The common effect of value on prioritized memory and category representation

Joshua Knobe[1] and Fiery Cushman[2]

[1]*Yale University,* [2]*Harvard University*

The way we represent categories depends on both the frequency and value of the category's members. Thus, for instance, prototype representations can be impacted both by information about what is statistically frequent and by judgments about what is valuable.. Notably, recent research on memory suggests that prioritized memory is also influenced by both statistical frequency and value judgments. Although work on conceptual representation and work on prioritized memory have thus far proceeded almost entirely independently, the patterns of existing findings provide evidence for a link between these two phenomena. In particular, these patterns provide evidence for the hypothesis that the impact of value on conceptual representation arises from its co-dependent relationship with prioritized memory.

1. The roles of frequency and value in category representation

It has long been established that statistical frequency plays an important role in shaping people's category representations [1]. Thus, people's representation of the category of birds will be influenced more by house sparrows than by condors, and their representation of the category of cognitive scientists will be skewed more towards bookish types than towards fashionable dandies.

Strikingly, research suggests that wherever statistical frequency plays a role in category representation, value does the same [2–5]. People's representation of the category of sunsets is not simply shaped by the statistical distribution of sunsets they observe; it is skewed toward sunsets that are especially striking and beautiful. Similarly, people's representation of the category of cognitive scientists will not simply reflect the statistical distribution of cognitive scientists; it will show an outsized influence of those individuals who are seen as especially outstanding scientists.

As we will see, existing studies suggest that the role of value judgments in category representation closely mirrors the role of statistical frequency. Across a range of different aspects of category representation, the impact of being seen as high-value appears to be similar to the impact of statistical frequency. We therefore face a larger theoretical question: Why is it that value judgments exert an impact on category representation that is so similar to the impact of statistical frequency?

Here we highlight a potential clue. It takes the form of a striking but unremarked correspondence. Quite separately from research on how we think about categories, another body of research asks *how we learn about and remember individuals*—say, the name of the cognitive scientist who presented in last week's seminar. Here, too, it has long been recognized that we remember things better the more frequently we encounter them [6,7], but recent research indicates that value judgments play an important role in prioritized memory as well [8–11]. Specifically, people tend to selectively learn and remember information about the features of objects they regard as good, while devoting less cognitive capacity to less good objects. Thus, for instance, we are more likely to remember the names of cognitive scientists whose presentations we especially enjoyed.

Within existing work, research on these two different effects has proceeded almost entirely separately. The aim of the present review is to bring the two together. Such a connection might help to elucidate the relationship between category representation and prioritized memory, while also helping to explain the impact of value judgments on both.

Correspondence should be directed to Joshua Knobe, joshua.knobe@yale.edu

2. The role of value in category representations

People's representations of a category—what is normal, typical, or representative of it—are influenced not only by information about which features of that category are common, but also by information about which features of that category are good.

Work on the representation of categories has been conducted across a broad array of different research areas, including everything from psychological research on concepts to linguistics research on gradable adjectives. Yet, these disparate research traditions have converged on a remarkably consistent pattern of findings regarding category representations in their various guises: Effects of value judgment and statistical frequency that coincide.

One of the most well-known of these effects involves representations of *prototypicality*. To illustrate, consider the category of grandmothers. It is possible to imagine two people such that both clearly count as grandmothers but people would tend to think that one of them is more of a "prototypical grandmother" than the other [1,12] (Rosch & Mervis, 1975; Armstrong, et al., 1983). What exactly accounts for these intuitions about prototypicality? Numerous studies show that one relevant factor is statistical frequency [1]. It is frequent for grandmothers to have certain features (e.g., being old) and infrequent for grandmothers to have other features (e.g., participating in boxing matches). Studies show that objects that have more statistically frequent features will be seen as more prototypical category members than those that have less frequent features. Strikingly, however, a number of studies show that value judgments can also play a role [2,3,13]. It is considered good for grandmothers to have certain features (e.g., skilled at baking cookies) and bad to have others (e.g., too busy to take care of their grandchildren). Even controlling for effects of statistical frequency, objects that have features that are regarded as good tend to be seen as more prototypical category members.

