

Determining Relevance: Close Enough is Good Enough

Sheldon J. Chow (schow55@uwo.ca)

Department of Philosophy, Stevenson Hall, The University of Western Ontario
London, Ontario, N6A 5B8 Canada

Abstract

Humans exhibit characteristic success in considering what is relevant in their cognitive tasks. Yet understanding how relevance is determined in cognition remains a problem. This paper seeks to make headway on this problem. The relevance problem is first introduced. Sperber and Wilson's influential theory of relevance is then discussed, but dismissed as inadequate. Some conditions are identified that an adequate definition of relevance might reasonably be expected to satisfy. A novel way to conceive of relevance is suggested which proves to be useful in understanding human cognitive performance.

Keywords: Relevance; concepts; file model; cognitive architecture.

The Relevance Problem

A longstanding problem in philosophy and cognitive science is understanding how we determine what is relevant to our cognitive tasks. The cognitive systems paradigmatically responsible for general reasoning and decision-making—so-called central systems—admittedly allow for free exchange of information. A dream of a snake biting its own tail, for example, bore in interesting and important ways on Kekulé's theorizing of the benzene molecule. A consequence of such a free exchange of information, however, is that, provided an appropriate set of background beliefs, any representation held by an agent can *in principle* be relevant to a given cognitive task in central cognition. Who won tonight's football game is *prima facie* irrelevant to whether there is beer in your friend's fridge. But if you believe that your friend's favourite football team played tonight, and that your friend usually overindulges in beer consumption whenever his favourite team wins, then your belief about who won tonight's game actually *is* relevant to your belief about beer in your friend's fridge. Since relevance can be determined only with respect to one's background beliefs, there is no way to circumscribe *a priori* the subset of representations that are relevant in a given occasion of reasoning or inference. Let us express this problem thus:

The relevance problem How a cognitive system considers only what is (mostly) relevant, or equivalently, how a cognitive system knows what is relevant.

Despite the fact that the relevance problem introduces a computational problem of sifting through heaps of information to decide what bears on the task at hand, humans seem to determine what is relevant in their cognitive tasks quickly and rather easily. This is not to say that we always know and consider what is relevant. For we often fail to do so, especially

when cognitive demands are high and/or when cognitive resources are low. Nevertheless, humans characteristically exhibit reasonable levels of success at identifying representations which are relevant to the task at hand. Such reasonable levels of success cannot be ascribed to chance or luck. Therefore, we are left to explain how humans (seem to) solve the relevance problem, short of considering the totality of one's beliefs (Fodor, 1987, 2000).

However, the inquiry into how relevance is determined is constrained by how relevance is defined. Indeed, whether and the extent to which relevant representations are picked out and brought to bear on a cognitive task will depend on what property we are concerned with. Yet, defining relevance is not an easy task. In this paper I provide a cursory overview of what is arguably the most influential account of relevance to date, namely Sperber and Wilson's *Relevance Theory*. I will then be able to use their Relevance Theory as a basis from which to propose an understanding of relevance that is supported by a view of concepts and cognition which draws on current theories in cognitive science, as well as Fred Dretske's information-theoretic epistemology.¹ This will allow me to show how relevance, or something like it, can be understood as naturally arising from human cognitive architecture, thus enabling the characteristic performance we observe in human reasoning.

Sperber and Wilson's Relevance Theory

Sperber and Wilson (1986/1995) developed their Relevance Theory in the context of communication and pragmatics. Humans tend to have an easy time communicating with each other, despite the fact that the meanings of utterances are enormously underdetermined. A simple example: Alice says to Bob, "Isn't that cute?" while nodding toward a chipmunk scurrying up a tree; Bob knows that by "that" Alice is referring to the chipmunk, and not to the birds in the other branches, the tree itself, or whatever else was within his perceptual field at the time of her utterance. According to Sperber and Wilson (SW henceforth), Bob understands that Alice was referring to the chipmunk because the stimulus of the

¹It should be kept in mind throughout the discussion that I am not intending to develop a full theory of relevance. I therefore do not offer any claim of completeness. Such a task deserves a more extensive treatment than what I can provide in this short paper. My contribution might be understood as a useful preliminary discussion of how the problem of relevance might be stated, and an identification of some conditions that an adequate definition of relevance might reasonably be expected to satisfy. Nevertheless, I shall suggest a potential way to conceive of relevance which proves to be useful in understanding human cognitive performance.

chipmunk running up the tree was *relevant* (or at least more so than any other present stimulus).

