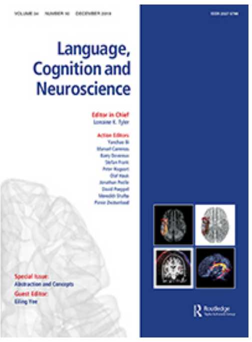

Language, Cognition and Neuroscience

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

Abstraction and concepts: when, how, where, what and why?

Eiling Yee

To cite this article: Eiling Yee (2019) Abstraction and concepts: when, how, where, what and why?, Language, Cognition and Neuroscience, 34:10, 1257-1265, DOI: 10.1080/23273798.2019.1660797

To link to this article: https://doi.org/10.1080/23273798.2019.1660797

Published online: 02 Oct 2019.

LANGUAGE, COGNITION AND NEUROSCIENCE
2019, VOL. 34, NO. 10, 1257–1265
https://doi.org/10.1080/23273798.2019.1660797

Routledge
Taylor & Francis Group

INTRODUCTION

Abstraction and concepts: when, how, where, what and why?

Eiling Yee [a,b]

[a]Department of Psychological Sciences, University of Connecticut, Storrs, CT, USA; [b]Connecticut Institute for the Brain and Cognitive Sciences, University of Connecticut, Storrs, CT, USA

ABSTRACT
It is increasingly apparent that sensorimotor information is a constitutive part of conceptual knowledge. Yet all concepts, even highly concrete ones (e.g. *dog*) include information that is abstracted across individual episodes of experience, departing somewhat from direct sensory or motor input. This process of *abstraction* is the essence of conceptual structure. This Special Issue brings together developmental, experimental, computational and cognitive neuroscientific perspectives on abstraction. The contributions address questions like: When during development do our concepts become less directly tied to sensory or motor knowledge? How (and where) in the brain does the process of abstraction happen? And what is the role of a concept's label in abstraction? In answering these questions, the contributions highlight that context–the developmental contexts of our first episodic experiences, and the linguistic contexts that accompany the development of conceptual knowledge in both children and adults–is at the root of conceptual knowledge.

ARTICLE HISTORY
Received 23 July 2019
Accepted 13 August 2019

KEYWORDS
Concepts; semantic memory; abstraction; abstract concepts; conceptual development

Over the past two decades, there have been an increasing number of demonstrations that sensory and motor (or *embodied*) information is a functional part of conceptual knowledge (for discussion, see Matheson & Barsalou, 2018). Most of this evidence comes from adults, but a role for sensory and motor information in conceptual knowledge is also consistent with evidence that young children's conceptual systems are acquired and organised via their sensory and motor experiences (for review, see Wellsby & Pexman, 2014).

At the same time, however, it is apparent that not all of conceptual knowledge can be directly attributed to simple sensory or motor input. At one end of the spectrum, there are many concepts that do not have any obvious sensory or motor correlates (e.g. the concept *concept*). But there are also intermediate cases, such as the concept *joy,* which not only correlate more with internal bodily states than with external sensory information but these internal sensory correlates can vary significantly (consider an exuberant joy vs. a calm joy; for review, see Winkielman, Coulson, & Niedenthal, 2018). In fact, even the concept *dog,* insofar as it maps on to animals of a large variety of shapes and sizes, arguably should have a relatively indistinct sensorimotor representation (for discussion see Barsalou, 1999). Furthermore, there is evidence that competent language users can derive (at least some) meaning from language without relying directly upon sensory and motor systems (for discussion, see Louwerse, 2011). For these reasons, conceptual knowledge must be more than simply sensorimotor – it seems that it also involves some "abstraction" away from direct sensory or motor information.

In this article, we use the term *abstraction* to refer to the process that, supported by sensitivity to statistical regularities in our environment (e.g. regularities across sensory input and our motor responses to that input), allows us to form and store (or, in exemplar/retrieval-based models, allows us to compute at retrieval) semantic information gleaned across our experiences. Via this process of abstraction we aggregate across individual episodes of direct sensory or motor experience, such that what is most common across our experiences with a given object or event gains prominence, and idiosyncratic properties, like a leaf sticking to your dog's coat, carry little weight. It is through this kind of abstraction that we are able to discern what various objects and events have in common, and group them together into concepts. We use the term *abstracted* to refer to the conceptual knowledge (which is dynamic in the sense that it is continually updated via experience) that is the output of this process of abstraction. We favour the term *abstracted* here because it accentuates that there was some process, operating across experience, from which the knowledge arose.