A second case is judgments of *normality*. For many categories, people seem to have a sense that a certain type of thing is normal: a normal lunch, a normal marriage, a normal way to treat one's students. Existing work suggests that such perceptions of normality have important downstream effects on a wide range of other judgments and behaviors [14–18]. So, how do people determine what is normal? Unsurprisingly, the more statistically frequent a feature is for a given category, the more it will be seen as normal for that category. But importantly, studies also show that there is a role for value judgment. Controlling for statistical frequency, there is a tendency such that the more an object is seen as good, the more it will be seen as normal [4].

A third case is that of the *standards* people use in evaluating gradable adjectives. Clearly, the size a table has to reach before it can be considered a "large table" is different from the size that a building has to reach before it can be considered a "large building." Within the existing literature, the usual way of thinking about this phenomenon is that we cannot apply an adjective like "large" to an object unless it exceeds a certain threshold - a standard - and this standard depends on the category we are discussing [19,20]. So then, how do people determine which standard is the right one for a given category? Research consistently finds that statistical frequency plays an important role [21–23]. For example, if you came to believe that it was more frequent for tables to have large sizes, the standard you used for applying "large" to tables would be higher. But, importantly, research also shows that value judgments can play a role [24]. Thus, if you came to believe that it was good for tables to have large sizes, then, even controlling for your judgments of statistical frequency, this judgment would lead your standard to be higher.

A final case is *what comes to mind*. Suppose we simply ask you to imagine a size that a table could be. There are no right or wrong answers here: you can just imagine any size - whichever size first that comes to mind. In completing this task, you seem to be sampling from some distribution. But what exactly is that distribution? Here again, the answer seems to be that both statistical frequency and value judgments play a role [5]. So if you think that it is very frequent for category members to have certain features, you will be more likely to think of an object that has those features, but also, controlling for statistical frequency, if you think that it is good to have certain features, you will be more likely to think of an object that has those features.

Looking at these four different effects of value judgment, it would at least be possible to suggest that each effect has a separate explanation. Yet, the effects observed in these very different areas all seem to show a very similar pattern. Wherever we observe an effect on frequency, we also observe an effect of value. Moreover, the effect is always in the same direction. When people see certain objects as good, the impact of this value judgment is always in the same direction as if those subjects were high in frequency, and conversely when people see certain objects as bad the impact of that value

Correspondence should be directed to Joshua Knobe, joshua.knobe@yale.edu

judgment is always in the same direction as if those objects were low in frequency. There is, therefore, at least some prima facie reason to prefer the strategy of seeking a single broader explanation. That is the strategy we pursue.

3. The role of value in prioritized memory

The role that value plays in representing categories has certain intriguing parallels with its role in how we learn about and remember individuals. Suppose you observe a collection of individual token objects, and you receive information about the features of each of these objects. For example, it might be that you observe a number of different scientists and receive information about each of them (where they live, how many hours they work per day, etc.). In most ordinary situations of this type, the limits on your cognitive capacities mean that you cannot possibly remember the features of all of the objects. Thus, you face a problem. If you can't remember the features of all of the objects, which should you prioritize remembering?

A large and growing body of research suggests that people are especially likely to remember things that are of high value [8–11,25]. (In much of the relevant empirical work, value is manipulated by establishing task-related financial or other rewards.) Presumably this is because, all else being equal, the valuable things are the ones we profit most from being able to remember and reason about—i.e., because this makes rational use of a limited cognitive resource [26,27]. Similarly, the fact that we are more likely to remember things when we encounter them more often can be understood, in part, as mnemonic prioritization of things we are more likely to encounter often in the future [6,7,28,29].

Current research indicates value-based prioritization in at least three phases of the memory process. First, people prioritize encoding high value targets. That is, when people believe that an object is good, they devote more cognitive capacity to accurately encoding information about it [10,11,30,31]. Second, even if people learn that an object is good only *after* the time that they observe it, comparable effects still arise, with people devoting more cognitive capacity to memory consolidation for objects they regard as good [8]. Finally, there is evidence for an effect of value on memory retrieval [9,27]. This effect is particularly strong in domains where valuable targets are more important to retrieve—e.g., when one retrieves memories of objects in order to use them, and higher-value objects have high-value uses. However, even when people believe that they will not be able to use the information in this way, the effect remains at a diminished level. In other words, people exhibit prioritized retrieval of high-value targets even when not task-relevant, and show still greater prioritization of high-value targets when justified by the task [9].

