Words, images, and concepts

Daniel A. Weiskopf

A central problem for any theory of cognition is to specify what the medium or vehicle of thought is. Historically, answers divide into roughly three categories. Some claim that the words of natural language are the medium of thought. Others claim that we think in sensory or perceptual images. Dual-coding approaches in psychology have sometimes attempted to split the difference between these (Paivio, 1990, 2007). A final possibility is that thought is neither verbal nor imagistic, but that it takes place in a special-purpose amodal conceptual medium. On this view, concepts may be similar in some respects to language and images, but they nevertheless belong to a distinct system with its own representational capacities and processes.

In *Words and Images*, Christopher Gauker develops a novel dual-coding approach to cognition. He proposes that conceptual thought should be identified with a kind of inner language use, while the medium of nonconceptual thought is imagistic. To represent an object, event, or scenario imagistically is to assign it to a location in perceptual similarity space. The dimensions of this space correspond to the objective qualities in the world that can be detected and processed by a creature's perceptual apparatus. A mark in perceptual space is a (possibly partial) vector that represents a certain set of values on these dimensions; distance in the space is therefore a measure of perceptual similarity between the objects responsible for the marks. An *image* is any sort of representation that makes use of these and only these resources.

The question is how many cognitive capacities can be explained in terms of this comparatively austere basis. A brief list culled from Gauker's text includes "imaginative understanding" (p. 150) of certain counterfactual scenarios, the ability to use this understanding

to represent goals and make plans (p. 225), the ability to reason about "imagistic causation" (pp. 161-163), the ability to formulate hypotheses (p. 226) and elaborate on images in coherent ways (p. 227), and the ability to use simple labels instrumentally and as devices to coordinate behavior (pp. 174-183).

This is an impressive list. At the outset, however, it isn't obvious that this model has sufficiently rich representations to capture elementary features of animal behavior, particularly those that seem to involve computing over abstract representations of space and time (Gallistel, 2009). Consider the ability to navigate through space—that is, not merely to avoid some present feature of the world but to move in a directed way towards a goal. This ability is phylogenetically ancient, appearing even in insects such as desert ants and honeybees, as well as in mammals such as gerbils. Foraging ants may leave the nest, which is a fixed starting point, and travel widely in search of food, returning successfully despite the fact that the origin of their voyage is no longer visible, and doing so without wastefully retracing every step along the way. Ants that are displaced in transit also follow return paths that would have taken them back to their origin point from their pre-displacement location.

There is widespread agreement that one mechanism that enables them to dead reckon so successfully is a form of path integration: as they travel, they update their position in allocentric space relative to the nest. This lets them calculate a direct return path by vector addition (Collett & Graham, 2004). This relatively standard model of dead reckoning requires assuming that ants have an odometer or accumulator that logs total distance traveled, as well as recording the direction of each leg of the journey. In walking insects, the accumulator is probably updated based on leg movements, while in flying insects such as honeybees it may be updated based on

optic flow plus bodily sensations such as the decrease in the volume of nectar being carried in their stomachs, which is a measure of total energy expended (Gould, 2004).

The relevance of this basic form of navigation is that it appears to draw on representations of space that go beyond those posited in the imagistic model, and also to involve representation of abstract quantities such as total distance traveled. The spatial model these insects use is not a purely egocentric one that draws only on distances to objects seen from their own visual perspective. Instead, they keep track of locations in the world relative to one another, even when none of those locations are currently being perceived. This form of spatial representation is not completely independent of perception, since insects also appear to take sensory "snapshots" that are associated with particular landmarks along their path. These snapshots can help them to recalibrate the path integration system on long voyages. Here perceptual similarity to stored images undoubtedly plays a role. But the spatial maps that are used in navigation do not themselves appear to be perceptual maps, but instead something more abstract and perspective-independent.

So the first challenge for the imagistic model is to account for capacities like allocentric spatial navigation, which seems to draw on a different set of representations than those used in perception. This is not a matter of simply adding dimensions to the same perceptual space, but of mapping locations in a different *kind* of space. Putting the point a little differently, allocentric navigation may involve spatial images, but not perceptual images. This actually points to a slightly wider moral, which is that focusing on a division of all mental representations into *either* perceptual or conceptual is probably too simple. There is a wide range of nonconceptual representations that may be used by specialized nonperceptual subsystems such as the dead reckoning device.

