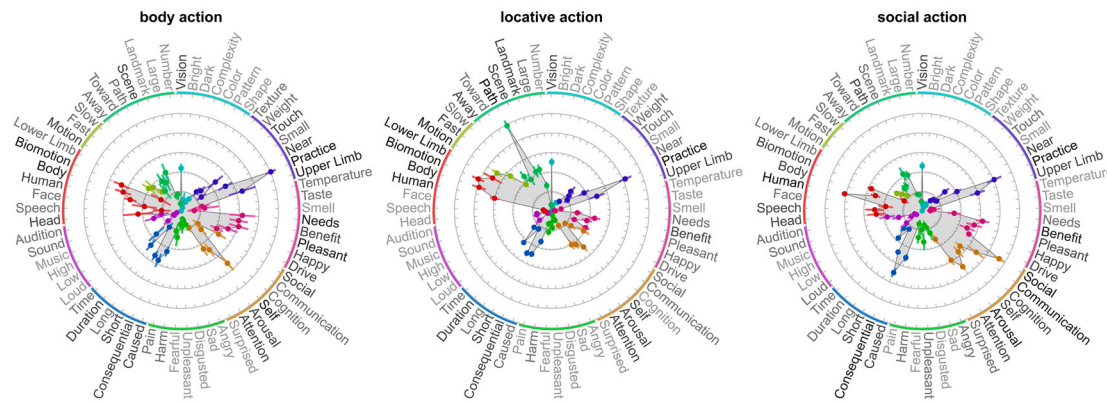| F#	EV	Interpretation	Highest-Loading Attributes (all > .35)
1	12.81	Vision/Touch	Pattern, Shape, Texture, Colour, Weight, Small, Vision, Dark, Bright
2	10.71	Negative	Disgusted, Unpleasant, Sad, Angry, Pain, Harm, Fearful, Consequential
3	6.93	Communication	Communication, Social, Head
4	6.86	Audition	Sound, Audition, Loud, Low, High, Music, Attention
5	6.18	Ingestion	Taste, Head, Smell, Toward
6	6.08	Self	Needs, Near, Self, Practice
7	5.86	Motion in space	Path, Fast, Away, Motion, Lower Limb, Toward, Slow
8	5.33	Human	Face, Body, Speech, Human, Biomotion
9	5.08	Surprise	Short, Surprised
10	4.52	Place	Landmark, Scene, Large
11	4.50	Upper limb	Upper Limb, Touch, Music
12	4.49	Reward	Benefit, Needs, Drive
13	4.07	Positive	Happy, Pleasant, Arousal, Attention
14	3.22	Time	Duration, Time, Number
15	2.37	Luminance	Bright
16	1.16	Slow	Slow

Note: Factors are listed in order of variance explained. Attributes with positive loadings are listed for each factor. F# = factor number. EV = rotated eigenvalue.

Figure 1. Mean attribute vectors for 6 concrete object noun categories. Attributes are grouped and colour-coded by general domain and similarity to assist interpretation. Blackness of the attribute labels reflects the mean attribute value. Error bars represent standard error of the mean. [To view this figure in colour, please see the online version of this Journal.]

Figure 2. Mean attribute vectors for 3 event noun categories and 3 abstract noun categories. [To view this figure in colour, please see the online version of this Journal.]
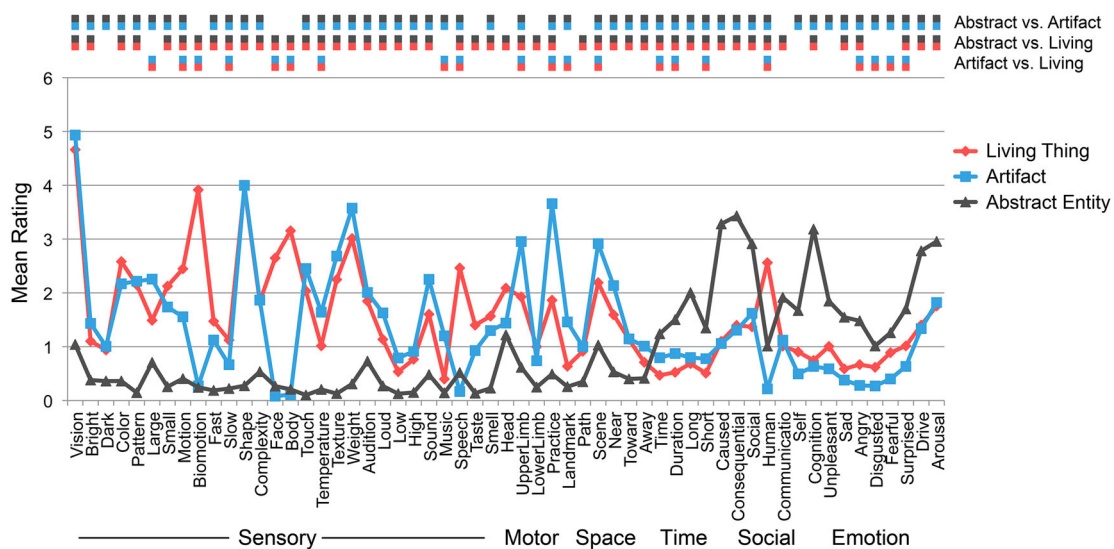
Figure 3. Mean attribute vectors for 3 verb categories. [To view this figure in colour, please see the online version of this Journal.]

the attributes specifically associated with animals and people, such as Biomotion, Face, Body, Speech, and Human, for which Artefacts received ratings as low as or lower than those for Abstract Entities. Conversely, Abstract Entities received higher ratings than the concrete categories on attributes related to temporal and causal experiences (Time, Duration, Long, Short, Caused, Consequential), social experiences (Social, Communication, Self), and emotional experiences (Unpleasant, Sad, Angry, Disgusted, Fearful, Surprised, Drive, Arousal). In all, 57 of the 65 attributes distinguished abstract from concrete entities.

Living Things were distinguished from Artefacts by the aforementioned attributes characteristic of animals and people. In addition, Living Things on average received higher ratings on Motion and Slow, suggesting they tend to have more salient movement

qualities. Living Things also received higher ratings than Artefacts on several emotional attributes (Angry, Disgusted, Fearful, Surprised), although these ratings were relatively low for both concrete categories. Conversely, Artefacts received higher ratings than Living Things on the attributes Large, Landmark, and Scene, all of which probably reflect the over-representation of buildings in this sample. Artefacts were also rated higher on Temperature, Music (reflecting an over-representation of musical instruments), Upper Limb, Practice, and several temporal attributes (Time, Duration, and Short). Though significantly different for Artefacts and Living Things, the ratings on the temporal attributes were lower for both concrete categories than for Abstract Entities. In all, 21 of the 65 attributes distinguished between Living Things and Artefacts.

Figure 4. Attributes distinguishing Artefacts, Living Things, and Abstract Entities. Pairwise differences significant at $p < .0001$ are indicated by the coloured squares above the graph. [To view this figure in colour, please see the online version of this Journal.]
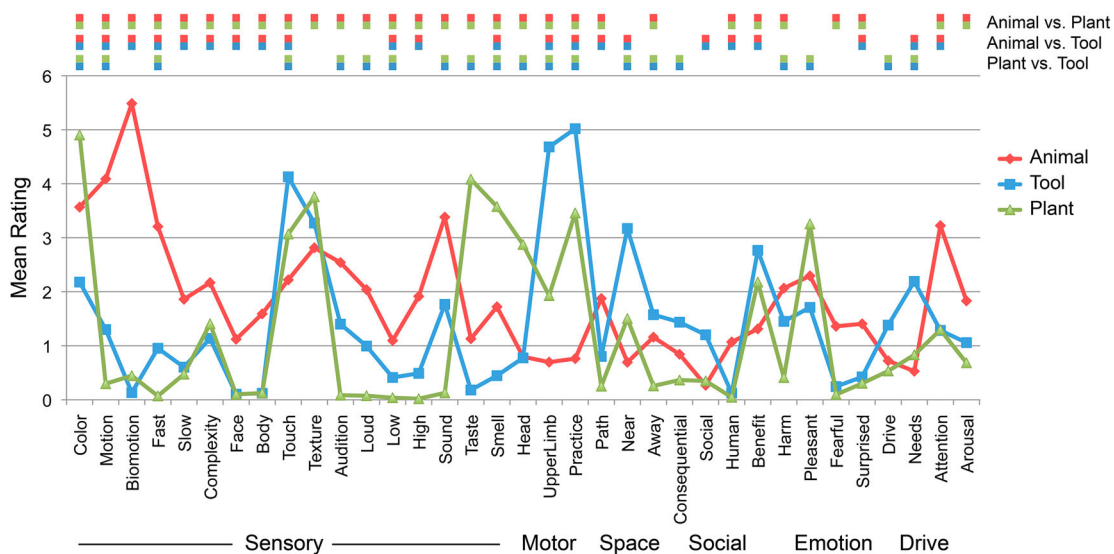
Figure 5 shows comparisons between the more specific categories Animals ($n = 30$), Plants ($N = 30$), and Tools ($N = 27$). Relatively selective impairments on these three categories are frequently reported in patients with category-related semantic impairments (Capitani, Laiacona, Mahon, & Caramazza, 2003; Forde & Humphreys, 1999; Gainotti, 2000). The data show a number of expected results (see Gainotti et al., 2009; Gainotti et al., 2013; Hoffman & Lambon Ralph, 2013, for previous similar comparisons between these categories), such as higher ratings for Animals than for either Plants or Tools on attributes related to motion; higher ratings for Plants on Taste, Smell, and Head attributes, related to their prominent role as food; and higher ratings for Tools and Plants on attributes related to manipulation (Touch, Upper Limb, Practice). Ratings on sound-related attributes were highest for Animals, intermediate for Tools, and virtually zero for Plants. Colour also showed a three-way split, with Plants > Animals > Tools. Animals, however, received higher ratings on visual Complexity.

In the spatial domain, Animals were viewed as having more salient paths of motion (Path), and Tools were rated as more close at hand in everyday life (Near). Tools were rated as more consequential and more related to social experiences, drives, and needs. Tools and Plants were seen as more beneficial than Animals, and Plants more pleasant than Tools. Animals and Tools were both viewed as having more potential for ha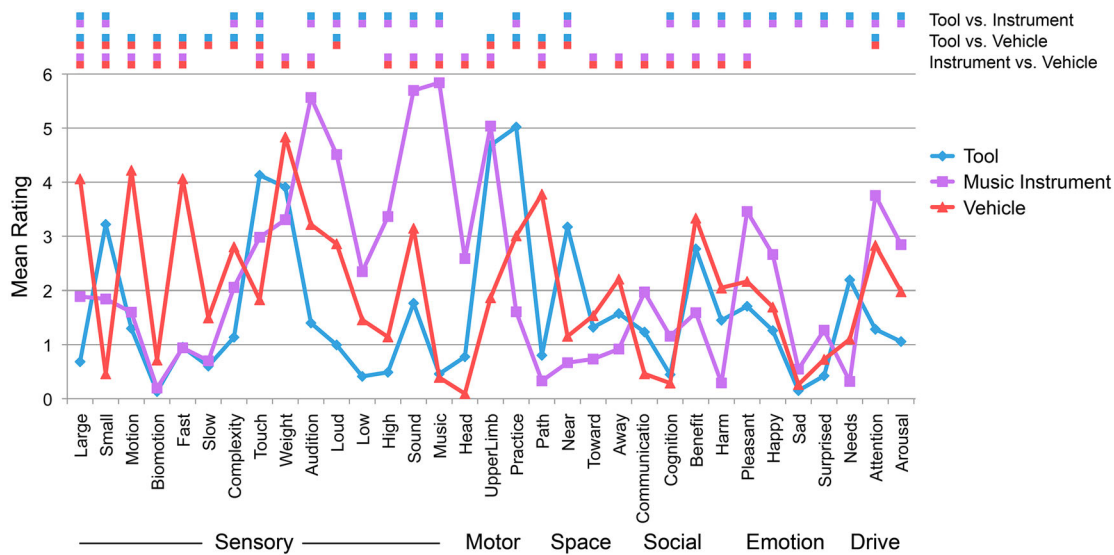rm than Plants, and Animals had higher ratings on Fearful, Surprised, Attention, and Arousal attributes, suggesting that animals are viewed as potential threats more than are the other categories. In all, 29 attributes distinguished between Animals and Plants, 22 distinguished Animals and Tools, and 20 distinguished Plants and Tools.

Figure 6 shows a three-way comparison between the specific artefact categories Musical Instruments ($N = 20$), Tools ($N = 27$), and Vehicles ($N = 20$). Again there were several expected findings, including higher ratings for Vehicles on motion-related attributes (Motion, Biomotion, Fast, Slow, Path, Away) and higher ratings for Musical Instruments on the sound-related attributes (Audition, Loud, Low, High, Sound, Music). Tools were rated higher on attributes reflecting manipulation experience (Touch, Upper Limb, Practice, Near). The importance of the Practice attribute, which encodes degree of personal experience manipulating an object or performing an action, is reflected in the much higher rating on this attribute for Tools than for Musical Instruments, whereas these categories did not differ on the Upper Limb attribute. The categories were distinguished by size in the expected direction (Vehicles > Instruments > Tools), and Instruments and Vehicles were rated higher than Tools on visual Complexity.

In the more abstract domains, Musical Instruments received higher ratings on Communication ("a thing or action that people use to communicate") and Cognition ("a form of mental activity or function of the mind"), suggesting stronger associations with mental

Figure 5. Attributes distinguishing Animals, Plants, and Tools. Pairwise differences significant at p < .0001 are indicated by the coloured squares above the graph. [To view this figure in colour, please see the online version of this Journal.]

Figure 6. Attributes distinguishing Musical Instruments, Tools, and Vehicles. [To view this figure in colour, please see the online version of this Journal.]

activity, but also higher ratings on Pleasant, Happy, Sad, Surprised, Attention, and Arousal, suggesting stronger emotional and arousal associations. Tools and Vehicles had higher ratings than Musical Instruments on both Benefit and Harm attributes, and Tools were rated higher on the Needs attribute ("someone or something that would be hard for you to live without"). In all, 22 attributes distinguished between Musical Instruments and Tools, 21 distinguished Musical Instruments and Vehicles, and 14 distinguished Tools and Vehicles.