Overall, then, the pattern of existing results does not suggest that the role of value in prioritized memory is due to some one specific cognitive process, such that if that one process could be removed the entire effect would disappear. Instead, it appears to be a pervasive tendency found across a wide variety of different cognitive processes. In all of these different processes, we appear to be finding an effect whereby more cognitive resources are devoted to information about individual objects that are seen as good than about individual objects that are not seen as good.

4. The relationship between category representation and prioritized memory

Thus far, we have been reviewing evidence for two kinds of effects. First, within research on category representation, representations that are influenced by statistical frequency are also influenced by value judgments. Moreover, this evidence indicates that these effects consistently appear in the same direction. Second, within research on the memory of individual token objects, there is enhanced memory for objects that are regarded as good than for objects that are regarded as bad.

What is the relationship between these two kinds of effects? Possibly, the relationship is coincidental: it might be that value judgments just happen to play a role in each of these separate effects but that the two effects do not relate in any way to each other. For example, it might be that the impact of value on category representation has nothing to do with prioritized memory and is instead the results of a completely separate process involving conceptual coherence [32–34]. At the moment, there is no concrete empirical evidence against this hypothesis, and it therefore remains very much an open possibility.

However, another natural hypothesis would be that the pattern we have been describing arises because of a deeper link between conceptual representation and prioritized memory. More specifically, it seems that conceptual representation and prioritized memory are intimately connected–indeed, each might ordinarily

Correspondence should be directed to Joshua Knobe, joshua.knobe@yale.edu

Box 1. Predicting the functional form of the impact of value on category representation

Suppose we assume that the impact of value on conceptual representation does indeed go in large part via the impact of value on remembering individual token objects. Drawing on this assumption, we can derive specific quantitative predictions about the functional form of this impact. Recent studies provide some intriguing initial support for these predictions.

We begin by modeling prioritized memory as a process that encounters a series of unique objects belonging to some category, each of which has a given feature *f* with probability *P(f)*. Then, by Bayes theorem, the expected proportion of remembered objects possessing feature *f* will be the *product* of the frequency of *f* among observed objects *P(f)* and the probability of an object being remembered given that is has that feature *P(r|f)*, normalized by the probability of remembering any encountered object *P(r)*:

$$p(f|r) \ = \ p(f) \bullet p(r|f)/p(r)$$

This is quite intuitive: An object with some feature is remembered just if it is both *encountered* and then, based on its value, *remembered*.

Suppose now that the probability of remembering feature *f* each time it is encountered tends to be higher when *f* is regarded as good than when *f* is regarded as bad. We can then use the equation above to generate a more specific prediction about the functional form of the impact of value on the probability of remembering a feature. Specifically, we should predict a value x frequency interaction, which follows from the multiplicative relationship between frequency *P(f)* and value *P(r|f)*. In other words, we should observe that the impact of value will be larger when the probability of encountering the feature is higher both in prioritized memory and, therefore, in conceptual representation.

Preliminary evidence, while limited, is encouraging. In one study participants were shown a series of novel objects that varied in length, with the probability distribution over length varied between participants. Some participants were told that long objects were the highest in value, while for others short objects were the highest in value. Finally, participants were asked to indicate the length of the first object that spontaneously came to mind.

Strikingly, the results did in fact show the predicted value x frequency interaction [5]. When a number of different possible models were compared, the one with the best fit was one in which the probability of a given feature *f* coming to mind is proportional to $p(f) \bullet \ e^{v(f)\tau}$, where $p(f)$ is the frequency of encountering feature f, $v(f)$ is the value of feature f and τ is an inverse temperature parameter. Interestingly, this is exactly the prediction we obtain if we model the impact of the value of a feature on the probability of remembering the feature by transforming the value into a probability in the most straightforward way using a softmax function, $p(r|f) \propto e^{v(f)\tau}$. Substituting this quantity into the equation above, we get that:

$$p(f|r) \propto \ p(f) \ \bullet \ e^{v(f)\tau}$$

Which is precisely the model that emerged independently from empirical inquiry on the relevant conceptual representation.

depend on the other, with bidirectional influences. Thus, the effect of value on conceptual representation and the effect of value on prioritized memory might not reflect two distinct psychological processes. Instead, these might just be two different symptoms of what is ultimately the very same psychological phenomenon.