Turn now to the role of conceptual representations. The primary feature that Gauker ascribes to conceptual thought is the capacity to judge that a particular belongs to a kind. As the book states up front: "Language is necessary for every kind of thought that involves judging, of some particular, that it belongs to some kind" (Gauker, 2013, p. 13). Predication requires representations that have the requisite structure: they need to have argument places to which representations of individuals can be bound. In his critique of Locke, for instance, Gauker argues that the ideas that empiricists say are given to us in perception do not come with such argument places built in, and therefore cannot be the source of our concepts (pp. 32-33). And while he entertains the thought that elementary atomic judgments may be captured by similarity space representations, he argues that judgments involving nested quantification cannot (pp. 108-109). Perceptions and perceptual similarities are unarticulated, and so cannot form the basis for conceptual thought.

The notion of a "kind" is not explicitly unpacked in Gauker's text. However, in the literature on natural kind terms and concepts, one prominent supposition is that kinds are ontological categories whose membership is not determined purely by appearance properties. To be a sample of a chemical or biological kind requires more than looking the right way, it requires having the right history, internal structure, and so on. Species membership might be determined by a creature's genetic makeup, its ability to interbreed with other species members, or its participation in a historically delimited clade, none of which are overtly visible traits of that particular creature. While these properties are typically referred to as "kind essences," we needn't adopt any such heavy-duty metaphysical construal here. All that matters is that the genuinely membership-determining properties for a kind typically either are or include properties that are not perceivable.

Not all concepts are kind concepts, since that would bar the existence of concepts picking out the various dimensions of perceptual quality space (e.g., color concepts), as well as many others that aren't designed to track kind groupings at all. But it might be a requirement on having concepts that a creature *can* represent kinds; or, more specifically, that a creature has perception-transcendent representational capacities. This implies the ability to represent categories whose boundaries cross-cut those that are established by the creature's perceptual systems. In this sense, kinds may group together individuals that are perceptually dissimilar and may distinguish between individuals that are perceptually similar. In doing this, kind concepts draw on some form of information that goes beyond what is given by the perceptual systems themselves, since if perception were sufficient to distinguish the category's membership perfectly it would not be a kind in the present sense of the term.

Putting these claims together, having concepts requires two things. First, it requires having a predicatively structured system of representations, typically one that supports thoughts having a degree of logical complexity up to embedded quantification. Second, it requires the ability to represent kind categories, or perception-transcendent groupings more generally.

For Gauker, learning a concept requires learning the use of a term, which is minimally to acquire a set of production and acceptance rules. These rules determine the circumstances under which a creature will be disposed to produce or accept atomic sentences. These atomic sentences are "labeling uses" of terms: they consist of an occurrence of a demonstrative plus a monadic predicate. In the learning situation, a person is confronted with a scenario registered as a mark in perceptual similarity space that falls a certain distance from other marks that record previous perceptual encounters. Some of these previous encounters will also have been associated with utterances of atomic labeling sentences, which are linked to the marks themselves. This can be

seen as a type of exemplar memory system that stores linguistically augmented perceptual records of encounters with objects or situations (Sloutsky & Fisher, 2004).

Equipped with such memory traces, a person is disposed to produce a sentence "this is F" in the event that the new mark is between the two nearest marks that were also labeled with "this is F" and is not between the nearest two marks that were labeled with "this is not F". More specifically, the new mark must fall between two nearest neighbor marks labeled with "this is F" that belong to a *cluster*, where a cluster is defined as a set of marks all of whose members are between at least two other members. The clustering requirement ensures that each mark is sufficiently similar to at least two other marks that receive the same linguistic label, though any two co-labeled marks chosen at random may be quite dissimilar from one another, so long as there is a chain of similar marks connecting them. Labels group together marks within an intrinsically boundaryless similarity space. Since a set of labeled marks may include dissimilar items, this may provide a way of representing categories that transcend perception. Appropriately distributed labels give a possible mechanism for representing kinds.

