

**Situational systematicity: A role for schema in understanding the differences between
abstract and concrete concepts**

Abstract

Abstract concepts differ from concrete concepts in a number of ways. Here, we focus on what we refer to as *situational systematicity*: The objects and relations that constitute an abstract concept (e.g., *justice*) are more dispersed through space and time than are the objects and relations that typically constitute a concrete concept (e.g., *chair*); a larger set of objects and relations might potentially constitute an abstract concept than a concrete one; and exactly *which* objects and relations constitute a concept is likely more context-dependent for abstract than for concrete concepts. We thus refer to abstract concepts as having low situational systematicity. We contend that situational systematicity, rather than abstractness *per se*, may be a critical determinant of the cognitive, behavioral, and neural phenomena typically associated with concepts. We also contend that investigating concepts through the lens of schema provides insight into the situation-based dynamics of concept learning and representation, and into the functional significance of the interactions between brain regions that make up the *schema control network*.

Keywords: concepts, semantic memory, episodic memory, abstract concepts, schema

Introduction

Abstract concepts like *idea*, *integrity*, *good*, and *pride* are central to human thought. They allow us to comprehend relationships among people and between people and their environments (e.g., “customers value the *pride* we have in our work”), to make inferences about things and thoughts that are not observable through the senses, or perhaps not observable at all (e.g., “her *integrity* shone through”), and, with the appropriate labels, to communicate about patterns of information in the world (e.g., “science is about *ideas*”). In this paper, we consider how abstract concepts relate to the theoretical notions of situation, episode, and schema. Thinking of abstract concepts—and in fact, any concept—as schemas affords the possibility of inheriting the neurobiological brain mechanisms, and associated empirical consequences, that have been proposed to underpin schema-based knowledge (see Gilboa & Marlatte, 2017, for review).

While the idea that concepts are schemas is not new (e.g., Rumelhart & Ortony, 1977; for discussion, see Murphy, 2002), there has been a gradual drift in the neurobiological literature to a distinction (which we suggest is unhelpful) between concepts as knowledge of category membership and schemas as knowledge about spatiotemporal and causal relations between category members in a given situation (e.g., Ghosh & Gilboa, 2014; but see Gureckis & Goldstone, 2010 and Barsalou et al., 2018 for a more unified approach). We take *situations* (following Barwise & Perry, 1983) to consist of objects in the world, their properties, and their relations to one another (*episodes* are situations that are grounded in space, time, and specific contexts; see Tulving, 1983, i.e., they are instances of situations). We view *object knowledge* as schema knowledge (see also Barsalou, 1999)—knowledge of the typical situations in which an object is found, of its typical perceptual correlates, and of the typical spatiotemporal and causal relations between that object and others (similar to *affordances*; Gibson, 1979; Glenberg, 1997).

Abstract (concept) knowledge is also schema knowledge to the extent that it applies to situations and constitutes knowledge not only of typical (and atypical) events that may accompany those situations, but also of the relations and interactions among the objects taking part in those events. Thus, like others before us, we take the view that concepts, both concrete and abstract, are schemas.

More broadly, schemas can be construed as semantic knowledge of situations and the typical events they are composed of (Bartlett, 1932). Some schemas refer to situations that are relatively specific and similar, in respect of the details of each situation or event (going to a restaurant—tables, chairs, menus, servers, food, etc.), whereas others refer to quite dissimilar situations (playing a game—a game of solitaire versus a board game versus a football game, etc.). These differences can be thought of as reflecting a continuum of *situational systematicity*—the extent to which the same objects and relations constitute the concept (i.e., schema) regardless of the situation. Bread and (optionally) butter tend to be constitutive of the concept *toast*, which has high situational systematicity—these same objects and relations will take part in different instances of toast. But *justice* has low situational systematicity, in that the constitutive objects and relations are not the same across situations to which the concept refers. Different instances of *justice* may involve quite different situations, objects, and relations (similar to different instances of *games*; Wittgenstein, 1953).

We, like others (e.g., Galbraith & Underwood, 1973; Hoffman et al., 2013; Schwanenflugel, 1991), claim that abstract concepts tend to be low in situational systematicity (others have used terms such as *context availability* and *semantic diversity* to capture similar ideas, but see the following sections for how our proposal differs). This has the consequence that abstract concepts, unlike concrete concepts, cannot so easily be learned or identified by analogy

to similar situations. Yet despite their differences, we propose that taking a schema-based perspective of both abstract and concrete concepts has the potential to contribute to our understanding of the cognitive, behavioral, and neurobiological phenomena that accompany the differences between the two.

In the sections that follow, we first briefly review accounts of abstract concepts that, like the one developed here, focus on their situational systematicity (or lack thereof). We then draw on the neurobiological literature on schema knowledge to identify a neurobiological process that differentially constrains learning and processing of concrete and abstract concepts. These differential constraints are afforded by the low systematicity typical of abstract concepts. They concern the ways in which the objects and relations that constitute the situational content of a concept are grounded in actual space, time, and experience (i.e., are bound to the episodic context during the actual experience). And because systematicity is graded, these constraints are graded also.