A key task for a future work will be to develop more concrete theories that describe this psychological phenomenon in detail. To pursue this task, we can turn to existing models that relate memory for token objects to category representations. We can then spell out the core idea of the present hypothesis by adding memory prioritization to those existing models. Here, we illustrate two possibilities, but these are neither mutually exclusive nor exhaustive.

One possible view is that value affects category representations simply as a result of its influence on our memory of particular token objects. People's representation of a category appears to be shaped in some way by the statistical distribution of the objects they actually encounter, but one might think that the category representation will be shaped to a greater extent, or perhaps exclusively, by those objects that people actually remember. Suppose now that people

Correspondence should be directed to Joshua Knobe, joshua.knobe@yale.edu

Box 2. Memory models often depend on category-level representations

How might representations of a category influence the way people remember individual items? It is well known that people "fill in" missing details in memory by drawing on general background knowledge [35]. We focus on two ways to formalize this idea: One drawn from Bayesian models of reconstructive memory [36,37], and another drawn from a widely-used variety of deep neural net [38]. In both cases, a key insight for our purposes is that models of how we remember individual things (people, objects, events, etc.) imply an explicit representation of the categories to which those things belong.

If a person observes object or event x and encodes an imperfect or incomplete memory trace z, then the task of "remembering" can be construed as a form of rational inference. One wishes to derive $p(x|z)$: a probability distribution over the true object or event, given one's memory trace. By Bayes rule this can be computed from the product of two quantities, normalized: $p(x|z) \propto p(z|x)p(x)$. The first quantity is given by a generative model of how objects and events give rise to their associated memory traces—i.e., the likelihood $p(z|x)$. It is the second term, however, that is particularly relevant to theories of conceptual representation. This term is the prior $p(x)$— i.e., the prior probability distribution over objects or events themselves. It has the effect of filling in missing information in memory, or correcting flawed information, by drawing on what is generally expected for the category. For instance, if your memory trace for the name of a friend's grandmother is equally likely given the true names "Alice" and "Alistair", generic knowledge about the typical names of grandmothers will help you select "Alice" as the most probable This style of model implies that in order to remember a specific thing, we must also represent information about the general category to which it belongs, establishing an important point of connection between memory and category representation.

A similar architectural motif arises in variational autoencoders (VAEs), a popular class of deep neural nets. VAEs encode high dimensional inputs (such as images) into a low dimensional space. This can be useful for efficient storage, and VAEs also encourage the organization of this space around interpretable features. Specifically, a deep neural network ("encoder network") is trained to project high-dimensional inputs into a low-dimensional latent representation. A second network ("decoder network") then attempts to reconstruct the original image, and the networks are trained to minimize loss during reconstruction. This architecture shows an especially striking ability to "fill in" missing details of an input—for instance, knowing that an incomplete image depicting an apparent cat would most likely be completed with whiskers. Of special importance for our proposal, VAEs organize their latent representations as probability distributions. Thus, one can sample novel individuals from a distribution defined over a category—akin to saying, "construct some new cat images"—and the sample distribution will approximate the category representation. Here, again, we see an architecture in which a system designed to encode, infer, and reconstruct features of individuals ends up also encoding information about categories, including a representation of what is typical for that category.

Although these are just two examples of specific computational models, they help to illustrate several broader themes. In order to reconstruct information from incomplete or noisy memories, we can "fill in" or adjust details drawn from category-level representations. Thus, rational memory systems come to represent not only individuals, but also categories. This may explain why effects of value on prioritized memory would exert a downstream influence on the way we represent and reason about categories.

have a higher probability of remembering an object to the extent that they represent it as good. Then the distribution of the category members they actually remember will be shaped both by the statistics of which objects they tend to encounter and by their value judgments about how good each of these objects are. This model generates a specific prediction about the way frequency and value should interact in determining category representations, which we outline in Box 1.