Although SW's Relevance Theory is mainly concerned with verbal/ostensive communication and comprehension in particular, they claim that their theory can be extended to the inference processes of ordinary thinking (1986/1995, p. 67; p. 75). This is because SW ground their account of relevance in a fundamental and general view of human cognition. Specifically, SW claims that cognition always tends toward efficiency (i.e., maximizing gains and minimizing costs), and furthermore that human cognition succeeds in increasing its efficiency by having a tendency toward *maximizing relevance*. According to their Relevance Theory, the relevance of an input to a process is a function of *cognitive effect* on the one hand, and *processing effort* on the other.²

It should be obvious that relevance, assessed in terms of cognitive effect and processing effort, comes in degrees. In addition, some input may yield greater cognitive effects on some occasions and less effects on others, or, depending on circumstances related to fatigue or stress, the same input may be more or less easy to process at different times. Thus, relevance is a relative property—relative to an individual and to a time. SW therefore provide the following two definitive conditions for relevance:

1. *Ceteris paribus*, the greater the cognitive effects achieved by an individual by processing some input, the more relevant that input is to that individual at that time.
2. *Ceteris paribus*, the greater the effort expended by an individual by processing some input, the less relevant that input is to that individual at that time.

When SW claim that human cognition has a tendency toward maximizing relevance—which they assert as The Cognitive Principle of Relevance³—they mean that human cognition is geared toward allocating cognitive resources to processing available inputs (from the environment or from memory) so as to maximize expected cognitive effects for the least expected effort.

Sperber and Wilson (1995) make the point that an individual is not interested in cognitive effects *per se*, but only insofar as such cognitive effects contribute to achieving certain cognitive goals, or otherwise fulfilling its functions. For

²SW develop their account in terms of constructing an appropriate context within which to process inputs. A context is simply just a set of assumptions within which input can be processed. For simplicity, I will pass over this aspect of their account; no harm is done to the central points of the present paper.

³SW's Relevance Theory makes claims about human cognition in general, but an important consequence of their Cognitive Principle, and one that is the basis for their work in pragmatics, is the Communicative Principle of Relevance: *Every ostensive stimulus conveys a presumption of its own optimal relevance*. This Communicative Principle of Relevance is really the centerpiece of SW's Relevance Theory. However, since communication and pragmatics are not the focus of this paper, the Communicative Principle will not be discussed any further.

indeed, there may be cognitive effects that are not worth having, or that contribute negatively to the individual's cognitive functions or goals. Thus, SW define a *positive cognitive effect* as a cognitive effect that contributes to the cognitive goals or functions of the individual in a positive way. This is typically achieved by alterations in the individual's beliefs. In SW's words, positive cognitive effects produce an “epistemic improvement” (Sperber & Wilson, 1995, p. 266); they make a “worthwhile difference” to the cognitive functioning of the individual (Wilson & Sperber, 2006, p. 608).

Thus far I have given a brief outline of SW's Relevance Theory. Nevertheless, there are a number of ways in which Relevance Theory is inadequate. For brevity, I mention here only two inadequacies, but these are most damaging. First, it is grossly unclear how to quantify cognitive effect and effort so as to make sense of the ratio between the two in SW's definition of relevance. Second, SW's conception of cognition is strictly in terms of deduction—processes performed by a “deductive device” that is supposed to “model the system used by human beings in spontaneous inference” (Sperber & Wilson, 1986/1995, p. 94). Of course, there is wide consensus now that spontaneous inference is not strictly deductive, but is to a large extent non-demonstrative. SW's account of cognition completely excludes such non-demonstrative cognition, as well as other aspects of cognition such as perception.⁴

In light of these shortcomings, the remainder of this paper is devoted to offering the beginnings of a more adequate account of relevance in cognition. I will begin by outlining a psychologically plausible theory of cognition. We shall see that this will allow a more naturalistic understanding of SW's notion of “positive cognitive effects”, but more importantly it has the potential to deliver a desired account of relevance.

