

Eiling Yee

Fluid semantics: Semantic knowledge is experience-based and dynamic

Abstract: Is our internal notion of, e. g., the object *lemon*, static? That is, do we have stable semantic representations that remain constant across time? Most semantic memory researchers still (at least tacitly) take a static perspective, assuming that only effects that can be demonstrated across a variety of tasks and contexts should be considered informative about the architecture of the semantic system. This chapter challenges this perspective by highlighting studies showing that the cognitive and neural representations of object concepts are fluid, changing as a consequence of the context that each individual brings with them (e. g., via current goals, recent experience, long-term experience, or neural degeneration). These findings support models of semantic memory in which rather than being static, conceptual representations are dynamic and shaped by experience, whether that experience extends over the lifetime, the task, or the moment.

Introduction

Lemon. We can consider its shape, approximate size, color, taste, texture, weight, etc. How is this information organized in the semantic/conceptual system? How does information about one concept relate to information about other concepts? How does an object concept's representation map onto the corresponding real world object? And what are some of the ways in which representations might change over time?

At first glance, conceptual representations (such as our internal notion of the object *lemon*) seem static. That is, we have the impression that there is something that lemon means (a sour, yellow, rugby ball-shaped, citrus fruit) and that

Eiling Yee, Department of Psychological Sciences, University of Connecticut; Basque Center on Cognition, Brain & Language

DOI 10.1515/9783110422658-012

this meaning does not vary. Research in semantic memory traditionally takes this static perspective. Yet, according to more recent accounts of semantic/conceptual knowledge, knowledge of object properties (e. g., color, shape and smell) is distributed (in part) across brain regions that underlie sensory and motor processing (e. g. Allport, 1985; Barsalou, 1999; Damasio, 1989), and concepts are multi-dimensional representations across those substrates. Thus, according to these “sensorimotor-based” distributed accounts, the “meaning” of a lemon is not an indivisible whole, but is distributed across a range of featural dimensions (cf. McRae, de Sa & Seidenberg, 1997; Rogers & McClelland, 2004; Tyler & Moss, 2001; Vigliocco, Vinson, Levis & Garrett, 2004).

In this chapter, we focus on some of the things that such distributed, sensorimotor-based accounts of semantic representations allow for. For example, sensorimotor-based models predict that sensory and motor features are among those that make up semantic representations, and they make predictions about the organization of semantic representations. Specifically, if a brain area that perceives a given feature is the same region that represents it, then if two concepts share a perceptual feature, their representations must overlap in that brain region. Also, if representations are distributed, this means that an entire concept would not have to be activated at once—its various semantic features (e. g., its shape, or the purpose for which it is used) could be activated at different rates or, be more activated in some circumstances than others. Furthermore, if representations are sensorimotor-based, then disrupting or interfering with a sensorimotor brain area could interfere with conceptual access.

Although some of the predictions that sensorimotor-based distributed models make about how we represent and access semantic knowledge may seem, on the surface, surprising, in the sections that follow, we describe evidence supporting these predictions. In section one, we suggest that the semantic features over which objects are represented can include not only more “abstract” features, but also sensory and motor features. We also argue that different features may have different time courses of activation. In section two, we describe evidence that semantic representations of object concepts can overlap in sensorimotor and multimodal cortices. In the third section, we argue that for object concepts, sensorimotor activity is part of (rather than peripheral to) their semantic representations. In the fourth and fifth sections, we suggest that prior experience and individual abilities modulate semantic access. Finally, we conclude that semantic representations are much more fluid than they may seem at first glance.

1 Which semantic features are included in object representations, and what is the time course over which these features are activated?

Since distributed models of semantic memory assume that concepts that share features have overlapping patterns of representation, activating a particular concept should also partially activate other concepts that share its features. The semantic priming effect, wherein identifying a target word is facilitated when it is preceded by a (conceptually) related prime word (e. g., Meyer & Schvaneveldt, 1971), can therefore be interpreted as support for distributed models. However, semantically related objects are often category coordinates, which are related in multiple ways, e. g., *crayon* and *pencil* are both thin, oblong, used for marking paper, and grasped with the thumb and the second and third fingers (cf. Kellenbach, Wijers & Mulder, 2000). As a result, unless the features of overlap are examined separately, it is not clear which are responsible for the facilitation.

Identifying which features are responsible for the facilitation has implications for theories of semantic memory. Specifically, if words (referring to concepts) that are related via *sensorimotor* features partially activate each other, this would suggest that such features constitute part of the representation of concepts. An increasing number of semantic priming studies have explicitly manipulated the semantic relationship between primes and targets. For example, in an auditory semantic priming paradigm, Myung, Blumstein & Sedivy (2006) observed that lexical decisions on target words were speeded if the prime was an object that was manipulated similarly for use (e. g., *key* primes *screwdriver* because using each involves twisting the wrist; we discuss this study further in section 3).

Other semantic priming studies have explored whether semantic priming is obtained when primes and targets have the same shape (e. g., coin-button), or are related via a more abstract dimension similar to function/purpose of use (e. g., apple-banana, or stapler-paperclip). Broadly speaking, priming has been observed for both shape similarity (Schreuder, Flores D'Arcais, & Glazenborg, 1984; Flores d'Arcais, Schreuder, & Glazenbor., 1985; Pecher, Zeelenberg and Raaijmakers, 1998; Taylor, 2005) and function similarity (Schreuder et al., 1984; Flores d'Arcais et al., 1985; Taylor, 2005). However, findings across experiments have been varied, and one explanation put forth for the differences is that different features may become active at different times during semantic activation (Schreuder et al., 1984). Below we describe studies using the “visual world”

eyetracking paradigm (Cooper, 1974; Tanenhaus, Spivey-Knowlton, Eberhard & Sedivy, 1995) that suggest that this is indeed the case.