As suggested earlier, it may be that all conceptual knowledge can, in a sense, be described as abstracted. We therefore use the term in a graded fashion, referring

CONTACT Eiling Yee eiling.yee@gmail.com

to knowledge that is "more abstracted" (e.g. the knowledge of what an *animal* is) and "less abstracted" (e.g. knowledge of what a *dog* is, or even of who a particular dog, *Zoey*, is). We use the term *abstract* to refer to knowledge that does not have *any* obvious sensory or motor basis (e.g. knowledge about the concepts *concept* or *truth*) – and, on our view, can therefore only be attained by relying heavily on language or on abstracting across diverse episodes of experience. Thus, we consider abstract knowledge to be conceptual knowledge from the more abstracted end of the continuum. *Abstract concepts* are composed, predominantly, of abstract knowledge, but may have sensorimotor correlates (e.g. the concept *joy*; for related discussion, see Barsalou, Dutriaux, & Scheepers, 2018).

But why have concepts in the first place? That is, what is the benefit of (via the process of abstraction) developing this conceptual knowledge? Most people who study concepts would agree that concepts allow us to make sense of and organise the world – they allow us, when encountering something new, not to have to learn from scratch what it does, how we should interact with it, and how it might change. In this sense, having concepts gives us the ability to *generalize* – which we define here as the ability to apply knowledge that we have learned to new situations. Thus, when we use concepts to make inferences or predictions about what will happen when we encounter something new (e.g. any new instance of a dog, or a leaf, or a pencil), we are generalising (these definitions of generalisation and abstraction follow the conventions in the statistical learning literature; see Altmann, 2017).[1] The notions of abstraction and generalisation recur repeatedly throughout this Special Issue. This is appropriate because (1) we have concepts in order to *use them* and (2) there can be no generalisation without first developing/abstracting the knowledge that we use to make generalisations.

The point that we use concepts to organise the world relates to the notion of different *levels of abstraction*, i.e. differences in degree of detail or precision, where the less detail there is, the greater the level of abstraction/the more abstracted the concept (e.g. *goldendoodle* vs. *dog* vs. *animal*). The idea is that the abstraction *process* (as described above) can produce these different levels of abstraction (for discussion of this relationship, see Burgoon, Henderson, & Markman, 2013). That is, all abstractions minimise the idiosyncratic details of the experiences from which they are derived (the leaf on my goldendoodle's coat, and even the length of her coat are unlikely to become important parts of my abstracted concept *dog*). But we can perform even more abstraction, to create the concept *animal*, which does not contain detail about the size or the shape of the creature. In fact, the different levels in hierarchical descriptions of semantic knowledge (e.g. Collins & Quillian, 1969) can be thought of as degrees of abstraction, with superordinate (i.e. more inclusive) levels corresponding to more abstraction, and subordinate levels corresponding to less (Rosch, 1978). And insofar as these more abstracted representations of concrete things (e.g. *animal*) lack clear sensory or motor correlates, they are similar to abstract concepts, as defined above.[2]

The major questions discussed in this Special Issue focus on exactly this central property of concepts – that all concepts, even concrete ones (e.g. *dog*) require abstraction away from their sensorimotor correlates. This process of abstraction is the essence of conceptual structure. It is from within this perspective that the papers in this Special Issue ask:

(1) Given that young children's conceptual systems are acquired and organised via their sensory and motor experiences, when, and how, during development does abstraction away from direct experiences occur (Sheya & Smith, 2018; Pexman, 2017; Sloutsky & Deng, 2017; Gentner & Asmuth, 2017)?
(2) What is the role of language, and in particular, a concept's label, in abstraction (Pexman, 2017; Sloutsky & Deng, 2017; Gentner & Asmuth, 2017; Connell, 2018; Lupyan & Lewis, 2017; Jones, 2018; Davis & Yee, 2018)?
(3) How do we generalise from (apply) this abstracted knowledge (Sheya & Smith, 2018; Pexman, 2017; Gentner & Asmuth, 2017; Connell, 2018; Lupyan & Lewis, 2017; Jones, 2018; Davis & Yee, 2018)?
(4) What is the underlying neurobiology that supports abstraction and generalisation – how/where are such processes grounded in brain function (Sheya & Smith, 2018; Gentner & Asmuth, 2017; Jones, 2018; Davis & Yee, 2018)?
(5) How are abstract concepts (e.g. *truth* and *fairness*), which seem to be devoid of physical and perceptual instantiation, learned and represented (Pexman, 2017; Gentner & Asmuth, 2017; Connell, 2018; Lupyan & Lewis, 2017; Davis & Yee, 2018)?