Concepts and Cognition

A number of cognitive scientists and philosophers have converged on an idea about the nature and function of mental representations (e.g., Evans, 1982; Fodor, 2008; Lawlor, 2001). The idea is a *file model* of cognition (Kahneman, Treisman, & Gibbs, 1992; Pylyshyn, 2003; Treisman, 1982; Treisman & Schmidt, 1982). There is no standard doctrine accepted by everyone who endorses the file model, but there is some common ground that can be identified.

According to the file model of cognition, one has a mental “file” for each object that one has beliefs about. Each mental file has a “label” that both picks out the file, and refers to the

⁴By the second edition of *Relevance* (1995), Sperber and Wilson had disavowed such a central systems architecture in favour of the massive modularity hypothesis. Though I note this only in passing, I believe that certain complications arise for their Relevance Theory within a massively modular architecture. This is a matter to be discussed on some other occasion, however. Sperber (2005) has recently suggested that cognitive efficiency, in terms of maximizing relevance, is achieved biologically—specifically, by optimally allocating energy in the brain. His proposal is to think of maximizing cognitive effects and minimizing effort in terms of noncognitive physiological processes. However, this entails an account of relevance entirely different from SW's Relevance Theory, and Sperber has not developed such an account in any detail.

object that is associated with the file. Further, each file holds a number of “notes” or “memos” that contain various kinds of information or beliefs, or representations more generally, about the given object.

It is important for the present purposes to understand that the file model can be conceived more particularly to be a theory of the structure of *concepts*. On this reading, a concept names a file that contains representations about the things in the concept’s extension. When thinking about cats, for instance, one calls upon one’s CAT file, which may contain such representations as *is a living thing*, *is a furry creature*, *is a quadruped*, *is grandma’s favourite animal*, etc.⁵ Notice now that what is contained in the file—what notes there are in the file—name other concepts: LIVING THING, FURRY CREATURE, QUADRUPED, GRANDMA’S FAVOURITE ANIMAL, etc. This is not classical decomposition. Rather, concept files contain notes that convey information about the concept in question. Hence, unpacking the notes of one’s CAT concept file is unpacking one’s beliefs and other stored knowledge about cats, the representations of which invoke further concepts. Moreover, whatever is inferred from the unpacked representations may excite further concepts, thereby bringing to mind files and representational content of their own. Thus, such a conceptual system forms a vast network among the beliefs and other representations which are contained in concept-files.

A very important feature of this account is that concepts do not exist independently of one another. Rather, this is a view of concepts wherein vast connections exist between them in virtue of their content. Every representation is a part of an interconnected network, where activation can spread through the system depending on the strengths and the organization of the connections between representations in the network. In principle, the structure allows for any content to activate any other content, which is a hallmark of central cognition (cf. Viger, 2006b).

We might therefore see that the file model of cognition allows for the following sort of picture to underwrite thought and inference. When one is presented with a given cognitive task, the nature and content of the task initially activates and primes⁶ a focused set of concepts and contents. Suppose, for instance, one is presented with the task of estimating whether it is likely or unlikely that there will be a plane crash within the month. This task would activate conceptual content concerning *planes*, *crashes*, *likelihoods*, *estimation*, *months*, *timeframes*, and more. Certainly, various parameters will constrain what conceptual content gets activated. Some parameters include those that are given by the language used (e.g., using “plane” rather than “jet”), as well as those that are suggested by the nature of the task (e.g., the task elicits a coarse-grained subjective likelihood assignment to a future event as opposed to, say, a fine-grained numerical subjec-

⁵The convention I adopt here is to represent concepts in small caps, contents in italics, and uses in quotes.