Situational Systematicity

Many types of experience, including linguistic, emotional, social, and assorted sensorimotor experiences, are associated with understanding abstract concepts (for review, see Binder et al., 2016; Borghi et al., 2017; for meta-analysis, see Desai et al., 2018; Wang et al., 2010). The degree to which each of these systems is active in processing abstract concepts, however, varies both across concepts as well as across situations—the properties associated with the *goodness* of something depend on whether that thing is, for example, a person, a fruit, or a machine. Even within these categories, the goodness of a machine, for example, would be assessed differently depending on its intended function, the situation in which it is placed, and so on (e.g., Noppeney & Price, 2004; for review, see Hoffman, 2016).

The notion that situational systematicity is critical to the processing and representation of abstract concepts was first advanced by Schwanenflugel (e.g., Schwanenflugel & Shoben, 1983; see also Schwanenflugel, 1991) under the *context availability hypothesis*, which suggests that it is difficult to generate plausible contexts in which abstract concepts are used, and thus their meaning is more often constrained by the linguistic context. On this view, processing advantages for concrete over abstract concepts are at least in part due to differences in context availability, and by extension, situational systematicity. In recent years, this approach has been augmented by Hoffman and colleagues (for review, see Hoffman, 2016), who demonstrate that words referring to abstract concepts are higher on a metric called *semantic diversity*, which measures the semantic variability of the different linguistic contexts in which a word appears (a proxy for experience). Abstract concepts tend to occur in highly variable contexts—consider that an *idea* can refer to both the *idea* to jump into a water-filled quarry at the encouragement of a group of rowdy friends, and the *idea* to write a paper about abstract concepts. A *chair*, on the other hand, typically occurs in contexts related to sitting, often in the presence of tables, and occasionally in the context of reaching (i.e., standing on one); these contexts are far more constrained than the ones that accompany ideas, or lies (see Barsalou et al., 2018, for further discussion).

Heightened semantic diversity has consequences for how concepts are processed in context, and these consequences are captured in the *controlled semantic cognition* framework (Hoffman et al., 2018; Lambon Ralph et al., 2017). Hoffman et al. (2015) found that anterior temporal regions (thought to be a *hub* region for semantic representation; Rogers et al., 2004) are responsive when concepts are processed following an informative context, while left inferior frontal gyrus (IFG; involved in the selection of situationally appropriate information; Thompson-Schill et al., 1997) is more active following an irrelevant context, as well as for abstract words in

particular. The finding of greater left IFG activation for abstract concepts is consistent with the majority of work investigating neural processing differences between abstract and concrete concepts (see Wang et al., 2010), and Hoffman et al. suggest that it may reflect the important role of top-down control mechanisms in processing abstract concepts: Because abstract words are semantically diverse (i.e., they can be used in a number of semantically distinct contexts), a degree of *control* (i.e., retrieval of the appropriate information given the context; see e.g., Badre & Wagner, 2005; Thompson-Schill et al., 1998; Thompson-Schill, 2003; Wagner et al., 2001) is necessary to suppress irrelevant information and facilitate retrieval of the appropriate meaning (Hoffman et al., 2015). That is, under the controlled semantic cognition framework, not only are there differences in representational *content* between abstract and concrete concepts, but there are also differences in the systems-level dynamics (i.e., semantic control) that facilitate comprehension.

Notably, much of this work on abstract concept processing has been on recognition or processing of *words* that refer to abstract concepts rather than on recognition of abstract concepts directly from the situations they denote (but see McRae et al., 2018). As noted earlier, semantic diversity, for example, has been operationalized in terms of words in their linguistic contexts, as observed in linguistic corpora. This focus on words reflects the fact that accessing an abstract concept non-linguistically is not so straightforward. Whereas an image of a chair, or of a person throwing a punch, can be recognized as instances of those things (e.g. Potter, 1976; Hafri et al., 2013), it is less clear what image one could present (other than an image of letters spelling out JUSTICE) to elicit the concept *justice*—*justice* relies on diverse kinds of information extracted across a range of situations before an instance of it can be recognized in the absence of language.

For this reason, abstract concepts are commonly viewed as being more language-dependent¹ than concrete concepts (e.g., Barsalou et al., 2008; Borghi & Binkofski, 2014; Hoffman et al., 2015; Paivio, 1971, 1991; for review, see Dove, 2016, 2018).

While we agree that language is important for abstract concepts, here we aim to bring the explanatory focus to the experiences driving their representations, rather than to the linguistic contexts from which such concepts might be activated. We therefore use the term “situational systematicity” (rather than “semantic diversity”) to characterize the systematicity (i.e., similarity) of the *real-world* contexts in which a concept might be acquired and recognized. This emphasis motivates focusing on the memory systems that are sensitive to systematicity in the environment and corresponding real-world events and situations, as well as on the intimate relationship between concepts and their contexts (see also Yee & Thompson-Schill, 2016). As we shall argue below, this focus on memory systems suggests an account, grounded in neurobiology, of the observed differences between concrete and abstract concepts.