Sheya and Smith (2018) begin by arguing for the importance of the interactions between the body and its environment on cognitive, and hence, conceptual development. They highlight several ways the child's sensory input changes in concert with the child's actions. For example, they point out that via seeing one's hand move, the input to visual brain regions can be linked to output from motor regions, coupling activity in these regions. Such sensory and motor couplings, they argue, will both affect the infant brain's functional connectivity in the short term, and may also, given repeated occurrences over the long-term, affect anatomical connectivity (for related discussion, see Pulvermüller, 2018). Moreover, as the infant grows, the inputs to her perceptual and motor systems change (e.g. what she can reach, and see, changes), which in turn influences her behaviour. Thus, the particular child and her specific environment structure each other. Sheya & Smith's developmental perspective makes clear that cognition emerges from the interplay of input and perception (see also Lupyan & Lewis, 2017; Pexman, 2017). These kinds of associations across cognitive faculties, and consequent functional and structural changes, are what allow for the development of conceptual knowledge.

Picking up many of these developmental threads, Pexman (2017) links them directly to conceptual representations. For instance, Pexman suggests that children's conceptual representations change as their physical development produces changes in how they interact with objects. Moreover, she reviews theories and evidence suggesting that there are changes across the lifespan in the extent to which sensory or motor information supports conceptual access in some tasks. In particular, for some conceptualisation tasks, the role of sensorimotor information may be more prominent for children than for adults. This developmental change may be related to object labels having a larger role in conceptual representations as development progresses, e.g. as children's language develops they can rely more upon labels to help them detect what is common across situations (in other words, labels contribute to the ability to develop abstractions across situations).

Pexman (2017) also makes the intriguing claim that examining development may shed light on how, in a sensorimotor-based framework, abstract concepts such as *truth* can be represented. She points out that if, as other authors have suggested, abstract concepts are grounded through language, emotion, and/or metaphor (see also Gentner & Asmuth, 2017) then children's abstract concepts (and their reliance on these systems) will necessarily change as their systems for language, emotion, and metaphor understanding develop. In other words, abstract concepts may rely on different systems depending in part on the developmental phase of the child. Overall, through highlighting the impact of cognitive development on the development of concepts, Pexman emphasises that concepts are dynamic and context-dependent.

Sloutsky and Deng (2017) also discuss conceptual development but focus on the role of language. On their view, labels are more than just one part of a conceptual representation – rather, Sloutsky & Deng define concepts as *lexicalized* classes of entities. They suggest that lexicalisation, i.e. having a label for a concept, is particularly important because it makes it possible to acquire knowledge about that concept via language. However, they point out that while some concepts originate in a "top-down" manner (i.e. via language) and are grounded later (e.g. *germ*), others originate in a "bottom-up" manner (i.e. via interactions with the world) and only become lexicalised later. The acquisition of bottom-up concepts is constrained by experiential and developmental considerations – they tend to be objects that are commonly experienced by the infant, and that can be acted upon (see also Pexman, 2017; Sheya & Smith, 2018).

How easily concepts can be acquired in a bottom-up or top-down manner may also be affected by another dimension on which, according to Sloutsky and Deng (2017), categories differ: the extent to which the members of the category have correlated features. Dense categories have members with many correlated features (e.g. race cars), whereas sparse categories have members that share few relevant features (e.g. red things). Thus, dense categories can be learned by generalising membership to things that are highly perceptually similar overall, but learning sparse categories seems to require the ability to selectively attend only to specific, relevant features. And, because the ability to selectively attend develops gradually throughout childhood, the ability to learn sparse categories should take time to develop as well. Sloutsky & Deng describe evidence that this is true, and also suggest that selective attention allows us to focus on one particular feature – the label, which is likely to be particularly predictive of what the other features of an object will be.

But language does more than label objects – language in context can also help learners infer various aspects of a word's meaning when its referent is not present and its meaning is unknown (see also Connell, 2018). Sloutsky and Deng (2017) describe a language-experience based account for the phenomenon whereby children appear to prefer to group things based on thematic (i.e. co-occurrence) relationships, such as grouping *monkey* and *banana*, whereas adults appear to prefer to group things based on taxonomic (i.e. categorical) relationships, such as grouping *monkey* and *dog*. In particular, they suggest that simple word co-occurrence (in linguistic parlance, *syntagmatic* relationships) supports learning thematic relations, and it is only after longer-term experience with language that learners accumulate enough information to develop sensitivity to when words play the same roles in sentences (i.e. *paradigmatic* relationships), which supports learning taxonomic relations. Thus, Sloutsky & Deng contend that conceptual development "is greatly amplified (if not transformed) by language … "

