A Theory of How Heuristics Work in Cognition

Sheldon J. Chow (sheldonjchow@gmail.com)
Institut für Philosophie, Leibniz Universität Hannover
Im Moore 21, 30167 – Hannover, Germany

Abstract

There is relatively little work in cognitive science that investigates the nature of heuristics and the mechanics of their operations. This paper outlines a theory of how heuristics work in cognition. A cognitive architecture is described which facilitates satisficing procedures which utilize little cognitive resources, i.e., heuristics. An interesting feature of this account is that heuristics are not assumed to operate with representational content, but with higher-order information (or knowledge) implicit in informationally rich representational systems and relations among them.

Keywords: Heuristics; satisficing; cognitive architecture; representations; representational systems; inference

Introduction

Heuristics are commonly understood as economical shortcut reasoning procedures that may not lead to optimal or correct results, but will generally produce outcomes that are in some sense satisfactory or "good enough". Since the seminal work of Kahneman and Tversky in the 1970s, there has been an explosion of research that tries to uncover various heuristics that humans typically rely on in problem-solving or decision-making. Yet there has been comparatively little research on the general nature of heuristics and how they operate in cognition.

An exception is the recent incarnation of heuristics research in psychology, exemplified by Gigerenzer's "fast and frugal heuristics" research program. Gigerenzer and his colleagues make various attempts to explain how heuristics work by proposing computational models and exposing the information that a given heuristic is sensitive to. However, it is not clear whether human cognition actually implements the proposed computational models (Ayal & Hochman, 2009; Dougherty, Franco-Watkins, & Thomas, 2008; Hilbig, 2010; Hilbig & Pohl, 2008; Oppenheimer, 2003), and moreover, those models are abstract with little care for how they would fit within a general account of cognition.

Even less philosophical attention has been paid to the nature of heuristics and heuristic reasoning. Yet philosophy as a discipline possesses valuable resources to advance investigations into these important matters, such as accounts of cognitive architecture, theories of representation, theories of concepts, and theories of reason and cognition.

In this paper I outline a general theory of how heuristics work in human cognition which enjoys inspiration from philosophical reflections on cognitive archtiecture. While heuristics research in the psychological literature focuses on determining what specific information or knowledge is used by given heuristics, I will here reorient the problem by discussing architectural properties of the mind that facilitate heuristics generally. More specifically, I will investigate cognitive architectural features that can constrain heuristic processes in such a way that ensures that they satisfice, operate systematically, and require little cognitive resources for their recruitment and execution (Chow, in press; cf. Shah & Oppenheimer, 2008). My suggestion is that heuristics exploit informationally rich representational systems in cognition. However, contrary to what is typically assumed in the psychological literature, this account maintains that heuristics generally don't operate over representational content, but over higher-order information (or knowledge) implicit in certain structures embodied by representational systems. In this way, heuristics *indirectly* exploit our representational systems, and this, I maintain, is what enables heuristics to mitigate cognitive costs presented by cognitive tasks.

Lessons from Philosophy

Let us begin with some informative remarks from Dennett to guide our discussion. In his paper, "Cognitive Wheels", he writes:

> Even if you have excellent knowledge (and not mere belief) about the changing world, how can this knowledge be represented so that it can be efficaciously brought to bear? . . . A walking encyclopedia will walk over a cliff, for all its knowledge of cliffs and the effects of gravity, unless it is designed in such a fashion that it can find the right bits of knowledge at the right times, so it can plan its engagements with the real world. (Dennett, 1984, pp. 140-141)

Dennett is speaking here about the frame problem, but his point can be generalized. For human cognition is confronted with, but manages to solve, a general problem of achieving and maintaining an organization of knowledge so as to enable access to the appropriate, relevant information in its cognitive tasks. Dennett's remarks suggest that humans must possess certain cognitive structures that exhibit the requisite organization—that appropriately represent knowledge—to ensure access to the right information for successful planning and action in a complex world. Thus, in general, without the right kind of cognitive architecture we would be unable to make the inferences that we in fact do, and which underlie much of our cognition.