For instance, we have found evidence that during visual object identification, information about the form of an object (e. g., that knives are oblong) becomes available sooner than information about its function (e. g., that they are used for cutting; Yee Huffstetler & Thompson-Schill, 2011). Specifically, in a visual world eyetracking study we found that when participants were briefly exposed (for 1 second) to an array of four objects and asked to click on the object corresponding to a heard word, they were sensitive to the fact that one of the other objects in the display was similar in shape (at the conceptual level) to that of the named object. For example, when they heard the target word “Frisbee”, they looked at a slice of pizza (another object that can be round). Importantly, shape similarity was not apparent in the visual depictions (e. g., a slice of pizza is triangular, a shape that a Frisbee cannot take, see Figure 1); hence, preferential fixations on the shape-related object were attributable to activation of conceptual shape information (and not to the current input to the senses).

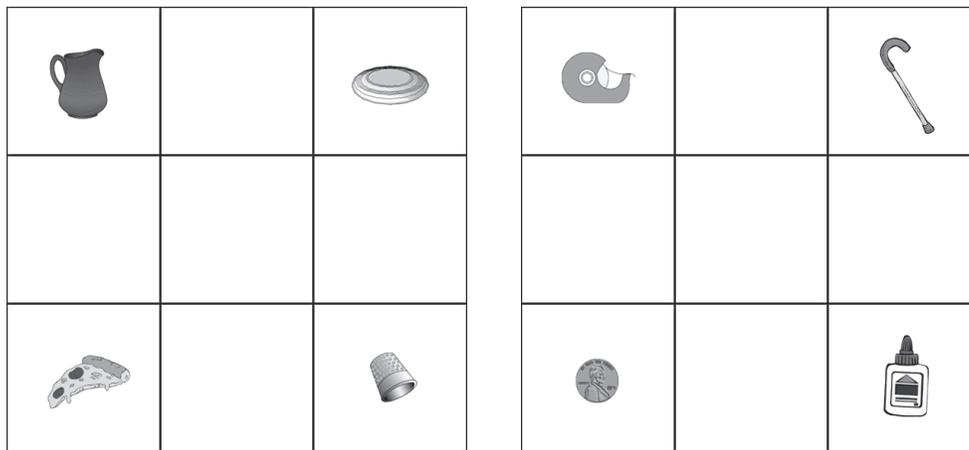


Figure 1: Shape (left panel) and function (right panel) related pairs from Yee et al. (2011).

However with the same 1-second exposure to the array, we found no preferential fixations on objects related in function to the named object. For example, when participants heard the word “tape”, they did not preferentially fixate on a bottle of glue. Yet when exposure to the array was lengthened to 2 seconds, we observed the opposite pattern—participants preferentially fixated on the function-related object, but not the shape-related object. Thus, these findings suggest that semantic activation during visual object recognition is a dynamically unfolding process in which function follows form.

Other studies using the visual world paradigm have also found that function information becomes active after other, arguably more perceptually grounded, aspects of conceptual knowledge. For instance, knowledge about the thematic relationships that an object can participate in (i. e., the knowledge that a broom is often paired with/seen with a dustpan, or that a steak is paired/seen with a knife), and knowledge about an object's structural characteristics (its shape, size and volume) appear to become available more rapidly and more transiently than information about its function (see Kalénine, Mirman, Middleton & Buxbaum, 2012 for thematic relationships and Lee, Middleton, Mirman, Kalénine & Buxbaum, 2013 for structural relationships).

Future work will be needed to learn the cause(s) of this chronology—in which not only does the activation of function information follow the activation of form, but form information seems to be more transiently activated.¹ With respect to function information becoming active later, one possibility is that because information about the purpose for which an object is used is not directly available via the senses, accessing this information requires more processing (and hence more time) than accessing perceptual information. Another possibility is that in the studies reviewed above, perceptual information has special status because the visual world paradigm task—essentially a visual search task—*requires* attending to perceptual information. A third (compatible) possibility is that presenting stimuli in the visual modality (regardless of whether the task is visual search) places emphasis on visual information.

Studies using written words provide a hint that the modality of the stimulus, rather than the visual search task *per se*, contributes to the earlier activation of perceptual information. For example, when reading sentences referring to objects (e. g., a calculator), information that is directly available from perceiving the object (e. g., the grasp that one would use to pick up a calculator to move it) becomes less active over time, while information that requires more abstract knowledge about the purpose of the object (e. g., the finger poke that one would use to operate a calculator) becomes more active over time (Bub & Masson, 2010).

Along the same lines, the semantic priming studies referred to at the beginning of this section that did obtain evidence of shape priming all used written words (Flores d'Arcais, et al., 1985; Schreuder et al., 1984; Pecher et al., Experiments 4 and 6). Moreover, as alluded to earlier, they also found evidence that priming based on function emerges more reliably at long, rather than short interstimulus

¹ One might speculate that form information rapidly decays (or is inhibited) in favor of function information because although form is needed for object recognition, once the object is recognized, other attributes, such as what it is used for, are typically more relevant.

intervals, whereas priming based on shape relatedness was larger at short than long interstimulus intervals (Flores d'Arcais et al., 1985; Schreuder et al., 1984). In other words, the chronology parallels what has been found in the visual world paradigm, even for a task that does not require visual search.

Although these findings rule out visual search *per se* as the sole reason that perceptual features become active earlier than more abstract features, is it possible that perceptual information becomes active early because both written words and visual objects place emphasis on the visual modality? Or does perceptual information become active first during conceptual activation, regardless of context? As we describe next, when stimuli are not visually presented, perceptual information does not always become active first, suggesting that stimulus modality (e. g., visual vs. auditory) does indeed affect the dynamics of featural activation.

In a cross-modal semantic priming study in which primes were auditory words (targets were written), priming for visually related targets was observed, but it was delayed in time relative to priming for targets related to the typical use of the prime² (Moss, McCormick & Tyler, 1997). Another study that presented object names auditorily obtained analogous results: Information that is directly available from viewing the object (here, information about how one grasps an object to move it) became active later (and for a shorter duration) than information about how an object is manipulated in order to use it – i. e., its function (Bub & Masson, 2012). Thus, in contrast to when primes are visual, when the primes (or the sole presented words) are auditory, activation of perceptual information can appear later than functional information (see Garcea & Mahon, 2012 for related work).