Overall, these category-related differences are intuitively sensible, and a number of them have been reported previously (Cree & McRae, 2003; Gainotti et al., 2009; Gainotti et al., 2013; Hoffman & Lambon Ralph, 2013; Tranel, Logan, Frank, & Damasio, 1997; Vigliocco, Vinson, Lewis, & Garrett, 2004). They highlight the underlying complexity of taxonomic category structure and illustrate how the present data can be used to explore this complexity.

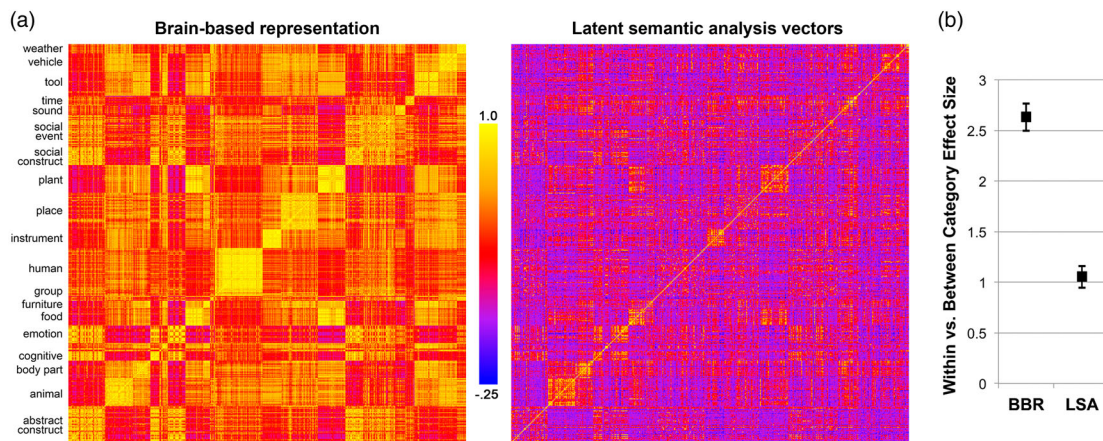### Category effects on item similarity: Brain-based versus distributional representations

Another way in which a priori category labels are useful for evaluating the representational scheme is by enabling a comparison of semantic distance between items in the same category and items in different categories. Figure 7A shows heat maps of the pairwise cosine similarity matrix for all 434 nouns in the sample, constructed using the brain-

based representation and a representation derived via latent semantic analysis (LSA; Landauer & Dumais, 1997) of a large text corpus (Landauer, Kintsch, & the Science and Applications of Latent Semantic Analysis (SALSA) Group, 2015). Vectors in the LSA representation were reduced to 65 dimensions for comparability with the brain-based vectors, though very similar results were obtained with a 300-dimensional LSA representation. Concepts are grouped in the matrices according to a priori category to highlight category similarity structure. The matrices are modestly correlated ($r = .31$, $p < .00001$). As the figure suggests, cosine similarity tended to be much higher for the brain-based vectors (mean = .55, SD = .18) than for the LSA vectors (mean = .19, SD = .17; $p < .00001$). More importantly, cosine similarity was generally greater for pairs in the same a priori category than for between-category pairs, and this difference, measured for each word using Cohen's d effect size, was much larger for the brain-based (mean = 2.64, SD = 1.09) than for the LSA vectors (mean = 1.09, SD = 0.85; $p < .00001$; Figure 7B). Thus, both the brain-based vectors and LSA vectors capture information reflecting a priori category structure, and the brain-based vectors show significantly greater effects of category co-membership than the LSA vectors.

Data-driven categorization of the words

Cosine similarity measures were also used to identify "neighbourhoods" of semantically close concepts

Figure 7. (A) Cosine similarity matrices for the 434 nouns in the study, displayed as heat maps, with words grouped by superordinate category. The matrix at left was computed using the brain-based vectors, and the matrix at right was computed using latent semantic analysis vectors. Yellow indicates greater similarity. (B) Average effect sizes (Cohen's d) for comparisons of within-category versus between-category cosine similarity values. An effect size was computed for each word, then these were averaged across all words. Error bars indicate 99% confidence intervals. BBR = brain-based representation; LSA = latent semantic analysis. [To view this figure in colour, please see the online version of this Journal.]

without reference to a priori labels. Table 7 shows some representative results (a complete list is available for download; see link prior to the References section). The results provide further evidence that the representations capture semantic similarity across concepts, although direct comparisons of these distance measures with similarity decision and priming data are needed for further validation.

A more global assessment of similarity structure was performed using k-means cluster analyses, with the mean attribute vectors for each word as input. A series of k-means analyses was performed, with the cluster parameter ranging from 20–30, reflecting the approximate number of a priori categories in Table 4. Although the results were very similar across this range, results in the range from 25–28 seemed to afford the most straightforward interpretations.

Results from the 28-cluster solution are shown in Table 8. It is important to note that we are not claiming that this is the "true" number of clusters in the data. Conceptual organization is inherently hierarchical, involving degrees of similarity, and the number of categories defined depends on the type of distinction one wishes to make (e.g., animal vs. tool, canine vs. feline, dog vs. wolf, hound vs. spaniel). Our aim here was to match approximately the "grain size" of the k-means clusters with the grain size of the a priori superordinate category labels so that these could be meaningfully compared.

Several of the clusters that emerged from the ratings data closely matched the a priori category assignments outlined in Table 4. For example, all 30 Animals clustered together, 13 of the 14 Body Parts, and all 20 Musical Instruments. One body part ("hair") fell in the Tools and

Table 7. Representational similarity neighbourhoods for some example words.

Word	Closest neighbours
actor	artist, reporter, priest, journalist, minister, businessman, teacher, boy, editor, diplomat
advice	apology, speech, agreement, plea, testimony, satire, wit, truth, analogy, debate
airport	jungle, highway, subway, zoo, cathedral, farm, church, theater, store, bar
apricot	tangerine, tomato, raspberry, cherry, banana, asparagus, pea, cranberry, cabbage, broccoli
arm	hand, finger, shoulder, leg, muscle, foot, eye, jaw, nose, lip
army	mob, soldier, judge, policeman, commander, businessman, team, guard, protest, reporter
ate	fed, drank, dinner, lived, used, bought, hygiene, opened, worked, barbecue, played
avalanche	hailstorm, landslide, volcano, explosion, flood, lightning, cyclone, tornado, hurricane
banjo	mandolin, fiddle, accordion, xylophone, harp, drum, piano, clarinet, saxophone, trumpet
bee	mosquito, mouse, snake, hawk, cheetah, tiger, chipmunk, crow, bird, ant
beer	rum, tea, lemonade, coffee, pie, ham, cheese, mustard, ketchup, jam
belch	cough, gasp, scream, squeal, whine, screech, clang, thunder, loud, shouted
blue	green, yellow, black, white, dark, red, shiny, ivy, cloud, dandelion
book	computer, newspaper, magazine, camera, cash, pen, pencil, hairbrush, keyboard, umbrella

Table 8. K-means cluster analysis of the entire 535-word set, with $k = 28$.

Animals	Body parts	Human types	Neutral human roles	Violent human roles	Plants and foods	Large bright objects
$n = 30$	$n = 13$	$n = 7$	$n = 37$	$n = 6$	$n = 44$	$n = 3$
chipmunk	arm	woman	diplomat	mob	apricot	cloud
hawk	finger	girl	mayor	terrorist	asparagus	sun
duck	hand	boy	banker	criminal	tomato	bonfire
moose	shoulder	parent	editor	victim	broccoli	
hamster	eye	man	minister	army	spaghetti	
horse	leg	child	guard	riot	jam	
camel	nose	family	businessman		cheese	
cheetah	jaw		reporter		cabbage	
penguin	foot		priest		cucumber	
pig	muscle		journalist		tangerine	

Non-social places	Social places	Tools and furniture	Musical instruments	Loud vehicles	Quiet vehicles	Festive social events
$n = 22$	$n = 20$	$n = 43$	$n = 23$	$n = 6$	$n = 18$	$n = 15$
field	office	spatula	mandolin	plane	boat	party
forest	store	hairbrush	xylophone	rocket	carriage	festival
bay	cathedral	umbrella	fiddle	train	sailboat	carnival
lake	church	cabinet	trumpet	subway	scooter	circus
prairie	lab	corkscrew	saxophone	ambulance	van	musical
apartment	park	comb	bugle	(fireworks)	bus	symphony
fence	hotel	window	banjo		cab	theater
bridge	embassy	shelves	bagpipe		limousine	parade
island	farm	stapler	clarinet		escalator	rally
elm	cafeteria	rake	harp		(truck)	celebrated

Verbal social events	Negative natural events	Nonverbal sound events	Ingestive events and actions	Beneficial abstract entities	Neutral abstract entities	Causal entities
$n = 14$	$n = 16$	$n = 11$	$n = 5$	$n = 20$	$n = 23$	$n = 19$
debate	cyclone	squeal	fed	moral	hierarchy	role
oration	hurricane	screech	ate	trust	sum	legality
testimony	hailstorm	gasp	drank	hope	paradox	power
negotiated	storm	whine	dinner	mercy	irony	law
advice	landslide	scream	barbecue	knowledge	(famous)	theory
apology	flood	clang		optimism	clue	agreement
plea	tornado	(loud)		intellect	whole	treaty
speech	explosion	cough		(spiritual)	(ended)	truce
meeting	avalanche	belch		peace	verb	fate
dictation	stampede	thunder		sympathy	patent	motive

Negative emotion entities	Positive emotion entities	Time concepts	Locative change actions	Neutral actions	Negative actions and properties	Sensory properties
$n = 30$	$n = 22$	$n = 16$	$n = 14$	$n = 23$	$n = 16$	$n = 19$
jealousy	joviality	morning	crossed	used	damaged	green
ire	gratitude	evening	left	bought	dangerous	dusty
animosity	fun	day	went	fixed	blocked	expensive
denial	joy	summer	ran	gave	destroyed	yellow
envy	liked	new	approached	helped	injured	blue
grievance	satire	spring	landed	found	broke	used
guilt	awe	era	walked	met	accident	white
snub	joke	year	marched	played	lost	(ivy)
sin	happy	winter	flew	wrote	feared	red
fallacy	(theme)	night	hiked	took	stole	empty

Note: Numbers below the cluster labels indicate the number of words in the cluster. The first 10 items in each cluster are listed in order from closest to farthest distance from the cluster centroid. Parentheses indicate words that do not fit intuitively with the interpretation.

Furniture cluster, and three additional sound-producing artefacts ("megaphone", "television", and "cellphone") clustered with the Musical Instruments. Ratings-based clustering did not distinguish edible Plants from Manufactured Foods, collapsing these into a single Foods cluster that excluded larger inedible plants ("elm", "ivy", "oak", "tree"). Similarly, data-driven clustering did not distinguish Furniture from Tools, collapsing these into a single cluster that also included most of the miscellaneous manipulated artefacts ("book", "door", "computer",

"magazine", "medicine", "newspaper", "ticket", "window") as well as several smaller non-artefact items ("feather", "hair", "ice", "stone").

Data-driven clustering also identified a number of intuitively plausible divisions within categories. Basic human biological types ("woman", "man", "child", etc.) were distinguished from human occupational roles, and human individuals or groups with negative or violent associations formed a distinct cluster. Compared to the neutral occupational roles, human type

concepts had significantly higher ratings on attributes related to sensory experience (Shape, Complexity, Touch, Temperature, Texture, High, Sound, Smell), nearness (Near, Self), and positive emotion (Happy; Table 9). Places were divided into two groups, which we labelled Non-Social and Social, that corresponded roughly to the a priori distinction between Natural Scenes and Places/Buildings, except that the Non-Social Places cluster included several manmade places ("apartment", fence", "bridge", "street", etc.). Social Places had higher ratings on attributes related to human interactions (Social, Communication, Consequential), and Non-Social Places had higher ratings on several sensory attributes (Colour, Pattern, Shape, Touch, and Texture; Table 10). In addition to distinguishing vehicles from other artefacts, data-driven clustering separated loud vehicles from quiet vehicles (mean Loud ratings 5.30 vs. 1.96, respectively). Loud vehicles also had significantly higher ratings on several other auditory attributes (Audition, High, Sound). The two vehicle clusters contained four unexpected items ("fireworks", "fountain", "river", "soccer"), probably due to high ratings for these items on Motion and Path attributes, which distinguish vehicles from other nonliving objects.

In the domain of event nouns, clustering divided the a priori category Social Events into two clusters that we labelled Festive Social Events and Verbal Social Events (Table 11). Festive Events had significantly higher ratings on a number of sensory attributes related to visual shape, visual motion, and sound and were rated as more Pleasant and Happy. Verbal Social Events had higher ratings on attributes related to human cognition (Communication, Cognition,

Table 9. Attributes that differ significantly between human types and neutral human roles (occupations).

Attribute	Human types	Neutral human roles
Bright*	**0.80 (0.28)**	0.37 (0.19)
Small*	**1.73 (1.19)**	0.63 (0.29)
Shape*	**3.96 (0.81)**	2.45 (0.55)
Complexity	**2.58 (0.50)**	1.78 (0.38)
Touch*	**2.22 (0.83)**	0.50 (0.25)
Temperature	**0.69 (0.37)**	0.34 (0.14)
Texture*	**1.69 (1.01)**	0.64 (0.24)
High*	**1.69 (1.12)**	0.39 (0.22)
Sound*	**2.73 (0.58)**	1.30 (0.52)
Smell*	**1.27 (0.61)**	0.36 (0.28)
Near*	**2.91 (0.77)**	0.84 (0.54)
Self	**2.71 (1.23)**	1.01 (0.77)
Happy	**3.50 (0.81)**	1.76 (0.91)

Note: Columns give mean for each rating, with standard deviations in parentheses. The larger value is indicated in bold font.
Differences significant at $p < .0001$. *$p < .00001$.

Table 10. Attributes that differ significantly between non-social places and social places.

Attribute	Non-social places	Social places
Colour*	**2.71 (1.09)**	0.99 (0.51)
Pattern*	**2.92 (0.82)**	1.05 (0.57)
Shape*	**3.89 (0.91)**	2.50 (0.78)
Touch*	**1.95 (1.03)**	0.47 (0.37)
Texture*	**2.76 (1.07)**	0.92 (0.41)
Time	0.36 (0.42)	**1.34 (0.92)**
Consequential*	0.65 (0.45)	**1.89 (1.00)**
Social*	0.84 (0.52)	**3.31 (0.96)**
Communication*	0.26 (0.19)	**1.38 (0.79)**
Cognition	0.20 (0.13)	**1.03 (0.84)**
Disgusted	0.08 (0.06)	**0.49 (0.38)**

Note: Columns give mean for each rating, with standard deviations in parentheses. The larger value is indicated in bold font.
Differences significant at $p < .0001$. *$p < .00001$.