Gentner and Asmuth (2017) discuss a closely related notion – the importance of language for learning abstract, relational concepts such as *cause* and *agree,* or even simpler terms such as *tall,* and *above.* Such relational concepts are important for cognition – they include not only most verbs and prepositions, but also many nouns (e.g. *response, goal, anchor, brother*). Gentner & Asmuth point out that learning relational concepts from non-linguistic experience alone (or what Sloutsky & Deng [2017] would call in a bottom-up manner) is challenging because relations are not obvious in the world. Therefore, it is useful to have a label to guide one's attention to the commonalities between things (for related comments on labels, see also Davis & Yee, 2018; Lupyan & Lewis, 2017; Pexman, 2017; Sloutsky & Deng, 2017).

Most abstract relational concepts arise, Gentner and Asmuth (2017) propose, via analogy to more concrete concepts, a process they refer to as the Career of Metaphor. For example, the abstract, relational meaning of *anchor* (anything that prevents something else from moving in an undesirable direction) may have come about through processing figurative language like "their religion is like an anchor". The idea is that processing such figurative language involves focusing on the characteristics that, e.g. the concrete concept *anchor* and *religion* have in common (i.e. performing *structural alignment* on the two representations), which renders these common characteristics more salient. Thus, repeated experience with such language allows us to gradually abstract the meaning of the metaphor in that the common relational structure becomes more prominent than the unshared concrete features. Gentner & Asmuth's account thus helps explain how metaphorical language may provide a path for the emergence of more abstract concepts from more concrete, embodied concepts.

A common thread across the four contributions that focus on conceptual development is that concepts change as the input to the child changes. It is worth recognising, however, that although they emphasise changes in infancy and childhood, the perspectives in these contributions are all compatible with similar changes occurring through adulthood (none of them suggest that after childhood, concepts become fixed). In other words, given that even as adults, the input that we receive changes as our experiences change in both the long and short term (e.g. via changes in our cultural, physical, or bodily environments, or in our knowledge or goals) our concepts should continue to change as well (for recent evidence of conceptual evolution due to changes in the body and knowledge, respectively, see Chrysikou, Casasanto, & Thompson-Schill, 2017; and Clarke, Pell, Ranganath, & Tyler, 2016; for review see Yee & Thompson-Schill, 2016).

Likewise, these authors' emphasis on the importance of the input is in harmony with the idea that concepts can also vary *across* individuals due to differences in experiences (e.g. due to differences in cultural, physical, or bodily environments, or in our knowledge or goals; e.g. Hoffman, McClelland, Ralph, & Matthew, 2018; for reviews, see Casasanto & Lupyan, 2015; Yee, Jones, & McRae, 2018). Such experience-based variability within and across individuals is, in fact, an inescapable consequence of an embodied or sensorimotor-based theory of concepts (e.g. Allport, 1985; Barsalou, 1999; for discussion, see Prinz, 2002; for a description and an implementation, respectively, of a mechanism that would produce such variability [a recurrent architecture based on Elman's (1990) simple recurrent network] see Yee & Thompson-Schill, 2016 and Hoffman et al., 2018).

Moving to adults, Connell's (2018) contribution addresses yet another aspect of conceptual variability. She suggests that in some circumstances, when we activate a concrete concept, we do not strongly activate the sensorimotor knowledge that we have about that concept. In other words, Connell suggests that when the task does not require a full activation of the concept (e.g. when the task is to simply judge whether *car* is related to *park*), we may rely on information derived from language, rather than simulating all of our experience-based knowledge about it. One reason for this, Connell suggests, is that linguistically derived information can typically be activated faster and less effortfully than simulated (i.e. embodied or grounded) information. Thus, task-permitting, we may rely more heavily on (computationally cheap) linguistic information when resources are limited or motivation is low. This works, in part, because we use language to talk about the world, and so the information that we derive about concepts from the distributional statistics of language will include much of the information that we derive from the distributional statistics of the world.

Connell (2018) also suggests, however, that rather than being simply a cheaper, imperfect proxy for sensorimotor information, linguistic information can also carry unique information (see also Davis & Yee, 2018; Gentner & Asmuth, 2017; Lupyan & Lewis, 2017; Pexman, 2017; Sloutsky & Deng, 2017). For example, Connell points out that because we use language not only to describe the world, but also to "question, analyze, interpret, abstract, and predict it … linguistic information may capture qualitatively different aspects of conceptual information" than simulated information – aspects which may be particularly important for abstract concepts.