Hence, the ebb and flow of different features seems to be influenced by the relationship between the modality of the stimulus and the specific feature: When the stimulus is presented visually, visual features may become active earlier than more abstract knowledge. In contrast, when the stimulus is presented auditorily, the reverse pattern is observed.³ The dynamics of featural activation during

² Perceptually related targets were primarily visible parts of the prime or what the prime is typically made of (e. g., *blouse-button* or *sandal-leather*), while use-related targets typically denoted the primary purpose of the prime or the location in which the prime is used (e. g., *blouse-wear*; *satchel-school*; *radio-music*).

³ Task may also interact with which features become available when: Rogers and Patterson (2007) have shown that when the task is categorization (e. g., judging whether the word “canary” [or on other trials, “bird”, or “animal”] correctly identifies a picture of a canary), information that can distinguish among objects in the same semantic neighborhood (e. g., the property “yellow” distinguishes canaries from robins) becomes available later than more general information (information that does not help in distinguishing amount objects in the same neighborhood; e. g., the property “can fly” is shared by most birds).

semantic activation, therefore, appear to depend upon the modality (i. e., the context) in which the concept is presented, and may also depend on which features are typically most relevant. Importantly, the fact that some types of featural information become active earlier than others during semantic activation is consistent with distributed models of semantic representation because in these models different features can be differentially activated.

With respect to the “sensorimotor-based” claim of sensorimotor-based distributed models, it is important to note that priming has been observed for both perceptual (e. g., shape) and motor (e. g., manipulation) features, *as well as* more abstract features (e. g., function). While perceptual overlap is clearly predicted on sensorimotor models, such models do not require overlap on more abstract features (such as an object’s function/purpose of use) as they cannot be directly perceived via any individual sensory modality. As a result, a *purely* sensorimotor-based model would not be sufficient to accommodate these patterns. Instead, a model is required in which higher order similarity can be represented—perhaps by abstracting across similarity in the contexts in which things are used. We will return to this idea in the next section.

In sum, the results described in this section are consistent with distributed semantic representations that are at least partially sensorimotor-based, but that also include higher-order information, such as the purpose for which an object is used.

2 In which brain regions does representational overlap occur?

The semantic relatedness effects described in the prior section suggest that *some-where* in the brain, the representations of objects overlap such that they can partially activate one another, but they do not address the nature of this neural representation (e. g., *where* in the brain the representations overlap). In this section we turn to studies using functional magnetic resonance imaging (fMRI) to address this question.

A large body of fMRI studies have indicated that the different components of semantic knowledge about an object (e. g., its color, action or sound) activate neural systems that are close to, or overlap with those involved in perceiving those sensory features, or producing those actions (e. g., for color: Martin, Haxby, Lalonde, Wiggs & Ungerleider, 1995; Simmons et al., 2007 and Hsu, Kraemer, Oliver, Schlichting & Thompson-Schill, 2011; for action: Chao and Martin, 2000; for sounds: Kiefer, Sim, Herrnberger, Grothe & Hoenig, 2008).

Findings like these provide support for sensorimotor-based models because they are consistent with the idea that the brain regions that are active when we *perceive and interact with* an object are the same ones that represent it. And recall that, because of this, sensorimotor-based models predict that when two objects share a sensory feature such as shape, their representations would overlap in brain regions involved in perceiving shape. Similarly, if objects are manipulated similarly (e. g., *key* and *screwdriver*) their representations should overlap in brain regions involved in performing object-related actions (i. e., in the “dorsal stream”). In two fMRI studies, we examined the neural encoding of two sensorimotor-based features (shape and manipulation) as well as a more abstract feature (function).

The paradigm we used takes advantage of the fact that repeated presentation of the same visual or verbal stimulus results in reduced fMRI signal levels in brain regions that process that stimulus, either because of neuronal “fatigue” (e. g., firing-rate adaptation) or because the initial activation of a stimulus’ representation is less neurally efficient than its subsequent activation (see Grill-Spector, Henson & Martin, 2006 for a review). In a typical fMRI-adaptation experiment, stimuli are presented which are either identical (which produces an adaptation/reduced hemodynamic response) or completely different (producing a recovery response). However, using stimuli pairs that are semantically related, rather than identical, also produces an adaptation effect (e. g., Kotz, Cappa, von Cramon & Friederici, 2002; Rissman, Eliassen & Blumstein, 2003; Matsumoto, Iidaka, Haneda, Okada & Sadato, 2005; Bedny, McGill & Thompson-Schill, 2008). Hence, the paradigm’s sensitivity to similarity means that it can be used to detect which brain regions encode different conceptual features.

In our studies, we used an fMRI-adaptation paradigm to obtain a neural metric of similarity between objects. We found that in regions involved in guiding actions (i. e., premotor cortex and intraparietal sulcus), the degree of fMRI-adaptation to a pair of objects is correlated with the degree of similarity in the actions we use to interact with them. For example, a key and a screwdriver (which we use with similar hand and wrist motions) have similar representations in action regions, just as they should if representations are sensorimotor based. We also found several regions in which degree of adaptation is correlated with similarity in function. Two of these regions (medial temporal lobe and posterior middle temporal gyrus) are thought to be involved in integrating information; the activation in these regions is consistent with the idea that encoding more abstract information (such as an object’s function) requires integrating information from multiple modalities (Yee, Drucker & Thomson-Schill, 2010). Unexpectedly, we found no evidence at all of representational overlap based on shape similarity, even when performing exploratory analyses that did not use stringent corrections for multiple statistical comparisons.

Recalling that semantic priming and eye-tracking studies suggest that the activation of information about an object's form may be transient, we hypothesized that our null effect for shape similarity could be due to the timing of the presentation of stimuli. In particular, we speculated that shape information about the prime was no longer active by the time the target appeared. We therefore conducted a second fMRI-adaptation study using a shorter interstimulus interval (ISI). However, even with an ISI similar to that used in the semantic priming studies that did observe a priming effect for shape, we found no evidence of adaptation for shape, anywhere in the brain (Yee, Musz & Thompson-Schill, 2012). Nor did we observe behavioral priming for shape-related pairs (we did observe behavioral priming for manipulation-related pairs). One possible reason for this null result is that the task (concreteness judgment)⁴ focused attention away from shape information.