Consequential) and, interestingly, were rated as both more Unpleasant and more Beneficial than Festive Events. A cluster labelled Negative Natural Events corresponded roughly to the a priori category Weather Events, with the addition of several non-meteorological negative or violent events ("explosion", "stampede", "battle", etc.). A cluster labelled Nonverbal Sound Events corresponded closely to the a priori category of the same name. Finally, a small cluster labelled Ingestive Events and Actions included several social events and verbs of ingestion, linked by high ratings on the attributes Taste, Smell, Upper Limb, Head, Toward, Pleasant, Benefit, and Needs.

Correspondence between data-driven clusters and a priori categories was less clear in the domain of abstract nouns. Clustering yielded two abstract

Table 11. Attributes that differ significantly between festive and verbal social events.

Attribute	Festive social events	Verbal social events
Vision*	**4.56 (0.93)**	1.41 (1.13)
Bright	**1.50 (0.98)**	0.10 (0.07)
Large*	**3.18 (1.38)**	0.53 (0.72)
Motion*	**3.60 (1.00)**	0.84 (0.85)
Fast*	**1.51 (0.63)**	0.38 (0.28)
Shape*	**1.40 (0.71)**	0.24 (0.21)
Complexity*	**2.89 (1.17)**	0.48 (0.61)
Loud*	**3.89 (0.92)**	1.49 (1.09)
Sound	**3.89 (1.15)**	1.96 (1.03)
Music	**3.21 (1.77)**	0.52 (0.78)
Head	1.85 (1.30)	**3.98 (1.09)**
Consequential*	1.61 (0.85)	**3.83 (0.82)**
Communication*	2.20 (1.14)	**5.33 (0.39)**
Cognition*	1.15 (0.44)	**3.70 (0.77)**
Benefit*	2.50 (0.53)	**3.75 (0.58)**
Pleasant*	**4.41 (0.85)**	1.78 (0.90)
Unpleasant*	0.38 (0.25)	**1.62 (0.59)**
Happy*	**4.26 (0.88)**	1.63 (0.92)
Angry*	0.34 (0.36)	**1.24 (0.61)**

Note: Columns give mean for each rating, with standard deviations in parentheses. The larger value is indicated in bold font.
Differences significant at $p < .0001$. *$p < .00001$.

entity clusters mainly distinguished by perceived benefit and association with positive emotions (Table 12), labelled Beneficial and Neutral, which comprise a variety of abstract and mental constructs. A third cluster labelled Causal Entities comprises concepts with high ratings on the Consequential and Social attributes (Table 13). Two further clusters comprise abstract entities with strong negative and positive emotional meaning or associations (Table 14). The final cluster of abstract entities corresponds approximately to the a priori category Time Periods, but also includes properties ("new", "old", "young") and verbs ("grew", "slept") associated with time duration concepts.

Three clusters consisted mainly of verbs. These included a cluster labelled Locative Change Actions, which corresponded closely with the a priori category of the same name and was defined by high ratings on attributes related to motion through space (Table 15). The second verb cluster contained a mixture of emotionally neutral physical and social actions. The final verb-heavy cluster contained a mix of verbs and adjectives with negative or violent associations.

The final cluster emerging from the k-means analysis consisted almost entirely of adjectives describing physical sensory properties, including

Table 12. Attributes that differ significantly between beneficial and neutral abstract entities.

Attribute	Beneficial abstract entities	Neutral abstract entities
Consequential	**3.42 (0.83)**	2.22 (0.95)
Self*	**3.61 (1.03)**	1.19 (1.02)
Cognition	**4.03 (1.38)**	2.19 (1.36)
Benefit*	**4.68 (0.70)**	2.31 (1.29)
Pleasant*	**3.84 (0.93)**	1.45 (0.78)
Happy*	**3.90 (1.05)**	1.32 (0.76)
Drive*	**3.76 (0.82)**	1.70 (0.60)
Needs*	**3.52 (1.16)**	1.30 (0.99)
Arousal*	**3.17 (0.94)**	1.76 (0.83)

Note: Columns give mean for each rating, with standard deviations in parentheses. The larger value is indicated in bold font.
Differences significant at p < .0001. *p < .00001.

Table 13. Attributes that differ significantly between causal entities and neutral abstract entities.

Attribute	Causal entities	Neutral abstract entities
Consequential*	**4.47 (0.67)**	2.22 (0.95)
Social*	**3.94 (1.21)**	1.80 (0.91)
Drive*	**3.46 (0.83)**	1.70 (0.60)
Arousal*	**2.95 (0.59)**	1.76 (0.83)

Note: Columns give mean for each rating, with standard deviations in parentheses. The larger value is indicated in bold font.
Differences significant at p < .0001. *p < .00001.

Table 14. Attributes that differ significantly between negative and positive emotion entities.

Attribute	Negative emotion entities	Positive emotion entities
Vision	0.75 (0.49)	**1.52 (0.81)**
Bright*	0.06 (0.06)	**0.62 (0.50)**
Dark	**0.56 (0.41)**	0.14 (0.16)
Pain*	**2.09 (1.05)**	0.19 (0.21)
Music	0.10 (0.14)	**0.57 (0.54)**
Consequential*	**4.42 (0.72)**	2.58 (0.80)
Self*	1.01 (0.56)	**2.52 (1.20)**
Benefit*	0.75 (0.65)	**3.37 (0.93)**
Harm*	**3.81 (0.93)**	0.46 (0.55)
Pleasant*	0.28 (0.25)	**4.71 (0.84)**
Unpleasant*	**4.80 (0.72)**	0.29 (0.37)
Happy*	0.23 (0.35)	**4.80 (0.92)**
Sad*	**3.46 (1.12)**	0.41 (0.58)
Angry*	**3.79 (1.06)**	0.29 (0.40)
Disgusted*	**2.70 (0.84)**	0.24 (0.37)
Fearful*	**2.59 (1.06)**	0.26 (0.27)
Needs*	0.37 (0.47)	**2.09 (0.96)**

Note: Columns give mean for each rating, with standard deviations in parentheses. The larger value is indicated in bold font.
Differences significant at p < .0001. *p < .00001.

colour, surface appearance, tactile texture, temperature, and weight.

Hierarchical clustering

A hierarchical cluster analysis of the 335 concrete nouns (objects and events) was performed to examine the underlying category structure in more detail. The major categories were again clearly evident in the resulting dendrogram. A number of interesting subdivisions emerged, as illustrated in the dendrogram segments shown in Figure 8. Musical Instruments, which formed a distinct cluster, subdivided further into instruments played by blowing (left subgroup, light blue, colour online) and instruments played with the upper limbs only (right subgroup), the latter of which contained a somewhat segregated group of instruments played by striking. "Piano" formed a distinct branch, probably due to its

Table 15. Attributes that differ significantly between locative change and abstract/social actions.

Attribute	Locative change actions	Neutral actions
Motion*	**3.81 (0.71)**	1.98 (0.81)
Biomotion*	**4.64 (0.67)**	2.87 (0.78)
Fast*	**3.23 (1.27)**	1.42 (0.76)
Body	**4.47 (0.98)**	2.74 (1.06)
Lower Limb*	**3.86 (2.08)**	1.11 (0.96)
Path*	**5.10 (0.84)**	2.05 (1.43)
Communication	1.01 (0.58)	**2.88 (1.51)**
Cognition	0.96 (0.31)	**2.53 (1.28)**

Note: Columns give mean for each rating, with standard deviations in parentheses. The larger value is indicated in bold font.
Differences significant at p < .0001. *p < .00001.

higher rating on the Lower Limb attribute. People concepts, which also formed a distinct cluster, showed evidence of a subgroup of "private sector" occupations (large left subgroup, green, colour online) and a subgroup of "public sector" occupations (middle subgroup, purple, colour online). Two other subgroups consisted of basic human types and people concepts with negative or violent associations (far right subgroup, magenta, colour online).

Animals also formed a distinct cluster, though two ("ant" and "butterfly") were isolated on nearby branches. The animals branched into somewhat distinct subgroups of aquatic animals (left-most subgroup, light blue, colour online), small land animals (next subgroup, yellow, colour online), large animals (next subgroup, red, colour online), and unpleasant animals (subgroup including "alligator," magenta, colour online). Interestingly, "camel" and "horse"— the only animals in the set used for human transportation—segregated together despite the fact that

there are no attributes that specifically encode "used for human transportation". Inspection of the vectors suggested that this segregation was due to higher ratings on both the Lower Limb and Benefit attributes than for the other animals. Association with lower limb movements and benefit to people, that is, appears to be sufficient to mark an animal as useful for transportation. Finally, Plants and Foods formed a large distinct branch, which contained subgroups of beverages, fruits, manufactured foods, vegetables, and flowers.

A similar hierarchical cluster analysis was performed on the 99 abstract nouns. Three major branches emerged. The largest of these comprised abstract concepts with a neutral or positive meaning. As shown in Figure 9, one subgroup within this large cluster consisted of positive or beneficial entities (subgroup from "attribute" to "gratitude," green, colour online). A second fairly distinct subgroup consisted of "causal" concepts associated with a high likelihood of consequences (subgroup from "motive" to "fate,"

Figure 8. Dendrogram segments from the hierarchical cluster analysis on concrete nouns. Musical instruments, people, animals, and plants/foods each clustered together on separate branches, with evidence of sub-clusters within each branch. [To view this figure in colour, please see the online version of this Journal.]
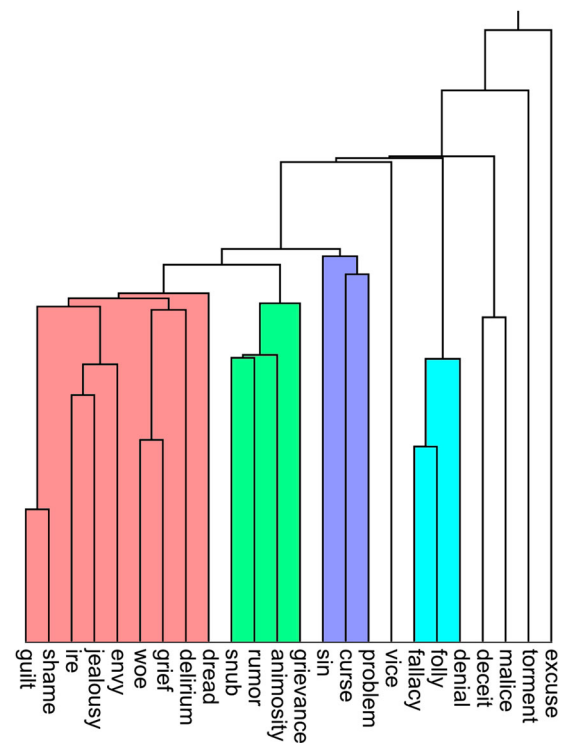
blue, colour online). A third subgroup included concepts associated with number or quantification (subgroup from "noun" to "infinity," light blue, colour online). A number of pairs (advice/apology, treaty/truce) and triads (irony/paradox/mystery) with similar meanings were also evident within the large cluster.

The second major branch of the abstract hierarchy comprised concepts associated with harm or another negative meaning (Figure 10). One subgroup within this branch consisted of words denoting negative emotions (subgroup from "guilt" to "dread," red, colour online). A second subgroup seemed to consist of social situations associated with anger (subgroup from "snub" to "grievance," green, colour online). A third subgroup was a triad (sin/curse/problem) related to difficulty or misfortune. A fourth subgroup was a triad (fallacy/folly/denial) related to error or misjudgment.

The final major branch of the abstract hierarchy consisted of concepts denoting time periods (day, morning, summer, spring, evening, night, winter).

Correlations with word frequency and imageability

Frequency of word use provides a rough indication of the frequency and significance of a concept. We predicted that attributes related to proximity and self-needs would be correlated with word frequency. Ima-

Figure 10. Negative abstract branches of the abstract noun dendrogram. [To view this figure in colour, please see the online version of this Journal.]

geability is a rough indicator of the overall salience of sensory, particularly visual, attributes of an entity, and thus we predicted correlations with attributes related to sensory and motor experience. An alpha level of

Figure 9. Neutral abstract branches of the abstract noun dendrogram. [To view this figure in colour, please see the online version of this Journal.]

.00001 (r = .1673) was adopted to reduce family-wise Type 1 error from the large number of tests.

Word frequency was positively correlated with Near, Self, Benefit, Drive, and Needs, consistent with the expectation that word frequency reflects the proximity and survival significance of concepts. Word frequency was also positively correlated with attributes characteristic of people, including Face, Body, Speech, and Human, reflecting the proximity and significance of conspecific entities. Word frequency was negatively correlated with a number of sensory attributes, including Colour, Pattern, Small, Shape, Temperature, Texture, Weight, Audition, Loud, Low, High, Sound, Music, and Taste. One possible explanation for this is the tendency for some natural object categories with very salient perceptual features to have far less salience in everyday experience and language use, particularly animals, plants, and musical instruments.

As expected, most of the sensory and motor attributes were positively correlated (p < .00001) with imageability, including Vision, Bright, Dark, Colour, Pattern, Large, Small, Motion, Biomotion, Fast, Slow, Shape, Complexity, Touch, Temperature, Weight, Loud, Low, High, Sound, Taste, Smell, Upper Limb, and Practice. Also positively correlated were the visual–spatial attributes Landmark, Scene, Near, and Toward. Conversely, many non-perceptual attributes were negatively correlated with imageability, including temporal and causal components (Duration, Long, Short, Caused, Consequential), social and cognitive components (Social, Human, Communication, Self, Cognition), and affective/arousal components (Unpleasant, Sad, Angry, Disgusted, Fearful, Surprised, Drive, Arousal).

Correlations with non-semantic lexical variables

The large size of the dataset afforded an opportunity to look for unanticipated relationships between semantic attributes and non-semantic lexical variables. Previous research has identified various phonological correlates of meaning (Lynott & Connell, 2013; Schmidtke, Conrad, & Jacobs, 2014) and grammatical class (Farmer, Christiansen, & Monaghan, 2006; Kelly, 1992). These data suggest that the evolution of word forms is not entirely arbitrary, but is influenced in part by meaning and pragmatic factors. In an exploratory analysis, we tested for correlation between the 65

semantic attributes and measures of length (number of phonemes), orthographic typicality (mean position-constrained bigram frequency, orthographic neighbourhood size), and phonologic typicality (phonologic neighbourhood size). An alpha level of .00001 was again adopted to reduce family-wise Type 1 error from the large number of tests and to focus on very strong effects that are perhaps more likely to hold across other random samples of words.