Situational Systematicity and the Activation of Conceptual Knowledge

The role of non-linguistic events and situations in abstract concepts remains relatively unexplored (but see Barsalou, 1999 and Barsalou et al., 2018, for theoretical discussion, Desai et al., 2018 for meta-analysis, and Wilson-Mendenhall et al., 2013, for evidence that the same brain regions that underpin imagining what other people might be thinking also subserve understanding of words like *convince*). Little empirical work has investigated, for example, whether abstract concepts actually *activate* situation knowledge and *vice versa*. In the one such

¹ In fact, the greater activity in left IFG (and other left-lateralized frontal and temporal regions) often observed for abstract compared to concrete concepts in functional MRI studies has typically been interpreted in support of dual-coding theory and the role of language in representing abstract concepts (e.g., Binder et al., 2005; for meta-analyses, see Binder et al., 2009; Wang et al., 2010), rather than as evidence of abstract concepts requiring greater top-down control (e.g., Hoffman et al., 2015).

study to date (McRae et al., 2018), real-world pictures depicting events with an abstract concept at play (e.g., two girls *sharing* a cob of corn) effected faster response times to related (e.g., *share*) as compared to unrelated (e.g., *convocation*) abstract words. Words referring to abstract concepts also activated situations—that is, response times were faster for related situations following an abstract word. Thus, abstract concepts can activate situations, and appropriate situations can activate abstract concepts.

Knowing that abstract concepts can indeed activate situations and *vice versa* is important. It demonstrates that abstract concepts are grounded in real-world experiences of situations and events. However, work in this area is in its early stages, and there remain many open questions. For example, the mutual activation observed by McRae et al. (2018) only tells us that abstract concepts and situations that *clearly* depict those concepts prime each other (the scenes were normed such that a separate group of participants reliably listed the target abstract word in response to the picture). It can be relatively easy to point to an instance of sharing (depending on the kind of sharing), or even to define the range of situations in which sharing might be happening. But the same is not true for concepts such as *justice*, in part because the timeframe over which *sharing* unfolds (during which various entailments might be verifiable) is typically shorter than that over which *justice* unfolds. We return to this below.

If abstract concepts tend to be associated with situations which are more diverse, on average, than those associated with concrete concepts, they may less strongly activate any particular situation as compared to concepts which are associated with a more restricted range of situations. And conversely, any particular situation may less strongly activate that concept—for example, concepts such as *chair* may more strongly activate the situations in which chairs occur as compared to concepts such as *justice* (Schwanenflugel & Shoben, 1983; see also

Schwanenflugel, 1991; Pulvermüller, 2013 and Barsalou et al., 2018). The point here is that the associative “strength” between a real-world situation and an abstract concept—a schema—will vary depending on the concept and its situational systematicity. This in turn has consequences for the role that the concept can play in directing attention towards the appropriate constitutive elements within the situation(s) to which the concept applies.

Attentional Control and Spatiotemporal Diversity

Low situational systematicity may make it less clear what elements of a situation need to be attended to, and which should not be attended to, in order to facilitate the interpretation of that situation (or set of situations) as constitutive of a particular concept or schema. For example, certain concepts may require, for their activation, more or less of the perceptual elements of a situation to be recognized: the concepts *chair* and *restaurant* are both reasonably concrete, and yet the concept *chair* likely requires recognition of a smaller subset of the perceptual elements in a particular scene as compared to the concept *restaurant*, which requires recognition over a larger spatial window (making *restaurant* less situationally systematic than *chair*). Relatedly, because of its lower situational systematicity, a concept such as *justice* cannot rely solely on situated elements that tend to appear in *every* situational instance of justice, unlike the meanings of *chair* or *restaurant*, which can rely on such elements (reflecting their relatively greater situational systematicity).

Abstract concepts may rely more on elements *across* scenes and situations—that is, across space and time—than do concrete objects, which rely more on elements within a scene or situation (see also Barsalou, 1999). Such spatiotemporal properties of concepts may also impact on the associative strength between a situation and the concepts associated with that situation, at least in part because it will be more difficult to detect systematicities within an extended

spatiotemporal window. In this sense, abstract concepts rely less on the configuration of elements in any *single* given scene than do concrete concepts, and may therefore require more inhibition of irrelevant parts of the situation. Thus, we return to the idea that semantic control may be an important process in respect of the activation of abstract concepts (and see the section below on the schema control network).

The control problem is made all the more challenging because of an added dimension of variability: As we have argued, not only is there spatiotemporal variability *across* concepts, but there is also considerable variability *within* concepts; there tends to be more within-concept variability for abstract than concrete concepts. However, in the discussion thus far of situational systematicity, we have conflated what are potentially (but perhaps not in actual fact) two independent dimensions along which concepts—and abstract concepts in particular—can vary. First, the extent to which the concept denotes a range of possible situations (e.g., *justice* in its many forms, or *sharing* a cob of corn versus a car); and second, the extent to which a given concept denotes a range of possible situations that vary with respect to the extent to which they are spread *across time* (e.g., for *justice*: arrest, imprisonment, determination of mitigating circumstances, subsequent release). In other words, there may also be variability in the *temporal window* over which the component elements and their relations associated with a concept are recognized, with the size of the window required for a given concept depending on the situation.

Consider the concept of *sharing*: sharing a cob of corn can be accomplished by two children biting on it simultaneously—a single discernable situation that can be recognized from a single “snapshot.” Or, one child can nibble on it, then walk around the table, then hand it to the other child, who then nibbles on it (a sequence across time of individual episodes that, taken together, entail an act of sharing). But imagine the case where one child nibbles on it, leaves it on

a plate, and then leaves; the other child now comes up and nibbles on it. Is this *sharing*, or is this *stealing*? To be recognized as an instance of sharing would require information (e.g., about the intentions of the first child) from situations *before* the actual nibbling. And while both the simultaneous and the sequential “sharing” events might be construed as instances of sharing, they are different and spread across different timespans.

