



# Typicality: Stable structures and flexible functions

Michael A. Dieciuc<sup>1</sup> · Jonathan R. Folstein<sup>1</sup>

Published online: 27 November 2018  
© Psychonomic Society, Inc. 2018

## Abstract

Some research suggests typicality is stable, other research suggests it is malleable, and some suggests it is unstable. The two ends of this continuum—stability and instability—make somewhat contradictory claims. Stability claims that typicality is determined by our experience of decontextualized feature correlations in the world and is therefore fairly consistent. Instability claims that typicality depends on context and is therefore extremely inconsistent. After reviewing evidence for these two claims, we argue that typicality’s stability and instability are not contradictory but rather complementary when they are understood as operating on two different levels. Stability reflects how information gets encoded into semantic memory—what we call *structural typicality*. Instability reflects the task-dependent recruitment of semantic knowledge—what we call *functional typicality*. Finally, we speculate on potential factors that may mediate between the recruitment of structural or functional typicality.

**Keywords** Typicality · Categorization

Our conceptual knowledge of the world is both a reflection of the world’s inherent structure *and* a distortion of it. It would surprise no one to say that trees and bushes share more physical properties with each other than trees and cars. Thus, it shouldn’t surprise anyone to find that trees and bushes are *conceptually* more similar to each other than trees and cars. In this way, the structure of our experienced environment impresses itself onto the conceptual structure of our mind/knowledge. However, our needs, goals, and contexts may require us to flexibly reorganize our conceptual information. Despite sharing very few physical properties, trees and cars may temporarily become more conceptually similar to one another than trees and bushes if I am looking for something heavy to tie a rope to. How do these differences in conceptual organization relate to one another? We will be examining this question via the lens of typicality research. Throughout this paper, we aim to do three things: one, point out that there is a tension in typicality research; two, offer a theoretical framework and terminology for understanding and potentially reconciling the tension; and three, speculate on

mechanisms that may be responsible for this tension. Before elaborating on our framework, we provide a brief introduction to typicality research at large.

## What is typicality?

In a very real sense, everything in our environment is unique. Fortunately, we can overcome the “blooming, buzzing, confusion” of uniqueness by categorizing the things we encounter according to some perceived equivalence. When people categorize, they take two or more things and place them into the same group despite their differences. In other words, categorizing minimizes the differences between things while maximizing their similarities. By default, the act of categorizing nonidentical things according to some perceived equivalence *necessarily* means that there will be variability between members. As a result of this variability, we recognize that some members of a category are better than others. For instance, sparrows, crows, penguins, and ostriches are all birds, but they are also quite clearly nonidentical. We might generally agree that sparrows and crows are more representative of birds in general than are penguins and ostriches. Typicality is simply the recognition that exemplars differ in how representative they are of the category and of their fellow members. Some exemplars are good examples, and some are not. In more formal terms, typical exemplars are highly representative of their category, and atypical exemplars are not.

---

✉ Michael A. Dieciuc  
dieciuc@psy.fsu.edu

<sup>1</sup> Department of Psychology, Florida State University, 1107 W. Call St., Tallahassee, FL 32306, USA

There are a number of different ways in which the typicality of an exemplar is assessed. The most common method involves presenting participants with an exemplar and a category and having them make “goodness-of-example” (also known as goodness-of-fit) ratings. For example, participants could be presented with a penguin and asked, “On a scale of 1 to 7, is this a good example of birds?” (Rosch & Mervis, 1975). The higher the rating, the more typical the exemplar. Another related method is to give participants a category and ask them to produce exemplars: “List as many birds as you can.” Typical members are produced more often and earlier than atypical members (Mervis, Catlin, & Rosch, 1976; Mervis & Rosch, 1981). Another method is to give participants a category (e.g., bird) and have them list features of that category (wings, beak, feathers, etc.). Typical members will have more features that overlap with that list, whereas atypical members have fewer features (Rosch & Mervis, 1975). Both production and feature-listing tasks predict goodness-of-example ratings. Generally speaking, exemplars that are produced more often, sooner, and share a higher number of features will also have higher goodness-of-example ratings.

Differences in typicality are important because they affect virtually all aspects of categorization. Relative to atypical exemplars, typical exemplars are learned better (Mervis & Pani, 1980; Rosch, Simpson, & Miller, 1976b), verified faster (McCloskey & Glucksberg, 1978; Smith, 1978), produced more often (Barsalou, 1983, 1985; Mervis et al., 1976), and generalized further during reasoning (Osherson, Smith, Wilkie, Lopez, & Shafir, 1990; Rips, 1975). Furthermore, these typicality effects are pervasive. They appear in different kinds of categories: abstract (Hampton, 1981), natural kinds (Osherson et al., 1990), feature based (Goldwater, Markman, & Stilwell, 2011), role governed (Rein, Goldwater, & Markman, 2010), ad hoc (Barsalou, 1983), and goal derived (Barsalou, 1985).

## Structure and function

As explained below, the totality of typicality research can be categorized into two different threads. There are those who argue that typicality is relatively stable because it reflects a structure of shared “features,” most of which are available from our experiences of the world, and there are those who argue that typicality is radically unstable because it is relative to the idiosyncratic context in which it is instantiated. For example, when Lambon Ralph, Sage, Jones, and Mayberry (2010), say that “the same core information is activated each time an entity is encountered” (p. 2719), they are implying that typicality is relatively static because the same core information determines the graded structure of the category each time. Conversely, Lebois, Wilson-Mendenhall, and Barsalou (2015) say that “concepts have no cores at all” (p. 1773), implying that typicality is not static. Similarly, Roth and

Shoben (1983) argue that “[typicality] undergoes a complete restructuring once context is introduced” (p. 369).

In this paper, we suggest that this disagreement over whether typicality is stable or unstable actually arises because the two camps are referring to slightly different (but related) theoretical constructs. We suggest that the relative stability of typicality is a reflection of how consistent correlations between features largely available from the world become encoded into long-term memory (what we call *structural typicality*) and that the radical instability of typicality is a reflection of how subsets of information in long-term memory are recruited in a task-dependent and context-specific manner in working memory (what we call *functional typicality*). Typicality ratings driven by structural typicality result from easy access to exemplars in long-term memory that share many features with fellow category members and few features with contrast categories—in other words, exemplars with high “cue validity.” In contrast, typicality ratings driven by functional typicality result from a good “fit” between an exemplar and the goals of the observer with respect to a particular situation. Our observation that concepts may involve two different components is not necessarily novel. It has been alluded to by many others (Barsalou, 1985, 1987, 2009; Lambon Ralph, Jefferies, Patterson, & Rogers, 2016; Rogers & McClelland, 2004). However, we believe that our terms *structural typicality* and *functional typicality* are a parsimonious way of framing the issue that helps highlight what the issue is, where the differences lie, and how the two may be reconciled.

A note of clarification is in order before we continue. Throughout this paper, we discuss the role the environment plays on our conceptual knowledge. To be clear: When we speak of the environment, we are speaking of one’s *experienced* environment. Obviously, the feature correlations of the arctic cannot influence my conceptual structure if I have never visited or studied the arctic. More importantly, the differences between environments will account for some of the differences in conceptual structure between people. It would be uninteresting to point out that someone living in Manhattan had a different conceptual structure of birds than someone living in the Amazonian jungle. Thus, some of the variability in typicality research could be due to the different feature correlations of different environments. However, what we hope to show is that there is extensive research suggesting that there is considerable consistency but, importantly, also considerable variability in the conceptual structure of people who share the *same* environment.

## Stable structures: Continued support for the Roschian view

Having briefly outlined our framework, we now turn to evidence in support of it. In this section, we review research

which suggests that typicality is stable and anchored in the observer's experienced world rather than in the goals of the observer. As a reminder, by stable we mean that typicality is relatively impervious to change; it is relatively consistent across individuals, time, and contexts.

## Rosch and family resemblance

The most prominent and well-known approach to typicality is that of family resemblance (Rosch & Mervis, 1975; Hampton, 1979, 1995).<sup>1</sup> This approach stems from the belief, expressed by Rosch and Mervis (1975), that our conceptual structure mirrors our experience of the structure of our environment:

The basic category cuts in the world are those which separate the information- rich bundles of attributes which form natural discontinuities. . . . The present study has shown that formation of prototypes of categories appears to be likewise nonarbitrary. The more prototypical a category member, the more attributes it has in common with other members of the category and the less attributes in common with contrasting categories. Thus, prototypes appear to be just those members of the category which most reflect the redundancy structure of the category as a whole. That is, categories form to maximize the information-rich clusters of attributes in the environment and, thus, the cue validity of the attributes of categories. (p. 602).