A more complex view would say that the impact of value on category representation does not necessarily have to go via an impact on people's memory for individual objects. At the heart of this second approach is the idea that people's ability to remember individual objects within a category is closely tied to more abstract representations of the category itself (see Box 2). For example, people's ability to remember individual songs is not just a matter of separately memorizing the sequence of notes in each separate song; it is in large part a matter of building up more abstract representations of different types of music. This abstract representation then makes it possible to 'fill in' information about the sequence of notes in each individual song. Thus, in computational models of category representation, there may be no sharp distinction between representations of the individual objects in a category and abstract representations of the category itself. For example, the very same neural net can serve both of these functions, as summarized in Box 2.

This second, more complex view gives us another way of understanding the impact of prioritized memory for high-value objects on category representation. As we noted above, this impact might arise in part because of a process that takes place separately on each occasion when people observe an object (with people devoting more computation to accurately remembering each high-value object than to remembering each low-value object). However, it is also possible that value judgments more directly impact the abstract category representation. If it is especially important for people to accurately remember high-value objects, and if people remember objects in part by having the right sort of abstract category representation, there might be processes whereby people's value judgments more directly impact their abstract category representation in ways that serve to facilitate memory for high-value objects.

In sum, we have been considering the hypothesis that the similar patterns we observe for conceptual representation and prioritized memory arise because of a deeper link between these two phenomena, and in this section, we pointed to two different ways of implementing that hypothesis in more concrete detail. Further work should continue to examine evidence for and against the hypothesis itself and also to explore these more detailed implementations. We outline several potential directions in Outstanding Questions.

5. Explaining category representation in terms of memory prioritization

The framework proposed here has the potential to give us a deeper understanding of why value has an impact on conceptual representation in the first place, from a functional perspective. Suppose we ignore prioritized memory and just focus on conceptual representation. It might then seem puzzling that value judgment has the sort of impact it does. Why exactly would it be helpful for prototypicality judgments, normality judgments, standards for gradable adjectives and what comes to mind to all be determined by a mixture of statistical frequency and value?

By contrast, when we turn to prioritized memory, the functionalist explanation is straightforward. We can easily see why it might be helpful under certain conditions to have especially accurate representations of high-value objects [6,28], and existing research has already developed precise formal accounts of how to do so optimally [27].

We have been suggesting that there is a tight, bidirectional link between conceptual representation and prioritized memory at a mechanistic level. But, if this suggestion turns out to be correct, it opens the door to a possible deeper explanation of the effect of value on conceptual representation at a functional level. Specifically, it might be that people have an array of psychological processes that are needed in order to facilitate especially accurate representations of high-value objects, and these psychological processes then generate an impact of value on conceptual representation. In this case, the importance of value for prioritized memory explains its effect on conceptual representation.

As an illustration, consider people's ordinary way of thinking about artists. If we ignore prioritized memory and just think about conceptual representation, it might seem puzzling that people's conceptual representation of the category of artists should be impacted by value. But suppose instead we focus on the psychological capacities people use to accurately represent individual artists. It makes sense that people might have a greater need to accurately represent great artists than to accurately represent mediocre artists. But then, if there is a tight link between conceptual representation and prioritized memory, the psychological processes that enable people to more accurately represent great artists will themselves lead to a representation of the category of artists that is impacted by value.

This approach might then give us insight into the precise notion of "value" that plays a role in conceptual representation. Existing work shows that there is some sense in which conceptual representation is biased toward things that are 'good,' but difficult questions arise about precisely which notion of goodness is relevant here [13]. The present hypothesis has some potential to enable further insight into this complex question.