There is undoubtedly something plausible about this picture. Words are powerful indicators of category membership for young children. In studies by Sandra Waxman and her colleagues, for example, labels enable 12-month old children to co-classify perceptually dissimilar objects at the superordinate level, a task which they find challenging in the absence of verbal cues (Waxman & Markow, 1995). In the developmental literature more generally, there is a distinction between models of word learning that are driven by perceptual associations and those that are (for lack of a better term) inferentially driven. The main explanatory apparatus of Gauker's model appears closer to the associationist proposal. Such models have been extensively

developed by Linda Smith and her colleagues (Smith, Jones, & Landau, 1996; Smith, 2000). However, they aren't without their challenges.

For instance, a set of studies with 1.5 and 2 year olds showed that how children choose to extend new words is a function of the ontological category that they are presumed to map onto, not just a matter of perceptual appearance (Booth, Waxman, & Huang, 2005). These children were shown a simple shape that was labeled (as a "dax" or a "riff") and described in terms of different characteristic properties: "this dax has a mommy and daddy who love it very much" versus "this dax was made by an astronaut to do a special job on her spaceship", for example. These suggested a construal of the object as either an animate creature or an artifact. They were then shown triads of objects and asked to find another dax. The triads varied along dimensions of similarity such as shape, size, and texture. Both animate and artifact labels were extended to same-shaped objects (but not different-shaped objects) and could be extended to different sized objects; however, only artifact labels were extended to objects having different texture. Since the same objects were used in each condition, any differences must be in how the object was construed by the child, and in the properties that are relevant to membership in that general category. Similar results were found in 3-year olds when they were presented with objects that had the appearance of animate entities (e.g., with big cartoonish eyes) but that were described in terms of artifact function (Booth & Waxman, 2002). Information about artifact function could override perceptual appearance in these children's projection of words to new instances.

Causal information pertaining to different biological roles also governs how words are extended. In studies with kindergarten-age children, line drawings of juvenile bug larvae were described as either the offspring of a pair of bugs, or as the prey of that pair (Opfer & Bulloch, 2007). Where a bug was characterized as an "offspring," its kind membership depended on the

child's beliefs about the biological species of its parents, rather than its perceptual similarity to the previously displayed offspring. On the other hand, when the same item was characterized as "prey," its kind membership was determined by its perceptual resemblance to previously named prey insects. When all of these biological facts are removed, however, children revert to classifying exemplars simply by perceived similarity. Information about different causal roles can govern judgments about classification even at a young age (mean in this study=5.7 years old). Importantly, this holds even when no verbal labels for the exemplars are presented.

While it is sometimes said that these studies show the influence of "conceptual information" on early classification, in the present context that would obviously be a contentious description. What does seem clear, however, is that word learning is influenced by more than perceptual similarities and labels. Quite specific causal expectations are also at play from a very early stage, making human language learning much richer than it is depicted in Gauker's model. Another way to put this point is that children have a set of early developing kind expectations that govern the sorts of causal structure they associate with particular categories, as well as how they categorize and make inductive inferences.

Gauker does propose that imagistic representation can support forms of causal reasoning. If I wonder what will happen if I poke a turtle with my finger, I may answer my own question by imagining poking it. This prompts recall of a past episode in which a similarly poked turtle withdrew its head, and recalling this leads me to think that the same thing will happen this time as well. Possibly I even grasp the causal generalization that poked turtles will, in general, withdraw. But, says Gauker, "I do not think I need to know any such rule, written out in my mind in words or in concepts" (p. 150). Nor is such a causal rule inscribed in images. This general fact about poked turtles is one that I know *implicitly*. It is not represented in an imagistic

format, nor in a non-imagistic one. Rather, it is represented by the fact that imagining the first event leads to retrieving or generating an image of the next event, in virtue of the former's similarity to past imagistic episodes.