In a vein similar to Dennett's remarks, some theorists believe that a modular architecture enables fast and frugal heuristics (e.g., Gigerenzer, 2000; Carruthers, 2006). There are many ways that one can characterize a module, but there are two core features that are generally believed to facilitate quick and computationally cheap reasoning (i.e., the sort of reasoning characteristic of heuristics). These features are *domain-specificity* and *informational encapsulation*.[1]

Informational encapsulation is supposed to confine the amount of information that can be surveyed to a highly restricted proprietary database, which significantly reduces the computational burden associated with information search. Nevertheless, there are at least two reasons to doubt that heuristics owe their ability to be fast and frugal to informational encapsulation. First, heuristics are pervasive in central cognition, and yet central systems don't exhibit modular encapsulation: a widely recognized feature of central systems is that they allow for the free exchange of information,[2] which is antithetical to encapsulation (cf. Samuels, 2005). Second, if a system is informationally encapsulated, there would be little need for heuristics: information search would be sufficiently restricted, and the role for heuristics would be superfluous.[3]

On the other hand, perhaps it is domain-specificity that enables heuristics to operate fast and frugally. Indeed, the potential benefits of domain-specificity and its role in cognition are especially apparent. As Samuels explains, "if a mechanism is sufficiently domain specific, then it becomes possible to utilize a potent strategy for reducing computational load, namely, to build into the mechanism substantial amounts of information about the domain in which it operates" (Samuels, 2005, p. 111). The suggestion is that heuristics can be fast and frugal in virtue of the domain-specific information built into the systems in which they operate.

Nevertheless, if we are to take the lessons from Dennett seriously, then we must add to this that the information built into a system must be encoded in a highly organized fashion. A mechanism can take advantage of generous amounts of domain-specific information encoded in a system, and thereby enable quick and efficient cognition, but if the information is structured or organized in specific ways, it would enable quicker and more efficient cognition. Contrariwise, if such information is not structured or organized in specific ways, it would retard the speed and efficiency of search and processing. Characteristic human performance on many cognitive tasks suggests that the structures heuristics exploit are distinctly organized to produce systematicity and robustness in reasoning and inference.

Moreover, it would seem that a system must not only have a specific internal structure, but also bear specific intra-system relations, which would further facilitate quick and efficient reasoning. Reasoning within one domain often bears in a number of ways on other domains. Without such connections between bodies of knowledge, we would not be able to make the rich heuristic inferences we characteristically do.

But for all this, there is no need to refer to these organized systems as "domain-specific". It seems rather that what is really doing the work in facilitating fast and frugal cognition is not domain-specificity *per se*, but the manner in which the specific kinds of information is encoded and organized.[4] For, similar to what was recently observed, it is possible that one can have an *unorganized* system of lots of domain-specific information, but it will be doubtful that a cognitive mechanism would exhibit the same speed and efficiency operating with this unorganized body of knowledge as it would operating with a highly structured system. Moreover, it is possible that one can have a domain-specific body of knowledge that is quite impoverished; and although a cognitive mechanism would likely be able to operate quickly and efficiently with such a system, owing primarily to the little information that ever gets considered, the inferences made would not be as robust as those made by operations over richer systems.

In summary, a plausible feature of cognitive architecture that would enable the fast and frugal operations of heuristics is the possession of numerous highly-organized epistemic systems that allow easy access to certain information in certain circumstances. Since neither informational encapsulation nor domain-specificity really play a role here, it would be improper to call these systems modules. But what, then, could these systems be? It turns out that this philosophically-inspired analysis converges on some important psychological models of cognition.