⁶For simplicity, I will often use “active”/“activated” to refer to both active/activated and primed content.

tive probability). Other parameters will have to do with factors affecting long-term memory recall, such as which concepts an individual possesses, the relative strengths at which conceptual content is stored in long-term memory, the ease with which conceptual content is activated (perhaps based on past activations), and the existence and relative strength of connections between concepts and conceptual content established by past inferences. In addition to these constraints, limits on time and cognitive resources will restrict what and how many conceptualizations occur. Yet, even with these constraints, a considerable amount of conceptual content may still get activated.

This has been a brief overview of the file model of cognition. Before I continue, I should note that the file model is best viewed as a useful analogy for thinking about concepts and the relations among them and their contents. The model is advanced here merely as a gesture toward a possible cognitive conceptual architecture. If the file model is to become a viable basis for a definition of relevance, it will need much further development and theoretical refinement—a task that is beyond the scope of this paper. Nevertheless, the main proposal here does not depend on the truth of the file model, specifically. It requires only that there exist highly organized systems of knowledge or representations in cognition which bear certain connections to one another. The file model gives us a tentative idea of what these systems might be. Yet, that a number of cognitive scientists have converged on it to model various aspects of cognition is suggestive of its plausibility.⁷

Reinterpreting Relevance

Let us now turn back to relevance. There is something intuitively right about the idea that relevance has something to do with what SW call “positive cognitive effects”. For indeed, it seems as if processing irrelevant information would not generally yield anything positive for the cognitive system. Yet, we would need a more precise understanding of positive cognitive effect which avoids the pitfalls of SW’s account, but which is somehow tied to relevance. The account of concepts and cognition in the previous section delivers this.

To see this, let us begin by drawing some lessons from Fred Dretske (1981). In developing his information-theoretic epistemology, Dretske commented that there is “a *relativization* of the information contained in a signal because *how much*

⁷Moreover, the file model lends itself to being interpreted within certain current neurological theories of conceptual cognition. For example, Barsalou’s (1999) *perceptual symbols systems* view of cognition, Patterson, Nestor, and Rogers’ (2007) theory of a *semantic hub*, and Damasio and Damasio’s theory of *convergence zones* (A. R. Damasio, 1989; H. Damasio, Tranel, Grabowski, Adolphs, & Damasio, 2004; Tranel, Damasio, & Damasio, 1997), each posits the existence of neurological sites that store and enact a “code” which specifies specific neural patterns to be (re)activated during mental representation. My preferred (albeit tentative) interpretation is to understand concept-files to be roughly analogous to the posited sites where the codes are stored; opening files can then be understood as instantiating the appropriate codes to (re)activate neural patterns for representation; and file labels can be understood as lexical terms that name concepts in natural language. (This last point was suggested to me by Chris Viger.)

information a signal contains, and hence *what* information it carries, depends on what the potential receiver already knows about the various possibilities that exist at the source" (p. 79). Although we might not know (or even cannot know) absolute measures associated with the amount of information generated by an event or carried by a signal, Dretske believes that *comparisons* can be made, "in particular comparisons between the amount of information generated by the occurrence of an event and the amount of information a signal carries about that event" (p. 54).

According to Dretske, our concepts are cognitive structures that enable us to extract and exploit certain information in the environment. Learning or enriching a concept provides us with the ability to encode (or digitalize) certain aspects of our (analog) sensory experience in such a way that we are able to cognitively respond in certain ways which we would not have been able to otherwise. Having the concept DAFFODIL, for example, enables one to see a daffodil as a daffodil (as opposed to a flower, or some green and white object), and thus allows one to have daffodil-thoughts and to cognize daffodil-stimuli in certain ways. Importantly, then, what information one can encode from stimuli crucially depends on what one already knows about the objects of the stimuli, or in the terms of the present account, on what and the way in which conceptual contents are coded in one's conceptual architecture.⁸