A challenge for future research is to determine what aspects of language may be particularly relevant for abstract concepts (for discussion, see Dove, 2018, 2019; Pulvermüller, 2013). Some candidates include that learning abstract concepts that refer to mental states or social situations may be difficult without receiving verbal explanations of them (e.g. Borghi & Binkofski, 2014). Another possibility is that understanding abstract concepts requires integrating information over a larger spatiotemporal window than does understanding concrete concepts (e.g. consider *chair* vs. *democracy* – the former can be recognised in the here and now, whereas the latter requires observation of an entire society and over an extended time window; see also Barsalou, 1999; Davis, Altmann, & Yee, 2019) and that language supports the memory and attentional demands necessary to integrate information over such larger windows.

Like most of the other contributors, Lupyan and Lewis (2017) also highlight some mechanisms by which language may shape our conceptual knowledge.[3] They suggest that words do more than activate embodied (or, in Connell's [2018] terms, *situated*) mental states, but they also help construct those mental states. For example, they point out that *distributional* models of semantics, which derive all of their knowledge from the distributional statistics of words in large corpora of language, can not only extract word pair similarity, but also, "by calculating a direction vector between one word … and another and then translating it onto another word, the network's learned representations are able to form a kind of analogical reasoning" (see also Gentner & Asmuth, 2017). Importantly, however, this feat should not be taken as evidence that experience of the world is unnecessary for conceptual knowledge – for the language that goes into these models, of course, comes from humans whose language includes information about their sensory experiences in the world (see also Connell, 2018).

Another important role for language, suggest Lupyan and Lewis (2017), is that applying a verbal label to a perceptual representation can emphasise features of that representation that are particularly relevant to the category, and separate (or facilitate abstraction of) the representation from irrelevant aspects of the specific experience (see also Connell, 2018; Davis & Yee, 2018; Gentner & Asmuth, 2017; Pexman, 2017; Sloutsky & Deng, 2017). This works, they claim, because most of the time, labels are entirely "unmotivated" (i.e. arbitrary) in the sense that the label is divorced from the sensorimotor correlates of the concept. Thus, the label is "uncertain" in a way that other features of a concept are not (i.e. hearing something described as a "dog" is uncertain in that it does not tell you the size or shape of the dog), and this uncertainty promotes (or reflects, depending on one's perspective) abstraction over idiosyncratic variability across specific instances of experience.

Although Lupyan and Lewis (2017) do not explicitly discuss this, another crucial aspect of the label is that, unlike other features, it is invariant. This invariance not only allows it to have Lupyan & Lewis' property of uncertainty, but may also promote its ability to act as a kind of hub, or anchoring point around which other, correlated features can coalesce – the intuition being that it should be easier to bind variable features to an invariant one than to bind variable features to other variable features.

In this last role, the label does not itself encode any semantic information. However as further evidence that language structure *can* capture significant semantic information, Jones (2018) points out that computational models that "learn" (or produce) distributed representations of word meaning from statistical redundancies that appear in very large corpora (i.e. distributional models) have been extremely successful at fitting aspects of human behaviour, e.g. responses in semantic priming, or word association tasks. Intriguingly, many of these models derive a word's meaning via a process of abstraction (usually referred to as *dimensionality reduction*) whereby a large corpus of text is processed, a word's most stable dimensions are identified and retained, and other information is discarded. Yet this dimensionality reduction process has limitations. Perhaps most problematically, Jones suggests, the human categorisation literature suggests that rather than being discarded, information about individual episodes is retained.[4] For this reason, Jones suggests, exemplar-based distributional semantic models (also called *retrieval-based models* or *memory models*) may be a better fit to human performance – in these models, all instances are retained and stored (subject to degradation through noise and decay), and prototype or abstraction-like effects are emergent artifacts of a retrieval process operating on this stored information. The success of these models at fitting human performance casts doubt on the need for a separate (abstracted) semantic memory store. Importantly, however, the idea that there may not be separate semantic and episodic memory stores does not mean that semantic knowledge about concepts does not exist – instead, according to these models semantic memory is a process, rather than a structure.