Word length (number of phonemes) was negatively correlated with Near and Needs, with a strong trend (p < .0001) for Practice, suggesting that shorter words are used for things that are near to us, central to our needs, and frequently manipulated. Similarly, orthographic neighbourhood size was positively correlated with Near, Needs, and Practice, with strong trends for Touch and Self, and phonological neighbourhood size was positively correlated with Near and Practice, with strong trends for Needs and Touch. Together, these correlations suggest that concepts referring to nearby objects of significance to our self needs tend to be represented by shorter or simpler morphemes with commonly shared spelling patterns. Position-constrained bigram frequency was not significantly correlated with any of the attributes, however, suggesting that semantic variables mainly influence segments larger than letter pairs.

Potential applications

The intent of the current study was to outline a relatively comprehensive brain-based componential model of semantic representation and to explore some of the properties of this type of representation. Building on extensive previous work (Allport, 1985; Borghi et al., 2011; Cree & McRae, 2003; Crutch et al., 2013; Crutch et al., 2012; Gainotti et al., 2009, 2013; Hoffman & Lambon Ralph, 2013; Lynott & Connell, 2009, 2013; Martin & Caramazza, 2003; Tranel et al., 1997; Vigliocco et al., 2004; Warrington & McCarthy, 1987; Zdrazilova & Pexman, 2013), the representation was expanded to include a variety of sensory and motor subdomains, more complex physical domains (space, time), and mental experiences, guided by the principle that all aspects of experience are encoded in the brain and processed according to information content. The utility of the representation ultimately depends on its completeness and veridicality, which are determined in no small part by the specific

attributes included and details of the queries used to elicit attribute salience values. These choices are almost certainly imperfect, and thus we view the current work as a preliminary exploration that we hope will lead to future improvements.

The representational scheme is motivated by an underlying theory of concept representation in the brain, and its principal utility—and main source of validation—is likely to be in brain research. A recent study by Fernandino et al. (Fernandino et al., 2015) suggests one type of application. The authors obtained ratings for 900 noun concepts on the five attributes Colour, Visual Motion, Shape, Sound, and Manipulation. Functional MRI scans, performed while participants read these 900 words and performed a concreteness decision, showed distinct brain networks modulated by the salience of each attribute, as well as multi-level convergent areas that responded to various subsets or all of the attributes. Future studies along these lines could incorporate a much broader range of information types both within the sensory domains and across sensory and non-sensory domains.

Another likely application is in the study of category-related semantic deficits induced by brain damage. Patients with these syndromes show differential abilities to process object concepts from broad (living vs. artefact) or more specific (animal, tool, plant, body part, musical instrument, etc.) categories. The embodiment account of these phenomena posits that object categories differ systematically in terms of the type of sensory–motor knowledge on which they depend. Living things, for example, are cited as having more salient visual attributes, whereas tools and other artefacts have more salient action-related attributes (Farah & McClelland, 1991; Warrington & McCarthy, 1987). As discussed by Gainotti (Gainotti, 2000), greater impairment of knowledge about living things is typically associated with anterior ventral temporal lobe damage, which is also an area that supports high-level visual object perception, whereas greater impairment of knowledge about tools is associated with frontoparietal damage, an area that supports object-directed actions. On the other hand, counterexamples have been reported, including patients with high-level visual perceptual deficits but no selective impairment for living things and patients with apparently normal visual perception despite a selective living thing impairment (Capitani et al., 2003).

The current data, however, support more recent proposals suggesting that category distinctions can arise from much more complex combinations of attributes (Cree & McRae, 2003; Gainotti et al., 2013; Hoffman & Lambon Ralph, 2013; Tranel et al., 1997). As exemplified in Figures 3–5, concrete object categories differ from each other across multiple domains, including attributes related to specific visual subsystems (visual motion, biological motion, face perception), sensory systems other than vision (sound, taste, smell), affective and arousal associations, spatial attributes, and various attributes related to social cognition. Damage to any of these processing systems, or damage to spatially proximate combinations of them, could account for differential category impairments. As one example, the association between living thing impairments and anterior ventral temporal lobe damage may not be due to impairment of high-level visual knowledge per se, but rather to damage to a zone of convergence between high-level visual, taste, smell, and affective systems (Gainotti et al., 2013).

The current data also clarify the diagnostic attribute composition of a number of salient categories that have not received systematic attention in the neuropsychological literature, such as foods, musical instruments, vehicles, human roles, places, events, time concepts, locative change actions, and social actions. Each of these categories is defined by a unique subset of particularly salient attributes, and thus brain damage affecting knowledge of these attributes could conceivably give rise to an associated category-specific impairment. Several novel distinctions emerged from a data-driven cluster analysis of the concept vectors (Table 8), such as the distinction between social and non-social places, verbal and festive social events, and three other event types (negative, sound, and ingestive). Although these data-driven clusters probably reflect idiosyncrasies of the concept sample to some degree, they raise a number of intriguing possibilities that can be addressed in future studies using a wider range of concepts.

Application to some general problems in componential semantic theory

Apart from its utility in empirical neuroscience research, a brain-based componential representation might provide alternative answers to some

longstanding theoretical issues. Several of these are discussed briefly below.

Feature selection

All analytic or componential theories of semantic representation contend with the general issue of feature selection. What counts as a feature, and how can the set of features be identified? As discussed above, standard verbal feature theories face a basic problem with feature set size, in that the number of all possible object and action features is extremely large. Here we adopt a fundamentally different approach to feature selection, in which the content of a concept is not a set of verbal features that people associate with the concept, but rather the set of neural processing modalities (i.e., representation classes) that are engaged when experiencing instances of the concept. The underlying motivation for this approach is that it enables direct correspondences to be made between conceptual content and neural representations, whereas such correspondences can only be inferred post hoc from verbal feature sets, and only by mapping such features to neural processing modalities (Cree & McRae, 2003). A consequence of defining conceptual content in terms of brain systems is that those systems provide a closed and relatively small set of basic conceptual components. Although the adequacy of these components for capturing fine distinctions in meaning is not yet clear, the basic assumption that conceptual knowledge is distributed across modality-specific neural systems offers a potential solution to the longstanding problem of feature selection.
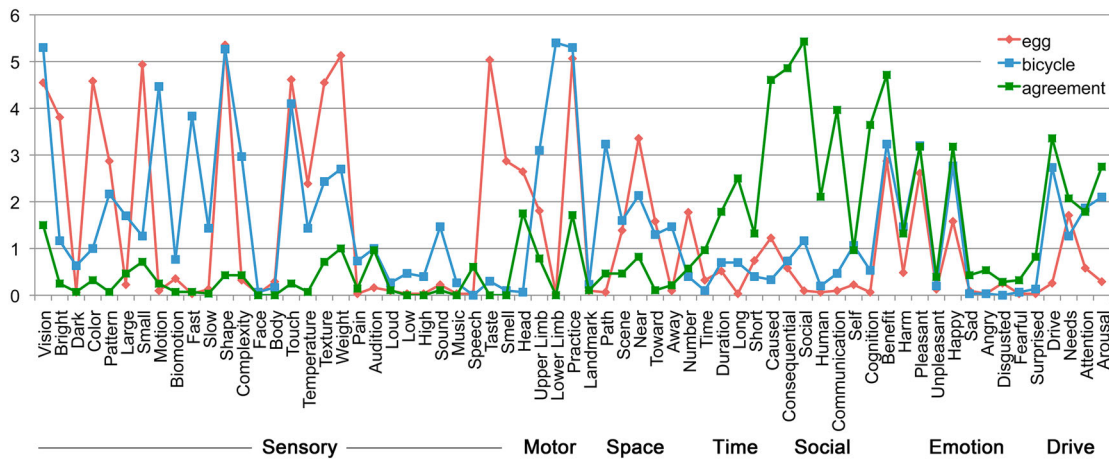
Feature weighting

Standard verbal feature representations also face the problem of defining the relative importance of each feature to a given concept. As Murphy and Medin (Murphy & Medin, 1985) put it, a zebra and a barber pole can be considered close semantic neighbours if the feature "has stripes" is given enough weight. In determining similarity, what justifies the intuition that "has stripes" should not be given more weight than other features? Other examples of this problem are instances in which the same feature can be either critically important or irrelevant for defining a concept. Keil (Keil, 1991) made the point as follows:

polar bears and washing machines are both typically white. Nevertheless, an object identical to a polar bear in all respects except colour is probably not a polar bear, while an object identical in all respects to a washing machine is still a washing machine regardless of its colour. Whether or not the feature "is white" matters to an item's conceptual status thus appears to depend on its other properties—there is no single weight that can be given to the feature that will always work. Further complicating this problem is the fact that in feature production tasks, people often neglect to mention some features. For instance, in listing properties of dogs, people rarely say, "has blood" or "can move" or "has DNA" or "can reproduce"—yet these features are central to our conception of animals.

Our theory proposes that attributes are weighted according to statistical regularities. If polar bears are always experienced as white, then colour becomes a strongly weighted attribute of polar bears. Colour would be a strongly weighted attribute of washing machines if it were actually true that washing machines are nearly always white. In actuality, however, washing machines and most other artefacts usually come in a variety of colours, so colour is a less strongly weighted attribute of most artefacts. A few artefacts provide counter-examples. Stop signs, for example, are always red. Consequently, a sign that was not red would be a poor exemplar of a stop sign even if it had the word "Stop" on it. School buses are virtually always yellow, thus a grey or black or green bus would be unlikely to be labelled a school bus. Semantic similarity is determined by overall similarity across all attributes rather than by a single attribute. Although barber poles and zebras both have salient visual patterns, they strongly differ in salience on many other attributes (shape complexity, biological motion, body parts, and motion through space) that would preclude them from being considered semantically similar.

Attribute weightings in our theory are obtained directly from human participants by averaging salience ratings provided by the participants. Rather than asking participants to list the most salient attributes for a concept, a continuous rating is required for each attribute. This method avoids the problem of attentional bias or pragmatic phenomena that result in incomplete representations using the feature listing procedure (Hoffman & Lambon Ralph,

Figure 11. Mean ratings for the concepts *egg*, *bicycle*, and *agreement*. [To view this figure in colour, please see the online version of this Journal.]

2013). The resulting attributes are graded rather than binary and thus inherently capture weightings. An example of how this works is given in Figure 11, which includes a portion of the attribute vectors for the concepts *egg* and *bicycle*. Given that these concepts are both inanimate objects, they share low weightings on human-related attributes (biological motion, face, body part, speech, social, communication, emotion, and cognition attributes), auditory attributes, and temporal attributes. However, they also differ in expected ways, including stronger weightings for *egg* on brightness, colour, small size, tactile texture, weight, taste, smell, and head action attributes, and stronger weightings for *bicycle* on motion, fast motion, visual complexity, lower limb action, and path motion attributes.

Abstract concepts

It is not immediately clear how embodiment theories of concept representation account for concepts that are not reliably associated with sensory–motor structure. These include obvious examples like *justice* and *piety*, for which people have difficulty generating reliable feature lists and whose exemplars do not exhibit obvious coherent structure. They also include somewhat more concrete kinds of concepts where the relevant structure does not clearly reside in perception or action. Social categories provide one example: categories like *male* encompass items that are grossly perceptually different (consider infant, child, and elderly human males, or the males of any given plant or animal species, which are certainly

more similar to the female of the same species than to the males of different species); categories like *Catholic* or *Protestant* do not appear to reside in sensory–motor structure; concepts like *pauper*, *liar*, *idiot* and so on are important for organizing our social interactions and inferences about the world but refer to economic circumstances, behaviours, and mental predispositions that are not obviously reflected in the perceptual and motor structure of the environment. In these and many other cases it is difficult to see how the relevant concepts are transparently reflected in the feature structure of the environment.

The experiential content and acquisition of abstract concepts from experience has been discussed by a number of authors (Altarriba & Bauer, 2004; Barsalou & Wiemer-Hastings, 2005; Borghi et al., 2011; Bretherton & Beeghly, 1982; Crutch et al., 2013; Crutch & Warrington, 2005; Crutch et al., 2012; Kousta et al., 2011; Lakoff & Johnson, 1980; Schwanenflugel, 1991; Vigliocco et al., 2009; Wiemer-Hastings & Xu, 2005; Zdrazilova & Pexman, 2013). Although abstract concepts encompass a variety of experiential domains (spatial, temporal, social, affective, cognitive, etc.) and are therefore not a homogeneous set, one general observation is that many abstract concepts are learned by experience with complex situations and events. As Barsalou (Barsalou, 1999) points out, such situations often include multiple agents, physical events, and mental events. This contrasts with concrete concepts, which typically refer to a single component (usually an object or action) of a situation. In both cases concepts are learned through experience, but abstract concepts

differ from concrete concepts in the distribution of content over entities, space, and time. Another way of saying this is that abstract concepts seem "abstract" because their content is distributed across multiple components of situations.

Another aspect of many abstract concepts is that their content refers to aspects of mental experience rather than to sensory–motor experience. Many abstract concepts refer specifically to cognitive events or states (*think*, *believe*, *know*, *doubt*, *reason*) or to products of cognition (*concept*, *theory*, *idea*, *whim*). Abstract concepts also tend to have stronger affective content than do concrete concepts (Borghi et al., 2011; Kousta et al., 2011; Troche et al., 2014; Vigliocco et al., 2009; Zdrazilova & Pexman, 2013), which may explain the selective engagement of anterior temporal regions by abstract relative to concrete concepts in many neuroimaging studies (see Binder, 2007; Binder et al., 2009, for reviews). Many so-called abstract concepts refer specifically to affective states or qualities (*anger*, *fear*, *sad*, *happy*, *disgust*). One crucial aspect of the current theory that distinguishes it from some other versions of embodiment theory is that these cognitive and affective mental experiences count every bit as much as sensory–motor experiences in grounding conceptual knowledge. Experiencing an affective or cognitive state or event is like experiencing a sensory–motor event except that the object of perception is internal rather than external. This expanded notion of what counts as embodied experience adds considerably to the descriptive power of the conceptual representation, particularly for abstract concepts.