Imagistic causality has two main characteristics: it is Michottean and Humean. As infant studies suggest, certain perceived interactions "look causal" to us. Consider the sight of a wrecking ball making contact with a wall, followed by its disintegration into a pile of bricks, mortar and dust. Gauker only mentions apparent contact as a key feature of the stimuli, but we can add containment and support relations, both of which are often implicated in the perception of causal relations. Token imagistic causality is registered when two contact, containment, or support events occur against an appropriate background. To generalize from instances of Michotte causation to a broader causal hypothesis involves a Humean process of matching the observed imagistic cause to similar remembered ones, and predicting a similar effect. This process may also be driven by temporal contingencies between these events even in the absence of Michottean contact.

Humans, however, use forms of causal reasoning and representation that are richer than the imagistic sort. For example, one prominent model of human causal inference is Patricia Cheng's Power PC model (Cheng, 1997). This model posits that human causal inference involves the assumption that observed contingencies among events are produced by the interactions of the unobserved causal powers possessed by objects. While the strength of these causal powers can be estimated from various probabilities, the powers themselves are independent of the observed frequencies, and the model treats them as interacting to generate these frequencies. Causal reasoning on this view involves attempting to learn the causal powers

possessed by objects in the environment, their strength, and how they interact to generate or inhibit events.

In a similar vein, a substantial body of research suggests that people tend to expect causes to produce effects by means of unobserved mechanisms. What transforms a probabilistic regularity between a *cue* and an *outcome* into a relationship between a *cause* and an *effect* is the presence of some "system of connected parts that operate or interact to make or force an outcome to occur" (Ahn & Kalish, 2000, p. 201). Where these mechanisms are assumed to be present, they organize and explain covariation information, and distinguish causal covariations from accidental ones. The domain-specific expectation of mechanistic structure is another factor that distinguishes human from Humean causal reasoning.

For a final example, consider research on the causal status effect in categorization (Ahn & Kim, 2001). The effect is constituted by the fact that people expect many of the features that belong to members of a category to have some cause, rather than to be merely accidentally possessed by category members, and they give this cause a high degree of precedence in determining category membership. The greater the degree of causal "fundamentality" possessed by a feature, the more significance it has in categorization, even if it is outweighed in strictly numerical terms by other features (Hayes & Rehder, 2012). This involves both the assumption of a certain generic causal structure possessed by many categories, and a default way in which this structure is processed. The preference for classifying in terms of causes rather than effects also extends to other tasks, such as judgments of similarity; for instance, people who display the same observable symptoms but have a different underlying cause are judged to be less similar than those who have the same underlying cause but different symptoms.

Causal reasoning, then, is characterized by a cluster of properties: a concern with tracking unobserved causal powers, the tendency to posit mechanisms to transform covariations into causal regularities, and the use of specific sorts of causal structures to represent categories. This cluster seems to exceed the explanatory reach of imagistic causality, and it is absent from most (if not all) nonhuman causal cognition. Its presence in human beings must have a different origin.

It is hard to see how the language faculty, all by itself, can be an explanation for this suite of cognitive capacities. The learning model Gauker sketches gives a mechanism for grouping together otherwise dissimilar perceptual scenarios, but not for representing these more distant abstractions. But more importantly, language itself is only a medium. Its structure makes available new forms of processing, but does not dictate any particular one of them. The challenge in understanding causal cognition is uncovering these highly specific inferential rules and processes. While there is no consensus on precisely how these processes operate, they need to be considerably more powerful than the contact-based, associationistic ones that are part of image-based thought. The addition of language to an image-based cognitive system does not suffice to explain the additional reasoning abilities of humans versus other nonlinguistic creatures; for that, a causal reasoning system is necessary.

To be clear, there is nothing *per se* in the model Gauker sketches that prohibits him from proposing that human cognition differs from its nonhuman relatives in more ways than just the addition of language. Perhaps inherently more powerful causal reasoning processes are also part of our cognitive phenotype. My hesitancy on this point is due to the fact that in his discussion of intrapersonal discourse, Gauker seems averse to positing this sort of machinery. He allows that "[t]here may be some truth in the idea that language is the medium by which we abstract general

truths" (p. 276), but goes on to add that he does not see the utility of such explicitly formulated generalizations: "I am not sure why the guidance provided by such generalizations could not be provided equally well or better by the devices of imagistic cognition" (p. 277). Causal models of the sort described here are paradigmatic ways of capturing these abstract generalizations, but there is little reason to think they are grounded solely in imagistic cognition.