Certain features tend to naturally co-occur in the environment: animals that have beaks tend to have feathers; things that are hard tend to be heavy. These correlations of features form natural “clusters” of things in our environment: trees, birds, cars. Long-term experience with these natural clusters shapes our conceptual structure. Early theories of category learning proposed that an observer's repeated exposure to these feature correlations forms a central tendency in their mind, a prototype, that is the extraction/summary of their average experience with each cluster of features (e.g., Posner & Keele, 1968). These prototypes are then used as comparison points for categorization. If we see a new dog (e.g., a dachshund or a golden retriever), we compare it to the prototype of all previously experienced dogs. The more features it has in common with the prototype, the more typical that dog is of other dogs.

Exemplar theories (Medin & Shaffer, 1978) and connectionist models (Rumelhart & Todd, 1993) soon emerged to challenge prototype theories, but even though they did not explicitly represent prototypes, they too predicted that that ease of categorization should often be sensitive to overall

similarity, that is, family resemblance, to the central tendency (e.g., Nosofsky & Johansen, 2000; Nosofsky & Zaki, 2002; Rogers & McClelland, 2004). In this context, similarity refers to the proportion of feature overlap and can be broken down into two different but interrelated components: within-category-similarity and between-category-similarity (Rosch & Mervis, 1975). These two components determine an exemplar's degree of typicality. For example, a rabbit is a highly typical mammal because it shares many features found in most, but not all, mammals: fur, warm-blooded, produces milk. In contrast, a hairless rabbit is an atypical mammal because it has low within-category similarity: It is missing a feature (fur) that is often found in other mammals. On the other hand, a platypus is an atypical mammal because it has high between-category similarity: Its webbed feet and duck bill are features often found in other animals (amphibians and birds). Notice that this approach emphasizes that typicality has to do with the correlation and interrelation of features found in the world and has very little to do with the observer's goals or context.

The notion that the structure of our mind mirrors our experience of the structure of the environment is supported by fMRI research. Jordan, Greene, Beck, and Fei-Fei (2016) used a technique in fMRI known as multivoxel pattern analysis (MVPA) to demonstrate that typicality is actually a principle of neural organization. MVPA is a technique that measures distributed patterns of activity across the brain and then allows the patterns to be compared with one another. Jordan et al. (2016) showed participants images from eight common taxonomic categories (e.g., cats, birds, cars, planes), recorded the pattern of activity across the brain for each category, created a prototype for each category by averaging the voxel patterns of all the stimuli, and then compared the patterns of typical and atypical exemplars with that of the prototype. They found that typical exemplars were more similar to the prototype than were atypical exemplars. Conversely, they found that, compared with typical exemplars, atypical exemplars were more similar to the prototype of other categories. The authors interpret this as neural evidence for Rosch's original distinction between within-category similarity (feature overlap with category members) and between-category similarity (feature overlap with members of other categories).

In short, this study is consistent with the notion that typicality in object selective areas of the brain (occipitotemporal cortex) is a reflection of relatively stable typicality differences in the world. One reason for this is that the neural activity was measured from brain areas representing perceptual features. The stability of perceptual representations across contexts is of course a matter of degree (e.g., Folstein, Palmeri, & Gauthier, 2013; Li, Ostwald, Giese, & Kourtzi, 2007), but it can at least be said that perceptual features are closely related to features in the world even if they can be distorted (e.g., Freedman, Riesenhuber, Poggio, & Miller, 2003). Thus,

<sup>1</sup> Please note, we use the terms *family resemblance*, *prototype*, and *central tendency* interchangeably.

evidence of family resemblance in perceptual cortex is particularly compelling because perceptual cortex is not as susceptible to reflecting differences in goals. A second reason this research is consistent with our notion of stable typicality is that the behavioral typicality ratings that predicted neural typicality were elicited without reference to any particular context,<sup>2</sup> so it is parsimonious to assume that context was not a critical variable in driving the typicality ratings. This assumption has occasionally been challenged, and we revisit it below, but the close relationship between typicality ratings and perceptual features suggests to us that it holds to some degree.

### Semantic dementia and the anterior temporal lobe

In addition to the visual areas (occipitotemporal regions of cortex) mapped out by Jordan et al. (2016), research in semantic dementia suggests that the anterior temporal pole is also critically involved in representing stable semantic structures which, as Roschian typicality effects predict, are organized according to similarity. Semantic dementia is a neurodegenerative disorder associated with atrophy of the anterior temporal pole, especially ventrolateral regions, which results in systematic and specific deficits in semantic memory (e.g., Mayberry, Sage, & Lambon Ralph, 2011). Individuals with semantic dementia present difficulties on both verbal and nonverbal tasks that recruit knowledge of categories. These tasks include reading and spelling fluency, word–nonword lexical decisions, object naming, object sorting, feature generation, and drawing (Patterson, 2007). Importantly, the pattern of performance on all these tasks is sensitive to typicality. As the severity of semantic dementia increases, typical exemplars are more robust against impairment than atypical exemplars.

Collectively, these studies suggest that the anterior temporal lobe may be structurally organized according to our experience of relatively stable feature correlations of our environment (Lambon Ralph et al., 2010). For example, Mayberry et al. (2011) found that individuals with semantic dementia simultaneously made undergeneralization and overgeneralization errors that were sensitive to typicality during a simple categorization task. Atypical stimuli that should have been categorized were not (e.g., a penguin was not categorized as a bird) and “pseudotypical” nonmembers were incorrectly categorized (e.g., a butterfly was categorized as a bird). This pattern of results was replicated in a nonverbal picture-drawing task (Bozeat et al., 2003). When asked to draw an exemplar from memory, patients with semantic dementia made systematic errors. For instance, when drawing atypical

exemplars, the distinctive features were omitted (e.g., a rhino with no horn, a camel with no hump), while the properties of typical exemplars were overextended (e.g., a duck with four legs). Importantly, Rogers, Patterson, Jefferies, and Lambon Ralph, (2015) show that typicality in these varied tasks affects performance over and above other factors that are also known to influence semantic processing (i.e., familiarity and specificity). Finally, these typicality effects have also been found in healthy controls. Woollams (2012) used repetitive transcranial magnetic stimulation (rTMS) to create a temporary “virtual lesion” in the anterior temporal lobe of healthy controls and then had them perform picture-naming tasks. Participants were slower to name atypical exemplars than typical exemplars after rTMS but not before, a pattern that matches those of individuals with semantic dementia.

The various studies just reviewed demonstrate that the pattern of deficits in semantic dementia is systematically related to typicality: Typical exemplars and features are more robust against impairment than atypical exemplars and features. This suggests three broad conclusions: One, the anterior temporal pole is important in processing semantic knowledge; two, the organization of semantic knowledge is sensitive to typicality; and three, typicality is a relatively stable feature of semantic memory that influences cognitive processing across different tasks and modalities

Evidence from ERPs also supports the notion that category structure, and therefore typicality structure, is stably stored in long-term memory. The N400 is a negative polarity event-related potential that occurs around 300–400 ms and is increased when access to semantic information is difficult (Kutas & Federmeier, 2011). Highly typical members elicit a smaller N400 amplitude and highly atypical members elicit a greater N400 amplitude (Fujihara, Nageishi, Koyama, & Nakajima, 1998; Heinze, Munte, & Kutas, 1998; see also Federmeier & Kutas, 1999).

Consistent with the idea that structural typicality is a reflection of the experienced world, some studies have shown that typicality does not change in response to changes in the observer’s mindset. For instance, Kim and Murphy (2011) found that asking participants to adopt different mental perspectives altered judgments of ideals but did *not* alter judgments of typicality, suggesting that typicality was probably driven by central tendency. Similarly, Hampton, Dubois, and Yeh (2006) found that, for most categories, the ability of acontextual typicality judgments to predict how people would categorize was unaffected when participants were instructed to take various mental perspectives that should have altered their behavior. These studies demonstrate that central tendency is a powerful determinant of typicality, one strong enough to overcome the mindset of the observer.