Turning to Psychology and Neuroscience

There is a long tradition in psychology which describes human representational systems in terms of networks. These models have roots in semantic processing theories, according to which information about words and their meanings are stored and organized in complex webs of nodes. Nodes are understood to be specific representations—variously conceived as concepts, features, or lexical items—with connections that reflect associative and/or semantic relations. (See Collins and Loftus (1975) for the most popular account of such a model, which was inspired by Quillian's (1962) theory of semantic memory for "an understanding machine".) A key idea behind these network models is that a given node can be activated by some stimulus (e.g., someone uttering the word "red" activates the node that represents *red*) and activation spreads from this node-of-origin to other nodes along existing connections (e.g., to the *fire engine*, *apple*, *rose* nodes, and perhaps nodes representing other colours, depending on existing connections). An initial activation will be specified

[1] These two properties were, of course, those (among others) that Fodor (1983) had originally used to define the input-output modules that he argued subserve peripheral systems.

[2] This is a feature acknowledged by those who deny central systems modularity (e.g., Fodor, 1983, 2000; Samuels, 2005) as well as those who advocate for it (e.g., Carruthers, 2006; Sperber, 1994).

[3] Of course, there may be other reasons why an encapsulated system might deploy heuristics other than information search (e.g., pattern matching), but the claim that I'm trying to make here is that informational encapsulation and heuristics can be divorced in such a way that warrants resisting the idea that heuristics satisfice and require little cognitive resources *in virtue of* informational encapsulation.

[4] That a system of information is dedicated to a specific domain no doubt contributes to the extent to which the system is organized, since there are natural relations among items of information within a domain. However, domain-specificity itself does not seem to be *necessary* for a body of knowledge to be structured and organized in ways conducive to quick and efficient exploitation.

by a certain degree of activation which will in turn be divvied out to other nodes through spreading; those downstream activations will thus be correspondingly weaker, and activation will be divvied up and spread from them, exciting yet weaker activations of further nodes, and so on until activation peters out. On some models, the connections are directional and weighted (according to contextual features, such as recency or repeated activation) such that activation is more easily facilitated along certain connections and less easily along others.

Such network models have been used as paradigms to investigate conceptual cognition, memory recall and organization, semantic priming, and the properties of lexical systems (Lucas, 2000), and they are successful in explaining and predicting a variety of these phenomena (Baronchelli, Ferrer-i-Cancho, Pastor-Satorras, Chater, & Christiansen, 2013). All of this can be seen as corroborating the philosophically-inspired architectural view sketched above. Networks may be understood to embody the organization of representational information required for fast and frugal cognition, where heuristic search is guided by spreading activation.

A Cognitive Architecture

Nevertheless, the usual characterizations of network nodes are insufficient for an account of cognition. As recently stated, nodes are usually understood variously as concepts, features, or lexical items; and *qua* nodes, they are generally conceived to be unstructured units. Now, lexical items really seem to be unstructured units, and as such their natures and roles in cognition may very well be captured by network models. However, the *embodiment* view of concepts or semantics maintains that semantic representations (at least partly) consist of sensory-motor representations (e.g., Barsalou, 1999; Meteyard, Rodriguez Cuadrado, Bahrami, & Vigliocco, 2012; Patterson, Nestor, & Rogers, 2007; Thompson-Schill, 2003). If the embodiment view is correct—and there is a growing body of evidence favouring it—concepts and other conceptual representations (e.g., features) are not unstructured, but rather are constituted by numerous multimodal representations.[5]

As I am advancing a theory of how heuristics work, these considerations urge me to recast network models for spreading activation in such a way that nodes are conceived to be structured units consisting of numerous multimodal representations (including affective and lexical information; Meteyard et al., 2012). In a way, nodes might be understood as networks unto themselves where activation of a given node is spread among the various representations belonging to it, depending on contextual features. For example, one's *apples* node may consist of various representations related to its semantic content (pertaining to colour, shape, taste, use, category, etc.). If a stimulus activates this node, a subset of multimodal representations[6] are initially tokened and activation would spread from there. For instance, representations concerning the shape and (typical) colour of apples may be initially activated, and activation would spread to representations concerning the taste and texture of flesh of apples. Activations from those representations would then spread from the *apples* node outward to, for instance, representations of the shapes and colours of other fruits, and perhaps of other small roundish objects, such as baseballs.