The fact that we did not observe an adaptation or priming effect for shape suggests an interesting potential difference between shape and manipulation information. Although neither shape *nor* manipulation information is required to perform a concreteness judgment, we observed evidence that manipulation knowledge was accessed, and no evidence that shape knowledge was. This may suggest that (at least for the objects that we tested) manipulation is a more routinely, or more strongly accessed feature of conceptual knowledge than is shape. One possibility is that this difference exists because shape is more context-dependent than manipulation; for example, the shape in which a book appears depends on whether, in a given instance, you are looking at it from the top or the side, and its shape can change depending upon whether it is open or closed. In contrast, how one manipulates an object for use tends to be more stable. The difference also highlights that different features can be activated independently, which provides more evidence that different features are separable components of distributed semantic knowledge.

These fMRI findings converge with the studies described in the first section to suggest that objects that are manipulated similarly, or that have similar functions have overlapping representations. Moreover, they provide information about *where* in the brain the representational overlap occurs. Specifically, objects that are manipulated similarly have overlapping representations in action regions, and objects that have similar functions have representational overlap in regions involved in integrating information. These findings suggest that semantic knowledge consists of both sensorimotor (e.g., manipulation) *and* abstract

⁴ Note that that in the behavioral studies described above, shape priming was most robust when the task was naming (Pecher et al., 1998).

(e. g., function) knowledge. Therefore although semantic memory may have sensorimotor information at its base, it must also incorporate a way for more abstract knowledge to be represented.

3 Is sensorimotor activity part of conceptual knowledge?

As described above, there is good evidence that semantic knowledge about object concepts is distributed across multimodal and sensorimotor brain regions. However, a frequently raised question that the studies described above do not address is whether the neural activity that is observed in sensorimotor regions during semantic access is *part* of the semantic knowledge being accessed, or rather, is incidental to it. That is, on some theoretical models, sensorimotor activity could, in principle, be incidental to the activation of an amodal (or ‘disembodied’) concept, rather than part of the concept (for discussion see Mahon & Caramazza 2008, Anderson & Spivey 2009, Chatterjee, 2010).

One often-cited fact that has been pointed to as evidence that such activations are only incidental is that there exist patients with motor or sensory deficits who, despite having difficulty performing, e. g., object-related actions, can retain the ability to name, and may also be able to describe the use of objects with strongly associated actions. These abilities have been taken as evidence that such individuals have intact conceptual knowledge and thus, that sensory or motor information is not part of conceptual knowledge (see Negri et al., 2007).

However, distributed models of semantic memory posit that conceptual representations include many different components (e. g., visual, auditory, and olfactory as well as action-oriented and multi-modal) that are distributed across cortex. Moreover, there is evidence that conceptual information is represented at multiple levels of abstraction, and consequently, depending upon the context, conceptual activation may involve the activation of some levels more than others (for discussion, see Thompson-Schill, 2003; Binder & Desai, 2011). For instance, in an fMRI study, Hsu, et al. (2011) asked participants to judge which of two objects a third object most resembled in color. When the three objects were all from the same color category (e. g., butter, egg yolk, and school bus – all are yellow), and so the task context required retrieving detailed color knowledge, the neural response overlapped more with brain regions involved in color perception than when two of the three objects were from different color categories (e. g., one red and two yellow objects) and therefore less detailed color knowledge was necessary. This finding implies that a task that requires a high degree of perceptual

resolution involves perceptual areas more than a task that can be performed on the basis of more categorical (or abstracted) knowledge.

For these reasons (as we and others—e. g., Taylor & Zwaan, 2009—have argued previously), having difficulty accessing part of a representation would not be expected to result in catastrophic conceptual loss (although depending upon the task, some impairment may be detectable). In other words, “brain damage leading to problems performing an action with a particular object does not entail difficulty recognizing that object . . . the object may be recognizable on the basis of other aspects of its representation (and the extent to which there are other aspects to rely upon may vary across individuals)” (Yee, Chrysikou, Hoffman, & Thompson-Schill, 2013, pp. 917–918).

In fact, not only may object recognition (e. g., of a typewriter) remain successful for an individual who has an impairment in e. g., accessing knowledge about object-directed action, but even access to information that is specifically related to the affected modality may not be *entirely* lost (e. g., access to knowledge about how one moves one’s fingers to press the keys of a typewriter may be delayed, but not lost). To illustrate, we have examined participants with ideomotor apraxia, which is a neurological impairment, typically caused by stroke that causes difficulty performing object-related actions. Our starting point was the finding that unimpaired participants exhibit priming for objects that are manipulated similarly—both standard auditory word priming (as described in section 1), as well as manipulation relatedness effects in the visual world paradigm (Myung, Blumstein & Sedivy, 2006). When we examined apraxic participants, we found that they had abnormally delayed access to manipulation information about objects (Myung, et al., 2010), and moreover, the amount of delay was correlated with *how much* difficulty they had performing object related actions (see also Lee, Mirman & Buxbaum, 2014). This finding shows that damage to a brain region supporting object-related action can hinder access knowledge about how objects are manipulated, even if that information is not entirely lost.

However, while our study with apraxic individuals demonstrated that problems performing object related actions can cause problems accessing manipulation knowledge about an object, it was not designed to test whether such manipulation knowledge is *part of* an object’s semantic representation. To answer the latter question, it is necessary to determine whether problems accessing manipulation knowledge about an object can interfere with thinking about that object more generally, e. g., even when access to manipulation information about it is not required. This is what we asked in a subsequent study. Specifically, we asked whether performing a concurrent manual task that is incompatible with how a given object is acted upon can interfere with thinking about that object. We found evidence that it can: If, while naming pictures, participants had to concurrently

perform an unrelated sequence of hand motions, picture naming was more disrupted for objects that are typically interacted with manually than objects that are less frequently interacted with manually (e. g., there was relatively more interference for pencils vs. tigers; Yee et al., 2013; see also Witt, Kemmerer, Linkenauger & Culham, 2010).

Thus, the context of a concurrent manual task interfered with people's ability to think about objects that are frequently manipulated. This demonstrates that 1) manipulation information is part of the representation of frequently manipulated objects, and 2) that our ability to think about a given object depends on the match between our mental representation of that object's meaning and what we are doing at the moment.