A last point concerns the nature of linguistic representation itself. In his final chapter, Gauker attempts to disavow any commitment to problematic forms of Whorfianism. This threat arises because if conceptual thought is inner speech, and if languages are substantially diverse in their morphological, syntactic, and semantic profiles, then conceptual structure will presumably also vary across speakers of different languages. Given massive linguistic diversity, the disunity of human conceptual thought would seem to follow. Wishing to avoid this conclusion, Gauker commits himself to a strong claim about the universal sameness of all the world's languages, which he says are "enough alike that people who think by means of any of them think basically alike" (p. 16; see also p. 270).

This may seem surprising, since how much commonality there is among languages is one of the most highly contested issues in linguistics. In traditional approaches to typology, such as Bernard Comrie's classic survey, for example (Comrie, 1989), we find a relatively firm commitment to universals. On the other hand, there are also significant dissenters from this consensus (see Evans & Levinson, 2009 for extensive discussion). Even the claim that all languages have "recursive structure" has been challenged, both on general theoretical grounds (Pullum & Scholz, 2010) and based on investigation of particular languages such as Pirahã (Everett, 2010). So the assumption that all of the world's 5000-7000 languages are at some basic level the same is hardly obvious.

But it is also unclear what "the same" means in this context. It cannot mean that languages have the same overt structure, since this isn't true, given the manifest variation in the their surface characteristics. It might mean what some advocates of universal grammar have in mind: that there is a common set of principles applicable to all (possible) human languages. For all languages to be the same is for them to be generated or constrained in the same way. This does not imply that the languages are the same in terms of their distinctive internal descriptions, however, or that they are represented and processed in the same way by their speakers. It means only that grammarians can find a compact set of general parameters or features that account for all of the possible variation in languages. As noted, this claim is contentious, but it also does not seem relevant here, since the fact that grammarians can produce such a list of universal characteristics does not address the issue of whether these superficially different languages are mentally represented as the same.

The sense of sameness that appears to matter for Gauker's purposes is one on which superficially diverse languages share a level of representation that abstracts away from their differences. For example, this would be an inner format in which SVO and VSO languages share a common ordering, and one on which polysynthetic and isolating languages share a common type of morphemic decomposition. A common level of representation at which these differences disappear to be replaced by a common encoding would give psychological substance to the idea that "all languages are the same" in a way that would block the worst Whorfian consequences. It is also, however, less clear that this is a view on which natural language itself is the medium of thought. On this view, thought takes places several levels removed from the familiar sounds and structures we read, write, speak, and hear. It is no coincidence that proposals about such a level

of abstract representation have typically considered it to be more appropriately called a semantic or conceptual level (Jackendoff, 1983).

I conclude with the following suggestion. Suppose that the sense in which all languages are the same is one in which they map onto a common underlying format of this sort, which has enough structure to support predicative judgments and contains the resources for representing kinds. Suppose also that there are causal reasoning mechanisms that can access this format and enable it to be used in reasoning tasks that go beyond what can be achieved using imagistic understanding of causality. And suppose furthermore that this system of representations is autonomous from overt language production, so that it can survive the kind of destruction of the language faculty that takes place in profound aphasia (Monti, Osherson, Martinez, & Parsons, 2007; Monti & Osherson, 2012; Varley, Klessinger, Romanowski, & Siegal, 2005), a speculation that Gauker himself at one point entertains (p. 170).

An abstract, autonomous, structured system of representations of this sort would be close enough to what many people—myself included—mean by a conceptual system (Weiskopf, 2014). On this view, the medium of thought is neither words nor images, but concepts. Since these modifications seem to be well-motivated, either from within Gauker's project or by the empirical phenomena, it is not clear how much distance there is between Gauker's view, suitably enriched, and his opponent's. It is an open question at this point whether there are further grounds for distinguishing them.

References

Ahn, W., & Kalish, C. W. (2000). The role of mechanism beliefs in causal reasoning. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and Cognition* (pp. 199–225). Cambridge: MIT Press.