By highlighting the role of episodic information in semantic memory, Jones (2018) contribution underscores an important point – although a recurrent theme throughout this Special Issue is that the process of abstraction is essential, none of the contributions suggest that during this process, individual instances are *entirely* discarded. Moreover, views of semantic memory as experience-based and context dependent make it difficult to meaningfully distinguish concepts from the contexts in which they occur (for recent discussion, see Casasanto & Lupyan, 2015; Connell & Lynott, 2014; Lebois, Wilson-Mendenhall, & Barsalou, 2015; Yee & Thompson-Schill, 2016). Thus, many current perspectives on semantic memory fit in well with exemplar/retrieval based models of semantic knowledge.

Finally, Davis and Yee's (2018) contribution presumes that theoretical approaches to conceptual representation can be informed by considering how such representations (and the processes that contribute to their development and deployment) might be constrained by the neurobiology of the brain. Their contribution focuses on two types of generalisation – generalisation based on taxonomic (or categorical) similarity, and generalisation based on thematic (or event-based) similarity. They propose *why* the two brain regions that have been suggested to support taxonomic and thematic generalisation, i.e. the anterior temporal lobe, and angular gyrus, respectively, are able to do so.

In particular, Davis and Yee (2018) suggest that the angular gyrus' strong reciprocal connectivity with posterior hippocampus supports its sensitivity to spatial and temporal episodic detail (i.e. event structure) from which thematic relations are generalised. They also propose that one of the factors that allows the anterior temporal lobe to support taxonomic generalisation is its connectivity to frontal language regions which support retrieval of labels, which play several roles in supporting taxonomic generalisation. Finally, they speculate that the statistical abstractions that that map individual episodic instances onto accumulated experience may be the same for thematic and taxonomic generalisations but that the abstractions supporting thematic generalisations are derived over larger spatiotemporal windows than are taxonomic (an idea also mentioned in Hoffman et al., 2018; for related discussion, see Mirman, Landrigan, & Britt, 2017).

A potentially significant issue that Davis and Yee (2018) only briefly discuss, however, is whether the role that labels play for conceptual processing may differ depending on the concept's level of abstraction. Labels may play a particularly important role for *putting things together* into taxonomic categories (especially during development) at higher levels of abstraction and for abstract concepts and relations (see also Connell, 2018; Davis & Yee, 2018; Gentner & Asmuth, 2017; Jones, 2018; Lupyan & Lewis, 2017; Pexman, 2017; Sloutsky & Deng, 2017), whereas labels may aid in *keeping things apart* at lower levels of abstraction (see also Davis & Yee, 2018). For instance, for superordinate categories, like *animal* or *tool*, and for abstract concepts such as *causation* (concepts for which integration across multiple features or contexts may be needed to identify the commonalities among instances) labels may be essential in that without them, the commonalities would go unnoticed (for developmental review, see Markman, 1990). In contrast, for subordinate categories, where there is a very high level of perceptual similarity among exemplars, labels may play a quite different role – rather than highlighting the commonalities among exemplars, subordinate level labels (e.g. adding a model name to a car) may provide an additional, discriminating feature that can help distinguish between highly similar entities (e.g. Gilbert, Regier, Kay, & Ivry, 2006).

Relatedly, it is possible that one reason that superordinate and subordinate level categories tend to be later to develop than basic level categories[5] is that these levels are particularly dependent on the ability to focus on specific aspects of a representation, including the label, and that this ability depends on prefrontal cortex, the brain region that takes longest to develop (for discussion see Sloutsky, 2010). It will be interesting to discover whether the development of prefrontal cortex indeed influences the extent to which children rely on labels versus sensorimotor information when learning new semantic knowledge.

Importantly, at both the superordinate and the subordinate levels, the very same mechanism underlies the label's ability to perform its different roles. At both levels, the label "creates for the learner a new realm of perceptible objects (the associated tags or labels) upon which to target her more basic capacities of statistical and associative learning." (p. 371; Clark, 2006) – the difference simply lies in the statistics that are associated with the label at the different levels.

The learner's sensitivity to statistical regularities may also help us gain some insight into the connection between thematic and taxonomic relations: to use an example concerning structure in language (Elman, 1990; see also Sloutsky & Deng, 2017), nouns are nouns (i.e. are members of the taxonomic category *nouns*) because they come before verbs, and potentially come after determiners and adjectives. Likewise, adjectives are a taxonomic category defined by coming after determiners and before nouns (this is an over-simplification, but it makes the point). And the *relationships* between nouns, verbs, and adjectives are essentially thematic relations (in linguistic terms, the relationship between the nouns and the verb in a sentence defines the thematic structure of that sentence). The insight to be gained here is that things are members of a *taxonomic* category by virtue of *sharing* the kinds of things they can co-occur with. And of course the flip side of this same coin is that *it is these very co-occurrences* that define the *thematic* relationships that each member of the category can enter into (and hence, they define the generalisations that can be made based on category membership).