Moreover, because activity in motor areas influences semantic retrieval, findings such as this one suggest that motor area activity is *more than a* “peripheral” part of conceptual knowledge—it is *part* of conceptual knowledge. More broadly, these findings from unimpaired participants converge with patient and fMRI work in supporting the idea that activity in sensorimotor brain regions can be part of an object's concept.

4 Can what we have recently been doing affect semantic activation?

The findings described above also have another implication: They suggest that our ability to access semantic representations is dynamic in the sense that it can change depending on what exactly we may be doing at the moment. This raises another question. Can what we have *recently* been doing also affect how we access semantic knowledge?

We addressed this question using a different approach than in our studies examining manipulation knowledge. Rather than testing whether accessing a particular semantic feature can be made more difficult (e. g., via brain damage or a concurrent incompatible task), we instead asked whether it is possible to *enhance* our ability to access a concept (or, a particular aspect of a concept) through a recent activity. We used the semantic feature of color as a test case (Yee, Ahmed & Thomson-Schill 2012). Color is particularly interesting: In the visual object recognition literature, a point of contention has been whether color is, or is not a part of an object's representation (Biederman & Ju, 1988). There is evidence that color only becomes part of the representation of objects for which it is both consistent (e. g., lemons are normally yellow) *and* for which it is important for distinguishability (e. g., color is necessary for distinguishing lemons from limes; Tanaka & Presnell,

1999). The idea that consistency is important is clearly compatible with sensorimotor-based models—the more consistently associated a color is with an object, the more frequently it will be experienced in that color, and it is those experiences that cause color to become a part of the object’s representation. Moreover, if distinguishability is also important, this suggests that *attention* can influence the extent to which color becomes a part of an object’s representation—the idea being that the more important color is for distinguishing an object, the more that attention will be focused on color when that object is experienced, which will strengthen color’s involvement in its representation. An intriguing extension of this idea is that, for objects for which color is already part of the representation, attention may play a role in the extent to which color information is activated in a given episode.

Our examination of color suggests that its activation is indeed modulated by attention: We found that recent experience can influence the activation of color as a semantic feature in a subsequent, unrelated task. Specifically, we have found that although lemons and daffodils, for example, overlap on the dimension of color (both are yellow), and might therefore be expected to partially activate one another, the word “lemon” does not ordinarily activate, or prime, “daffodil”. However, it can prime “daffodil” if participants’ attention has been focused on color in a prior task involving unrelated items (e. g., color words in a Stroop task; Yee Ahmed & Thompson-Schill, 2012). Similar findings have been reported in other modalities. For example, Pecher, and colleagues (1998) observed shape priming (e. g., the word “coin” priming the word “button”) only when, prior to the priming experiment, participants made shape judgments about the objects to which the words referred. Thus, recent experience can linger long enough to affect conceptual activation in a subsequent, unrelated task.

Additional evidence that recent experience affects semantic activation comes from van Dantzig and colleagues (2008), who have shown that the modality to which attention is directed immediately prior to thinking about objects can affect conceptual activation: Between trials that required participants to make true-false judgments on sentences referring to object properties (e. g., *broccoli is green* or *soup is hot*), participants responded to either a visual light, an auditory noise, or a tactile vibration; property judgments were faster when the modality to which the sentence referred was the same as the preceding perceptual stimulus.

Along similar lines, but using a more implicit measure of conceptual activation, Bermeitinger, Wentura and Frings (2011) found that when an independent task directing attention to shape was interspersed with a semantic priming task, priming for words referring to natural kinds (for which shape is known to be a particularly important feature) was greater than priming for artifacts. In contrast, priming was greater for artifacts (for which action is known to be a particularly important feature) than for natural kinds when the interspersed task directed

attention to action. By inserting a task that directs attention to one modality or another, these two studies converge with our color-priming study (Yee et al., 2012) to show that directing attention to a particular modality changes subsequent conceptual activation such that information related to that modality is activated more easily. Several other behavioral studies have reported compatible results (e. g., Martens, Ansorge, & Kiefer, 2011; Van Dam, Rueschemeyer, Lindemann & Bekkering, 2010), and neural activation patterns are also consistent with this kind of attentionally modulated flexibility in semantic activation (Hoenig, Sim, Bochev, Herrnberger & Kiefer, 2008; Mummery, Patterson, Hodges, & Price, 1998; Phillips, Noppeney, Humphreys, Price, 2002; Thompson-Schill, Aguirre, D'Esposito, & Farah, 1999; Rogers, Hocking, Mechelli, Patterson, & Price, 2005; Van Dam, van Dijk, Bekkering & Rueschemeyer, 2012; for review, see Willems & Francken, 2012).

5 Is your lemon different from mine?

So far, we have described evidence supporting distributed, sensorimotor-based semantic representations. However, we have not yet considered (1) whether factors intrinsic to the individual, such as individual differences in cognitive abilities, might impact an individual's semantic processing in general, or (2) whether such factors might impact the extent to which a given individual's semantic activation is affected by what they are doing at the moment or what they have been recently doing. In this section, we speculate about some individual factors that we hypothesize may affect both semantic processing and the influence of concurrent and recent experience on such processing.

First, individual differences in processing preferences may impact conceptual activation: In the color-priming study described in the prior section, in which lemon only primed daffodil if participants' attention had been focused on color in a prior, ostensibly unrelated Stroop task, we also observed that individual differences in the ability to selectively focus on color in the Stroop task predicted the amount of priming (see Figure 2). This relationship could reflect differences in the degrees to which people attend to or perceive color (as in Hsu et al., 2011): Individuals who attend to color more would more strongly associate conceptual color with both color words in the Stroop task (e. g., green) and with the names of objects in the priming task (e. g., cucumber).

However, another compatible possibility relates to selective attention: The general ability to selectively attend to one dimension at the expense of others (e. g., to focus on a word's font color while ignoring its meaning) is an aspect of cognitive control (Posner & Snyder, 1975) and varies across individuals. Thus, a

high capacity for selective attention could manifest as enhanced selective attention to the features most relevant for the current task (in this case, judging animal status), and hence as less activation of other features (e. g., color).