To illustrate, consider two very different ways in which one might understand the idea that people's representation of tigers could be biased toward tigers that are 'good.' On one hand, one might think that the best tiger is the strongest, most ferocious one; on the other, one might think that the strongest, most ferocious tiger is highly dangerous and that the best tiger to encounter is the one that is weakest and most docile. Just thinking about conceptual representation in isolation, it might not be clear which of these two notions of

Correspondence should be directed to Joshua Knobe, joshua.knobe@yale.edu

goodness should play a role. The present hypothesis suggests a way forward. It says that the notion of goodness that plays a role in conceptual representation is the one that plays a role in prioritized memory. Thus, if the notion of goodness that plays a role in prioritized memory is the one according to which the best tiger is the strongest and most ferocious, the hypothesis predicts that conceptual representation of tigers would be biased toward the strongest and most ferocious tiger. Further work could continue to explore this issue by looking at which exact notion of goodness plays a role in prioritized memory and then asking whether that exact notion of goodness is the one that plays a role in conceptual representation.

Going even further in this direction, it seems likely that there will be cases in which people show prioritized memory not for objects that are good but rather for objects that are bad. Take the domain of illness. One might think that people would show better memory for illnesses that are especially bad (cancer, AIDS) than for illnesses that are less bad (the common cold). The present hypothesis would then predict that this same effect should emerge for conceptual representation. For example, cancer should be regarded as an especially prototypical illness, and should be especially likely to come to mind, relative to what one would expect based on the purely statistical considerations. Further research could test this prediction.

6. Concluding remarks

We have reviewed a surprising consilience between recent findings in the study of category representation and recent findings in the study of prioritized memory. Research in category representation finds an effect of value, with judgments that something is good generally having a very similar impact to something being statistically frequent. Research in prioritized memory also finds an effect of value, with greater encoding, consolidation and retrieval for things that are seen as good than for things that are not seen as good. A question now arises about the relationship between these two findings.

We have proposed a hypothesis according to which they reflect a common psychological phenomenon, grounded in the bidirectional influences between prioritized memory and category representation. The human capacity for remembering individual token objects shows prioritized memory for things people regard as good. But the human capacity for category representation is not a completely separate capacity; it is deeply entwined with the capacity for representing token objects. Thus, the functional importance of prioritizing memory for token objects that are regarded as good, together with the coupling of memory and category representation, may explain why objects that are regarded as good also play an outsized role in category representation.

References

1. Rosch, E. and Mervis, C.B. (1975) Family resemblances: Studies in the internal structure of categories. *Cogn. Psychol.* 7, 573–605
2. Foster-Hanson, E. and Rhodes, M. (2019) Is the most representative skunk the average or the stinkiest? Developmental changes in representations of biological categories. *Cogn. Psychol.* 110, 1–15
3. Lake, B. et al. (2015) Deep neural networks predict category typicality ratings for images. In *37th Annual Meeting of the Cognitive Science Society*. Cognitive Science Society.
4. Bear, A. and Knobe, J. (2017) Normality: Part descriptive, part prescriptive. *Cognition* 167, 25–37
5. Bear, A. *et al.* (2020) What comes to mind? *Cognition* 194, 104057
6. Anderson, J.R. and Schooler, L.J. (1991) Reflections of the Environment in Memory. *Psychol. Sci.* 2, 396–408
7. Nussenbaum, K. and Hartley, C.A. (2021) Developmental change in prefrontal cortex recruitment supports the emergence of value-guided memory. *eLife* 10, e69796
8. Braun, E.K. *et al.* (2018) Retroactive and graded prioritization of memory by reward. *Nat. Commun.* 9, 4886
9. Aka, A. and Bhatia, S. (2021) What I like is what I remember: Memory modulation and preferential choice. *J. Exp. Psychol. Gen.* 150, 2175–2184
10. Wittmann, B.C. *et al.* (2005) Reward-Related fMRI Activation of Dopaminergic Midbrain Is Associated with Enhanced Hippocampus-Dependent Long-Term Memory Formation. *Neuron* 45, 459–467
11. Hennessee, J.P. *et al.* (2019) Forget me not: Encoding processes in value-directed remembering. *J. Mem. Lang.* 106, 29–39
12. Armstrong, S.L. *et al.* (1983) What some concepts might not be. *Cognition* 13, 263–308
13. Barsalou, L.W. (1985) Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *J. Exp. Psychol. Learn. Mem. Cogn.* 11, 629–654