One prominent example of how the model enables representation of abstract concepts is by inclusion of attributes that code social knowledge (see also Crutch et al., 2013; Crutch et al., 2012; Troche et al., 2014). Empirical studies of social cognition have identified several neural systems that appear to be selectively involved in supporting theory of mind processes, "self" representation, and social concepts (Araujo et al., 2013; Northoff et al., 2006; Olson, McCoy, Klobusicky, & Ross, 2013; Ross & Olson, 2010; Schilbach et al., 2012; Schurz et al., 2014; Simmons, Reddish, Bellgowan, & Martin, 2010; Spreng et al., 2009; Van Overwalle, 2009; van Veluw & Chance, 2014; Zahn et al., 2007). Many abstract words *are* social concepts (*cooperate*, *justice*, *liar*, *promise*, *trust*) or have salient social content (Borghi et al., 2011). An example of how attributes related to social cognition play a role in representing abstract concepts is shown in Figure 11, which compares the attribute vectors for *egg* and *bicycle* with the attribute vector for *agreement*. As the figure shows, ratings on sensory–motor attributes are generally low for *agreement*, but this concept receives much higher ratings than the other concepts on social cognition attributes (Caused, Consequential, Social, Human, Communication). It also receives higher ratings on several temporal attributes (Duration, Long) and on the Cognition attribute. Note also the high rating on the Benefit attribute and the small but clearly non-zero rating on the head action attribute, both of which might serve to distinguish *agreement* from many other abstract concepts that are not associated with benefit or with actions of the face/head.

Although these and many other examples illustrate how abstract concepts are often tied to particular kinds of mental, affective, and social experiences, we do not deny that language plays a large role in facilitating the learning of abstract concepts. Language provides a rich system of abstract symbols that efficiently "point to" complex combinations of external and internal events, enabling acquisition of new abstract concepts by association with older ones. In the end, however, abstract concepts, like all concepts, must ultimately be grounded in experience, albeit experiences that are sometimes very complex, distributed in space and time, and composed largely of internal mental events.

Context effects and ad hoc categories

The relative importance given to particular attributes of a concept varies with context. In the context of moving a piano, the weight of the piano matters more than its function, and consequently the piano is likely to be viewed as more similar to a couch than to a guitar. In the context of listening to a musical group, the piano's function is more relevant than its weight, and consequently it is more likely to be conceived as similar to a guitar than to a couch (Barsalou, 1982, 1983). Componential approaches to semantic representation assume that the weight given to different feature dimensions can be controlled by context, but the mechanism by which such weight adjustments occur is unclear. The possibility of learning different weightings for different

contexts has been explored for simple category learning tasks (Minda & Smith, 2002; Nosofsky, 1986), but the scalability of such an approach to real semantic cognition is questionable and seems ill-suited to explain the remarkable flexibility with which human beings can exploit contextual demands to create and employ new categories on the spot (Barsalou, 1991).
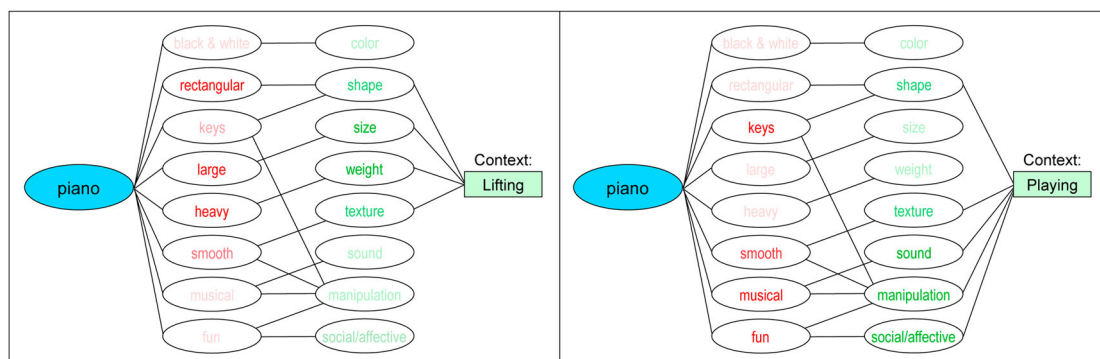
We propose that activation of attribute representations is modulated continuously through two mechanisms: top-down attention and interaction with attributes of the context itself. The context draws attention to context-relevant attributes of a concept. In the case of lifting or preparing to lift a piano, for example, attention is focused on action and proprioception processing in the brain, and this top-down focus enhances activation of action and proprioceptive attributes of the piano. This idea is illustrated in highly simplified schematic form on the left side of Figure 12. The concept of *piano* is represented for illustrative purposes by set of verbal features (first column of ovals, red, colour online), which are coded in the brain in domain-specific neural processing systems (second column of ovals, green, colour online). The context of lifting draws attention to a subset of these neural systems, enhancing their activation. Changing the context to playing or listening to music (right side of Figure 12) shifts attention to a different subset of attributes, with corresponding enhancement of this subset. In addition to effects mediated by attention, there could be direct interactions at the neural system level between the distributed representation of the target concept (piano) and the context concept. The concept of *lifting*, for example, is partly encoded in action and proprioceptive systems that overlap with those encoding *piano*. As discussed in more detail in the following section on conceptual combination, we propose that overlapping neural representations of two or more concepts can produce mutual enhancement of the overlapping components.

The enhancement of context-relevant attributes of concepts alters the relative similarity between concepts, as illustrated by the piano–couch–guitar example given above. This process not infrequently results in purely functional groupings of objects (e.g., *things one can stand on to reach the top shelf*) or "ad hoc categories" (Barsalou, 1983). The above account provides a partial mechanistic account of such categories: Ad hoc categories are formed when concepts share the same context-related attribute enhancement.

Conceptual combination

Language use depends on the ability to productively combine concepts to create more specific or more complex concepts. Some simple examples include modifier–noun (*red jacket*), noun–noun (*ski jacket*), subject–verb (*the dog ran*), and verb–object (*hit the ball*) combinations. Semantic processes underlying conceptual combination have been explored in numerous studies (Clark & Berman, 1987; Conrad & Rips, 1986; Downing, 1977; Gagné, 2001; Gagné & Murphy, 1996; Gagné & Shoben, 1997; Levi, 1978; Medin & Shoben, 1988; Murphy, 1990; Rips, Smith, & Shoben, 1978; Shoben, 1991; Smith & Osherson, 1984; Springer & Murphy, 1992). We begin here by

Figure 12. Schematic illustration of context effects in the proposed model. Neurobiological systems that represent and process conceptual attributes are shown in green, with font weights indicating relative amounts of processing (i.e., neural activity). Computations within these systems give rise to particular feature representations, shown in red. As a context is processed, this enhances neural activity in the subset of neural systems that represent the context, increasing the activation strength of the corresponding features of the concept. [To view this figure in colour, please see the online version of this Journal.]

attempting to define more precisely what kinds of issues a neural theory of conceptual combination might address. First among these is the general mechanism by which conceptual representations interact. This mechanism should explain how new meanings can "emerge" by combining individual concepts that have very different meanings. The concept *house*, for example, has few features in common with the concept *boat*, yet the combination *house boat* is nevertheless a meaningful and unique concept that emerges by combining these constituents. The second kind of issue that such a theory should address are the factors that cause variability in how particular features interact. The concepts *house boat* and *boat house* contain the same constituents but have entirely different meanings, indicating that constituent role (here modifier vs. head noun) is a major factor determining how concepts combine. Other combinations illustrate that there are specific semantic factors that cause variability in how constituents combine. For example, a *cancer therapy* is a therapy for cancer, but a *water therapy* is not a therapy for water. In our view it is not the task of a semantic theory to list and describe every possible such combination, but rather to propose general principles that govern underlying combinatorial processes.

As a relatively straightforward illustration of how neural attribute representations might interact during conceptual combination, and an illustration of the potential explanatory power of such a representation, we consider the classical feature *animacy*, which is a well-recognized determinant of the meaningfulness of modifier–noun and noun–verb combinations (as well as pronoun selection, word order, and aspects of morphology). Standard verbal feature-based models and semantic models in linguistics represent animacy as a binary feature. The difference in meaningfulness between (1) and (2) below:

(1) the angry boy
(2) *the angry chair

is "explained" by a rule that requires an animate argument for the modifier *angry*. Similarly, the difference in meaningfulness between (3) and (4) below:
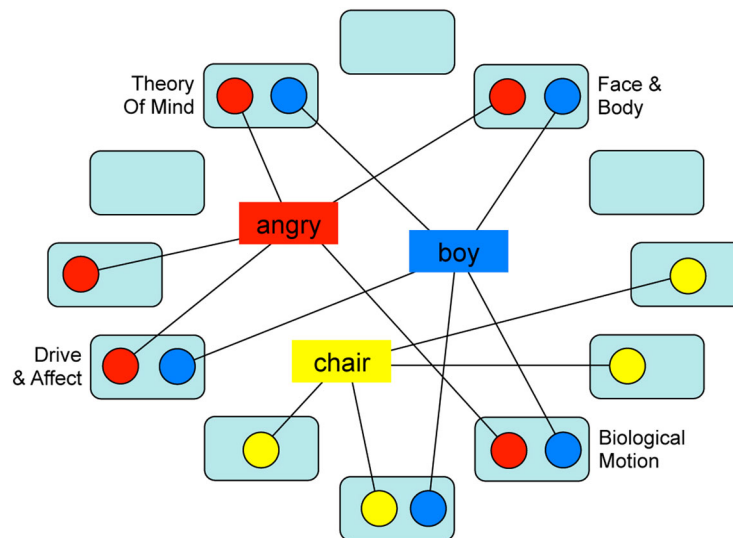
(1) the boy planned
(2) *the chair planned

is "explained" by a rule that requires an animate argument for the verb *plan*. The weakness of this kind of theoretical approach is that it amounts to little more than a very long list of specific rules that are in essence a restatement of the observed phenomena. Absent from such a theory are accounts of why particular concepts "require" animate arguments or even why particular concepts are animate. Also unexplained is the well-established "animacy hierarchy" observed within and across languages, in which adult humans are accorded a relatively higher value than infants and some animals, followed by lower animals, then plants, then non-living objects, then abstract categories (Yamamoto, 2006), suggesting rather strongly that animacy is a graded continuum rather than a binary feature.

From a developmental and neurobiological perspective, knowledge about animacy reflects a combination of various kinds of experience, including perception of biological motion, perception of causality, perception of one's own internal affective and drive states, and the development of theory of mind. Biological motion perception affords an opportunity to recognize from vision those things that move under their own power. Movements that are non-mechanical due to higher order complexity are simultaneously observed in other entities and in one's own movements, giving rise to an understanding (possibly mediated by "mirror" neurons) that these kinds of movements are executed by the moving object itself rather than by an external controlling force. Perception of causality, beginning with visual perception of simple collisions and applied force vectors, and progressing to multimodal and higher order events, provides a basis for understanding how biological movements can effect change. Perception of internal drive states and of sequences of events that lead to satisfaction of these needs provides a basis for understanding how living agents with needs effect changes. This understanding, together with the recognition that other living things in the environment effect changes to meet their own needs, provides a basis for theory of mind. On this account, the construct "animacy", far from being a binary symbolic property, arises from complex and interacting sensory, motor, affective, and cognitive experiences.

How might knowledge about animacy be represented and interact across lexical items at a neural level? As suggested schematically in Figure 13, we propose that interactions occur when two concepts

Figure 13. Schematic illustration of conceptual combination in the proposed model. Combination occurs when concepts share overlapping neural attribute representations. The concepts *angry* and *boy* share attributes that define animate concepts. [To view this figure in colour, please see the online version of this Journal.]

activate a similar set of modal networks, and they occur less extensively when there is less attribute overlap. In the case of animacy, for example, a noun that engages biological motion, face and body part, affective, and social cognition networks (e.g., "boy") will produce more extensive neural interaction with a modifier that also activates these networks (e.g., "angry") than will a noun that does not activate these networks (e.g., "chair").

To illustrate how such differences can arise even from a simple multiplicative interaction, we examined the vectors that result from multiplying mean attribute vectors of modifier–noun pairs with varying degrees of meaningfulness. Figure 14 (bottom) shows these interaction values for the pairs "angry*boy" and "angry*chair". As shown in the figure, *boy* and *angry* interact as anticipated, showing enhanced values for attributes concerned with biological motion, face and body part perception, speech production, and social and human characteristics, whereas interactions between *chair* and *angry* are negligible. Averaging the interaction values across all attributes gives a mean interaction value of 3.26 for angry*boy and 0.83 for angry*chair.

Figure 15 shows the average of these mean interaction values for 116 nouns, divided by taxonomic category. The averages roughly approximate the "animacy hierarchy", with strong interactions for nouns that refer to people (*boy, girl, man, woman, family, couple*, etc.) and human occupation roles

(*doctor, lawyer, politician, teacher*, etc.), much weaker interactions for animal concepts, and the weakest interactions for plants.

The general principle suggested by such data is that concepts acquired within similar neural processing domains are more likely to have similar semantic structures that allow combination. In the case of "animacy" (which we interpret as a verbal label summarizing a particular pattern of attribute weightings, rather than an a priori binary feature), a common semantic structure captures shared "human-related" attributes that allow certain concepts to be meaningfully combined. A chair cannot be angry because chairs do not have movements, faces, or minds that are essential for feeling and expressing anger, and this observation applies to all concepts that lack these attributes. The power of a comprehensive attribute representation constrained by neurobiological and cognitive developmental data is that it has the potential to provide a first-principles account of *why* concepts vary in animacy, as well as a mechanism for predicting which concepts should "require" animate arguments.