### Interim summary of stability

Behavioral, neuroimaging, semantic dementia, and ERP evidence all show that there is a systematic pattern to the way

<sup>2</sup> Throughout the paper, we refer to *acontextual typicality judgments*, or *typicality in isolation*. In a sense, no decision or judgment is made without a context (Casasanto & Lupyan, 2015), but what we mean is that typicality judgments in response to labels alone are highly unconstrained contexts. Examples of constraining contexts are given later in the paper, and we speculate on likely mechanisms of contextual constraint in the final section.

information is organized neurocognitively. That pattern is typicality: Neural representations form clusters of similar exemplars corresponding to categories, and exemplars near the center of those clusters are easiest to access. To the degree that typicality is driven by the number of features shared with fellow category members, it should be a fairly stable phenomenon and, as some studies suggest, relatively resistant to changes in task and context (e.g., Patterson, 2007; Rogers et al., 2015). This is because the correlation of features in the environment does not change drastically or quickly within a single observer.

The view that typicality relies on feature sharing implies that the role played by the observer and any observer-related variables should be relatively small. This is because the number of features an exemplar shares with its category's central tendency does not depend on the observer or their mindset. In the following sections, we review evidence that typicality does in fact depend on observer variables, supporting our claim that typicality can be driven by a separate set of cognitive structures. The most convincing cases, reviewed last, show that typicality can be manipulated within a single observer. First, however, we turn to demonstrations that different observers, often from different cultures, have different typicality structures and that these structures are not driven by central tendency but by goals and cultural significance.

### **Malleability: Experience with conceptual domains recruits ideals**

The previous section reviewed research showing that typicality is fairly stable because it reflects stable feature correlations in the world. The current section reviews research showing that a stable world is only part of the equation. The other part of the equation is our interaction and engagement with the world. Although two people may live in the same environment, their engagement with that environment can be extremely different. People may be driven by different goals, have different areas of expertise, or different attentional biases. These factors can alter long-term memory gradually over long periods of time and also affect how memories are retrieved in certain situations. Specifically, this section will review how typicality is affected by culture and expertise, both of which influence observers to shift away from the use of central tendency and toward representations called ideals as critical points of comparison for categorization. As a result of this shift, exemplars similar to ideals become more typical than exemplars similar to central tendency.

What precisely is an ideal? Ideals are exemplars possessing characteristics that facilitate the functional purpose of the category (Barsalou, 1985). Thus, it has been argued that ideals are representative, not because they share features with the most fellow category members and the least nonmembers

but because they fit well with a set of goals. This is most apparent in explicitly goal-derived categories, such as “things to take camping,” discussed in more detail in a later section. One likely reason a tent is typical camping equipment is not because it shares many features with a flashlight and a grill, but because it is portable and provides shelter, both of which are important for camping. A second likely reason is that a tent is almost always encountered in camping situations and is thus retrieved from long-term memory via association with a well-learned context or schema. By either explanation, a tent is a typical thing to take camping, not because it shares features with other things to take camping but because it fits with the camping scenario. Many taxonomic categories, like “athlete,” are also associated with goals and therefore have ideals that differ from central tendency. According to family resemblance, the typical athlete is the one that runs the average speed, jumps the average height, and has the average strength (see below for some important caveats to this). According to ideals, the typical athlete is the one that is fastest, most agile, and strongest because these characteristics will maximize success. Thus, it has been noted that ideals will often have extreme values rather than average values because extreme values are often best for achieving goals (Barsalou, 1985).

### **Ideals and culture**

The use of ideals is highly influenced by the culture of the observer. Ojalehto and Medin (2015) define culture as “a way of life, often equated with shared knowledge or what one needs to know to live successfully in a community” (p. 250). Ojalehto and Medin (2015) argue that conceptual structure is inseparable from culture; culture seeps into our experience, and our experience forms our concepts. Culture influences how a person engages with her environment, thereby influencing what her goals, and therefore her ideals, are.

As reviewed in the previous section, there is a large body of research supporting central tendency as the key determinant of typicality. However, most of this work has used biological categories (e.g., trees, fish, birds) and undergraduate participants who may have little actual experience with these biological categories. In contrast, research using experts and people of different cultures suggests a radically different conclusion: namely, that typicality appears to be organized around ideals that incorporate nonperceptual information (Medin & Atran, 2004). Bailenson, Shum, Atran, Medin, and Coley (2002) found that ideals structured typicality for U.S. bird experts and Mayan participants, who also had considerable experience with birds, whereas U.S. undergraduates relied on central tendency. Burnett, Medin, Ross, and Blok (2005) demonstrated that not only do cultures use ideals when processing important categories but—as one would expect if different cultures are associated with different goals—what counts as an ideal differs between cultures. They compared

the typicality structure for fishing experts from two different cultures: Native American Menominee Indians and European American Wisconsinites. For both kinds of experts, typicality was determined by ideals rather than by similarity to central tendency. For Menominee fisherman, sturgeon and trout were of greater value, and therefore most typical, whereas the European American valued gamefish more, rating them as the most typical.

Similarly, Lynch, Coley, and Medin (2000) found a comparable pattern when comparing novices (undergraduates) and three different kinds of experts (taxonomists, landscapers, and park maintenance personnel) on their typicality ratings of trees. Even though trees are a common category with fairly clear perceptible features, structures, and correlations, central tendency was not the primary determinant of typicality for either novices or experts. For novices, perceived typicality was based almost solely on familiarity. For experts, ideals had the most influence on typicality. Furthermore, the nature of the ideals shifted between experts. For maintenance personnel, perceived typicality was based on the ideal of weediness. For taxonomists and landscapers, perceived typicality was based on the ideal of weediness and height.

So far the data suggest that typicality can differ between novices and experts, that it can differ between people of different cultures, and that these differences can be driven by ideals rather than by central tendency. To complicate matters even further, Medin, Lynch, Coley, and Atran (1997) demonstrate that people can have multiple, different ways of organizing their concepts, and therefore typicality. Again, using landscapers, maintenance personnel, and taxonomists, Medin et al. (1997) found that some experts (i.e., maintenance personnel) used the same conceptual structures for sorting and reasoning while other experts (i.e., landscapers) switched between different conceptual structures for sorting and reasoning.

It is important to emphasize that the results above do not merely rely on the observation that cultures differ in what they consider most typical. Although many of the studies above compared groups that lived in similar geographical areas, raw differences in typicality across groups could still be accounted for by differences in physical environment, which would result in uninteresting differences in structural typicality. This was not a concern for the studies above because central tendency was measured separately for each culture, controlling for differences in physical environment. Most importantly, central tendency either did not predict or did not uniquely predict typicality ratings. Instead, when cultures had experience with the conceptual domain, typicality was predicted by usefulness or cultural significance. For instance, tree experts rated trees more typical if they were less “weedy,” and Menominee Indians rated fish as more typical when they were sacred (cultural significance) or frequently used for food

(usefulness). This is inconsistent with typicality driven by sharing features with other category members (structural typicality) and consistent with typicality driven by fit between exemplars and goals (functional typicality).

Research on nonundergraduate populations suggests that when observers have accumulated knowledge about a conceptual domain, typicality is driven by links with that knowledge rather than sharing features with fellow category members. Why might this be? One possible mechanism could be provided by links between exemplars and associated contextual information in long-term memory. In knowledgeable populations, exemplars become strong cues for linked information about usefulness and significance. This retrieved information could be used in judgments of representativeness and could even affect ease of conceptual access for the exemplars with which it is strongly associated.<sup>3</sup> We review evidence for links between concepts and contexts below.

### Interim summary of malleability

We have argued in this section that when it comes to determining typicality, central tendency is only one half of the equation. The other half of the equation comes from the observer and how she interacts with her environment. The fact that typicality was not correlated with central tendency demonstrates that typicality is not solely determined by the environment. Typicality is also determined by the *amount* of engagement, the *type* of engagement, and the *purpose* of engagement with the environment. This implies that typicality—and conceptual structure more generally—is a compromise between the affordances of the environment and the mental and practical needs of the observer (see Medin et al., 1997, for similar thoughts). This compromise can lead to differences. Typicality can differ between novices and experts (e.g., Lynch et al., 2000), between people of different cultures (e.g., Burnett et al., 2005), and can even differ within a person depending on the task at hand (Medin et al., 1997).