Conceiving network nodes in this way seems to be precisely what is needed for the philosophically-inspired analysis of cognitive architecture above. The main lesson there was that highly organized epistemic structures are needed to facilitate fast and frugal cognition. In light of what has just been discussed, the internal structure of nodes together with the structure of the embedding network appear to provide the right sort of structures to facilitate access to salient representations (to be used to make inferences and judgments) by means of spreading activation.

Moreover, it would be worthwhile to maintain the features of many network models mentioned above, namely that the connections between nodes are directed and weighted (Baronchelli et al., 2013). Directed connections determine the flow of information from one node to another. This means that, if the connection between node A and node B is directed from the first to the second, activation of A will spread to B but not necessarily the other way around, unless there also exists a directed connection from B to A. On the other hand, weighted connections can be roughly understood as a parameter that helps facilitate information flow between connected nodes, where heavily weighted connections allow for greater information flow than less-heavily weighted connections. If we assume that node activation can vary in strength, then greater information flow will affect the degree to which nodes are activated downstream through spreading. For example, if node A has directed connections to nodes B and C, but A's connection to B has a greater weight than its connection to C, activation will be spread more readily to B enabling it to be more strongly activated relative to C.

Consider an example. Activating representations of *cats* might almost always activate representations of *furriness* for an individual, but activating representations of *furriness* might activate representations of *cats* only in certain contexts (the directed connection is heavily weighted from *cats* to *furriness*, and the directed connection is lighter the other way). Furthermore, if this individual is very familiar with cats, representations of *furriness* might be more readily activated relative to representations of *purr*, even though both *furriness* and *purr* bear directed connections from *cats*. This will entail

[5]There has been significant discussion regarding the relation between lexical knowledge and semantic knowledge, and the implication this would have for network models. Treating these matters is beyond the scope of this paper. However, I take it as plausible that semantic information is not exhausted by semantic content (Meteyard et al., 2012), and that lexical networks are separate but embedded within semantic networks (Lucas, 2000; cf. Barsalou, Santos, Simmons, & Wilson, 2008).

[6]Henceforth, when I refer to representations I mean to refer to multimodal representations, even though I omit explicit mention of their multimodal character.

that, for this individual across contexts involving cat representations, cats' furriness will in general be more salient than their purring. In short, these features add to the organized structure of one's representational system that, as we saw, is so important to facilitate access to the right representations at the right times.[7]

How Heuristics Work

When faced with a cognitive task, the relations among nodes in a network embody a rich source of implicit (putative) knowledge concerning the active representations, but which is not contained within the representations themselves. Such implicit information includes which representations are associated with others, structural or hierarchical relation information among representations (cf. Lenat, 1982), the measure and strength of association among representations, and which representations appear to be salient. This information is there to be used in various ways by cognitive processes that have access to it. The present account suggests that heuristics exploit this implicit (putative) knowledge, not primarily the representational content of the representations concerned[8] (cf. Clark, 2002).

To illustrate, let us briefly consider Gigerenzer's Take the Best heuristic (Gigerenzer & Goldstein, 1999):

> (i) Search among cues; (ii) stop search when one cue discriminates between two alternatives; (iii) choose the object picked out by step (ii).

Gigerenzer's explication of Take the Best assumes that an individual has a subjectively ranked order of beliefs about cues that may discriminate between objects, and which are considered sequentially when making predictions about the object along certain dimensions. The highest ranked cue which discriminates is termed "the best", and the object that has a positive value with respect to this best cue will be predicted to possess some criterion. For example, suppose you had to choose which of a pair of cities is larger. You might take recognition as a cue that would discriminate and thereby inform your choice. However, if both cities are recognized, you move down your ordered list of cues to the next one that discriminates. Suppose that having a university is the next cue that discriminates on the criterion of city size, and you know that one city has a university while the other does not. Having a university would be "the best" cue on this occasion—you would thus use Take the Best and infer that the city with the university is the larger of the two (Gigerenzer & Goldstein, 1999).