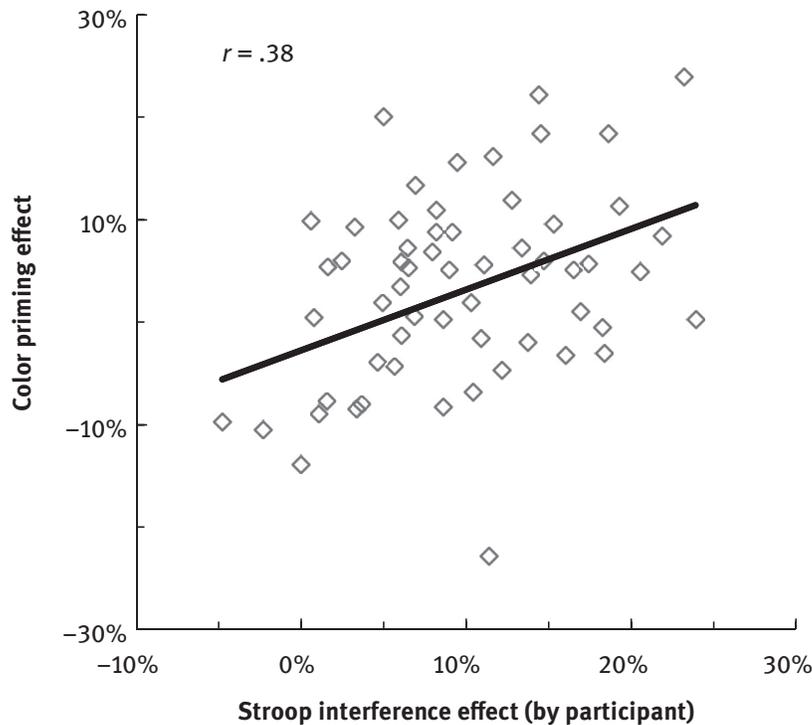


Figure 2: Scatter plot showing correlation between Stroop interference effect and color-priming effect (Yee et al., 2012).

This account is particularly interesting because it is consistent with the proposal that cognitive control regulates the ability to selectively attend to the task-relevant features of a concept in general (see Kan & Thompson-Schill, 2004a; Thompson-Schill, Bedny, & Goldberg, 2005). Evidence consistent with this account has come from a recent study demonstrating that inhibitory electrical stimulation over left prefrontal cortex (a brain region that supports cognitive control) interferes with the ability to categorize objects according to a specific attribute (e. g., “round or red things”) relative to categorizing objects at a more general level (e. g., “things that hold water”; Lupyan, Mirman, Hamilton, & Thompson-Schill, 2012; for related work, see also Chrysikou et al., 2013; Kan & Thompson-Schill, 2004b; Lupyan & Mirman, 2013).

If individual processing preferences or cognitive control abilities indeed affect the extent to which semantic activation is affected by context, this suggests that “meaning” must vary, not only from one context to another, but also within

a given context, from one person to another. Yet if this is true, then how is it that, in the face of different experiences, individuals are able to understand each other when they use conventionalized labels? In other words, if what is retrieved from semantic memory is so variable across time and individuals, then shouldn't communication be even more difficult than it is? Fortunately, there are significant commonalities in human experience (and especially within a given culture) that would lead different individuals' representations (and their labels) to be similar enough for most practical purposes. Moreover, often communication does not require that the interlocutors be activating the exact same conceptual representations (for further consideration of this point, see Casasanto & Lupyan, 2015; Connell & Lynott, 2014; Taylor & Zwaan, 2009), and in cases in which greater precision is desirable, communication often requires clarification (“no, no, no . . . that's not kale, it's chard!”).

Conclusions

Collectively, the studies reviewed above suggest that semantic knowledge about objects is instantiated as patterns of activation that are distributed across both sensorimotor and abstract features, with relationships between concepts captured by overlap in these patterns. Perhaps more surprisingly, they also suggest that semantic knowledge is fluid, changing not only as a function of our individual experiences with objects but, even more surprisingly, as a function of what we have recently been doing, and even as the process of object recognition unfolds. Thus, semantic memory is not static.

This fluidity highlights that semantic representations, like the word forms that refer to them, are highly context dependent. That is, just as the sequence of sounds that we produce to refer to something depends upon our long-term experiences (e. g., language, accent, whether we are toddlers or adults), short-term goals (e. g., register, emphasis) and the current articulatory context (e. g., surrounding words), so does the representation that is activated by those sounds.

That this kind of malleability exists in the mature semantic system has an important implication. It suggests that the same architecture that, in infancy, permits the semantic system to develop through interactions with the world,⁵

⁵ See Yee & Thompson-Schill (2016) for a discussion of one such architecture—an emergentist approach to cognitive development (e. g., Elman et al., 1996) in which context effects are a natural consequence of how concepts are acquired and represented in the first place, and in which change over time is naturally accommodated.

also allows the mature semantic system to be sensitive to, and change as a consequence of, the ever-richer contexts in which we, as humans, continue to develop. Thus, while semantics are fluid, there may be continuity across the lifespan in the underlying mechanisms of the system that allows for this fluidity.