<sup>3</sup> We noted above that ideals sometimes have extreme rather than average values along certain dimensions. Exemplars with extreme values can sometimes be advantaged (e.g., categorized more accurately and receive higher typicality ratings) for reasons related to feature sharing if they minimize similarity to salient contrast categories (Kim & Murphy, 2011; Levering & Kurtz, 2006; Palmeri & Nosofsky, 2001). Ideals of this type would therefore be predicted by structural typicality alone. There was a hint of this in the study of Lynch et al. (2000), who found that height predicted typicality in tree experts, but extreme values do not appear to be the most important factor driving typicality in these studies. Furthermore, Goldstone and colleagues found that ease of conceptual access to extreme exemplars could be manipulated while holding contrast categories constant by changing the degree to which categories were “interrelated” versus “isolated,” a phenomenon that could not be accounted for by feature sharing alone (Goldstone, 1996; Goldstone, Steyvers, & Rogosky, 2003).

## Flexible functions: Typicality can be rapidly altered within individuals

The previous section showed that long-term experience can gradually change typicality over time. Culture and expertise are acquired over long periods of time, suggesting that the effect of these experiential factors on typicality could be driven by the effect of learning on relatively stable long-term memory structures. The current section will review research which suggests that typicality is much more pliable and depends on goals, context, and perspectives that change rapidly from moment to moment.

### Ad hoc categories and typicality

There is an entire class of categories that cannot be explained by structural typicality yet still exhibits graded conceptual structure. Ad hoc categories are equivalence classes that are created on the spot in working memory (Barsalou, 1983, 1985, 1987). Examples of ad hoc categories include things like “ways to escape being killed by the mafia,” “places to look for antique desks,” and “things conquerors take as plunder” (Barsalou, 1983). These categories have two factors that preclude them from being explained by structural typicality. One, they are not directly stored in long-term memory. By definition, an ad hoc category is one that is created on the spot. Instead of being retrieved from memory, it is created to fit a very specialized and specific set of circumstances (e.g., it is not often we have to form a category such as “things to grab from a burning building”). Two, ad hoc categories often consist of exemplars that defy central tendency and family resemblance. For example, the list of “things to take camping” includes things like tent, grill, axe, and bug spray. These exemplars do not have many, if any, common perceptual features. It is not surprising, therefore, that ad hoc or goal-derived categories are sometimes influenced by ideals rather than central tendency (Barsalou, 1985), another indication that typicality for these categories may be calculated in working memory rather than long-term memory.

### Context

Perhaps ad hoc categories are the exception to the rule, and central tendency really is the most important determinant of typicality for common categories. Research comparing typicality ratings in the absence versus presence of an overt context and across different contexts suggests that even common categories are susceptible to changes or reversals in typicality. For instance, Freeman (2014) showed that some perceptual features can be considered atypical in one context and typical in another. Specifically, Freeman (2014) had participants categorize faces as male or female, with half of the participants seeing a majority of sex-typical targets (males

with short hair, females with long hair; the normative context) and other participants seeing a majority of sex-atypical targets (males with long hair, females with short hair; the counternormative context). He found that competition in participants’ responses, as measured via mouse tracking, was contingent on the context (i.e., whether the stimulus list contained predominantly sex-typical or sex-atypical stimuli). Sex-typical faces induced competition in the counternormative condition but not the normative condition, while sex-atypical faces induced competition in the normative condition but not the counternormative condition.

Similarly, Roth and Shoben (1983) asked participants to make typicality ratings to stimuli in isolation and compared that with typicality ratings done in scenarios (“the truck driver had the beverage and a donut at the truck stop”). They found that the typicality ranking of beverages changed depending on whether the specified context was a midmorning break (coffee, tea, milk) or a donut store at a trucker’s stop (coffee, milk, tea). Importantly, typicality rankings of the beverages in isolation were not predictive of typicality in context, and, inconsistent with cue validity accounts, typicality in context was not predicted by similarity to the best exemplar (e.g., coffee was most typical in the donut shop, but milk was more typical than tea in that context even though tea was more similar to coffee). Roth and Shoben (1983) argue that typicality judgments made in isolation have no bearing on typicality judgments made in context. They go on to make the more radical claim that context restructures typicality from the ground up and that there is no such thing as an invariant semantic space.

Yeh and Barsalou (2006) report more evidence that typicality in isolation is different from typicality in context. When they asked participants to make typicality judgments in isolation, the correlation between judgments was .45 between people and .81 within a person after a two week delay. However, when a context was specified, the correlation rose to .70 between people and to .88 within a person. Thus, when typicality is judged in isolation, the agreement between people is moderately low, suggesting that people are recruiting or utilizing different types of information in their judgments. When typicality is judged in context, the agreement between people substantially increases, suggesting that people are recruiting or utilizing similar information.

Absent a context, people’s typicality judgments are actually quite variable. This is supported by evidence that shows that the features people produce for a category are also quite variable (Yeh & Barsalou, 2006). When participants were asked to produce features for categories, on average only 40% of one person’s features appeared in another person’s list of features. This suggests that different people utilized different information for the same category. (This is reminiscent of the “things to grab from a burning building” example given above.) Similarly, when participants were brought back after 2 weeks to complete the same property generation task, only 67% of

properties overlapped. This suggests that the same person utilized different information for the same category on two different occasions. This within-person variability is particularly important because it helps rule out the possibility that differences in typicality are a product of comparing two people that have different but stable and consistent conceptual structure. The point of all this is to show that there is substantial variability, both between and especially *within* individuals, in the information we recruit to represent a category and in the typicality judgments we make of that category. Collectively, the results reported in Yeh and Barsalou (2006) suggest two things: One, typicality in isolation is different from typicality in context; two, stability might be an artifact that arises from averaging across groups of people.

The research reviewed so far has contributed to the view that conceptual structure should be construed not in terms of clusters of similar objects with shared features but in terms of fits between objects and situations. Situations act as cues for concepts that are typically found in them, and objects act as cues for their typical situations, with typicality judgments driven by the degree of match between object context (Barsalou, 2009).

This view is supported by the effect of context on perception, which some have argued plays an integral role in conceptual representations (Barsalou, 2003; Goldstone & Barsalou, 1998). In this research, visual context alters perception toward objects that are most typical for the depicted situation. As one example, objects are easier to recognize when seen against typical rather than atypical backgrounds (Davenport, 2007; Davenport & Potter, 2004; Palmer, 1975). Several examples also come from research on face categorization. Facial expression can influence perception of race (Hugenberg & Bodenhausen, 2004), and race can influence perception of facial expression (Hugenberg & Bodenhausen, 2003). Likewise, visually presented background scenes (Freeman, Ma, Barth, et al., 2013; Freeman, Ma, Han, & Ambady, 2013) as well as socially meaningful attire (Freeman, Penner, Saperstein, Schetz, & Ambady, 2011) have been shown to influence the perception of facial features. In both cases, the categorization of racially atypical faces—which is to say racially ambiguous faces—can be “pushed around” and influenced by context. Thus, a racially atypical face is more likely to be categorized as Asian if it is shown on an Eastern background (Freeman, Ma, Barth, et al., 2013; Freeman, Ma, Han, et al., 2013). This suggests that background context may “contaminate” the categorization of a face and influence the perception of features, even though the background information was uninformative in the context of the laboratory task, in which all backgrounds were equally associated with all races.

Interestingly, some kinds of context amplify rather than reduce the influence of structural typicality in semantic processing. In particular, verbal labels increase the effect of

typicality in picture-processing tasks. When labels were included in instructions to rate the typicality of a picture (e.g., a chair), the difference in typicality ratings between the most and least typical members increased (Lupyan, 2008). In another study, participants had to indicate whether a picture (e.g., of a dog) matched a preceding verbal label (“dog”) or various other kinds of nonlabels, including associated nonverbal sounds (a barking sound) and words for associated actions (“bark”). Typicality ratings were more strongly associated with faster reaction times in the label condition than in any other condition (Lupyan & Thompson-Schill, 2012). Finally, the label “triangle” increases typicality effects for several tasks related to drawing and recognizing triangles (Lupyan, 2017).<sup>4</sup>

## Perspectives

In addition to context, an observer’s state of mind—her mental perspective—also influences typicality. Perspectives frame what information is available, the meaningfulness of that information, and how that information will be used. Although observers are limited by their perspective, they are not limited *to* their perspective. One might take the perspective of an acquaintance with a particular political orientation when deciding what present to get his daughter for her birthday. Would a gender normative princess doll, a gender-neutral Play-Doh set, or a counternormative toy gun be considered a more ideal gift by the parents?