Thus, depending on the situation(s), the conditions for activating the same abstract concept may be *informationally* and *temporally diverse*, requiring different kinds of information, across different timeframes: Sharing a cob of corn may involve relations among objects and entities that can be apprehended within a single situation, and may or may not require knowledge about intentionality (depending on the precise situation), whereas sharing a car may involve relations that have to be apprehended across situations, that is, at greater spatiotemporal scales than are involved in apprehension of the components of sharing a cob of corn.

Concepts that are situationally diverse (i.e., that have low situational systematicity) in the various ways we have identified above require, for their identification, that we attend *more* to informative elements of a situation (or of a sequence of situations) and attend *less* to uninformative ones. While the semantic diversity account broadly predicts that this attentional control system inhibits competing (lexical-semantic) representations and selectively activates the appropriate (lexical-semantic) representation, in the following section we extend this notion by drawing on the neurobiology of schema processing to develop an account that attempts to explain the balance between inhibition and activation during the apprehension of conceptually relevant information from the *environment*. Specifically, we conjecture that this control system enables a complementary relationship between brain mechanisms sensitive to schema-congruent information in the environment on the one hand, and on the other, brain mechanisms that encode

the arbitrary associations that ground our everyday experience in the episodic contexts of those experiences.

Neurobiological Underpinnings: The Schema Control Network

A dynamic network of brain regions underpins the integration of information into schemas, and the way those brain regions interact depends on the degree of congruence between an object and a corresponding scene (van Kesteren et al., 2013). Specifically, medial prefrontal cortex (mPFC) interacts with hippocampus (HPC) such that when encoding is successful (as measured by a next-day recall task), medial prefrontal regions are more active in cases where an object is schematically consistent with a scene (e.g., coffee paired with a scene of a café), while hippocampal regions are more active in cases where an object is schematically inconsistent with a scene (e.g., a fish paired with a scene of a café). This suggests that mPFC helps to guide assimilation of stimuli into preexisting schematic representations, while HPC accommodates arbitrary, non-systematic information (see also van Kesteren et al., 2012). This is consonant with the episodic memory literature (e.g., Rugg et al., 2012; for review, see Davachi, 2006) showing that medial temporal activity reflects the degree to which (arbitrary) contextual information is encoded.

mPFC and HPC are functionally complementary in that mPFC essentially encodes (or is functionally connected to the substrates encoding) systematic relationships between objects and the contexts (including other objects) with which they typically co-occur (see schema), whereas HPC “blindly” encodes the relationships, systematic or arbitrary, between objects/contexts that co-occur within individual episodes (relational binding; Cohen & Eichenbaum, 1993). In fact, there is evidence that mPFC is more active when information is schema-congruent, and this deactivates HPC (we take *schema congruency* to reflect the relative match between the current

situation and stored information about the relevant schema). Conversely, when information is schema-incongruent, mPFC is less active while HPC is more active (van Kesteren et al., 2012). The complementarity between the two systems facilitates integration of episodic detail into pre-existing schema, to the extent that those details are schema-congruent, while also down-regulating the irrelevant relational associations (i.e., associated through co-occurrence).² We conjecture that by suppressing irrelevant associations, attention is essentially directed to the schema-relevant detail at the expense of the irrelevant detail. Thus, if a schema is partially activated by a particular episode, the mPFC–HPC system can partially suppress situational details irrelevant to that schema, essentially increasing the signal-to-noise ratio and the likelihood that additional relevant detail can be apprehended from the episode. This would allow the schema to be more strongly activated.

How does this complementarity between schema-relevant and schema-irrelevant information, and between mPFC and HPC, enrich our understanding of the distinction between concrete and abstract concepts? As argued above, our (and others') claim is that the constitutive situational elements of abstract concepts are more sparsely distributed, through space and time, than are the constitutive situational elements of concrete concepts (i.e., situational systematicity is lower for abstract concepts). Further, we have suggested that this low situational systematicity makes it more challenging to detect patterns of information in the environment that are congruent with stored information about abstract concepts. Thus, when a schema for a low situational systematicity concept is partially activated, the complementary relationship between mPFC and HPC allows attention to be selectively directed to the more relevant (according to top-down

² Although down-regulated, these associations are central to enabling the system to generate individuated “tokens” of experience (distinguishing knowledge about the type of a thing, whether an object or an event, from knowledge about the token thing—the actual individuated object or event as grounded in a particular space and time; Altmann & Ekves, 2019).

information from the partially activated schema) co-occurring elements at the expense of more spurious co-occurrences. Hence, the complementarity between mPFC and HPC supports the increased need for selective attention, across space and time, required for the learning and identification of more abstract concepts.

Thus, low situational systematicity requires “top-down” influence (i.e., pattern completion) from the schema to direct attention within or across situations. Where is the source of this top-down/schema information? There is consensus across both the schema (for review, see Gilboa & Marlatte, 2017) and concept literatures (for review, see Binder & Desai, 2011; for discussion, see also Davis & Yee, 2018, and Desai et al., 2018) that the angular gyrus (AG), with its dense connections to medial temporal regions, frontal control systems, and multimodal association areas (Geschwind, 1972), likely plays a dominant role in activating schematic event knowledge and constitutes the source of the top-down influence required by low situational systematicity. Further, because multiple schema may become active at once (though some may be more active than others and their relative activity should depend on the situation; see also Barsalou et al., 2018) we suggest that left IFG, via its role in selecting contextually appropriate information (e.g., Thompson-Schill et al., 1997; Thompson-Schill, 2003), modulates this top-down influence by up-regulating situationally appropriate schema.