To explain how this heuristic works in terms of the account outlined in the present paper, we begin by observing that a set of representations will be activated by the cognitive task. Network nodes that might be active in this task can include clusters of representations for *city*, *infrastructure*, *roads*, *traffic*, *buildings*, *university*, among others; the sorts of nodes and representations that get activated in a given task is open-ended depending on the representational and conceptual wherewithal of the individual. Now, it is important to understand that certain beliefs are implied about the cue upon which the choice in question is made. At the very least, we can infer that the said cue was believed to be "the best" on which to make the choice in question, and believing that cue to be "the best" implies certain things about the representational content of the cue (as possessed by the chooser), as well as certain things about how the cue fits within the chooser's representational system. For instance, if one believes that having a university is a good predictor of relative city size, one's university representations must be rich enough to support such a belief, or at least to support the belief that having a university is a better predictor of relative city size than some other cue (e.g., having a major highway). All of this knowledge is conceived as constituting a particular node as a cluster of university representations (as indicated above). Implicit information built into one's university node might include, in addition to various representations and conceptualizations that are required to understand what a university is, information about a university's function in society; how it houses a number of faculty, staff, and students; perhaps the relative social classes of members of a community typical of such faculty, staff, and students; maybe the ways in which having a university relates to the economy of a city; and probably much more. Indeed, it is such information that generally *makes one believe* that having a university is a good predictor of relative city size. In other words, this implicit knowledge will help guide one's reasoning to infer that the city with the university is a large city, or at least larger than the other. (We might recall here the lesson from Dennett discussed above.)

The diagram in Figure 1 is a rough characterization of a possible network that would be active in the Take the Best task we have been discussing. The structure exhibited suggests to the chooser that having a university, being a capital city, having an airport, and having a major highway run through the city are implicitly believed to indicate relative city size. Take the Best can thus be viewed as operating with the suggested structure to deliver an inference that the city in question is the larger of the two under consideration. In general, heuristics are conceived to exploit active structures made available through spreading activation (i.e., heuristics

[7] This sort of account of cognitive architecture is notably different from what is assumed by classical architectures such as ACT-R and Soar. The latter sorts of architectures assume that cognition consists of operations over propositions, and further that rules of inference (e.g., productions) are explicitly represented propositionally. On the present account, there is no assumption that cognition is propositional, even though the account is compatible with this assumption. Further, there is no assumption that heuristics are rule-like, even though (again) the account is compatible with rule-governed operations. Thus, the present account has certain advantages over classical architectures insofar as it is more general, and can accommodate non-classical architectures as well as multimodal representations. In addition, unlike what is needed for ACT-R or Soar, the present account does not assume that problems and problem-spaces are or must be well-defined.

[8] I say "primarily" since some representational content may figure into some heuristic processes. My point, nonetheless, is that representational content is not what heuristics generally operate with.

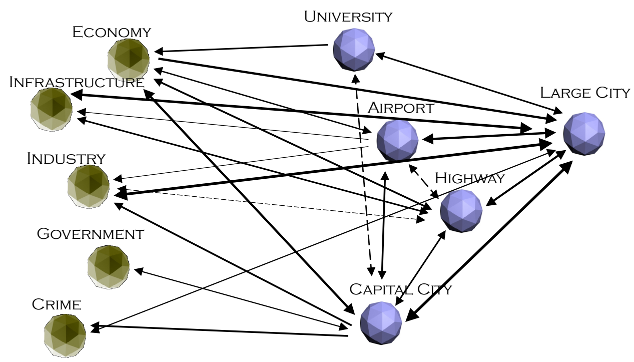

Figure 1: Proposed network structure arising from Goldstein and Gigerenzer's (1999) task of predicting which of two cities is larger. Each node is understood to be structured clusters of representations. Arrows represent directed connections and line thickness represents weights of connections.

don't operate over non-active connections). Note in the Figure the multiple, heavily-weighted connections pointing to the *large city* node—the Take the Best heuristic would exploit this information made available through the spreading activation process.