Ahn, W., & Kim, N. S. (2001). The causal status effect in categorization: an overview. In D. L. Medin (Ed.), *The Psychology of Learning and Motivation*, Vol. 40 (pp. 23-65). San Diego, CA: Academic Press.

Booth, A. E., & Waxman, S. R. (2002). Word learning is "smart": evidence that conceptual information affects preschoolers' extension of novel words. *Cognition*, *84*, B11–22.

Booth, A. E., Waxman, S. R., & Huang, Y. (2005). Conceptual information permeates word learning in infancy. *Developmental Psychology*, *41*, 491–505.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.

Collett, T. S., & Graham, P. (2004). Animal navigation: path integration, visual landmarks and cognitive maps. *Current Biology*, *14*, R475–7.

Comrie, B. (1989). *Language Universals and Linguistic Typology* (2nd ed.). Chciago: University Of Chicago Press.

Evans, N., & Levinson, S. C. (2009). The myth of language universals: language diversity and its importance for cognitive science. *The Behavioral and Brain Sciences*, *32*, 429–494.

Everett, D. L. (2010). You drink. You drive. You go to jail. Where's recursion? Posted online at http://ling.auf.net/lingBuzz/001141.

Gallistel, C. R. (2009). The Foundational Abstractions. In M. Piattelli-Palmarini, J. Uriagereka, & P. Salaburu (Eds.), *Of Minds and Language: A Dialogue with Noam Chomsky in the Basque Country* (pp. 58–73). Oxford, UK: Oxford University Press.

Gauker, C. (2013). *Words and Images: An Essay on the Origin of Ideas*. Oxford University Press.

Gould, J. L. (2004). Animal navigation. *Current Biology*, *14*, R221–4.

Hayes, B. K., & Rehder, B. (2012). The development of causal categorization. *Cognitive Science*, *36*, 1102–28.

Jackendoff, R. (1983). *Semantics and Cognition*. Cambridge, MA: MIT Press.

Monti, M. M., & Osherson, D. N. (2012). Logic, language and the brain. *Brain Research*, *1428*, 33–42.

Monti, M. M., Osherson, D. N., Martinez, M. J., & Parsons, L. M. (2007). Functional neuroanatomy of deductive inference: a language-independent distributed network. *NeuroImage*, *37*, 1005–16.

Opfer, J. E., & Bulloch, M. J. (2007). Causal relations drive young children's induction, naming, and categorization. *Cognition*, *105*, 206–17.

Paivio, A. (1990). *Mental Representations: A Dual Coding Approach*. New York: Oxford University Press.

Paivio, A. (2007). *Mind and Its Evolution: A Dual Coding Theoretical Approach*. New York: Psychology Press.

Pullum, G. K., & Scholz, B. C. (2010). Recursion and the infinitude claim. In Harry Van der Hulst (Ed.), *Recursion and Human Language* (pp. 113-137). Berlin: Mouton de Gruyter.

Sloutsky, V. M., & Fisher, A. V. (2004). Induction and categorization in young children: a similarity-based model. *Journal of Experimental Psychology. General*, *133*, 166–88.

Smith, L. B. (2000). Learning how to learn words: An associative crane. In R. M. Golinkoff, K. Hirsh-Pasek, L. Bloom, L. B. Smith, A. L. Woodward, N. Akhtar, M. Tomasello, & G. Hollich (Eds.), *Becoming a Word Learner: A Debate on Lexical Acquisition* (pp. 51–80). New York: Oxford University Press.

Smith, L. B., Jones, S. S., & Landau, B. (1996). Naming in young children: a dumb attentional mechanism? *Cognition*, *60*, 143–71.

Varley, R. A., Klessinger, N. J. C., Romanowski, C. a J., & Siegal, M. (2005). Agrammatic but numerate. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 3519–24.

Waxman, S. R., & Markow, D. B. (1995). Words as invitations to form categories: Evidence from 12-to 13-month-old infants. *Cognitive Psychology, 29*, 257-302.

Weiskopf, D. A. (2014). The architecture of higher thought. In M. Sprevak & J. Kellestrup (Eds.), *New Waves in Philosophy of Mind* (pp. 242–261). New York: Palgrave Macmillan.