Barsalou and Sewell (1984) provided evidence that taking another perspective systematically alters typicality. Participants rated the typicality of exemplars while adopting different international, domestic, or personal perspectives. Not only did typicality systematically change across perspectives, participants taking the same perspective had high agreement with one another. For example, people had high agreement that, from the perspective of “hippies,” vegetarian dinners are typical and meat dinners are atypical. Conversely, from a more Southern perspective, meat dinners are typical and vegetarian dinners are atypical. This systematic shift in typicality suggests that perspectives reorganize typicality in a coherent fashion, carving it into shape via top-down constraints. Furthermore, the flexibility with which someone can shift between their own perspective and someone else’s perspective is consistent with our notion that functional typicality is a highly dynamic process that occurs online in an individual’s working memory.

<sup>4</sup> Equilateral and isosceles triangles are more typical than other kinds of triangles, like right triangles. The author discusses multiple reasons why these triangles might have been most typical. One of his favored explanations involves minimizing similarity with salient contrast categories, which, as we have mentioned above, is consistent with access to structural typicality. We also believe it is plausible that isosceles triangles are near to the central tendency for the triangle category, although the author argues against this.



Hampton et al. (2006) found some additional evidence that perspectives alter conceptual structure. Participants categorized exemplars while taking perspectives corresponding to different purposes: either a “technical” purpose encouraging a strict scientific classification or a “pragmatic” purpose requiring a less strict classification. Interestingly, and contrary to the study’s prediction, technical perspectives induced classification decisions that were strongly correlated with acontextual typicality ratings, suggesting that categorization was driven by central tendency. Categorization from pragmatic perspectives was less correlated with acontextual typicality. The effect was most pronounced when the pragmatic perspective evoked a concrete situation, like a department store: categorization of furniture and tools was influenced by where one would expect to find them in the store, reducing the influence of central tendency. For instance, electrical things like televisions were less likely to be categorized as furniture because they belonged in the electronic section. Whereas the study supports the influence of perspectives on typicality (in this case measured by the probability of accepting an exemplar as a category member), it should be noted that the correlation between central tendency (measured by acontextual typicality judgments) and classification was quite strong for all perspectives, underscoring the strong influence of structural typicality.

Kim and Murphy (2011) demonstrated that perspectives also alter ideals which, as we have seen, can sometimes determine typicality. For example, consider what the ideal of a cigarette is. From the perspective of a tobacco company, the ideal cigarette is highly addictive. Conversely, from the perspective of a consumer, the ideal cigarette is nonaddictive. They asked participants to rate the idealness of positive or negative descriptions of objects (e.g., an nonaddictive vs. addictive cigarette) from different perspectives (e.g., a smoker or an antismoking kit manufacturer). Whereas perspectives altered which type of cigarette was most ideal, they did not alter which was rated as most typical. Thus, the study confirmed the importance of perspectives in semantic cognition but failed to support the hypothesis that perspectives reduced the influence of central tendency on typicality ratings.

### Interim summary of instability

The current section reviewed evidence supporting our notion of functional typicality. For one, ad hoc categories are mental structures that, one, do not, by definition, preexist in long-term memory, and, two, often defy central tendency in terms of common features. In addition, we reviewed evidence showing that small changes in context can have big changes in typicality, sometimes reversing judgments completely. This instability has led some researchers to argue that typicality is a dynamic and highly flexible judgment that is *always* constructed on the fly, in the moment, and for the moment. From this perspective, typicality is a fleeting context-dependent

judgment that emerges in response to the current nexus of internal and external constraints. If this is the case, then perhaps it does not even make sense to talk about what typicality means absent a specified context because there is no “invariant semantic space” (see Barsalou, 1985; Roth & Shoben, 1983). The notion that there is no “invariant semantic space” is a radical challenge that at first glance clashes head on with the notion that central tendency determines typicality.

### Structural and functional typicality are not mutually exclusive

So far we have reviewed literature showing the stability of typicality, the malleability of typicality, and the instability of typicality. Stability reflects the clusters of feature correlations in the experienced world, malleability reflects differential engagement with the world, and instability reflects pressing in the moment contextual demands. Are the two ends of this spectrum—stability and instability—mutually exclusive? Throughout this paper, we have suggested that this disagreement can be reconciled by breaking down typicality into two different constructs. On one level, typicality may reflect the way information gets encoded into semantic memory. We call this *structural typicality* because it reflects how long-term experience with the world builds the relatively stable similarity-based structure of information in our long-term memory (e.g., Rogers & McClelland, 2004). On a different level, typicality may reflect how information gets selectively recruited and used in working memory (Barsalou, 1983, 1987, 2003). This would be consistent with all the research showing that typicality is relative to, and dependent on, context. We call this *functional typicality* because it reflects how fleeting context and task demands transiently recruit a subset of the information in semantic memory to rapidly construct temporary conceptual structures.

Viewed in this manner, the apparent disagreement over typicality’s stability or instability is circumvented altogether by realizing that the two threads of research reflect different aspects of semantic cognition that both contribute to ease of access to concepts. Structural typicality reflects organization of information in semantic memory based on similarity. Functional typicality reflects the ability to construct, deconstruct, and reconstruct our concepts in multiple ways depending upon the current needs of the observer and the current affordances of the environment. From this perspective, structural typicality is fairly stable, functional typicality is fairly unstable, and there is no real contradiction between the two. This can be understood by way of analogy. There is a fairly systematic order to the way I arrange dishes in my kitchen cabinets. Small and large plates go together, cups and mugs go together, but how I use my dishes depends on what meal I am having. Some meals require large plates, some require small plates plus a

bowl. The structure by which my dishes are arranged in the cabinet is fairly stable. Conversely, the function of how I selectively recruit the dishes is fairly flexible (unstable). Big plates and small plates might go together in the cabinet (structurally), whereas small plates and bowls might go together when eating (functionally). In this same way, two things may be similar in structural typicality while simultaneously being dissimilar in functional typicality (or vice versa). If I was told to set the table for dinner but I did not know what kind of meal I was having, I might grab the plates and utensils that are most easily accessible. But if I knew we were having something specific, I would then reach for the dishes that were most relevant, even if they were less accessible.<sup>5</sup>

### Contextual constraint

Before concluding we would like to speculate on a potential mediating factor that may help explain when structural and functional typicality are recruited. To be clear, we do not think the two types of typicality operate in an “either/or” fashion. We already have evidence that the two operate simultaneously and interactively. Barsalou (1985) found evidence that both central tendency (structural typicality) and ideals (functional typicality) simultaneously explained a unique amount of variance in typicality judgments. In his words, “multiple determinants can simultaneously determine graded structure in a particular category” (p. 644). If both factors can be simultaneously active, the question becomes: What conditions favor (i.e., preferentially recruit) structural typicality? What conditions favor functional typicality? And how do the two types of typicality interact with one another?

Barsalou (1985) argued that context determines which factors dominate. Some contexts favor central tendency while other contexts favor ideals. Specifically, he argued that when a context made the ideals of a category apparent, then central tendency had little to no effect on judgments. This suggests that it is perhaps the degree of constraint within a context that mediates the degree to which structural and functional typicality are recruited (e.g., Roth & Shoben, 1983). It is plausible that less constrained contexts (e.g., “George saw the bird”) would encourage the recruitment of information that is most accessible in long-term memory. This would likely be the central tendency of the category (e.g., a robin in this case). Conversely, it is plausible that clearly specified, highly constrained contexts

(e.g., “George saw the bird walk across the yard at the farm”) would encourage the recruitment of information that is most relevant to the context. Here, information and schemas about farms and the types of birds found there (e.g., chickens) would likely be recruited and processed in working memory, resulting in a context-driven typicality judgment.

### Evidence, counterevidence, and considerations

Our dissociation between structural typicality, functional typicality, and the speculation that context mediates between them has some support but also some challenges. First, the evidence for structural typicality observed using fMRI and in semantic dementia patients demonstrated that typicality norms collected in the absence of context predicted neural activation patterns (Iordan et al., 2016) or consistent patterns of concept loss (Lambon Ralph et al., 2016). These studies suggest a relatively stable graded structure can be observed when context is unconstrained. Conversely, it goes without saying that the studies observing functional typicality effects provided clear constrained contexts within which exemplars should be categorized or judged for typicality (e.g., Barsalou & Sewell, 1984; Freeman, Ma, Barth, et al., 2013; Freeman, Ma, Han, et al., 2013; Freeman et al., 2011; Roth & Shoben, 1983).