But how is a child to know, as she builds a body of knowledge about a concept, what the range of possible co-occurrences (or more properly, contextual contingencies) is? This is where language may play yet another role. As a proxy for experience, labels may be important precisely because without them, we do not have the capacity to experience the full range of co-occurrences – the range of *contexts* – needed to generate the nuances of the category (for discussion see Louwerse, 2018; see also Connell, 2018; Lupyan & Lewis, 2017; Sloutsky & Deng, 2017) This may be especially true for abstract, relational concepts, e.g. *cause* (see also Gentner & Asmuth, 2017). And for categories that have few exemplars, labels in context may help make up for that sparseness,[6] in the extreme creating for the child categories of things for which there are no real-world exemplars (e.g. fairies).

Thus, for the child, language may step in and help fill the void that results from lack of experience with all those contexts. It is not surprising, then, that conceptual development proceeds apace with language development and, indeed cognitive development. Conceptual knowledge is, among other things, knowledge about the contexts in which things happen, and there are different routes to gaining that knowledge – labels, specifically, and language more generally, may be an essential stopgap during development. Hence, it is no coincidence that so much of this Special Issue about abstraction and concepts is devoted to development, and to language.

Notes

1. Some authors use the terms *generalization* or *induction* to refer to the process of deriving information across instances that may not be available from any individual instance – i.e., the process that we have called *abstraction.* When reading the literature (including this Special Issue) it is therefore essential to be sensitive to *how* authors are using terms to determine what they mean. Indeed, this definitional inconsistency relates to one of the challenges that has been discussed in the literature on abstract concepts – that what an abstract word (like the word *generalization*) refers to is strongly context dependent (e.g., Schwanenflugel, 1991).
2. Although historically, most of the literature on concepts has focused on what might be termed "simple concepts", that is, the kind of concrete object and action concepts that we learn early on during development, and that we use in our basic interactions with the world, abstract concepts are beginning to receive more attention. This is warranted because much of adult language – over half by some measures (Connell, 2018; for discussion see Lupyan & Winter, 2018) – is about concepts that have no obvious sensory or motor manifestation, e.g., *truth, fairness*. Furthermore, even simple concepts can contain features (e.g., the functions of many objects, such as *used to tell time*) that do not have obvious sensorimotor correlates.
3. Interestingly, despite the recent resurgence of research relating to linguistic relativity, the body of work focusing on the *mechanisms* by which language may shape our conceptual knowledge remains relatively small. The contributions to this Special Issue are an important step forward.
4. Note, however, that this does not mean that this information is retained veridically without loss of information. Rather it means that depending upon the retrieval cue, some aspects of individual episodes can be retrieved and contribute to the retrieved representation.
5. This is not to say that labels do not facilitate categorisation at the basic level (Lupyan & Lewis [2017] review evidence that they do, and for the same reason that they help at the superordinate level)—the idea is simply that they may be particularly useful at the superordinate and subordinate levels.
6. Sloutsky and Deng (2017) use *sparse* in a slightly different way--to refer to the number of features that distinguish members of the category from others.

Acknowledgments

Thank you to the authors for contributing to this special issue. Each of their papers is referenced with the year in which it first appeared online. Thanks to Katherine White and Gerry Altmann for helpful comments. This work was supported by NSF IGERT DGE-1144399 to the University of Connecticut.

Disclosure statement

No potential conflict of interest was reported by the author.

Funding

This work was supported by National Science Foundation (NSF) IGERT DGE-1144399 to the University of Connecticut.

ORCID

Eiling Yee http://orcid.org/0000-0001-6614-9025

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.

Altmann, G. T. M. (2017). Abstraction and generalization in statistical learning: Implications for the relationship between semantic types and episodic tokens. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*. doi:10.1098/rstb.2016.0060

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

Barsalou, L. W., Dutriaux, L., & Scheepers, C. (2018). Moving beyond the distinction between concrete and abstract concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1752). doi:10.1098/rstb.2017.0144

Borghi, A. M., & Binkofski, F. (2014). *Words as social tools: An embodied view on abstract concepts*. New York, NY: Springer.

Burgoon, E. M., Henderson, M. D., & Markman, A. B. (2013). There are many ways to see the forest for the trees: A tour guide for abstraction. *Perspectives on Psychological Science*, *8*(5), 501–520.