As an example of the interplay between these distinct components of the schema control circuitry, let us return to the relatively abstract concept *sharing*. Given that this concept can refer to a range of situations (from sharing custody of a child, to sharing a car, to sharing a sandwich) all of which might be activated via different cues, we would argue that *sharing* has low situational systematicity. Thus, if we observe someone at a fairground picnic table exclaim, “*Hey Bill, you should be sharing!*” AG will presumably partially activate potential *sharing* schemas

compatible with this situation.³ Because AG may activate several distinct sharing schemas, all of which are compatible with the situation, we conjecture that left IFG aids in upregulating the activation of the most appropriate schema(s). (In line with the well-established finding that left IFG is more active for abstract than concrete objects [e.g., Hoffman et al., 2015], concrete concepts, being more situationally systematic, are more constrained with respect to the mapping from situation to concept and hence there would be less demand for IFG-modulated schema selection.) Selective attention to the more situationally appropriate schema allows mPFC to direct attention to schema-appropriate elements in the environment—i.e., whatever is in Bill’s peripersonal space (given the experiential knowledge that in this situation, whatever will be shared is likely within that space). This would increase the likelihood that we would correctly attend to the likely object of the sharing, which would in turn facilitate pattern completion of an appropriate sharing schema for that kind of object (modulating activation of AG). We do not envision this process as discrete and sequential but rather as continuous and interactive.

However, the fairground is naturally distracting, with many co-occurrences irrelevant to the sharing schema. Here, there may be an advantage to the fact that activating a schema with low situational systematicity is a gradual process—we hypothesize that mPFC remains active throughout this process, resulting in sustained inhibition of HPC (given the complementary relationship between mPFC and HPC described above). This sustained inhibition suppresses the encoding of irrelevant co-occurrences, increasing the signal-to-noise ratio. This lessens the potential for distraction away from potential to-be-shared objects, thereby increasing the

³ We view as equivalent an approach in which a single concept such as *sharing* can refer to a range of situations and an approach in which *sharing* is a complex concept comprising several *kinds* of sharing (essentially corresponding to different, but overlapping, schemas); cf. *polysemy*.

likelihood that, when additional information is apprehended, we successfully understand the nature of *sharing* in this context.

To give an example of a concept with relatively higher situational systematicity, if we later hear someone call out, “*Bill, I have a chair for you!*” largely the same processes would unfold, but the demands on the components would differ. That is, AG would partially activate the *chair* schema which, in this case, will match a perceptible object *chair*, meaning that any mPFC-mediated direction of attention to schema-relevant elements for pattern completion will be rapidly accomplished, and mPFC need not remain active. Thus, there will be no *sustained* inhibition of HPC, leaving it free to encode arbitrary elements of the situation (e.g., the color of the chair).

In line with our claims about the role of the mPFC, a recent meta-analysis by Desai et al. (2018) found that both a temporal parietal region (containing the AG) and the very same ventral portion of mPFC that has been implicated in processing schema-congruent information (van Kesteren et al., 2013) are active when processing abstract concepts (presented as words). Moreover, this particular ventral portion of mPFC stands out as a region that emerges when performing activation likelihood estimation meta-analyses on tasks that target abstract but not concrete concepts (Desai et al., 2018, Figure 3).⁴

One corollary of our view is that if arbitrary episodic relations are suppressed in the case of abstract concepts (due to schema-based and mPFC-modulated suppression of irrelevant detail), and more so than in the case of concrete concepts, it will be harder to recall irrelevant episodic details for situations that activate abstract concepts. Indeed, we have observed evidence

⁴ Although most of the studies included in the meta-analysis conducted by Desai et al. (2018) did not directly compare concrete and abstract concepts, the finding that we refer to here (depicted in Figure 3 of Desai et al.) compares results of an ALE meta-analysis collapsing across tasks that target several domains of abstract concepts with an ALE meta-analysis that attempts to include only tasks and contrasts targeting concrete concepts.

that this may be true: in a recent study (Davis et al., 2019), arbitrary episodic detail was operationalized as an arbitrary context paired with some stimulus of interest, for instance, an arbitrarily colored box (e.g., red or green) surrounding a word, or different speakers presenting each word from a list (e.g., male or female). Memory for the context (whether boxes or voices) was worse for abstract compared to concrete concepts, and when we simply presented those concepts in either the same or a different context (here, box color) from that seen at encoding, people were *worse* at recognizing abstract (but not concrete) concepts when they were presented in the *same* context than when presented in a different one during a recognition phase. Abstract concepts therefore seem to be less effective cues to arbitrary episodic details, and in fact, appear to be suppressed in the context of arbitrary details.

One consequence of this last finding is that, if irrelevant episodic detail (arbitrary co-occurrence) is the basis for experiencing tokenized instances of a concept (see also Altmann & Ekves, 2019), it should be harder to encode and/or recall an *instance* of a situation associated with an abstract concept—more so than for a concrete concept—because the arbitrary episodic details associated with that instance will be harder to recall (see e.g. Schwanenflugel et al., 1992). Thus, although we have claimed that abstract concepts rely more on *broad*er episodic details than do concrete concepts, we end with an account that predicts that this increased reliance results in poorer encoding/recall of *individual* episodic details concerning the situation(s) to which an abstract concept applies. Further research is required to address this prediction and its implications for how abstract concepts are acquired in the first place.

