

What Is Conceptual Glue?

ERIC MARGOLIS

Department of Philosophy & Rice Cognitive Sciences Program — MS14, Rice University, P.O. Box 1892, Houston, TX 77251-1892, USA (E-mail: margolis@ruf.rice.edu)

Abstract. Conceptual structures are commonly likened to scientific theories, yet the content and motivation of the theory analogy are rarely discussed. Gregory Murphy and Douglas Medin’s “The Role of Theories in Conceptual Coherence” is a notable exception and has become an authoritative exposition of the utility of the theory analogy. For Murphy and Medin, the theory analogy solves what they call the problem of conceptual coherence or the problem of conceptual glue. I argue that they conflate a number of issues under these rubrics and that in each case either the problem to be solved isn’t subject to a general solution or the theory analogy is of little use. The issues I consider are: (1) what makes a concept efficient, useful, and informative, (2) what makes a concept refer to what it does, (3) what makes a set of objects form a single category, and (4) what makes concepts combine in one way rather than another.

Key words: category, concept, conceptual coherence, conceptual glue, mental theory.

Since the mid-1980s it has nearly become a platitude in cognitive psychology that many concepts should be understood on analogy with the terms of scientific theories.¹ At the same time, there has been little discussion of the content and motivation of the theory analogy. What exactly is the significance of the claim that concepts are embedded in mental theories? And why should we believe that this claim is true? Gregory Murphy and Douglas Medin’s influential paper, “The Role of Theories in Conceptual Coherence” (Murphy & Medin 1985), is one of the few places where anyone has explicitly addressed these questions, so it has become an authoritative expression of the view.² For Murphy and Medin, the chief motivation for introducing the theory analogy is what they call the problem of conceptual coherence. But, for all they say, it is not at all clear what the problem of conceptual coherence is or even whether there is a single problem that this notion circumscribes.

1. The Problem(s) of Conceptual Coherence

“Conceptual coherence” is a term of art that was introduced in Murphy & Medin (1985). They claim that previous theories of concepts (e.g., the Classical Theory and the Prototype Theory) fail to address the problem of conceptual coherence and that, as a consequence, these theories must be supplemented or abandoned. “The main thesis of this article is that current ideas, maxims, and theories concerning the structure of concepts are insufficient to provide an account of conceptual coherence” (Murphy & Medin 1985, p. 289). And, in their view, conceptual coherence offers a paramount constraint on a theory of concepts, a constraint that is met only by introducing the theory analogy. “The keystone of our explanation is



Minds and Machines 9: 241–255, 1999.

© 1999 Kluwer Academic Publishers. Printed in the Netherlands.

that people's theories of the world embody conceptual knowledge and that their conceptual organization is partly represented in their theories" (Murphy & Medin 1985, pp. 289–90). So what is the problem of conceptual coherence?

Despite Murphy and Medin's claim that the problem of conceptual coherence is of such great concern to a theory of concepts and their insistence that earlier theories of concepts fail to address the problem of conceptual coherence, they have little to say about what the problem is. Perhaps this is because they think that once it is pointed out that it is obvious. But, on the contrary, their remarks suggest a number of questions. Here is nearly all that they say by way of a positive characterization of the problem (Murphy & Medin 1985, p. 289, p. 290, p. 291, respectively):

- (1) Why is a given set of objects grouped together to form a category? That is, why is it that some groupings are informative, useful, and efficient, whereas others are vague, absurd, or useless? The current surge of interest in people's concepts has provided much information about conceptual structure and content. Yet, the central question of what makes a category seem coherent has only been sketchily addressed and incompletely answered.
- (2) [W]e do not propose a new model of conceptual representation. Rather, we present a theory of what the glue is that holds a concept together and an account of what sorts of concepts are easy to learn, use and remember, with the understanding that conceptual models must build appropriate structures to account for the facts discussed.
- (3) We have already hinted at what we mean by a coherent category. It is one whose members seem to hang together, a grouping of objects that makes sense to the perceiver. We do not give an operational definition of coherence because we do not wish to tie it to a particular theoretical framework. There are a number of measures that might reflect coherence, including how easily the concept is learned and used, and there may be others that are not known yet.

You will notice that these passages slip between conditions for *conceptual* coherence and conditions for *category* coherence. For Murphy and Medin, categories are the semantic values of concepts; they are the groups of objects, events, and so on, that concepts represent. Throughout their discussion, it is unclear whether Murphy and Medin are principally interested in the coherence of categories or the coherence of concepts, but presumably the two are interdefinable: A category is coherent just in case a concept that express it is coherent.