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

Chrysikou, E. G., Casasanto, D., & Thompson-Schill, S. L. (2017). Motor experience influences object knowledge. *Journal of Experimental Psychology: General*, *146*(3), 395–408.

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

Clarke, A., Pell, P. J., Ranganath, C., & Tyler, L. K. (2016). Learning warps object representations in the ventral temporal cortex. *Journal of Cognitive Neuroscience*, *28*(7), 1010–1023.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, *8* (2), 240–247.

Connell, L. (2018). What have labels ever done for us? The linguistic shortcut in conceptual processing. *Language, Cognition and Neuroscience*, *34*(10), 1308–1318.

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

Davis, C. P., Altmann, G. T. M., & Yee, E. (2019). Situational systematicity: A role for schema in understanding the differences between abstract and concrete concepts. *PsyArXiv*. doi:10.31234/osf.io/5qud8

Davis, C. P., & Yee, E. (2018). Features, labels, space, and time: Factors supporting taxonomic relationships in the anterior temporal lobe and thematic relationships in the angular gyrus. *Language, Cognition and Neuroscience*, *34*(10), 1347–1357.

Dove, G. (2018). Language as a disruptive technology: Abstract concepts, embodiment and the flexible mind. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373* (1752). doi:10.1098/rstb.2017.0135

Dove, G. (2019). More than a scaffold: Language is a neuroenhancement. *Cognitive Neuropsychology*. doi:10.1080/02643294.2019.1637338

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

Gentner, D., & Asmuth, J. (2017). Metaphoric extension, relational categories, and abstraction. *Language, Cognition and Neuroscience*, *34*(10), 1298–1307.

Gilbert, A. L., Regier, T., Kay, P., & Ivry, R. B. (2006). Whorf hypothesis is supported in the right visual field but not the left. *Proceedings of the National Academy of Sciences*, *103*(2), 489–494.

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

Jones, M. N. (2018). When does abstraction occur in semantic memory: Insights from distributional models. *Language, Cognition and Neuroscience*, *34*(10), 1338–1346.

Lebois, L. A., Wilson-Mendenhall, C. D., & Barsalou, L. W. (2015). Are automatic conceptual cores the gold standard of semantic processing? The context-dependence of spatial meaning in grounded congruency effects. *Cognitive Science*, *39*(8), 1764–1801.

Louwerse, M. M. (2011). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, *3*(2), 273–302.

Louwerse, M. M. (2018). Knowing the meaning of a word by the linguistic and perceptual company it keeps. *Topics in Cognitive Science*, *10*(3), 573–589.

Lupyan, G., & Lewis, M. (2017). From words-as-mappings to words-as- cues: The role of language in semantic knowledge. *Language, Cognition and Neuroscience*, *34*(10), 1319–1337.

Lupyan, G., & Winter, B. (2018). Language is more abstract than you think, or, why aren't languages more iconic? *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1752). doi:10.1098/rstb.2017.0137

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

Matheson, H. E., & Barsalou, L. W. (2018). Embodiment and grounding in cognitive neuroscience. *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*, *3*, 1–27.

Mirman, D., Landrigan, J. F., & Britt, A. E. (2017). Taxonomic and thematic semantic systems. *Psychological Bulletin*, *143*(5), 499–520.

Pexman, P. M. (2017). The role of embodiment in conceptual development. *Language, Cognition and Neuroscience*, *34*(10), 1274–1283.

Prinz, J. J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA: MIT Press.

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

Pulvermüller, F. (2018). Neural reuse of action perception circuits for language, concepts and communication. *Progress in Neurobiology*, *160*, 1–44.

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.

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

Sheya, A., & Smith, L. (2018). Development weaves brains, bodies and environments into cognition. *Language, Cognition and Neuroscience*, *34*(10), 1266–1273.

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

Sloutsky, V. M., & Deng, W. (2017). Categories, concepts, and conceptual development. *Language, Cognition and Neuroscience*, *34*(10), 1284–1297.

Wellsby, M., & Pexman, P. M. (2014). Developing embodied cognition: Insights from children's concepts and language processing. *Frontiers in Cognitive Science*, *5*, 506.

Winkielman, P., Coulson, S., & Niedenthal, P. (2018). Dynamic grounding of emotion concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1752). doi:10.1098/rstb.2017.0127

Yee, E., Jones, M. N., & McRae, K. (2018). Semantic memory. In J. T. Wixted, & S. L. Thompson-Schill (Eds.), *Stevens' Handbook of experimental Psychology and cognitive neuroscience* (4th ed., Vol. 3: Language and Thought, pp. 319–356). New York: Wiley.

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