Implications for Situated Conceptualization

While the present discussion has been couched in the divide between abstract and concrete concepts, the overarching goal is to explain these differences not in terms of

abstractness *per se*, but in terms of the systematicity of the situations in which we experience those concepts (i.e., schemas). Another recent proposal, based on the situated conceptualization framework (e.g., Barsalou, 2009), and referred to as the *brain as a situation processing architecture* (BASPA; Barsalou et al., 2018) has also sought to move past this ontological distinction between abstract and concrete by considering two different dimensions: *Situational elements* (which include external elements such as settings, agents, objects, actions, and outcomes, as well as internal elements such as emotions, motivations, and other internal states) and *situational integrations* which relate and integrate across situational elements in a given situation. These two dimensions—situational elements and integrators—provide a useful framework for characterizing situated conceptual processing, and like BASPA, our approach highlights that concrete concepts, as compared to abstract concepts, involve situations that are more systematic (or in BASPA terms, more “predictable”).

BASPA views concrete concepts as originating in the processing of external elements, and contrasts them with abstract concepts which are viewed as originating in the processing of internal elements and elements that must be integrated across. However, our view is that even external elements involve integration—for example, when considering *actions* and *outcomes* as external elements, what is an outcome, or indeed the action causing that outcome, if not an *integration* of elements (and their changing states; Altmann & Ekves, 2019) across time? And what are a concrete object’s affordances if not integrations, again, across time, space, and the elements (both internal and external) occupying that space and time? Thus, we suggest that all concepts, concrete and abstract, involve situational integrations, and that the difference between “more abstract” and “more concrete” concepts lies in the extent to which the to-be-integrated situation is more or less systematic.

Our approach also extends the situated conceptualization framework by providing a neurobiological mechanism to manage the interplay within any given episode between elements which are systematically related to the situated conceptualization and elements which are present only arbitrarily (and which contribute to the *episodic experience* of that instance of the concept). Specifically, we have argued that the likelihood of elements being down-regulated in any given instantiation of a concept increases as a function of (1) their arbitrariness with respect to the situated conceptualization and (2) the situational systematicity of the (abstract) concept (or situational integration) at play.

The idea that situational systematicity varies across (and within) concepts also implies variability in the *spatial* and *temporal windows* over which concepts are learned, employed, recognized, and situated. There is some evidence that increasingly abstract language is used to traverse increasingly large spatial and temporal distances: the abstractness of natural language use on Twitters increases as a function of the spatial distance between an individual tweeter and a referent location, and as a function of the temporal distance between the present and a referenced point in time (Sneffjella & Kuperman, 2015). This accords with work in other domains suggesting that abstraction functions to support “mental travel” across progressively greater spatial and temporal distances (see Gilead et al., 2019; Trope & Liberman, 2010). And if, as we conjecture, the recruitment of top-down control mechanisms in learning and processing abstract concepts has to do, at least in part, with the duration and variability of this spatiotemporal window, then the dynamic activity of the mPFC–HPC network should track a measure of this window in online situated conceptualization—with greater engagement of the network when recognizing a concept whose constituent features are usually more dispersed through space and/or time. For example,

⁵ This was operationalized as the average abstractness in tweets containing at least four words from Brysbaert et al.’s (2014) concreteness norms.

the constitutive features of the concept *bread* can be recognized within a small spatiotemporal window, but *justice* requires a much broader window, and *sharing* a window somewhere between the two. Even within concrete concepts, one might expect constitutive features of *chair* to be recognized in a more circumscribed spatial window than those of *restaurant* (in addition to requiring a smaller *number* of the perceptual elements in a particular scene, as described earlier). Critically, mPFC–HPC engagement should reflect prior experience of the apprehension of the concept’s constitutive features from the environment.

Experimentally manipulating these factors in a learning context would be an important next step: for example, to-be-learned concepts could be varied in respect of the degree to which to-be-integrated elements are spread over physical space and time. Not only would mPFC be critical in tracking systematicities that are more spread over space and time, but as the presence of arbitrary situational elements increases—which would inevitably happen as consequence of the constitutive features of a concept being more sparsely distributed—mPFC would down-regulate HPC and suppress encoding of those arbitrary elements.

Concluding Remarks

Thinking of concepts as schema, and of whether the situations to which they apply can be recognized purely bottom-up (based on sensory experience) or whether they more likely rely on top-down (previously stored) knowledge to direct attention within and across the relevant situations, provides testable predictions about the neural mechanisms by which they operate. We have proposed that these mechanisms reflect interactive activity within a schema control network, i.e., within brain regions critical to navigating situational systematicity. Navigating situational systematicity involves understanding the relationship between concepts, pre-existing schema knowledge, and the spatiotemporally distributed environmental cues that contribute to

the activation of situationally appropriate schema knowledge. The brain regions of the schema control network include AG (activation of schema knowledge), left IFG (selective attention toward situation-relevant schema knowledge), mPFC (encoding of, and hence direction of attention towards, schema-relevant relationships between objects and their contexts), and HPC (“blind” encoding of relationships, relevant or otherwise, between objects and their contexts).