A natural account of relevance follows from this picture in terms of the amount of information received from a source (such as a stimulus). More specifically, the greater the amount of information received, the greater the relevance of that information. For example, suppose that Alice and Bob are on a nature walk. Alice is a botanist, whereas Bob never cared for plant science. As Alice and Bob gaze upon the flora of the forest floor, they both cognitively extract a vast amount of information from their respective perceptual scenes. However, Alice's conceptual knowledge is so rich that she is able to extract more specialized information than Bob does or even can, having to do with the various kinds of plants that they come across. In this way, the perceptual scene carries more information for Alice than for Bob. Of course, they both process the *same* information, but Alice can *cognitively extract* more information. Whereas Bob simply sees a plant, Alice sees Blindwood ivy; whereas Bob simply sees flowers, Alice sees daffodils. Importantly, certain information in Alice's and Bob's perceptual scene is very *relevant* to Alice but not so relevant to Bob. And the information from the perceptual

⁸According to information theory, and thus according to Dretske, a signal carries 0 information if one already knows the message the signal carries. I disagree, since for a signal to carry 0 information, one would have to be unable to extract and conceptualize any information from the source in question (cf. Gabbay & Woods, 2003). And such a thing does not seem to be likely, or even possible, in human cognitive practice. Processing information one already knows may be *redundant*, but it may also be useful to strengthen or reaffirm one's beliefs. A corollary of this view is that a signal will always carry information for human cognitive systems; the *amount* of information in a signal will depend on what the receiver knows, or more specifically, on activated conceptual representations. More on this presently.

scene that Alice finds relevant is just that information she is able to extract via her conceptual knowledge. On the other hand, such information is not as relevant to Bob because he cannot represent the information in the same ways, since he lacks the conceptual wherewithal to do so.⁹

Therefore, I suggest that the relevance of a stimulus to a given cognitive system (or agent) depends on the amount of information received from that stimulus. Since the amount of information received depends on one's conceptual wherewithal to attend to and code specific information in certain ways, whether and the extent to which something is relevant is dependent on the informational content of one's concepts. But this is not the entire story, since relevance will also depend on the context and cognitive task. Suppose, for instance, that both Alice and Bob are botanists, but Alice is interested in finding a rare flower while Bob is interested in seeing a specific species of ivy. Both Alice and Bob could code the same information in the same ways, but because of their different goals and cognitive tasks, Alice will find flower information more relevant than Bob, and Bob will find ivy information more relevant than Alice.

This can be easily accounted for, however, once it is understood that, in setting up one's goals and preparing for one's cognitive task, one requisitely activates a number of concepts with specific conceptualizations and representations, as described above. This will in a sense serve as a filtering mechanism for focusing attention. Alice will thus be (cognitively/conceptually) geared to attend to specific information related to a specific flower whereas Bob will be (cognitively/conceptually) geared to attend to specific information related to ivies, as each will have prepared a set of concepts and representations upon embarking on their respective cognitive tasks. The set of concepts and representations that gets activated when one prepares for a given cognitive task will tend to be relevant, although this may not always be the case. What gets activated will depend on previous experience, past activations, and the extant relations and connections among the activations. These relations will constrain and guide inference.

Thus, on this view, the degree to which information is relevant is a matter of the *informativeness* of information. In other words, relevance is conceived to be a measure of how much information gets encoded given one's cognitive wherewithal and preparation (i.e., the activations and connective relations among one's conceptual representations). The same information can be more or less informative, and hence relevant.

⁹We might consider here a situation in which Bob is trying to learn some botany, in which case the information received from the stimuli given by the forest floor would seem to be more relevant to Bob than Alice, since Alice already knows what Blindwood ivy is and looks like, etc. However, this situation is a shift in context from the scenario described in the main text. Specifically, Bob's cognitive task has changed, and therefore so has his cognitive preparation. The difference in cognitive preparation, along with, say, Alice's instruction, will entail a difference in the kinds and amounts of information extracted from the environment, and hence a difference in relevance. See below.

vant, depending on the agent and cognitive task. Understood this way, relevance is not just a property of information *per se*—not just a matter of what information gets processed—but a matter of *how information gets processed*.