In many ways, our distinction between structural and functional typicality parallels the conceptual deficits found in patients with semantic dementia and semantic aphasia (Jefferies & Lambon Ralph, 2006; Rogers et al., 2015). This is perhaps not surprising since our distinction between structural and functional typicality is admittedly highly influenced by the hub and the control network in the Controlled Semantic Cognition framework (Lambon-Ralph et al., 2016; see also Folstein & Dieciuc, 2018). As discussed earlier, patients with semantic dementia have compromised anterior temporal lobes (Mayberry et al., 2011), a region considered by some to be a hub of conceptual information (Lambon Ralph, 2014). These patients tend to have conceptual impairments that are consistent across modalities (e.g., words, sounds, drawings) and across tasks (e.g., naming, matching, categorizing). Furthermore, this pattern of deficits is sensitive to typicality (Patterson, 2007; Rogers et al., 2015). Atypical exemplars—because they are not as robustly encoded in long-term memory—deteriorate before typical exemplars (Bozeat et al., 2003). Both of these points—poor performance across tasks, modalities, and contexts and the deterioration of atypical exemplars before typical exemplars—are consistent with our notion of structural typicality. However, this literature also provides support for our notion of functional typicality. Patients with semantic aphasia tend to have intact anterior temporal lobes but have compromised temporoparietal and frontal regions (Jefferies & Lambon Ralph, 2006), regions involved in the task-dependent recruitment of conceptual information (Lambon Ralph et al., 2016). Unlike patients with

<sup>5</sup> The Controlled Semantic Cognition framework (Lambon Ralph et al., 2016, see Box 4) takes a very similar approach. However, the authors imply that *typicality* is a term reserved for how information is organized in a “hub” located in the anterior temporal lobe—what would be the equivalent of our structural typicality. In this paper, we have tried to argue that both forms of typicality—the one reflecting the structure of long-term memory and the one that is fleetingly computed in working memory—are equally important and equally deserving of the term *typicality*. In a companion paper (Folstein & Dieciuc, 2018) we elaborate on the relationship between the Controlled Semantic Cognition framework and the structural/functional distinction.

semantic dementia, these patients do better on some tasks and worse on others (Jefferies & Lambon Ralph, 2006). This suggests that it is not the similarity-based semantic structures supporting structural typicality that are impaired but the ability to flexibly recruit that information in a task-dependent manner that is impaired. It is the latter ability that supports functional typicality effects.

The research on cultural differences in typicality provides both some support but also some challenges for our framework. Recall that Medin et al. (1997) showed that the very same individual could flexibly alter their category structure (and the ideals that were organizing it) depending on the task they were performing. This is consistent with our notion of functional typicality reflecting on the spot computations relative to the current contextual demands of the observer and their goals. However, this line of research also presents a potential challenge. In these studies, categorization and typicality rating tasks do not include highly constrained contexts. If our earlier notion that less constrained contexts are more likely to reflect structural typicality is true, then why is typicality observed to be related to ideals in this case rather than central tendency? One plausible explanation for this is that participants with large amounts of experience with a category (e.g., fish) come to retrieve highly consistent contexts (e.g. sport fishing) in association with them, even when no context is given. These contexts would be processed in working memory and associated with ideals. Recent experience with relevant contexts could also exert an influence on typicality rating tasks that would normally be weakly constrained.

Nonexpert participants might also *sometimes* retrieve contexts with concepts when no context is provided (e.g., Aminoff, Kveraga, & Bar, 2013), resulting in functional typicality judgments, but use structural typicality for other concepts. This would be consistent with the correlations reported by Yeh and Barsalou (2006). They found that correlations of typicality judgments were actually quite low between people (.45) and within a person (.81) when no context was specified, possibly suggesting inconsistent use of structural and functional typicality. One way to explicitly test for this might be to have participants make categorization judgments while performing a second task (e.g., subtraction by threes). Presumably, the dual task would interfere with their ability to consciously impose a mental context and therefore should lead to judgments more reflective of structural typicality. We would expect this type of dual-task set-up to increase the correlations between people's judgments. Another way to test for this would be to explicitly ask participants how they are performing the task by using a verbal report protocol where participants "speak aloud" and explain their thought process (Ericsson & Crutcher, 1991).

Overall, there is ample evidence that context changes typicality (e.g., Barsalou & Sewell, 1984; Freeman, 2014; Roth & Shoben, 1983; Yeh & Barsalou, 2006; for general reviews on

context effects, see Casasanto & Lupyan, 2015, and Yee & Thompson-Schill, 2016), consistent with our notion of functional typicality. One important study, however, does appear to contradict our speculation that contextual constraint mediates between structural and functional typicality. We would predict that highly constrained contexts would recruit processes that emphasize selection of particular highly congruent, "functionally typical" exemplars, thereby decreasing the influence of long-term category structure (structural typicality) on ease of semantic access. But data from Federmeier and Kutas (1999) found just the opposite. Participants in that study read context sentences prior to reading the final word, which was manipulated—for example, "They wanted to make the hotel look more like a tropical resort. So along the driveway, they planted rows of . . ." The expected final word, *palms*, elicited a lower amplitude N400 ERP component than unexpected endings, reflecting more difficult semantic access context-incongruent endings and broadly consistent with contextually constrained typicality. However, when the sentence final word did not share a category with the expected ending (e.g., *tulips*), the N400 was greatly increased, but when the final word shared a category with the expected word (e.g., *pinos*), the N400 had an intermediate value. Critically, the authors were able to compare this effect in highly constraining versus less constraining sentences. Highly constraining sentences have been shown to decrease expectations for anything but the sentence final word (Schwanenflugel & Shoben, 1985), predicting an increased N400 for the unexpected category member. Instead, the opposite was observed—N400 amplitude to the unexpected category member (*pinos*) was lower in the highly constraining than less constraining sentences, suggesting that highly constraining contexts made fellow category members easier to access semantically even when their fit to the context was poor.<sup>6</sup> Rather than overriding the influence of stable category structure, contextual constraint increased it.

Whereas this result clearly constitutes a challenge to any simplistic view of contextual constraint's effect on semantic access, it should be emphasized that it, along with the rest of the N400 literature, generally supports the claim that ease of semantic access is modified by context. The amplitude of the N400 is sensitive to both typicality in isolation (Fujihara, Nageishi, Koyama, & Nakajima, 1998; Heinze et al., 1998) and semantic constraint provided by a preceding sentence (Kutas & Federmeier, 2011), supporting the distinction between functional and structural typicality. Regarding how the findings of Federmeier and Kutas (1999) might fit within our framework, we noted above that single words appear to provide a particularly strong, possibly privileged, link to the long-term similarity based semantic structures that underlie structural typicality effects (Lupyan, 2012).

<sup>6</sup> Note that a robust difference between expected words and unexpected noncategory members was robust for both sentence types, ruling out floor effects due to overly weak contextual constraint.

Even though it is clear that sentence context constrains meaning, it might not be surprising that the very strong links between single words and semantic structure interact with sentence context in counterintuitive ways during sentence processing.

## Underlying mechanisms of contextual constraint

In the previous section, we proposed that strongly versus weakly constraining contexts mediated access to functional versus structural typicality, respectively. In this final section, we briefly address the question of what makes a context strongly or weakly constraining.

The first variable is the *specificity* with which a context is described or experienced. Specific contexts have many details that are described explicitly or indirectly cued. The single word *camping*, for example, cues a very specific schema that includes a campground, a fire, and a forest. “Winter camping” and “beach camping” cue different sets of details that differently constrain the category of “things to take camping”. Several of the studies reviewed above (e.g. Barsalou, 1983, 1985; Federmeier & Kutas, 1999; Roth & Shoben, 1983) also use sentences or phrases to explicitly provide contextual details. Nonverbal cues can also be efficient ways to specify details. These include pictures of scenes (e.g., Davenport & Potter, 2004; Freeman, Ma, Barth, et al., 2013; Freeman, Ma, Han, et al., 2013), but can also include more simple cues like sounds. Edmiston and Lupyan (2015) showed that sound cues facilitated access to specific exemplars, whereas labels facilitated access a broader set of category members. For instance, a low bark facilitated a matching judgment for larger dogs more than smaller dogs because larger dogs have lower barks.

The second variable is the strength of *associations* between concepts and contexts that are formed by mere concurrence. Barsalou (1985, 1987) proposed that some goal-derived categories, possibly including “things to take camping” are well learned and encoded in long-term memory, while other, truly ad hoc categories, such as “ways to escape being assassinated by the mafia” must be constructed on the fly for the first time. Bar and colleagues (e.g., Aminoff et al., 2013) have also shown that some objects are stronger cues for contexts than others. We speculated above that many years of experience with a conceptual domain might form very strong associations between concepts and contexts so that contexts are retrieved automatically and affect typicality judgments. This view is consistent with parallel distributed processing models, which can account for many different semantic phenomena using principles of association (e.g., Rogers & McClelland, 2004).