We have focused here on animacy because it is a strong and at the same time a relatively complex and poorly understood factor influencing conceptual combination. Similar principles might conceivably be applied, however, in other difficult cases. For the *cancer therapy* versus *water therapy* example given above, the difference in meaning that results from the two combinations is probably due to the fact
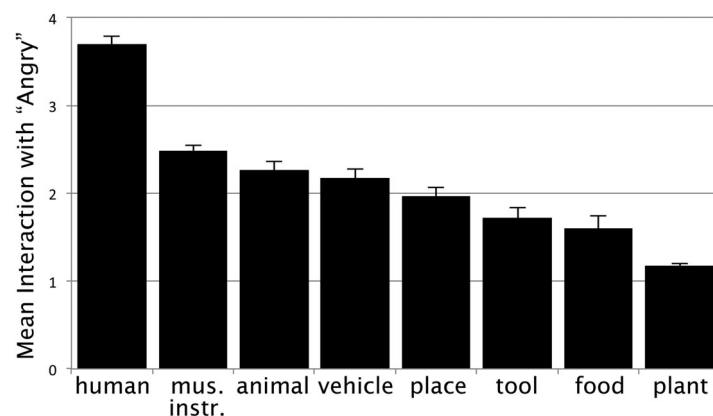
Figure 14. Top: Mean attribute ratings for *boy*, *chair*, and *angry*. Bottom: Interaction values for the combinations *angry boy* and *angry chair*. [To view this figure in colour, please see the online version of this Journal.]

that cancer is a negatively valenced concept, whereas water and therapy are usually thought of as beneficial. This difference results in very different constituent relations (therapy *for* vs. therapy *with*). A similar example is the difference between *lake house* and *dog house*. A lake house is a house located at a lake, but a dog house is not a house located at a dog. This difference reflects the fact that both houses and lakes have fixed locations (i.e., do not move and can be used as topographic landmarks), whereas this is not true of dogs. The general principle in these cases is that the meaning of a combination is strongly determined by attribute congruence. When Concepts A and B have opposite congruency relations with Concept C,

the combinations AC and BC are likely to capture very different constituent relationships. The content of these relationships reflects the underlying attribute structure of the constituents, as demonstrated by the locative relationship *located at* arising from the combination *lake house* but not *dog house*.

At the neural level, interactions during conceptual combination are likely to take the form of additional processing within the neural systems where attribute representations overlap. This additional processing is necessary to compute the "new" attribute representations resulting from conceptual combination, and these new representations tend to have added salience. As a simple example involving unimodal

Figure 15. Average mean interaction values for combinations of *angry* with a noun, for eight noun categories. Interactions are higher for human concepts (people and occupation roles) than for any of the non-human concepts (all $p < .001$). Error bars represent standard error of the mean.

modifiers, imagine what happens to the neural representation of *dog* following the modifier *brown* or the modifier *loud*. In the first case, there is residual activation of the colour processor by *brown*, which, when combined with the neural representation of *dog*, causes additional computation (i.e., additional neural activity) in the colour processor because of the interaction between the colour representations of *brown* and *dog*. In the second case, there is residual activation of the auditory processor by *loud*, which, when combined with the neural representation of *dog*, causes additional computation in the auditory processor because of the interaction between the auditory representations of *loud* and *dog*. As mentioned in the previous section on context effects, a similar mechanism is proposed (along with attentional modulation) to underlie situational context effects, because the context situation also has a conceptual representation. Attributes of the target concept that are "relevant" to the context are those that overlap with the conceptual representation of the context, and this overlap results in additional processor-specific activation. Seen in this way, situational context effects (e.g., lifting vs. playing a piano) are essentially a special case of conceptual combination.

Scope and limitations of the theory

We have made a systematic attempt to relate semantic content to large-scale brain networks and biologically plausible accounts of concept acquisition. The approach offers potential solutions to several long-standing problems in semantic representation, particularly the problems of feature selection, feature weighting, and specification of abstract word content. The primary aim of the theory is to better understand lexical semantic representation in the brain and, specifically, to develop a more comprehensive theory of conceptual grounding that encompasses the full range of human experience.

Although we have proposed several general mechanisms that may explain context and combinatorial effects, much more work is needed to achieve an account of word combinations and syntactic effects on semantic composition. Our simple attribute-based account of why the locative relation AT arises from *lake house*, for example, does not explain why *house lake* fails despite the fact that both nouns have fixed locations. A plausible explanation is that the syntactic roles of the constituents specify a head–modifier = ground–figure isomorphism that requires the modifier concept to be larger in size than the head noun concept. In principle, this behaviour could be captured by combining the size attribute ratings for the nouns with information about their respective syntactic roles. Previous studies have shown such effects to vary widely in terms of the types of relationships that arise between constituents and the "rules" that govern their lawful combination (Clark & Berman, 1987; Downing, 1977; Gagné, 2001; Gagné & Murphy, 1996; Gagné & Shoben, 1997; Levi, 1978; Medin & Shoben, 1988; Murphy, 1990). We assume that many other attributes could be involved in similar interactions.

The explanatory power of a semantic representation model that refers only to "types of experience" (i.e., one that does not incorporate all possible verbally generated features) is still unknown. It is far from clear, for example, that an intuitively basic distinction like *male/female* can be captured by such a representation. The model proposes that the concepts *male* and *female,* as applied to human adults, can be distinguished on the basis of sensory, affective, and social experiences without reference to specific encyclopaedic features. This account appears to fail, however, when *male* and *female* are applied to animals, plants, or electronic connectors, in which case there is little or no overlap in the experiences associated with these entities that could explain what is male or female about all of them. Such cases illustrate the fact that much of conceptual knowledge is learned directly via language and with only very indirect grounding in experience. In common with other embodied cognition theories, however, our model maintains that all concepts are ultimately grounded in experience, whether directly or indirectly through language that refers to experience. Establishing the validity, utility, and limitations of this principle, however, will require further study.

The underlying research materials for this article can be accessed at http://www.neuro.mcw.edu/resources.html.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

References

Albright, T. D., Desimone, R., & Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque. *Journal of Neurophysiology*, *51*, 16–31.

Allison, T., Puce, A., & McCarthy, G. (2000). Social perception from visual cues: Role of the STS region. *Trends in Cognitive Sciences*, *4*, 267–278.

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: Churchill Livingstone.

Altarriba, J., & Bauer, L. M. (2004). The distinctiveness of emotion concepts: A comparison between emotion, abstract, and concrete words. *The American Journal of Psychology*, *117*, 389–410.

Amorapanth, P. X., Widick, P., & Chatterjee, A. (2010). The neural basis for spatial relations. *Journal of Cognitive Neuroscience*, *22*, 1739–1753.

Anders, S., Eippert, F., Weiskopf, N., & Veit, R. (2008). The human amygdala is sensitive to the valence of pictures and sounds irrespective of arousal: An fMRI study. *Social, Cognitve and Affective Neuroscience*, *3*, 233–243.

Anders, S., Lotze, M., Erb, M., Grodd, W., & Birbaumer, N. (2004). Brain activity underlying emotional valence and arousal: A response-related fMRI study. *Human Brain Mapping*, *23*, 200–209.

Andersen, R. A., Snyder, L. H., Bradley, D. C., & Xing, J. (1997). Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual Review of Neuroscience*, *20*, 303–330.

Araujo, H. F., Kaplan, J., & Damasio, A. (2013). Cortical midline structures and autobiographical-self processes: An activation-likelihood estimation meta-analysis. *Frontiers in Human Neuroscience*, *7*, 548.

Aristotle. (1995/350 BCE). Categories (J. L. Ackrill, Trans.). In J. Barnes (Ed.), *The complete works of Aristotle* (pp. 3–24). Princeton: Princeton University Press.

Baayen, R. H., Piepenbrock, R., & Gulikers, L. (1995). *The CELEX lexical database (CD-ROM)* (2.5 ed.). Philadelphia: Linguistic Data Consortium, University of Pennsylvania.

Badre, D., & D'Esposito, M. (2007). Functional magnetic resonance imaging evidence for a hierarchical organization in the prefrontal cortex. *Journal of Cognitive Neuroscience*, *19*, 2082–2099.

Barlow, H. B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology. *Perception*, *1*, 371–394.

Barrett, L. F. (2006). Are emotions natural kinds? *Perspectives on Psychological Science*, *1*, 28–58.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory and Cognition*, *10*, 82–93.

Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, *11*, 211–227.

Barsalou, L. W. (1991). Deriving categories to achieve goals. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. *27*, pp. 1–64). San Diego, CA: Academic Press.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 577–660.

Barsalou, L. W., & Wiemer-Hastings, K. (2005). Situating abstract concepts. In D. Pecher & R. Zwaan (Eds.), *Grounding cognition: The role of perception and action in memory, language, and thought* (pp. 129–163). New York: Cambridge University Press.

Bartels, A., & Zeki, S. M. (2000). The architecture of the colour centre in the human visual brain: New results and a review. *European Journal of Neuroscience*, *12*, 172–193.

Baucom, L. B., Wedell, D. H., Wang, J., Blitzer, D. N., & Shinkareva, S. V. (2012). Decoding the neural representation of affective states. *Neuroimage*, *59*, 718–727.

Beauchamp, M. S., Haxby, J. V., Jennings, J. E., & DeYoe, E. A. (1999). An fMRI version of the Farnsworth-Munsell 100-hue test reveals multiple color-sensitive areas in human ventral occipitotemporal cortex. *Cerebral Cortex*, *9*, 257–263.

Belin, P., Zatorre, R. J., Lafaille, P., Ahad, P., & Pike, B. (2000). Voice-selective areas in human auditory cortex. *Nature*, *403*, 309–312.

Berlin, B., & Kay, P. (1969). *Basic color terms: Their universality and evolution*. Berkeley: University of California Press.

Binder, J. R. (2007). Effects of word imageability on semantic access: Neuroimaging studies. In J. Hart & M. A. Kraut (Eds.), *Neural basis of semantic memory* (pp. 149–181). Cambridge, UK: Cambridge University Press.

Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, *15*, 527–536.

Binder, J. R., Desai, R., Conant, L. L., & Graves, W. W. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, *19*, 2767–2796.

Bird, H., Franklin, S., & Howard, D. (2001). Age of acquisition and imageability ratings for a large set of words, including verbs and function words. *Behavior Research Methods, Instruments, & Computers*, *33*, 73–79.

Blos, J., Chatterjee, A., Kircher, T., & Straube, B. (2012). Neural correlates of causality judgment in physical and social context: The reversed effects of space and time. *Neuroimage*, *63*, 882–893.

Borghi, A. M., Flumini, A., Cimatti, F., Marocco, D., & Scorolli, C. (2011). Manipulating objects and telling words: a study on concrete and abstract words acquisition. *Frontiers in Psychology*, *2*, 15.

Boronat, C. B., Buxbaum, L. J., Coslett, H. B., Tang, K., Saffran, E. M., Kimberg, D. Y., & Detre, J. A. (2005). Distinctions between manipulation and function knowledge of objects: Evidence from functional magnetic resonance imaging. *Cognitive Brain Research*, *23*, 361–373.

Bowers, J. S. (2009). On the biological plausibility of grandmother cells: Implications for neural network theories in psychology and neuroscience. *Psychological Review*, *116*, 220–251.

Bretherton, I., & Beeghly, M. (1982). Talking about internal states: The acquisition of an explicit theory of mind. *Developmental Psychology*, *18*, 906–921.

Burgess, N. (2008). Spatial cognition and the brain. *Annals of the New York Academy of Sciences*, *1124*, 77–97.

Calvo-Merino, B., Glaser, D. E., Grezes, J., Passingham, R. E., & Haggard, P. (2005). Action observation and acquired motor skills: An fMRI study with expert dancers. *Cerebral Cortex*, *15*, 1243–1249.

Capitani, E., Laiacona, M., Mahon, B., & Caramazza, A. (2003). What are the facts of semantic category-specific deficits? A critical review of the clinical evidence. *Cognitive Neuropsychology*, *20*, 213–261.

Carota, F., Moseley, R., & Pulvermüller, F. (2012). Body part-specific representations of semantic noun categories. *Journal of Cognitive Neuroscience*, *24*, 1492–1509.

Carter, R. M., & Huettel, S. A. (2013). A nexus model of the temporal-parietal junction. *Trends in Cognitive Sciences*, *17*, 328–336.

Chow, H. M., Mar, R. A., Xu, Y., Liu, S., Wagage, S., & Braun, A. R. (2013). Embodied comprehension of stories : Interactions between language regions and modality-specific neural systems. *Journal of Cognitive Neuroscience*, *26*, 279–295.

Clark, E., & Berman, R. (1987). Types of linguistic knowledge: Interpreting and producing compound nouns. *Journal of Child Language*, *14*, 547–567.

Clark, J. M., & Paivio, A. (2004). Extensions of the Paivio, Yuille, and Madigan (1968) norms. *Behavior Research Methods, Instruments, & Computers*, *36*, 371–383.

Colibazzi, T., Posner, J., Wang, Z., Gorman, D., Gerber, A., Yu, S., … Peterson, B. S. (2010). Neural systems subserving valence and arousal during the experience of induced emotions. *Emotion*, *10*, 377–389.

Conrad, F., & Rips, L. (1986). Conceptual combination and the given-new distinction. *Journal of Memory and Language*, *25*, 255–278.

Conway, B. R., & Tsao, D. Y. (2009). Color-tuned neurons are spatially clustered according to color preference within alert macaque posterior inferior temporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *106*, 18034–18039.

Cortese, M. J., & Fugett, A. (2004). Imageability ratings for 3,000 monosyllabic words. *Behavior Research Methods, Instruments, and Computers*, *36*, 384–387.

Coslett, H. B., Saffran, E. M., & Schwoebel, J. (2002). Knowledge of the body: A distinct semantic domain. *Neurology*, *59*, 359–363.

Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, *132*, 163–201.

Crutch, S. J., Troche, J., Reilly, J., & Ridgway, G. R. (2013). Abstract conceptual feature ratings: the role of emotion, magnitude, and other cognitive domains in the organization of abstract conceptual knowledge. *Frontiers in Human Neuroscience*, *7*, Article 186.

Crutch, S. J., & Warrington, E. K. (2005). Abstract and concrete concepts have structurally different representational frameworks. *Brain*, *128*, 615–627.

Crutch, S. J., Williams, P., Ridgway, G. R., & Borgenicht, L. (2012). The role of polarity in antonym and synonym conceptual knowledge: Evidence from stroke aphasia and multidimensional ratings of abstract words. *Neuropsychologia*, *50*, 2636–2644.

Culham, J. C., Brandt, S. A., Cavanagh, P., Kanwisher, N. G., Dale, A. M., & Tootell, R. B. (1998). Cortical fMRI activation produced by attentive tracking of moving targets. *Journal of Neurophysiology*, *80*, 2657–2670.

Cunningham, W. A., Raye, C. L., & Johnson, M. K. (2004). Implicit and explicit evaluation: fMRI correlates of valence, emotional intensity, and control in the processing of attitudes. *Journal of Cognitive Neuroscience*, *16*, 1717–1729.

Dehaene, S. (1997). *The number sense: How the mind creates mathematics*. New York: Oxford University Press.

Denny, B. T., Kober, H., Wager, T. D., & Ochsner, K. N. (2012). A meta-analysis of functional neuroimaging studies of self- and other judgments reveals a spatial gradient for mentalizing in medial prefrontal cortex. *Journal of Cognitive Neuroscience*, *24*, 1742–1752.