If this is right, then positive cognitive effect can be understood in terms of the informativeness of information that is processed—processing information yields positive cognitive effects insofar as the results are informative. But this is just to refer to the degree of relevance of the information in question. This means that positive cognitive effect is yielded by processing relevant information. We therefore have our intuitive connection between positive cognitive effect and relevance borne out by the present account: Processed information has positive cognitive effect *because* it is informative; and the degree to which it is informative is the degree to which the information is relevant. Cognitive effects and effort therefore do not define relevance. Rather, concepts and the extant relations between their contents will facilitate cognitive effects and effort in processing information. Hence, it is structured concepts that deliver relevance, which in turn produces cognitive effects with little effort.

This account of relevance is definitely an improvement over SW's account. On the one hand, there is no problem of how to quantify cognitive effect and effort to make sense of the ratio between the two; instead, relevance is measured directly in terms of degree of informativeness. On the other hand, there is no conceiving cognition strictly in terms of deduction; the adopted view of cognition is amenable to deductive and non-demonstrative inference alike.

How Do We Determine What is Relevant?

We are now in a position to see how the foregoing account of concepts and cognition can help to explain our characteristic levels of success in our reasoning. Given the present account, we have more constraints on our reasoning than we may know. Specifically, a certain kind of relevance is determined by the extant relations among the contents of our concepts. The kind of relevance I have in mind is *de facto* relevance (cf. Gabbay & Woods, 2003), in which information appears to be relevant due to the architectural characteristics of cognition. More specifically, and to continue the line of reasoning in the previous section, something is (more or less) *de facto* relevant if it appears to be (more or less) informative when processed against a given set of activated concepts. In this way, *de facto* relevance cannot be determined *a priori*, as should be expected. Instead, it simply arises out of the nature and structures of our concepts.

The *de facto* relevance established by the extant relations within and between activated conceptual contents appears to be enough for humans to get by on. The kind of relevance that matters to the relevance problem spelled out in the first section of this paper, on the other hand, is a kind of *objective* relevance which exists independently of cognizers. Many examples of objective relevance come from science. For example, the motions of terrestrial objects are relevant to the

motions of the planets, and this is an objective fact, but we did not know this until Newton came along.¹⁰

There will certainly be times when we fail to process objectively relevant information, or when we process information that is not very objectively relevant at all. In some cases we may end up processing some objectively *irrelevant* information. Moreover, inevitably there will be cases in which we fail to process *de facto* relevant information, due to cognitive limitations, fatigue, stress, or some other extraneous factor. Under satisfactory conditions, however, our activated concepts, with their contents and extant relations, provide a network that informs cognition of what is *de facto* relevant, and constrains and guides its processing accordingly. The situation may not be ideal, but it is good enough for us to get by on—indeed, such is to be expected from satisficing organisms. On the other hand, when we enter into certain high-stake arenas, such as science or philosophy, we alter our standards, and *de facto* relevance is no longer good enough. In such circumstances, objective relevance is sought, and this is likely why progress and getting things right are much more difficult to achieve in these endeavours.

At the same time, however, we might understand the foregoing account of concepts and cognition as contributing to how humans manage to solve, not the relevance problem stated above, but a more fundamental problem. To see what I mean, let us consider Daniel Dennett's (1984) example of fixing a midnight snack. He noticed that such a mundane task requires copious amounts of knowledge: "We know trillions of things; we know that mayonnaise doesn't dissolve knives on contact, that a slice of bread is smaller than Mount Everest, that opening the refrigerator doesn't cause a nuclear holocaust in the kitchen" (p. 136). Dennett also noticed that we must possess a system of representing this knowledge in such a way that it is accessible on demand. This would require a system that is organized in such a way that achieves the efficient representation and access we observe in humans. In his words: "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" (pp. 140-1). From these considerations, we might say that humans solve a *representational relevance problem*:

The representational relevance problem How a cognitive system embodies the informational organization, and enables access to the relevant information, that seems to be required for human-like cognitive performance.