A third variable, which we refer to as *pragmatic constraints*, is difficult to disentangle from mere association. Pragmatic constraints are causal relationships within a situation that narrow the range of appropriate, functionally typical, concepts. These constraints are most clear in highly ad hoc categories, many of which

are used in everyday life, like “things used to prop open a door” or “things used to swat a fly other than a fly swatter” (Corbett, Jefferies, & Lambon Ralph, 2011). The alternative fly-swatting device must be easy to lift, wide enough to swat a fly, and available in one’s particular setting. A rolled-up newspaper fits these constraints well and is therefore likely to be functionally typical. Barsalou and colleagues have proposed that appropriate matches between objects and situations are calculated using “situated simulations” in which interactions between objects in a relevant situation are perceptually simulated in modality specific cortical areas (Barsalou, 1999, 2003, 2009). How and whether pragmatic constraints interact with association strength to affect encoding of fit between objects and situations (functional typicality), is not currently known.

Finally, Lupyan has argued that labels provide privileged access to semantic information (Lupyan, 2012), and we argued above, based on evidence that labels increase typicality effects, that labels bias semantic cognition towards stable, similarity-based, conceptual structures in long-term memory. By this view, labels can be considered special features of context that recruit structural rather than functional typicality. However, the status of labels as part of language creates an interesting tension between our claim that structural typicality is driven by features in the world and the idea that our view of the world is shaped by language (e.g., Whorf, 1956). Language is a feature of the observer’s culture and may thus be linked to the culture’s goals, which we have argued drive functional typicality.

Whereas a full treatment of this issue may be beyond the scope of this paper, we can offer brief speculations. First, labels will often correspond to clusters of similar objects, as was proposed by Rosch and colleagues (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976a). When this holds, labels simply reflect the structure of the environment and might be thought of as amplifying structural typicality as we have suggested. In cases where, for whatever reason, labels are applied in a way that does not respect similarity based on nonlinguistic features, we might speculate that, over time, this will alter long-term memory structures, making differently labeled exemplars less similar than before. This effect would be consistent with the label feedback hypothesis and with certain computational models of category learning (Rogers & McClelland, 2004), as well as with the finding that labels facilitate learning novel categories (Lupyan, Rakison, & McClelland, 2007). By this view, labels may play a special role in altering structural typicality, but would not influence functional typicality because they do not alter the fit between exemplars and situations in working memory.

## Conclusion

Throughout this paper we have argued that typicality’s stability and instability are not at odds with one another. Rather,

they are reflections of two different levels of semantic knowledge: one of which is structural and the other of which is functional. Others have hinted at this distinction between stable and unstable typicality in different terms (Barsalou, 1985, 1987, 2009; Lambon Ralph et al., 2016; Rogers & McClelland, 2004), but we argue that the novel terms *structural typicality* and *functional typicality* are useful terminology that serve as a simple yet clear way of dissociating and framing different threads of research that might on the surface appear to be irreconcilable. Structural typicality is a reflection of long-term interaction with one's experienced environment. This molds our base knowledge into a fairly stable shape based on the feature correlations we interact with. Functional typicality is a reflection of the short term, in the moment and for the moment, demands which dictate how some of our long-term knowledge gets transiently recruited. Viewed from this perspective, the notion of invariant semantic spaces (stability) and context-dependent judgments (instability) are not contradictory so much as they are complementary. Seen from our framework, the disagreement over typicality's stability becomes less about "Is typicality stable or unstable?" and more a question of emphasis and timing—for example, "When and under what circumstances is structural typicality more dominant than functional typicality (and vice versa)? And how do the two interact?" We speculate that contextual constraint may mediate between structural and functional typicality, but we encourage future research to investigate whether this speculation is right or misguided.

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## References

- Aminoff, E. M., Kveraga, K., & Bar, M. (2013). The role of the parahippocampal cortex in cognition. *Trends in Cognitive Sciences*, 17(8), 379–390.
- Bailenson, J. N., Shum, M. S., Atran, S., Medin, D. L., & Coley, J. D. (2002). A bird's eye view: Biological categorization and reasoning within and across cultures. *Cognition*, 84(1), 1–53.
- Barsalou, L. (2003). Situated simulation in the human conceptual system. *Language and Cognitive Processes*, 18(5/6), 513–562.
- Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, 11(3), 211–227.
- Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(4), 629.
- Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts reconsidered: The ecological and intellectual bases of categories*. New York: Cambridge University Press.
- Barsalou, L. W. (1999). Perceptions of perceptual symbols. *Behavioral and Brain Sciences*, 22(4), 637–660.
- Barsalou, L. W. (2009). Situating concepts. In P. Robbins & M. Aydede (Eds.) *The Cambridge handbook of situated cognition* (pp. 236–263). Cambridge: Cambridge University Press.
- Barsalou, L. W., & Sewell, D. R. (1984). *Constructing representations of categories from different points of view* (Tech. Report). Atlanta: Emory University.
- Bozeat, S., Lambon Ralph, M. A., Graham, K. S., Patterson, K., Wilkin, H., Rowland, J., ... Hodges, J. R. (2003). A duck with four legs: Investigating the structure of conceptual knowledge using picture drawing in semantic dementia. *Cognitive Neuropsychology*, 20, 27–47.
- Burnett, R. C., Medin, D. L., Ross, N. O., & Blok, S. V. (2005). Ideal is typical. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 59(1), 3.
- Casasanto, D., & Lupyan, G. (2015). All concepts are ad hoc concepts. In E. Margolis & S. Laurence (Eds.), *The conceptual mind: New directions in the study of concepts* (pp. 543–566). Cambridge: MIT Press.
- Corbett, F., Jefferies, E., & Lambon Ralph, M. A. (2011). Deregulated semantic cognition follows prefrontal and temporo-parietal damage: Evidence from the impact of task constraint on nonverbal object use. *Journal of Cognitive Neuroscience*, 23, 1125–1135.
- Davenport, J. L. (2007). Consistency effects between objects in scenes. *Memory & Cognition*, 35, 393–401.
- Davenport, J. L., & Potter, M. C. (2004). Scene consistency in object and background perception. *Psychological Science*, 15, 559–564.
- Edmiston, P., & Lupyan, G. (2015). What makes words special? Words as unmotivated cues. *Cognition*, 143, 93–100.
- Ericsson, K. A., & Crutcher, R. J. (1991). Introspection and verbal reports on cognitive processes—Two approaches to the study of thinking: A response to Howe. *New Ideas in Psychology*, 9(1), 57–71.
- Federmeier, K. D., & Kutas, M. (1999). A rose by any other name: Long-term memory structure and sentence processing. *Journal of Memory and Language*, 41, 469–495.
- Folstein, J. R., & Dieciuc, M. A. (2018). *Typicality is influenced by both stable and flexible representations: a review and synthesis of theoretical accounts*. Manuscript in preparation.
- Folstein, J. R., Palmeri, T. J., & Gauthier, I. (2013). Category learning increases discriminability of relevant object dimensions in visual cortex. *Cerebral Cortex*, 23, 814–823.
- Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2003). A comparison of primate prefrontal and inferior temporal cortices during visual categorization. *Journal of Neuroscience*, 23, 5235–5246.
- Freeman, J. B. (2014). Abrupt category shifts during real-time person perception. *Psychonomic Bulletin & Review*, 21(1), 85–92.
- Freeman, J. B., Ma, Y., Barth, M., Young, S. G., Han, S., & Ambady, N. (2013). The neural basis of contextual influences on face categorization. *Cerebral Cortex*, bht238.
- Freeman, J. B., Ma, Y., Han, S., & Ambady, N. (2013). Influences of culture and visual context on real-time social categorization. *Journal of Experimental Social Psychology*, 49(2), 206–210.
- Freeman, J. B., Penner, A. M., Saperstein, A., Scheutz, M., & Ambady, N. (2011). Looking the part: Social status cues shape race perception. *PLOS ONE*, 6(9), e25107.
- Fujihara, N., Nageishi, Y., Koyama, S., & Nakajima, Y. (1998). Electrophysiological evidence for the typicality effect of human cognitive categorization. *International Journal of Psychophysiology* 29, 65–75.
- Goldstone, R. L. (1996). Isolated and interrelated concepts. *Memory & cognition*, 24(5), 608–628.
- Goldstone, R. L., & Barsalou, L. W. (1998). Reuniting perception and conception. *Cognition*, 65(2/3), 231–262.
- Goldstone, R. L., Steyvers, M., & Rogosky, B. J. (2003). Conceptual interrelatedness and caricatures. *Memory and Cognition*, 31, 169–180.
- Goldwater, M. B., Markman, A. B., & Stilwell, C. H. (2011). The empirical case for role-governed categories. *Cognition*, 118(3), 359–376.