We have not attempted here to provide an account of the *content* of abstract concepts (see Wilson-Mendenhall et al., 2013, for a study that does, and Binder et al., 2016; Borghi et al., 2017 for relevant reviews); rather, we have explored the functional significance of some of the neurobiological mechanisms that underpin the interplay between one particular aspect of semantic memory (i.e., schema knowledge, as mediated by AG and mPFC) and one particular aspect of episodic memory (served by HPC, and which mediates the experience of the situations to which such schema-based knowledge can be applied). We have, in effect, translated this functional significance into expected functional differences between abstract and concrete concepts in respect of how they are grounded in actual, encountered situations. Further research is required to explore these functional differences, their relation to other associative mechanisms elsewhere in the brain implicated in conceptual knowledge (for discussion, see e.g., Barsalou et al., 2018; Pulvermüller, 2018), and their generalizability to different kinds of abstract concepts that appear to be served by different brain regions (for review, see e.g., Desai et al., 2018).

With traditional views on a dichotomy between abstract and concrete concepts fading, much work remains to be done to understand the neurocognitive mechanisms by which the situational and spatiotemporal dynamics underpinning conceptual knowledge operate, and how conceptual knowledge is learned in respect of these dynamics. Considering the situational systematicity of concepts—and the neurobiology of how information is integrated into

schemas—is a useful and productive alternative to the traditional perspective that distinguishes between “abstract” and “concrete” concepts.

References

- Altmann, G. T. M. & Ekves, Z., (2019). Events as intersecting object histories: A new theory of event representation. *Psychological Review*. doi:10.1037/rev0000154
- Badre, D., & Wagner, A. D. (2005). Frontal lobe mechanisms that resolve proactive interference. *Cerebral Cortex*, 15(12), 2003–2012.
- Barsalou, L. W. (1999). Perceptions of perceptual symbols. *Behavioral and Brain Sciences*, 22(04), 637–660.
- Barsalou, L. W. (2009). Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521), 1281–1289.
- 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), 20170144.
- Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In *Symbols, embodiment, and meaning* (pp. 245–283).
- Bartlett, F. C. (1932). *Remembering: An experimental and social study*. Cambridge, UK: Cambridge University.
- Barwise, J., & Perry, J. (1991). Situations and attitudes. *Noûs*, 25(5), 743–770.
- Binder, J. R., Conant, L. L., Humphries, C. J., Fernandino, L., Simons, S. B., Aguilar, M., & Desai, R. H. (2016). Toward a brain-based componential semantic representation. *Cognitive Neuropsychology*, 3294, 1–45.
- Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, 15(11), 527–536.

- Binder, J. R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, *19*(12), 2767–2796.
- Borghi, A. M., & Binkofski, F. (2014). *Words as social tools: An embodied view on abstract concepts*. New York, NY: Springer.
- Borghi, A. M., Binkofski, F., Castelfranchi, C., Cimatti, F., Scorolli, C., Borghi, A. M., ... Tummolini, L. (2017). The challenge of abstract concepts. *Psychological Bulletin*, *143*(3), 263–292.
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, *46*(3), 904–911.
- Cohen, N. J., & Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system*. Cambridge, MA: MIT Press.
- Davachi, L. (2006). Item, context and relational episodic encoding in humans. *Current Opinion in Neurobiology*, *16*(6), 693–700.
- 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*. doi:10.1080/23273798.2018.1479530
- Davis, C. P., Paz-Alonso P. M., Altmann, G. T. M., & Yee, E. (2019). Abstract concepts and the suppression of arbitrary episodic context. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (p. 1592). Montreal, QC: Cognitive Science Society.
- Desai, R. H., Reilly, M., & van Dam, W. (2018). The multifaceted abstract brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1752), 20170144.

- Dove, G. (2016). Three symbol ungrounding problems: Abstract concepts and the future of embodied cognition. *Psychonomic Bulletin & Review*, 23(4), 1109–1121.
doi:10.3758/s13423-015-0825-4
- Dove, G. (2018). Language as a disruptive technology: Abstract concepts, embodiment and the flexible mind. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170135.
- Galbraith, R. C., & Underwood, B. J. (1973). Perceived frequency of concrete and abstract words. *Memory & Cognition*, 1(1), 56–60.
- Geschwind, N. (1972). Language and the brain. *Scientific American*, 226(4), 76–83.
- Ghosh, V. E., & Gilboa, A. (2014). What is a memory schema? A historical perspective on current neuroscience literature. *Neuropsychologia*, 53, 104–114.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton-Mifflin.
- Gilead, M., Trope, Y., & Liberman, N. (2019). Above and beyond the concrete: The diverse representational substrates of the predictive brain. *Behavioral and Brain Sciences*, 1–63.
- Gilboa, A., & Marlatte, H. (2017). Neurobiology of schemas and schema-mediated memory. *Trends in Cognitive Sciences*, 21(8), 618–631.
- Glenberg, A. M. (1997). What memory is for. *Behavioral and brain sciences*, 20(1), 1–19.
- Gureckis, T., & Goldstone, R. L. (2010). Schema. In P. C. Hogan (Ed.), *The Cambridge encyclopedia of language science*. Cambridge, UK: Cambridge University Press.
- Hafri, A., Papafragou, A., & Trueswell, J. C. (2013). Getting the gist of events: Recognition of two-participant actions from brief displays. *Journal of Experimental Psychology: General*, 142(3), 880–905. doi:10.1037/a0030045