I believe that the present account of concepts and cogni-

¹⁰Notice that objective relevance also cannot be determined *a priori*; but rather than against a set of beliefs, relevance is determined against a background of facts and phenomena. The motions of terrestrial objects, for instance, are not objectively relevant *simpliciter*; it is objectively relevant with respect certain phenomena, such as planetary motion.

tion is precisely what enables humans to solve the representational relevance problem. For as I have illustrated, concepts are organized in such a way that the extant relations between their contents facilitates access to *de facto* relevant information. Such information may not be objectively relevant, but it will almost certainly be the kind of information that is needed to guide successful action, and this is all that is needed for human-like performance.

It is interesting to note, however, that it seems that much of the information that is *de facto* relevant turns out to be objectively relevant a lot of the time. This is evident from how humans get on in the world, and the success rate of many human inferential endeavours. I can only speculate why this is so: it is likely an outcome of some evolutionary process. This would explain our reasonable levels of success in bringing to bear objectively relevant information on our cognitive tasks. I admit that this is not a deep explanation. However, if we conceive our conceptual system to have evolved to track *things in the world*, then it should not be much of a mystery why our conceptual wherewithal reflects the organized structure of information in the world, including objective relevance relations. In this way, then, we can conceive *de facto* relevance to be built up by systems that track objective relevance. And, just like any cognitive system that tracks stuff in the world, sometimes things work out and sometimes things go awry; and sometimes cognitive systems track truths but not all the time (such as the perceptual systems; cf. Viger, 2006a). It seems, however, that *in the main* cognition tracks truths in the world, and is quite good at it. Thus, the *de facto* relevance embodied by the relations within and between concepts by and large reflects objective relevance in the world. This is what makes *de facto* relevance close enough to objective relevance; and close enough is good enough.

Acknowledgments

Thanks to Chris Viger for all his comments and suggestions. Thanks also to the anonymous referees whose comments have improved this paper.

References

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

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.

Dretske, F. I. (1981). *Knowledge and the flow of information*. Cambridge, MA.: The MIT Press.

Evans, G. (1982). *The varieties of reference*. Oxford: Oxford University Press.

Fodor, J. A. (1987). Modules, frames, fridges, sleeping dogs and the music of the spheres. In J. L. Garfield (Ed.), *Modularity in knowledge, representation and natural-language understanding* (p. 25-36). 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.

Fodor, J. A. (2008). *LOT 2: The language of thought revisited*. Oxford: Clarendon Press.

Gabbay, D., & Woods, J. (2003). *A practical logic of cognitive systems, Volume 1. Agenda relevance: A study in formal pragmatics*. Amsterdam: Elsevier.

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24, 175-219.

Lawlor, K. (2001). *New thoughts about old things: Cognitive policies as the ground of singular concepts*. New York: Garland Publishing.

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.

Pylyshyn, Z. W. (2003). *Seeing and visualizing: It's not what you think*. Cambridge: The MIT Press.

Sperber, D. (2005). Modularity and relevance: How can a massively modular mind be flexible and context-sensitive? In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Structure and contents* (p. 53-68). Oxford: Oxford University Press.

Sperber, D., & Wilson, D. (1986/1995). *Relevance: Communication and cognition*. Oxford: Blackwell (2nd edition, 1995).

Sperber, D., & Wilson, D. (1995). Postface to the second edition of *Relevance: Communication and Cognition*.

Tranel, D., Damasio, H., & Damasio, A. R. (1997). A neural basis for the retrieval of conceptual knowledge. *Neuropsychologia*, 35, 1319-1327.

Treisman, A. (1982). Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 194-214.

Treisman, A., & Schmidt, H. (1982). Illusory conjunction in the perception of objects. *Cognitive Psychology*, 14, 107-142.

Viger, C. (2006a). Is the aim of perception to provide accurate representations? a case for the “no” side. In R. J. Stainton (Ed.), *Contemporary debates in cognitive science* (p. 275-288). Malden, MA: Blackwell Publishing.

Viger, C. (2006b). Symbols: What cognition requires of representationalism. *Protosociology: The International Journal of Interdisciplinary Research*, 22, 40-59.

Wilson, D., & Sperber, D. (2006). Relevance theory. In L. R. Horn & G. Ward (Eds.), *The handbook of pragmatics* (p. 607-632). Malden, MA: Blackwell Publishing.