Derrfuss, J., Brass, M., Neumann, J., & von Cramon, D. Y. (2005). Involvement of the inferior frontal junction in cognitive control: Meta-analyses of switching and Stroop studies. *Human Brain Mapping*, *25*, 22–34.

Desai, R., Binder, J. R., Conant, L. L., & Seidenberg, M. S. (2009). Activation of sensory-motor and visual areas by sentence comprehension. *Cerebral Cortex*, *20*, 468–478.

Diekhof, E. K., Kaps, L., Falkai, P., & Gruber, O. (2012). The role of the human ventral striatum and the medial orbitofrontal cortex in the representation of reward magnitude—An activation likelihood estimation meta-analysis of neuroimaging studies of passive reward expectancy and outcome processing. *Neuropsychologia*, *50*, 1252–1266.

Downing, P. (1977). On the creation and use of English compound nouns. *Language*, *53*, 810–842.

Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, *293*, 2470–2473.

Duncan, J., & Owen, A. M. (2000). Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends in Neurosciences*, *23*, 475–483.

Eisenberger, N. I. (2012). The neural bases of social pain: Evidence for shared representations with physical pain. *Psychosomatic Medicine*, *74*, 126–135.

Ekman, P. (1992). An argument for basic emotions. *Cognition and Emotion*, *6*, 169–200.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*, 598–601.

Epstein, R. A., Parker, W. E., & Feiler, A. M. (2007). Where am I now? Distinct roles for parahippocampal and retrosplenial cortices in place recognition. *Journal of Neuroscience*, *27*, 6141–6149.

Etkin, A., Egner, T., & Kalisch, R. (2011). Emotional processing in anterior cingulate and medial prefrontal cortex. *Trends in Cognitive Sciences*, *15*, 85–93.

Evans, V. (2004). The structure of time: Language, meaning and temporal cognition. *Amsterdam: John Benjamins*.

Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of Experimental Psychology: General*, 120, 339–357.

Farmer, T. A., Christiansen, M. H., & Monaghan, P. (2006). Phonological typicality influences on-line sentence comprehension. *Proceedings of the National Academy of Sciences USA*, 103, 12203–12208.

Fernandino, L., Binder, J. R., Desai, R. H., Pendl, S. L., Humphries, C. J., Gross, W., … Seidenberg, M. S. (2015). Concept representation reflects multimodal abstraction: A framework for embodied semantics. *Cerebral Cortex*, published online March 5, 2015. doi:10.1093/cercor/bhv020.

Ferstl, E. C., & von Cramon, D. Y. (2007). Time, space and emotion: fMRI reveals content-specific activation during text comprehension. *Neuroscience Letters*, 427, 159–164.

Ferstl, E. C., Rinck, M., & von Cramon, D. Y. (2005). Emotional and temporal aspects of situation model processing during text comprehension: An event-related fMRI study. *Journal of Cognitive Neuroscience*, 17, 724–739.

Fonlupt, P. (2003). Perception and judgement of physical causality involve different brain structures. *Cognitive Brain Research*, 17, 248–254.

Forde, E. M. E., & Humphreys, G. W. (1999). Category-specific recognition impairments: a review of important case studies and influential theories. *Aphasiology*, 13, 169–193.

Friebel, U., Eickhoff, S. B., & Lotze, M. (2011). Coordinate-based meta-analysis of experimentally induced and chronic persistent neuropathic pain. *Neuroimage*, 58, 1070–1080.

Fugelsang, J. A., Roser, M. E., Corballis, P. M., Gazzaniga, M. S., & Dunbar, K. N. (2005). Brain mechanisms underlying perceptual causality. *Cognitive Brain Research*, 24, 41–47.

Gagné, C. L. (2001). Relation and lexical priming during the interpretation of noun-noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 236–254.

Gagné, C. L., & Murphy, G. L. (1996). Influence of discourse context on feature availability in conceptual combination. *Discourse Processes*, 22, 79–101.

Gagné, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 71–87.

Gainotti, G. (2000). What the locus of brain lesion tells us about the nature of the cognitive defect underlying category-specific disorders: A review. *Cortex*, 36, 539–559.

Gainotti, G., Ciaraffa, F., Silveri, M. C., & Marra, C. (2009). Mental representation of normal subjects about the sources of knowledge in different semantic categories and unique entities. *Neuropsychology*, 23, 803–812.

Gainotti, G., Spinelli, P., Scaricamazza, E., & Marra, C. (2013). The evaluation of sources of knowledge underlying different conceptual categories. *Frontiers in Human Neuroscience*, 7, Article 40.

Garrard, P., Lambon Ralph, M. A., Hodges, J. R., & Patterson, K. (2001). Prototypicality, distinctiveness and intercorrelation: Analyses of the semantic attributes of living and nonliving concepts. *Journal of Cognitive Neuroscience*, 18, 125–174.

Gerber, A. J., Posner, J., Gorman, D., Colibazzi, T., Yu, S., Wang, Z., … Peterson, B. S. (2008). An affective circumplex model of neural systems subserving valence, arousal, and cognitive overlay during the appraisal of emotional faces. *Neuropsychologia*, 46, 2129–2139.

Glasgow, K., Roos, M., Haufler, A., Chevillet, M., & Wolmetz, M. (2016). Evaluating semantic models with word-sentence relatedness. arXiv:1603.07253. Retrieved from http://arxiv.org/abs/1603.07253 [cs.CL]

Goldberg, R. F., Perfetti, C. A., & Schneider, W. (2006). Perceptual knowledge retrieval activates sensory brain regions. *Journal of Neuroscience*, 26, 4917–4921.

González, J., Barros-Loscertales, A., Pulvermüller, F., Meseguer, V., Sanjuán, A., Belloch, V., & Ávila, C. (2006). Reading cinnamon activates olfactory brain regions. *Neuroimage*, 32, 906–912.

Graziano, M. S., Yap, G. S., & Gross, C. G. (1994). Coding of visual space by premotor neurons. *Science*, 266, 1054–1057.

Grill-Spector, K., & Malach, R. (2004). The human visual cortex. *Annual Review of Neuroscience*, 27, 649–677.

Grondin, S. (2010). Timing and time perception: A review of recent behavioral and neuroscience findings and theoretical directions. *Attention, Perception, and Psychophysics*, 72, 561–582.

Grossman, E. D., & Blake, R. (2002). Brain areas active during visual perception of biological motion. *Neuron*, 35, 1167–1175.

Hamann, S. (2012). Mapping discrete and dimensional emotions onto the brain: controversies and consensus. *Trends in Cognitive Sciences*, 16, 458–466.

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

Harvey, B. M., Klein, B. P., Petridou, N., & Dumoulin, S. O. (2013). Topographic representation of numerosity in the human parietal cortex. *Science*, 6150, 1123–1128.

Hasson, U., Levy, I., Behrmann, M., Hendler, T., & Malach, R. (2002). Eccentricity bias as an organizing principle for human high-order object areas. *Neuron*, 34, 479–490.

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

Hendry, S. H. C., Hsiao, S. S., & Bushnell, M. C. (1999). Somatic sensation. In M. J. Zigmond, F. E. Bloom, S. C. Landis, J. L. Roberts, & L. R. Squire (Eds.), *Fundamental neuroscience* (pp. 761–789). San Diego: Academic Press.

Hoffman, P., & Lambon Ralph, M. A. (2013). Shapes, scents and sounds: Quantifying the full multi-sensory basis of conceptual knowledge. *Neuropsychologia*, 51, 14–25.

Horn, J. (1965). A rationale and test for the number of factors in factor analysis. *Psychometrika*, 30, 179–185.

Hsu, N. S., Frankland, S. M., & Thompson-Schill, S. L. (2012). Chromaticity of color perception and object color knowledge. *Neuropsychologia*, 50, 327–333.

Hsu, N. S., Kraemer, D., Oliver, R., Schlichting, M. L., & Thompson-Schill, S. L. (2011). Color, context, and cognitive style: Variations in color knowledge retrieval as a function of task and subject variables. *Journal of Cognitive Neuroscience*, 23, 1–14.

Jackendoff, R. (1990). *Semantic structures*. Cambridge, MA: MIT Press.

Jäncke, L., Koeneke, S., Hoppe, A., Rominger, C., & Hänggi, J. (2009). The architecture of the golfer's brain. *PLoS ONE*, 4, e4785.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17, 4302–4311.

Kassam, K. S., Markey, A. R., Cherkassky, V. L., Loewenstein, G., & Just, M. A. (2013). Identifying emotions on the basis of neural activation. *PLoS One*, 8, e66032–e66032.

Keil, F. (1991). The emergence of theoretical beliefs as constraints on concepts. In S. C. Gelman & R. Gelman (Eds.), *The epigenesis of mind: Essays on biology and cognition* (pp. 237–256). Hillsdale, NJ: Erlbaum.

Kellenbach, M. L., Brett, M., & Patterson, K. (2001). Large, colourful or noisy? Attribute- and modality-specific activations during retrieval of perceptual attribute knowledge. *Cognitive, Affective, and Behavioral Neuroscience*, 1, 207–221.

Kelly, M. H. (1992). Using sound to solve syntactic problems: The role of phonology in grammatical category assignments. *Psychological Review*, 99, 349–364.

Kemmerer, D. (2006). The semantics of space: Integrating linguistic typology and cognitive neuroscience. *Neuropsychologia*, 44, 1607–1621.

Kemmerer, D., & Tranel, D. (2008). Searching for the elusive neural substrates of body part terms: A neuropsychological study. *Cognitive Neuropsychology*, 25, 601–629.

Kiefer, M., & Pulvermüller, F. (2012). Conceptual representations in mind and brain: Theoretical developments, current evidence and future directions. *Cortex*, 48, 805–825.

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, 12224–12230.

Kober, H., Barrett, L. F., Joseph, J., Bliss-Moreau, E., Lindquist, K., & Wager, T. D. (2008). Functional grouping and cortical-subcortical interactions in emotion: A meta-analysis of neuroimaging studies. *Neuroimage*, 42, 998–1031.

Konkle, T., & Oliva, A. (2012). A real-world size organization of object responses in occipitotemporal cortex. *Neuron*, 74, 1114–1124.

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, 14–34.

Kragel, P. A., & LaBar, K. S. (2014). Advancing emotion theory with multivariate pattern classification. *Emotion Review*, 6, 160–174.

Kranjec, A., Cardillo, E. R., Schmidt, G. L., Lehet, M., & Chatterjee, A. (2012). Deconstructing events: The neural bases for space, time, and causality. *Journal of Cognitive Neuroscience*, 24, 1–16.

Kranjec, A., & Chatterjee, A. (2010). Are temporal concepts embodied? A challenge for cognitive neuroscience. *Frontiers in Psychology*, 1, 1–9.

Kuchinke, L., Jacobs, A. M., Grubich, C., Vo, M. L. H., Conrad, M., & Herrmann, M. (2005). Incidental effects of emotional valence in single word processing: An fMRI study. *Neuroimage*, 28, 1022–1032.

Kuiper, N. A., & Rogers, T. B. (1979). Encoding of personal information: self-other differences. *Journal of Personality and Social Psychology*, 37, 499–514.

Laiacona, M., Allamano, N., Lorenzi, L., & Capitani, E. (2006). A case of impaired naming and knowledge of body parts: Are limbs a separate category? *Neurocase*, 12, 307–316.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.

Lamm, C., Decety, J., & Singer, T. (2011). Meta-analytic evidence for common and distinct neural networks associated with directly experienced pain and empathy for pain. *Neuroimage*, 54, 2492–2502.

Landau, B., & Jackendoff, R. (1993). What and where in spatial language and spatial cognition. *Behavioral and Brain Sciences*, 16, 217–238.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211–240.

Landauer, T. K., Kintsch, W. & the Science and Applications of Latent Semantic Analysis (SALSA) Group (2015). *LSA applications*. Boulder, CO: University of Colorado. Retrieved from http://lsa.colorado.edu/

Leaver, A. M., & Rauschecker, J. P. (2010). Cortical representation of natural complex sounds: Effects of acoustic features and auditory object category. *Journal of Neuroscience*, 30, 7604–7612.

Lench, H. C., Flores, S. a., & Bench, S. W. (2011). Discrete emotions predict changes in cognition, judgment, experience, behavior, and physiology: A meta-analysis of experimental emotion elicitations. *Psychological Bulletin*, 137, 834–855.

Levi, J. (1978). *The syntax and semantics of complex nominals*. New York: Academic Press.

Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, 22, 1027–1038.

Lewis, J. W., Wightman, F. L., Brefczynski, J. A., Phinney, R. E., Binder, J. R., & DeYoe, E. A. (2004). Human brain regions involved in recognizing environmental sounds. *Cerebral Cortex*, 14, 1008–1021.

Lewis, P. A., Critchley, H. D., Rotshtein, P., & Dolan, R. J. (2007). Neural correlates of processing valence and arousal in affective words. *Cerebral Cortex*, 17, 742–748.

Liebenthal, E., Binder, J. R., Spitzer, S. M., Possing, E. T., & Medler, D. A. (2005). Neural substrates of phonemic perception. *Cerebral Cortex*, 15, 1621–1631.

Lindquist, K. A., Siegel, E. H., Quigley, K. S., & Barrett, L. F. (2013). The hundred-year emotion war: are emotions natural kinds

or psychological constructions? Comment on Lench, Flores, and Bench (2011). *Psychological Bulletin*, 139, 255–263.

Lindquist, K. A., Wager, T. D., Kober, H., Bliss-Moreau, E., & Barrett, L. F. (2012). The brain basis of emotion: A meta-analytic review. *Behavioral and Brain Sciences*, 35, 121–143.

Lingnau, A., Ashida, H., Wall, M. B., & Smith, A. T. (2009). Speed encoding in human visual cortex revealed by fMRI adaptation. *Journal of Vision*, 9, 3.

Longo, M., Azanon, E., & Haggard, P. (2010). More than skin deep: Body representation beyond primary somatosensory cortex. *Neuropsychologia*, 48, 655–668.

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

Lynott, D., & Connell, L. (2013). Modality exclusivity norms for 400 nouns: The relationship between perceptual experience and surface word form. *Behavior Research Methods*, 45, 516–526.

Maguire, E. A., Frith, C. D., Burgess, N., Donnett, J. G., & O'Keefe, J. (1998). Knowing where things are: Parahippocampal involvement in encoding object locations in virtual large-scale space. *Journal of Cognitive Neuroscience*, 10, 61–76.

Makin, T. R., Holmes, N. P., & Zohary, E. (2007). Is that near my hand? Multisensory representation of peripersonal space in human intraparietal sulcus. *Journal of Neuroscience*, 27, 731–740.