Correspondence should be directed to Joshua Knobe, joshua.knobe@yale.edu

14. Foster-Hanson, E. and Lombrozo, T. (2022) How "is" shapes "ought" for folk-biological concepts. *Cogn. Psychol.* 139, 101507
15. Goldring, M.R. and Heiphetz, L. (2020) Sensitivity to ingroup and outgroup norms in the association between commonality and morality. *J. Exp. Soc. Psychol.* 91, 104025
16. Henne, P. *et al.* (2021) Norms Affect Prospective Causal Judgments. *Cogn. Sci.* 45
17. Roberts, S.O. *et al.* (2019) The role of group norms in evaluating uncommon and negative behaviors. *J. Exp. Psychol. Gen.* 148, 374–387
18. Willemsen, P. and Kirfel, L. (2019) Recent empirical work on the relationship between causal judgements and norms. *Philos. Compass* 14
19. Kennedy, C. and McNally, L. (2005) Scale Structure, Degree Modification, and the Semantics of Gradable Predicates. *Language* 81, 345–381
20. Verheyen, S. *et al.* (2018) Subjectivity in gradable adjectives: The case of *tall* and *heavy*. *Mind Lang.* 33, 460–479
21. Lassiter, D. and Goodman, N.D. (2017) Adjectival vagueness in a Bayesian model of interpretation. *Synthese* 194, 3801–3836
22. Schöller, A. and Franke, M. (2017) Semantic values as latent parameters: Testing a fixed threshold hypothesis for cardinal readings of few & many. *Linguist. Vanguard* 3
23. Schmidt, L.A. *et al.* (2009) How tall is tall? Compositionality, statistics, and gradable adjectives. presented at the Proceedings of the 31st annual conference of the cognitive science society, 31, pp. 2759–2764
24. Egré, P. and Cova, F. (2015) Moral asymmetries and the semantics of "many." *Semant. Pragmat.* 8
25. Madan, C.R. *et al.* (2012) High Reward Makes Items Easier to Remember, but Harder to Bind to a New Temporal Context. *Front. Integr. Neurosci.* 6
26. Shohamy, D. and Adcock, R.A. (2010) Dopamine and adaptive memory. *Trends Cogn. Sci.* 14, 464–472
27. Mattar, M.G. and Daw, N.D. (2018) Prioritized memory access explains planning and hippocampal replay. *Nat. Neurosci.* 21, 1609–1617
28. Anderson, R.B. *et al.* (1997) Need probability affects retention: A direct demonstration. *Mem. Cognit.* 25, 867–872
29. Griffiths, T.L. *et al.* (2007) Google and the Mind: Predicting Fluency With PageRank. *Psychol. Sci.* 18, 1069–1076
30. Dodgson, D.B. and Raymond, J.E. (2020) Value associations bias ensemble perception. *Atten. Percept. Psychophys.* 82, 109–117
31. Rosenbaum, G.M. *et al.* (2022) Valence biases in reinforcement learning shift across adolescence and modulate subsequent memory. *eLife* 11, e64620
32. Marchant, N. *et al.* (2023) A Context-Dependent Bayesian Account for Causal-Based Categorization. *Cogn. Sci.* 47, e13240
33. Ponari, M. *et al.* (2020) The role of emotional valence in learning novel abstract concepts. *Dev. Psychol.* 56, 1855–1865
34. Rehder, B. and Kim, S. (2010) Causal status and coherence in causal-based categorization. *J. Exp. Psychol. Learn. Mem. Cogn.* 36, 1171–1206
35. Bartlett, F.C. (1932) *Remembering: A study in experimental and social psychology*, Cambridge University Press
36. Gershman, S.J. (2021) The rational analysis of memory. In *The Oxford handbook of human memory*
37. Hemmer, P. and Steyvers, M. (2009) A Bayesian Account of Reconstructive Memory. *Top. Cogn. Sci.* 1, 189–202
38. Kingma, D.P. and Welling, M. (2013) Auto-Encoding Variational Bayes. DOI: 10.48550/ARXIV.1312.6114