On this view, we might thus understand one's belief that having a university is a good predictor of relative city size to be implicit in metainformational structures that exist among representations comprising the network. In this way, beliefs about discriminating cues need not be explicitly represented. This is contrary to Gigerenzer's view that the belief that a cue discriminates is explicitly represented and searched. Also contrary to Gigerenzer's view is that heuristics are here understood to operate by generally ignoring representational content. Instead, heuristics utilize cognitive structures (viz. connections that exist among nodes and representations) that implicitly embody *higher-order* information (or knowledge) about active representations and their content[9] (cf. Clark, 2002).

In addition, on the present account, it will be very common for one to activate lots of conceptual content. This is because the cognitive tasks we typically face have a *high cognitive load*, and as such, navigating these tasks invoke large amounts of representations and informationally rich connections among them. This is antithetical to Gigerenzer's assumption that heuristic decision-making is often based on one reason (Gigerenzer & Goldstein, 1999; Gigerenzer, 2000). Indeed, it is unlikely that one will only invoke a handful of cues to decide, especially if one possesses significant knowledge about the task (cf. Hilbig, 2010). For instance, as illustrated in Figure 1, one might invoke numerous representations concerning not only universities, but also concerning whether a major highway runs through one or both cities in question; whether one or both cities have a professional sports team; whether one has heard one city in the news more often than the other; whether there is a famous museum in one or both cities; whether one knows of a major river in either of the two cities; whether one has any friends or relatives who have visited either city; and much more. And all the activated nodes and representations would bear various connections to one another, thus providing structures to be exploited by heuristics and other processes.

Importantly, despite the exploitation of generous amounts of information, none of this implies that heuristics are not frugal. Frugality is a key feature of heuristics, and moreover, heuristics are supposed to satisfice (Simon, 1957). And this entails that most information will not get considered in their processing (Gigerenzer & Gaissmaier, 2011); that is, satisficing and frugality go hand-in-hand. According to the account advanced here, one's representational wherewithal does most of the heavy lifting in cognition, not the heuristics themselves. That is, the cognitive architecture bears the informational burden. Heuristics are cognitively cheap solutions to problems for which an expensive cost has been requisitely paid.[10]

Conclusion

The theory outlined in this paper was advanced to explain how heuristics work in cognition. It is my hope that it will also serve to motivate researchers to take seriously the need to better understand what cognitive heuristics are and how they operate. Since heuristics are supposed to play a large role in our cognitive lives, our picture of human cognition will remain incomplete without such an understanding.

I acknowledge that more work needs to be done to corroborate the foregoing understanding of how heuristics work. I (all too briefly) illustrated how the Take the Best heuristic can be understood in terms of the theory outlined in this paper. Other heuristics in the literature can be similarly understood, but it is not possible to show this here. Further research must be done to explore more fully the nature of our representational structures and systems, the nature of the relations between representations and representational content, as well as

[9] I propose as a possibility that the information embodied by the wighted, directed connections among representations may be stored in convergence zones (A. R. Damasio, 1989; H. Damasio, Tranel, Grabowski, Adolphs, & Damasio, 2004), which is accessible to heuristics.

[10] van Rooij, Wright, and Wareham (2012) have recently (and compellingly) argued that the common conception that heuristics are solutions to computationally intractable (NP-hard) problems is mistaken. Roughly, they argue that, since intractable problems are by definition intractable, any solution (heuristic or otherwise) would be a solution to some other *tractable* problem (cf. Besold, 2013). I lack the space to fully discuss this matter. However, the theory offered here avoids the problem discussed by van Rooij et al. since, as noted, there is no assumption that there are prepackaged, well-defined, computationally intractable problems that are solved by heuristics (cf. footnote 7 above). Instead, the problems that heuristics solve are ones that arise from the (dynamic) patterns of spreading activation in the network. In this way, problem-spaces are sort of defined on the fly, and heuristics set about navigating *those* spaces. Importantly, the relevant problems can very well be understood as fully tractable. Of course, we are then faced with the (separate) task of producing a clear computational-level characterization of the problems that are solved.

to provide more detail on how heuristics are sensitive to and interact with the structures that comprise a network.

Acknowledgments

I would like to thank Richard Samuels, Chris Viger, and the anonymous referees of EAPCogSci2015 for helpful comments and suggestions.