References

- Allport, D. A. (1985). Distributed memory, modular subsystems and dysphasia. In S. K. Newman & R. Epstein (Eds.), *Current Perspectives in Dysphasia* (pp. 207–244). Edinburgh: Churchill Livingstone.
- Anderson, S. E., & Spivey, M. J. (2009). The enactment of language: Decades of interactions between linguistic and motor processes. *Language and Cognition*, 1, 87–111.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral Brain Science*, 22, 577–660.
- Bedny, M., McGill, M., & Thompson-Schill, S.L. (2008). Semantic adaptation and competition during word comprehension. *Cerebral Cortex* 18(11), 2574–2585.
- Bermeitinger, C., Wentura, D., & Frings, C. (2011). How to switch on and switch off semantic priming effects: Activation processes in category memory depend on focusing specific feature dimensions. *Psychonomic Bulletin & Review*, 18, 579–585.
- Biederman, I., & Ju, G. (1988). Surface versus edge-based determinants of visual recognition. *Cognitive psychology*, 20(1), 38–64.
- Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, 15, 527–536.
- Bub, D. N., & Masson, M. E. J. (2010). On the nature of hand-action representations evoked during written sentence comprehension. *Cognition*, 116, 394–408.
- Bub, D. N., & Masson, M. E. J. (2012). On the dynamics of action representations evoked by names of manipulable objects. *Journal of Experimental Psychology: General*, 141, 502–517.
- Casasanto, D., & Lupyan, G. (2015). All concepts are ad hoc concepts. In E. Margolis & S. Laurence (Eds.), *The conceptual mind: New directions in the study of concepts* (pp. 543–566). Cambridge: MIT Press.
- Chao, L.L., & Martin, A. (2000). Representation of manipulable man-made objects in the dorsal stream. *Neuroimage*, 12, 478–484.
- Chatterjee, A. (2010). Disembodying cognition. *Language and Cognition*, 2, 79–116.
- Chrysikou, E. G., Hamilton, R. H., Coslett, H. B., Datta, A., Bikson, M., & Thompson-Schill, S. L. (2013). Non-invasive transcranial direct current stimulation over the left prefrontal cortex facilitates cognitive flexibility in tool use. *Cognitive Neuroscience*, 4, 81–89.
- Connell, L. & Lynott, D. (2014). I see/hear what you mean: semantic activation in visual word recognition depends on perceptual attention. *Journal of Experimental Psychology: General*. 143(2), p. 527–533.
- Cooper, R. M. (1974). The control of eye fixation by the meaning of spoken language: A new methodology for the real-time investigation of speech perception, memory, and language processing. *Cognitive Psychology*, 6, 84–107
- Damasio, A. R. (1989). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation*, 1, 123–132.

- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press, Bradford Books.
- Flores d'Arcais, G. & Schreuder, R. & Glazenborg, G. (1985). Semantic activation during recognition or referential words. *Psychological Research* 47, 39–49.
- Garcea, F. E., & Mahon, B. Z. (2012). What is in a tool concept? Dissociating manipulation knowledge from function knowledge. *Memory & cognition*, 40(8), 1303–1313.
- Grill-Spector, K., Henson, R. & Martin, A., 2006. Repetition and the brain: neural models of stimulus-specific effects, *Trends in Cognitive Sciences* 10(1), 14–23.
- Hoenig, K., Sim, E.-J., Bochev, V., Herrnberger, B. & Kiefer, M. (2008). Conceptual flexibility in the human brain: Dynamic recruitment of semantic maps from visual, motion and motor-related areas. *Journal of Cognitive Neuroscience*, 20(10), 1799–1814.
- Hsu, N.S., Kraemer, D.J.M., Oliver, R.T., 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*, 29(3), 2544–2557.
- Kalénine, S., Mirman, D., Middleton, E. L., and Buxbaum, L. J. (2012). Temporal dynamics of activation of thematic and functional knowledge during conceptual processing of manipulable artifacts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(5), 1274–1295. DOI: 10.1037/a0027626.
- Kan, I. P. & Thompson-Schill, S. L. (2004a). Selection from perceptual and conceptual representations. *Cognitive, Affective, & Behavioral Neuroscience*, 4, 466–482. PMID: 15849891.
- Kan I. P., & Thompson-Schill, S. L. (2004b). Effect of name agreement on prefrontal activity during overt and covert picture naming. *Cognitive, Affective, & Behavioral Neuroscience*, 4, 43–57.
- Kellenbach, M. L., Wijers, A. A., & Mulder, G. (2000). Visual semantic features are activated during the processing of concrete words: Event-related potential evidence for perceptual semantic priming. *Cognitive Brain Research*, 10, 67–75.
- Kiefer, M., Sim, E.-J., Herrnberger, B., Grothe, J. & Hoenig, K. (2008). The sound of concepts for markers for a link between auditory and conceptual brain systems. *The Journal of Neuroscience*, 28, 12224–12230.
- Kotz, S.A., Cappa, S.F., von Cramon, D.Y., & Friederici, A.D., 2002. Modulation of the lexical-semantic network by auditory semantic priming: *An event-related functional MRI study*. *NeuroImage* 17, 1761–1772.
- Lee, C.-L., Middleton, E. L., Mirman, D., Kalénine, S., and Buxbaum, L. J. (2013). Incidental and context-responsive activation of structure- and function-based action features during object identification. *Journal of Experimental Psychology: Human Perception and Performance*, 39(1), 257–270. DOI: 10.1037/a0027533.
- Lee, C.-L., Mirman, D., and Buxbaum, L.J. (2014). Abnormal dynamics of activation of object use information in apraxia: Evidence from eyetracking. *Neuropsychologia*, 59, 13–26. DOI: 10.1016/j.neuropsychologia.2014.04.004.
- Lupyan, G. and Mirman, D. (2013). Linking language and categorization: Evidence from aphasia. *Cortex*, 49(5), 1187–1194. DOI: 10.1016/j.cortex.2012.06.006.
- Lupyan, G., Mirman, D., Hamilton, R., & Thompson-Schill, S. L. (2012). Categorization is modulated by transcranial direct current stimulation over left prefrontal cortex. *Cognition*, 124(1), 36–49.
- Mahon, B.Z. & Caramazza, A. (2008). A Critical Look at the Embodied Cognition Hypothesis & a New Proposal for Grounding Conceptual Content. *Journal of Physiology – Paris*, 102, 59–70.