Malach, R., Reppas, J. B., Benson, R. R., Kwong, K. K., Jiang, H., Kennedy, W. A., … Tootell, R. B. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proceedings of the National Academy of Sciences USA*, 92, 8135–8139.

Martin, A., & Caramazza, A. (2003). Neuropsychological and neuroimaging perspectives on conceptual knowledge: An introduction. *Cognitive Neuropsychology*, 20, 195–212.

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.

Mazlow, A. H. (1943). A theory of human motivation. *Psychological Review*, 50, 370–396.

Mazzola, L., Faillenot, I., Barral, F. G., Mauguiere, F., & Peyron, R. (2012). Spatial segregation of somato-sensory and pain activations in the human operculo-insular cortex. *Neuroimage*, 60, 409–418.

McGlone, F., & Reilly, D. (2010). The cutaneous sensory system. *Neuroscience and Biobehavioral Reviews*, 34, 148–159.

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

McRae, K., de Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, 126, 99–130.

Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, 20, 158–190.

Medina, J., & Coslett, H. B. (2010). From maps to form to space: Touch and the body schema. *Neuropsychologia*, 48, 645–654.

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

Minda, J. P., & Smith, J. D. (2002). Comparing prototype-based and exemplar-based accounts of category learning and attentional allocation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 275–292.

Miquée, A., Xerri, C., Rainville, C., Anton, J. L., Nazarian, B., Roth, M., & Zennou-Azogui, Y. (2008). Neuronal substrates of haptic shape encoding and matching: A functional magnetic resonance imaging study. *Neuroscience*, 152, 29–39.

Murphy, G. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language*, 29, 259–288.

Murphy, G., & Medin, D. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289–316.

Nakic, M., Smith, B. W., Busis, S., Vythilingam, M., & Blair, R. J. R. (2006). The impact of affect and frequency on lexical decision: The role of the amygdala and inferior frontal cortex. *Neuroimage*, 31, 1752–1761.

Nielen, M. M. A., Heslenfeld, D. J., Heinen, K., Van Strien, J. W., Witter, M. P., Jonker, C., & Veltman, D. J. (2009). Distinct brain systems underlie the processing of valence and arousal of affective pictures. *Brain and Cognition*, 71, 387–396.

Noordzij, M. L., Neggers, S. F. W., Ramsey, N. F., & Postma, A. (2008). Neural correlates of locative prepositions. *Neuropsychologia*, 46, 1576–1580.

Northoff, G., Heinzel, A., de Greck, M., Bermpohl, F., Dobrowolny, H., & Panksepp, J. (2006). Self-referential processing in our brain–a meta-analysis of imaging studies on the self. *Neuroimage*, 31, 440–457.

Nosofsky, R. M. (1986). Attention, similiarity and the identification-categorization relationship. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 115, 39–57.

Nyberg, L., Marklund, P., Persson, J., Cabeza, R., Forkstam, C., Petersson, K., & Ingvar, M. (2003). Common prefrontal activations during working memory, episodic memory, and semantic memory. *Neuropsychologia*, 41, 371–377.

O'Connor, B. P. (2000). SPSS and SAS programs for determining the number of components using parallel analysis and Velicer's MAP test. *Behavior Research Methods, Instruments, & Computers*, 32, 396–402.

Olson, I. R., McCoy, D., Klobusicky, E., & Ross, L. A. (2013). Social cognition and the anterior temporal lobes: A review and theoretical framework. *Social Cognitive & Affective Neuroscience*, 8, 123–133.

Peelen, M. V., Atkinson, A. P., & Vuilleumier, P. (2010). Supramodal representations of perceived emotions in the human brain. *Journal of Neuroscience*, 30, 10127–10134.

Pelphrey, K. A., Morris, J. P., Michelich, C. R., Allison, T., & McCarthy, G. (2005). Functional anatomy of biological motion perception in posterior temporal cortex: An fMRI study of eye, mouth and hand movements. *Cerebral Cortex*, 15, 1866–1876.

Peters, J., & Büchel, C. (2010). Neural representations of subjective reward value. *Behavioural Brain Research*, 213, 135–141.

Phan, K. L., Wager, T., Taylor, S. F., & Liberzon, I. (2002). Functional neuroanatomy of emotion: A meta-analysis of emotion activation studies in PET and fMRI. *Neuroimage*, 16, 331–348.

Piazza, M., Izard, V., Pinel, P., Bihan, D. L., & Dehaene, S. (2004). Tuning curves for approximate numerosity in the human intraparietal sulcus. *Neuron*, 44, 547–555.

Posner, J., Russell, J. A., Gerber, A., Gorman, D., Colibazzi, T., Yu, S., … Peterson, B. S. (2009). The neurophysiological bases of emotion: An fMRI study of the affective circumplex using emotion-denoting words. *Human Brain Mapping*, 30, 883–895.

Posner, J., Russell, J. A., & Peterson, B. S. (2005). The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and Psychopathology*, 17, 715–734.

Postle, N., McMahon, K. L., Ashton, R., Meredith, M., & de Zubicaray, G. I. (2008). Action word meaning representations in cytoarchitectonically defined primary and premotor cortices. *Neuroimage*, 43, 634–644.

Rees, G., Friston, K. J., & Koch, C. (2000). A direct quantitative relationship between the functional properties of human and macaque V5. *Nature Neuroscience*, 3, 716–723.

Rips, L., Smith, E., & Shoben, E. (1978). Semantic composition in sentence verification. *Journal of Verbal Learning and Verbal Behavior*, 17, 375–401.

Rogers, T. B., Kuiper, N. A., & Kirker, W. S. (1977). Self-reference and the encoding of personal information. *Journal of Personality and Social Psychology*, 35, 677–688.

Röhl, M., & Uppenkamp, S. (2012). Neural coding of sound intensity and loudness in the human auditory system. *Journal of the Association for Research in Otolaryngology*, 13, 369–379.

Rolls, E. T., & Grabenhorst, F. (2008). The orbitofrontal cortex and beyond: From affect to decision-making. *Progress in Neurobiology*, 86, 216–244.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, D. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382–439.

Ross, L. A., & Olson, I. R. (2010). Social cognition and the anterior temporal lobes. *Neuroimage*, 49, 3452–3462.

Rueschemeyer, S.-A., Pfeiffer, C., & Bekkering, H. (2010). Body schematics: On the role of the body schema in embodied lexical-semantic representations. *Neuropsychologia*, 48, 774–781.

Russell, J. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, 39, 1161–1178.

Said, C. P., Moore, C. D., Engell, A. D., & Haxby, J. V. (2010). Distributed representations of dynamic facial expressions in the superior temporal sulcus. *Journal of Vision*, 10, 1–12.

Satpute, A. B., Fenker, D. B., Waldmann, M. R., Tabibnia, G., Holyoak, K. J., & Lieberman, M. D. (2005). An fMRI study of causal judgments. *European Journal of Neuroscience*, 22, 1233–1238.

Saxe, R., & Wexler, A. (2005). Making sense of another mind: The role of the right temporo-parietal junction. *Neuropsychologia*, 43, 1391–1399.

Schilbach, L., Bzdok, D., Timmermans, B., Fox, P. T., Laird, A. R., Vogeley, K., & Eickhoff, S. B. (2012). Introspective minds: Using ALE meta-analyses to study commonalities in the neural correlates of emotional processing, social & unconstrained cognition. *PLoS One*, 7, e30920-e30920.

Schmidtke, D. S., Conrad, M., & Jacobs, A. M. (2014). Phonological iconicity. *Frontiers in Psychology*, 5, Article 80.

Schreiner, C. E., & Winer, J. A. (2007). Auditory cortex mapmaking: Principles, projections, and plasticity. *Neuron*, 56, 356–365.

Schurz, M., Radua, J., Aichhorn, M., Richlan, F., & Perner, J. (2014). Fractionating theory of mind: A meta-analysis of functional brain imaging studies. *Neuroscience and Biobehavioral Reviews*, 42, 9–34.

Schwanenflugel, P. (1991). Why are abstract concepts hard to understand? In P. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 223–250). Hillsdale, NJ: Erlbaum.

Schwarzlose, R. F., Baker, C. I., & Kanwisher, N. (2005). Separate face and body selectivity on the fusiform gyrus. *Journal of Neuroscience*, 25, 11055–11059.

Schwoebel, J., & Coslett, H. B. (2005). Evidence for multiple, distinct representations of the human body. *Journal of Cognitive Neuroscience*, 17, 543–553.

Serino, A., & Haggard, P. (2010). Touch and the body. *Neuroscience and Biobehavioral Reviews*, 34, 224–236.

Shaoul, C., & Westbury, C. (2013). A reduced redundancy USENET corpus (2005–2011). Edmonton, AB: University of Alberta. Retrieved from http://www.psych.ualberta.ca/~westburylab/downloads/usenetcorpus.download.html

Shoben, E. (1991). Predicating and nonpredicating combinations. In P. J. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 117–135). Hillsdale, NJ: Erlbaum.

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, 2802–2810.

Simmons, W. K., Reddish, M., Bellgowan, P. S., & Martin, A. (2010). The selectivity and functional connectivity of the anterior temporal lobes. *Cerebral Cortex*, 20, 813–825.

Smith, E. E., & Osherson, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive Science*, 8, 337–361.

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 174–215.

Spreng, R. N., Mar, R. A., & Kim, A. S. N. (2009). The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: A quantitative meta-analysis. *Journal of Cognitive Neuroscience*, 21, 489–510.

Springer, K., & Murphy, G. (1992). Feature availability in conceptual combination. *Psychological Science*, 3, 111–117.

Talmy, L. (1983). How language structures space. In H. Pick & L. Acredolo (Eds.), *Spatial orientation: Theory, research, and application* (pp. 225–282). New York: Plenum Press.

Tanaka, K., Saito, H., Fukada, Y., & Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *Journal of Neurophysiology*, 66, 170–189.

Tettamanti, M., Buccino, G., Saccuman, M. C., Gallese, V., Danna, M., Scifo, P., … Perani, D. (2005). Listening to action-related sentences activates fronto-parietal motor cicuits. *Journal of Cognitive Neuroscience*, 17, 273–281.

Tootell, R. B., Reppas, J. B., Kwong, K. K., Malach, R., Born, R. T., Brady, T. J., … Belliveau, J. W. (1995). Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *Journal of Neuroscience*, *15*, 3215–3230.

Tracy, J. L., & Randles, D. (2011). Four models of basic emotions: a review of Ekman and Cordaro, Izard, Levenson, and Panksepp and Watt. *Emotion Review*, *3*, 397–405.

Tranel, D., & Kemmerer, D. (2004). Neuroanatomical correlates of locative prepositions. *Cognitive Neuropsychology*, *21*, 719–749.

Tranel, D., Logan, C. G., Frank, R. J., & Damasio, A. R. (1997). Explaining category-related effects in the retrieval of conceptual and lexical knowledge for concrete entities: Operationalization and analysis of factors. *Neuropsychologia*, *35*, 1329–1339.

Troche, J., Crutch, S., & Reilly, J. (2014). Clustering, hierarchical organization, and the topography of abstract and concrete nouns. *Frontiers in Psychology*, *5*, Article 360.

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*, 474–486.

Van Overwalle, F. (2009). Social cognition and the brain: A meta-analysis. *Human Brain Mapping*, *30*, 829–858.

van Veluw, S. J., & Chance, S. A. (2014). Differentiating between self and others: An ALE meta-analysis of fMRI studies of self-recognition and theory of mind. *Brain Imaging and Behavior*, *8*, 24–38.

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

Vigliocco, G., Vinson, D. P., Lewis, W., & Garrett, M. F. (2004). Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive Psychology*, *48*, 422–488.

Vinson, D. P., Vigliocco, G., Cappa, S., & Siri, S. (2003). The breakdown of semantic knowledge: Insights from a statistical model of meaning representation. *Brain and Language*, *86*, 347–365.

Vytal, K., & Hamann, S. (2010). Neuroimaging support for discrete neural correlates of basic emotions: A voxel-based meta-analysis. *Journal of Cognitive Neuroscience*, *22*, 2864–2885.

Wallentin, M., Østergaarda, S., Lund, T. E., Østergaard, L., & Roepstorff, A. (2005). Concrete spatial language: See what I mean? *Brain and Language*, *92*, 221–233.

Warrington, E. K., & McCarthy, R. A. (1987). Categories of knowledge. *Further fractionations and an attempted integration. Brain*, *110*, 1273–1296.

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*, 829–853.

Wiemer-Hastings, K., & Xu, X. (2005). Content differences for abstract and concrete concepts. *Cognitive Science*, *29*, 719–736.

Wilson, M. D. (1988). The MRC Psycholinguistic Database: Machine Readable Dictionary, Version 2. *Behavior Research Methods, Instruments, & Computers*, *20*, 6–10.

Wilson-Mendenhall, C. D., Barrett, L. F., & Barsalou, L. W. (2013). Neural evidence that human emotions share core affective properties. *Psychological Science*, *24*, 947–956.

Woods, A. J., Hamilton, R. H., Kranjec, A., Minhaus, P., Bikson, M., Yu, J., & Chatterjee, A. (2014). Space, time, and causality in the human brain. *Neuroimage*, *92*, 285–297.

Wu, D. H., Morganti, A., & Chatterjee, A. (2008). Neural substrates of processing path and manner information of a moving event. *Neuropsychologia*, *46*, 704–713.

Wu, D. H., Waller, S., & Chatterjee, A. (2007). The functional neuroanatomy of thematic role and locative relational knowledge. *Journal of Cognitive Neuroscience*, *19*, 1542–1555.

Yamamoto, M. (2006). *Agency and impersonality: Their linguistic and cultural manifestations*. Amsterdam: J. Benjamins Pub. Co.

Zahn, R., Moll, J., Krueger, F., Huey, E. D., Garrido, G., & Grafman, J. (2007). Social concepts are represented in the superior anterior temporal cortex. *Proceedings of the National Academy of Sciences USA*, *104*, 6430–6435.

Zatorre, R. J., Belin, P., & Penhune, V. B. (2002). Structure and function of auditory cortex: Music and speech. *Trends in Cognitive Sciences*, *6*, 37–46.

Zdrazilova, L., & Pexman, P. M. (2013). Grasping the invisible: semantic processing of abstract words. *Psychonomic Bulletin and Review*, *20*, 1312–1318.

Zeki, S. M. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *The Journal of Physiology*, *236*, 549–573.