- Hampton, J. A. (1979). Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 18, 441–461.
- Hampton, J. A. (1981). An investigation of the nature of abstract concepts. *Memory & Cognition*, 9, 149–156.
- Hampton, J. A. (1995). Testing the prototype theory of concepts. *Journal of Memory and Language*, 34(5), 686.
- Hampton, J. A., Dubois, D., & Yeh, W. (2006). Effects of classification context on categorization in natural categories. *Memory & Cognition*, 34(7), 1431–1443.
- Heinze, H.-J., Muentz, T.-F., & Kutas, M. (1998). Context effects in a category verification task as assessed by event-related brain potential (ERP) measures. *Biological Psychology*, 47, 121–135.
- Hugenberg, K., & Bodenhausen, G. V. (2003). Facing prejudice implicit prejudice and the perception of facial threat. *Psychological Science*, 14(6), 640–643.
- Hugenberg, K., & Bodenhausen, G. V. (2004). Ambiguity in social categorization: The role of prejudice and facial affect in race categorization. *Psychological Science*, 15(5), 342–345.
- Jordan, M. C., Greene, M. R., Beck, D. M., & Fei-Fei, L. (2016). Typicality sharpens category representations in object-selective cortex. *NeuroImage*, 134, 170–179.
- Jefferies, E., & Lambon Ralph, M. A. (2006). Semantic impairment in stroke aphasia versus semantic dementia: A case-series comparison. *Brain*, 129, 2132–2147.
- Kim, S., & Murphy, G. L. (2011). Ideals and category typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(5), 1092.
- Kutas, M., Federmeier, K.D., (2011). Thirty years and counting: Finding meaning in the N400 component of the event-related brain potential (ERP). *Annual Review of Psychology* 62, 621–647.
- Lambon Ralph, M. A. (2014). Neurocognitive insights on conceptual knowledge and its breakdown. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1634), 20120392.
- Lambon Ralph, M. A., Jefferies, E., Patterson, K., & Rogers, T. T. (2016). The neural and computational bases of semantic cognition. *Nature Reviews Neuroscience*, 18, 43–55.
- Lambon Ralph, M. A., Sage, K., Jones, R. W., & Mayberry, E. J. (2010). Coherent concepts are computed in the anterior temporal lobes. *Proceedings of the National Academy of Sciences*, 107(6), 2717–2722.
- 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.
- Levering, K., & Kurtz, K. J. (2006). The influence of learning to distinguish categories on graded structure. Paper presented at the Proceedings of the 28th annual conference of the Cognitive Science Society. Vancouver
- Li, S., Ostwald, D., Giese, M., & Kourtzi, Z. (2007). Flexible coding for categorical decisions in the human brain. *Journal of Neuroscience*, 27, 12321–12330.
- Lupyan, G. (2008). From chair to “chair”: A representational shift account of object labeling effects on memory. *Journal of Experimental Psychology: General*, 137 (2), 348.
- Lupyan, G. (2012). Linguistically modulated perception and cognition: The label-feedback hypothesis. *Frontiers in psychology*, 3, 54.
- Lupyan, G. (2017). The paradox of the universal triangle: Concepts, language, and prototypes. *The Quarterly Journal of Experimental Psychology*, 70(3), 389–412.
- Lupyan, G., Rakison, D. H., & McClelland, J. L. (2007). Language is not just for talking: Redundant labels facilitate learning of novel categories. *Psychological Science*, 18, 1077–1083.
- Lupyan, G., & Thompson-Schill, S. L. (2012). The evocative power of words: Activation of concepts by verbal and nonverbal means. *Journal of Experimental Psychology: General*, 141(1), 170.
- Lynch, E. B., Coley, J. B., & Medin, D. L. (2000). Tall is typical: Central tendency, ideal dimensions, and graded category structure among tree experts and novices. *Memory & Cognition*, 28, 41–50.
- Mayberry, E. J., Sage, K., & Lambon Ralph, M. A. (2011). At the edge of semantic space: The breakdown of coherent concepts in semantic dementia is constrained by typicality and severity but not modality. *Journal of Cognitive Neuroscience*, 23(9), 2240–2251.
- McCloskey, M. E., & Glucksberg, S. (1978). Natural categories: Well defined or fuzzy sets? *Memory & Cognition*, 6, 462–472.
- Medin, D. L., & Atran, S. (2004). The native mind: biological categorization and reasoning in development and across cultures. *Psychological review*, 111(4), 960.
- Medin, D. L., Lynch, E. B., Coley, J. D., & Atran, S. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology*, 32, 49–96.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological review*, 85(3), 207.
- Mervis, C. B., Catlin, J., & Rosch, E. (1976). Relationships among goodness-of-example, category norms, and word frequency. *Bulletin of the Psychonomic Society*, 7, 283–284.
- Mervis, C. B., & Pani, J. R. (1980). The acquisition of basic object categories. *Cognitive Psychology*, 12, 496–522.
- Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology*, 32, 89–115.
- Nosofsky, R. M., & Johansen, M. K. (2000). Exemplar-based accounts of “multiple-system” phenomena in perceptual categorization. *Psychonomic Bulletin and Review*, 7(3), 375–402.
- Nosofsky, R. M., & Zaki, S. R. (2002). Exemplar and prototype models revisited: Response strategies, selective attention, and stimulus generalization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(5), 924.
- Ojalehto, B. L., & Medin, D. L. (2015). Perspectives on culture and concepts. *Annual Review of Psychology*, 66, 249–275.
- Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, 97, 185–200.
- Palmer, S. E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition*, 3, 519–526.
- Palmeri, T. J., & Nosofsky, R. M. (2001). Central tendencies, extreme points, and prototype enhancement effects in ill-defined perceptual categorization. *The Quarterly Journal of Experimental Psychology Section A*, 54, 197–235.
- Patterson, K. (2007). The reign of typicality in semantic memory. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362(1481), 813–821.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77(3p1), 353.
- Rein, J. R., Goldwater, M. B., & Markman, A. B. (2010). What is typical about the typicality effect in category-based induction?. *Memory & Cognition*, 38(3), 377–388.
- Rips, L. J. (1975). Inductive judgements about natural categories. *Journal of Verbal Learning & Verbal Behavior*, 14, 665–681.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge: MIT Press.
- Rogers, T. T., Patterson, K., Jefferies, E., & Ralph, M. A. L. (2015). Disorders of representation and control in semantic cognition: Effects of familiarity, typicality, and specificity. *Neuropsychologia*, 76, 220–239.
- Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976a). Basic objects in natural categories. *Cognitive Psychology*, 8, 382–439.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573–605.
- Rosch, E., Simpson, C., & Miller, R. (1976b). Structural bases of typicality effects. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 491–502.

- Roth, E. M., & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, *15*(3), 346–378.
- Rumelhart, D. E., & Todd, P. M. (1993). Learning and connectionist representations. *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 3–30). Cambridge: MIT Press.
- Schwanenflugel, P. J., & Shoben, E. J. (1985). The influence of sentence constraint on the scope of facilitation for upcoming words. *Journal of Memory and Language*, *24*(2), 232–252.
- Smith, E. E. (1978). Theories of semantic memory. In W.K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 6, pp. 1–56). Potomac, MD: Erlbaum.
- Whorf, B. L. (1956). *Language, thought, and reality: Selected writings of Benjamin Lee Whorf*. Cambridge: MIT Press.
- Woollams, A. M. (2012). Apples are not the only fruit: The effects of concept typicality on semantic representation in the anterior temporal lobe. *Frontiers in Human Neuroscience*, *6*, 85.
- Yee, E., & Thompson-Schill, S. L. (2016). Putting concepts into context. *Psychonomic Bulletin & Review*, *23*, 1015–1027.
- Yeh, W., & Barsalou, L. W. (2006). The situated nature of concepts. *The American Journal of Psychology*, *119*(3), 349–384.