- Martens, U., Ansorge, U., & Kiefer, M. (2011). Controlling the unconscious: Attentional task sets modulate subliminal semantic and visuo-motor processes differentially. *Psychological Science*, *22*, 282–291.
- 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.
- Matsumoto, A., Iidaka, T., Haneda, K., Okada, T., & Sadato, N., 2005. Linking semantic priming effect in functional MRI and event-related potentials. *NeuroImage* *24*(3), 624–634.
- 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.
- Meyer, D.E., & Schvaneveldt, R.W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*, 227–234.
- Moss, H.E., McCormick, S.F., & Tyler, L.K. (1997). The time course of activation of semantic information during spoken word recognition. *Language and Cognitive Processes*, *12*, 695–731.
- Mummary, C. J., Patterson, K., Hodges, J. R., & Price, C. J. (1998). Functional neuroanatomy of the semantic system: Divisible by what? *Journal of Cognitive Neuroscience*, *10*, 766–777.
- Myung, J. Y., Blumstein, S. E., & Sedivy, J. C. (2006). Playing on the typewriter, typing on the piano: manipulation knowledge of objects. *Cognition*, *98*(3), 223–243.
- Myung, J., Blumstein, S.E., Yee, E., Sedivy, J.C., Thompson-Schill, S.L., & Buxbaum, L.J. (2010). Impaired access to manipulation features in apraxia: Evidence from eyetracking and semantic judgment tasks. *Brain and Language*, *112*, 101–112.
- Negri, G. A. L., Rumiati, R. I., Zadini, A., Ukmar, M., Mahon, B. Z., & Caramazza, A. (2007). What is the role of motor simulation in action and object recognition? Evidence from apraxia. *Cognitive Neuropsychology*, *24*, 795–816.
- Pecher, D., Zeelenberg, R., & Raaijmakers, J. G. W. (1998). Does pizza prime coin? Perceptual priming in lexical decision and pronunciation. *Journal of Memory and Language*, *38*, 401–418. doi:10.1006/jmla.1997.2557
- Phillips JA, Noppeney U, Humphreys GW, Price CJ. (2002.) Can segregation within the semantic system account for category specific deficits? *Brain*. *125*, 2067–2080.
- Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 55–85). Hillsdale: Lawrence Erlbaum Associates.
- Rissman, J., Eliassen, J.C., & Blumstein, S.E., 2003. An event-related fMRI investigation of implicit semantic priming. *Journal of Cognitive Neuroscience* *15*(8), 1160–1175.
- Rogers, T. T. and McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press
- Rogers, T. T., & Patterson, K. (2007). Object categorization: Reversals and explanations of the basic-level advantage. *Journal of Experimental Psychology: General*, *136*, 451–469.
- Rogers, T. T., Hocking, J., Mechelli, A., Patterson, K. and Price, C. (2005). Fusiform activation to animals is driven by the process, not the stimulus. *Journal of Cognitive Neuroscience*, *17*, 434–445.
- Schreuder, R., Flores D'Arcais, G. B., & Glazenborg, G. (1984). Effects of perceptual and conceptual similarity in semantic priming. *Psychological Research*, *45*(4), 339–354.
- Simmons, W., Ramjee, V., Beauchamp, M., McRae, K., Martin, A., & Barsalou, L. (2007). A common neural substrate for perceiving and knowing about color. *Neuropsychologia*, *45*, 2802–2810.

- Tanaka, J. M., & Presnell, L.M. (1999). *Perception & Psychophysics*, *61*, 1140–1153.
- Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*, 1632–1634.
- Taylor, J.R. (2005). On The Perceptual Basis of Semantic Memory: Representation, process and attentional control revealed by behavior and event-related brain potentials. Unpublished Doctoral Dissertation.
- Taylor, L. J. & Zwaan, R. A. (2009). Action in cognition: The case of language. *Language and Cognition*, *1*, 45–58.
- Thompson-Schill, S. (2003). Neuroimaging studies of semantic memory: Inferring how from where. *Neuropsychologia*, *41*, 280–292.
- Thompson-Schill, S. L., Aguirre, G. K., D'Esposito, M., & Farah, M. J. (1999). A neural basis for category and modality specificity of semantic knowledge. *Neuropsychologia*, *37*, 671–676.
- Thompson-Schill, S. L., Bedny, M., & Goldberg, R.F. (2005). The frontal lobes and the regulation of mental activity. *Current Opinion in Neurobiology*, *15*, 219–224.
- Tyler, L.K., & Moss, H.E. (2001). Towards a distributed account of conceptual knowledge. *Trends in Cognitive Sciences*, *5*, 244–252.
- Van Dam, W. O., Rueschemeyer, S.-A., Lindemann, O., & Bekkering, H. (2010). Context effects in embodied lexical- semantic processing. *Frontiers in Psychology*, *1*, 150.
- van Dam, W. O., van Dijk, M., Bekkering, H., & Rueschemeyer, S.-A. (2012). Flexibility in embodied lexical-semantic representations. *Human Brain Mapping*, *33*, 2322–2333.
- van Dantzig, S., Pecher, D., Zeelenberg, R., & Barsalou, L.W. (2008). Perceptual processing affects conceptual processing. *Cognitive Science*, *32*, 579–590.
- 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.
- Willems, R. M. & Francken, J. C. (2012). Embodied Cognition: Taking the next step. *Frontiers in Psychology*, *3*, 582.
- Witt, J. K., Kemmerer, D., Linkenauger, S. A., & Culham, J. (2010). A functional role for motor simulation in naming tools. *Psychological Science*, *21*, 1215–1219.
- Yee, E., Ahmed, S., & Thompson-Schill, S.L. (2012). Colorless green ideas (can) prime furiously. *Psychological Science*, *23*(4), 364–369.
- Yee, E., Chrysikou, E., Hoffman, E., & Thompson-Schill, S.L. (2013). Manual Experience Shapes Object Representation. *Psychological Science*, *24* (6), 909–919.
- Yee, E., Drucker, D. M., & Thompson-Schill, S. L. (2010). fMRI-adaptation evidence of overlapping neural representations for objects related in function or manipulation. *NeuroImage*, *50*(2), 753–763. doi:10.1016/j.neuroimage.2009.12.036
- Yee, E., Hufstetler, S., & Thompson-Schill, S.L. (2011). Function follows form: Activation of shape and function features during object identification. *Journal of Experimental Psychology: General*, *140*, 348–363. doi: 10.1037/a0022840
- Yee, E., Musz, E., & Thompson-Schill, S.L. Mapping the Similarity Space of Concepts in Sensorimotor Cortex. *Poster presented at the Neurobiology of Language Conference*, San Sebastian. October, 2012.
- Yee, E., & Thompson-Schill, S.L. (2016). Putting concepts into context. *Psychonomic Bulletin and Review*, *23*(4), 1015–1027.