References

Ayal, S., & Hochman, G. (2009). Ignorance or integration: The cognitive processes underlying choice behavior. *Journal of Behavioral Decision Making*, *22*, 455-474.

Baronchelli, A., Ferrer-i-Cancho, R., Pastor-Satorras, R., Chater, N., & Christiansen, M. H. (2013). Networks in cognitive science. *Trends in Cognitive Science*, *17*, 348-358.

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

Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. D. Vega, A. M. Glenberg, & A. C. Graesser (Eds.), *Symbols, embodiment, and meaning* (p. 245-283). Oxford: Oxford University Press.

Besold, T. R. (2013). Formal limits to heuristics in cognitive systems. *Proceedings of the Second Annual Conference on Advances in Cognitive Systems (Poster Collection)*, 19-28.

Carruthers, P. (2006). *The architecture of the mind: Massive modularity and the flexibility of thought*. Oxford: Oxford University Press.

Chow, S. J. (in press). Many meanings of 'heuristic'. *The British Journal for the Philosophy of Science*.

Clark, A. (2002). Local associations and global reason: Fodor's frame problem and second-order search. *Cognitive Science Quarterly*, *2*, 115-140.

Collins, A. M., & Loftus, E. F. (1975). A spreading activation theory of semantic processing. *Psychological Review*, *82*, 407-428.

Damasio, A. R. (1989). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition*, *33*, 25-62.

Damasio, H., Tranel, D., Grabowski, T. J., Adolphs, R., & Damasio, A. R. (2004). Neural systems behind word and concept retrieval. *Cognition*, *92*, 179-229.

Dennett, D. C. (1984). Cognitive wheels: The frame problem of AI. In C. Hookway (Ed.), *Minds, machines, and evolution* (p. 129-152). Cambridge: Cambridge University Press.

Dougherty, M. R., Franco-Watkins, A. M., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review*, *115*, 199-213.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge: The MIT Press.

Fodor, J. A. (2000). *The mind doesn't work that way: The scope and limits of computational psychology*. Cambridge: The MIT Press.

Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the real world*. New York: Oxford University Press.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451-482.

Gigerenzer, G., & Goldstein, D. G. (1999). Betting on one good reason: The Take the Best heuristic. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (p. 75-95). New York: Oxford University Press.

Hilbig, B. E. (2010). Reconsidering "evidence" for fast and frugal heuristics. *Psychonomic Bulletin and Review*, *17*, 923-930.

Hilbig, B. E., & Pohl, R. F. (2008). Recognizing users of the recognition heuristic. *Experimental Psychology*, *55*, 394-401.

Lenat, D. B. (1982). The nature of heuristics. *Artificial Intelligence*, *19*, 189-249.

Lucas, M. (2000). Semantic priming without association: A meta-analytic review. *Psychonomic Bulletin & Review*, *7*, 618-630.

Meteyard, L., Rodriguez Cuadrado, S., Bahrami, B., & Vigliocco, G. (2012). Coming of age: A review of embodiment and the neuroscience of semantics. *Cortex*, *48*, 788-804.

Oppenheimer, D. M. (2003). Not so fast! (and not so frugal!): Rethinking the recognition heuristic. *Cognition*, *90*, B1-B9.

Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? the representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, *8*, 976-987.

Quillian, M. R. (1962). A revised design for an understanding machine. *Mechanical Translation*, *7*, 17-29.

Samuels, R. (2005). The complexity of cognition: Tractability arguments for massive modularity. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Structure and contents* (p. 107-121). Oxford: Oxford University Press.

Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, *134*, 207-222.

Simon, H. A. (1957). *Models of man, social and rational*. New York: John Wiley.

Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (p. 39-67). Cambridge: Cambridge University Press.

Thompson-Schill, S. L. (2003). Neuroimaging studies of semantic memory: Inferring "How" from "Where". *Neuropsychologia*, *41*, 280-292.

van Rooij, I., Wright, C. D., & Wareham, T. (2012). Intractability and the use of heuristics in psychological explanations. *Synthese*, *187*, 471-487.