- Hoffman, P. (2016). The meaning of 'life' and other abstract words: Insights from neuropsychology. *Journal of Neuropsychology*, 10(2), 317–343.
- Hoffman, P., Binney, R. J., & Lambon Ralph, M. A. (2015). Differing contributions of inferior prefrontal and anterior temporal cortex to concrete and abstract conceptual knowledge. *Cortex*, 63, 250–266.
- Hoffman, P., Lambon Ralph, M. A., & Rogers, T. T. (2013). Semantic diversity: A measure of semantic ambiguity based on variability in the contextual usage of words. *Behavior Research Methods*, 45(3), 718–730.
- Hoffman, P., McClelland, J. L., & Lambon Ralph, M. A. (2018). Concepts, control and context: A connectionist account of normal and disordered semantic cognition. *Psychological Review*, 125(3), 293–328.
- Lambon Ralph, M. A., Jefferies, E., Patterson, K., & Rogers, T. T. (2017). The neural and computational bases of semantic cognition. *Nature Reviews Neuroscience*, 18(1), 42.
- McRae, K., Nedjadrasul, D., Pau, R., Lo, B. P. H., & King, L. (2018). Abstract concepts and pictures of real-world situations activate one another. *Topics in Cognitive Science*, 1–15.
- Murphy, G. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- Noppeney, U., & Price, C. J. (2004). Retrieval of abstract semantics. *Neuroimage*, 22(1), 164–170.
- Paivio A. (1971). *Imagery and verbal processes*. New York, NY: Holt, Rinehart & Winston.
- Paivio, A. (1991). Dual coding theory: Retrospect and current status. *Canadian Journal of Psychology*, 45(3), 255–287.
- Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory*, 2(5), 509–522.

- 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.
- Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., & Patterson, K. (2004). Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review*, 111(1), 205–235.
- Rugg, M. D., Vilberg, K. L., Mattson, J. T., Sarah, S. Y., Johnson, J. D., & Suzuki, M. (2012). Item memory, context memory, and the hippocampus: fMRI evidence. *Neuropsychologia*, 50(13), 3070–3079.
- Rumelhart, D. E., & Ortony, A. (1977). The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, W. E. Montague (Eds.), *Schooling and the acquisition of knowledge* (pp. 99–135). Hillsdale, NJ: Lawrence Erlbaum Associates
- Schwanenflugel, P. J. (1991). Why are abstract concepts hard to understand? In P. J. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 223–250). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Schwanenflugel, P. J., Akin, C., & Luh, W. M. (1992). Context availability and the recall of abstract and concrete words. *Memory & Cognition*, 20(1), 96–104.
- Schwanenflugel, P. J., & Shoben, E. J. (1983). Differential context effects in the comprehension of abstract and concrete verbal materials. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(1), 82–102.

- Snefjella, B., & Kuperman, V. (2015). Concreteness and psychological distance in natural language use. *Psychological Science*, 26(9), 1449–1460. doi:10.1177/0956797615591771
- Thompson-Schill, S. L. (2003). Neuroimaging studies of semantic memory: Inferring “how” from “where”. *Neuropsychologia*, 41, 280–292.
- Thompson-Schill, S. L., D’Esposito, M., Aguirre, G. K., & Farah, M. J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: a reevaluation. *Proceedings of the National Academy of Sciences*, 94(26), 14792–14797.
- Thompson-Schill, S. L., Swick, D., Farah, M. J., D’Esposito, M., Kan, I. P., & Knight, R. T. (1998). Verb generation in patients with focal frontal lesions: a neuropsychological test of neuroimaging findings. *Proceedings of the National Academy of Sciences of the United States of America*, 95(26), 15855–15860.
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440–463.
- Tulving, E. (1983). Ecphoric processes in episodic memory. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 302(1110), 361–371.
- van Kesteren, M. T. R., Beul, S. F., Takashima, A., Henson, R. N., Ruiter, D. J., & Fernández, G. (2013). Differential roles for medial prefrontal and medial temporal cortices in schema-dependent encoding: From congruent to incongruent. *Neuropsychologia*, 51(12), 2352–2359.
- van Kesteren, M. T. R., Ruiter, D. J., Fernández, G., & Henson, R. N. (2012). How schema and novelty augment memory formation. *Trends in Neurosciences*, 35(4), 211–219.
- Wagner, A. D., Paré-Blagoev, E. J., Clark, J., & Poldrack, R. A. (2001). Recovering meaning: left prefrontal cortex guides controlled semantic retrieval. *Neuron*, 31(2), 329–338.

- Wang, J., Conder, J. A., Blitzer, D. N., & Shinkareva, S. V. (2010). Neural representation of abstract and concrete concepts: A meta-analysis of neuroimaging studies. *Human Brain Mapping, 31*(10), 1459–1468.
- Wilson-Mendenhall, C. D., Simmons, W. K., Martin, A., & Barsalou, L. W. (2013). Contextual processing of abstract concepts reveals neural representations of nonlinguistic semantic content. *Journal of Cognitive Neuroscience, 25*(6), 920–935.
- Wittgenstein, L. (1953/2010). *Philosophical investigations*. John Wiley & Sons.
- Yee, E., & Thompson-Schill, S. L. (2016). Putting concepts into context. *Psychonomic Bulletin & Review, 23*(4), 1015–1027.