These passages offer a number of criteria for conceptual coherence. (1) and (2) suggest that the coherence of a concept is constituted by its informativeness, its usefulness, and its efficiency, and the presence of "conceptual glue", which "hold[s] a concept together". (3) continues this characterization. A coherent concept (category?) is one "whose members seem to hang together", whose corresponding category is "a grouping of objects which make sense to the perceiver", and which is easily learned and used. The difference between (1) and (2), on the one hand, and (3), on the other, is that (1) and (2) present the conceptual glue of a concept

and its informativeness, usefulness, etc., on a par, while (3) reports the latter as an indication of the former. So if we take (3) as the definitive characterization of the problem of conceptual coherence, then the problem is just to explain the nature of conceptual glue.

‘Conceptual glue’ is an awful metaphor, hardly more illuminating than the initial rubric (“conceptual coherence”) which it is supposed to explicate. Moreover, the gloss Murphy and Medin give for conceptual glue is consistent with several distinct interpretations. We will explore these in some detail, but first we should consider the more tentative criteria for conceptual coherence that Murphy and Medin cite. The point I wish to emphasize is that these criteria don’t locate a problem of sufficient unity to warrant a general solution, at least not one with the depth that one would expect, given the present enthusiasm for the theory analogy.

2. The Initial Characteristics of Coherence

It would be nice to start with some examples of coherent and incoherent concepts and see exactly how the coherent ones are more useful, more informative, and so on. The difficulty is that Murphy and Medin have no examples, apart from a few cases where the apparent incoherence of a concept is said to be misleading, given further information. Still, I think we can generate some examples of our own. What they seem to have in mind is the difference between representations of ordinarily arbitrary collections of objects versus representations of ordinarily salient (or at least, non-arbitrary) collections. For example, the collection of birds, boats, kidneys, and triangles is, by any intuitive standard, arbitrary, while the collection of animals is generally salient. Correspondingly, the concept BIRDS, BOATS, KIDNEYS AND TRIANGLES is incoherent, while the concept ANIMALS is coherent; or, treating coherence as a matter of degree, the first is much less coherent than the second. The difference between the two concepts has to do with the collections of objects that they represent. In particular, don’t be misled by the fact that in referring to these concepts we’ve used a complex expression in one case and a simple expression in the other. These expressions are names for concepts, not structural descriptions. For all that’s been said, either concept may be simple or complex. The relevant difference between them is that one represents a group of objects that seems to be an arbitrary collection, while the other represents a group of objects that is commonly treated as a salient collection.³

From Murphy and Medin’s characterization of conceptual coherence, one would think that incoherent concepts impose a severe cognitive burden. They are uninformative, useless, vague, absurd, inefficient, difficult to learn and remember. My own image of a person who gets caught entertaining incoherent concepts is someone who is himself incoherent — a babler who can’t make it out of a room without injury for all the penalties his concepts render. But this is overstated. If the coherence of a concept is a function of the salience of its corresponding category, then even paradigmatic incoherent concepts needn’t impose a cognitive burden.

To see why this is so, consider what happens in mundane instances of language comprehension.

Language comprehension, like all perception, involves successive computations that result in judgments of the applications of concepts. In language comprehension, the important concepts include (at least) phonological ones, syntactic ones, and semantic ones. A person categorizes a linguistic signal as a token of a certain linguistic type by assigning it to representations of the right type. For instance, recovering the word “dog” from a token signal involves a procedure that results in the judgment that the token is token of the type THE WORD ‘DOG’. And recovering the structural analysis of the phrase “the dog” from a token signal involves a procedure that results in the judgment that the token is a token of the type [*NP* [*DET* THE] DOG]. Similarly, recovering the literal meaning of a sentence involves a procedure that results in a judgment about the application of a representation that expresses a content that the speaker has in mind. If you say “Tom is a cookie thief” (meaning that Tom is a cookie thief) and I understand you, then I will token a complex representation with the content *Tom is a cookie thief*. Invariably this will include the concept COOKIE THIEF.

These mundane facts have consequences for the notion of conceptual coherence. Language comprehension includes judgments of the applications of concepts that express the contents of the communicative intention of the speaker. And language comprehension is automatic, incredibly fast, and reliable. Yet — this is the important point — there is no bound on the arbitrariness of the communicative intentions that can be the subject of a linguistic exchange. Communicative intentions about arbitrary collections of objects are often perfectly suited to be encoded linguistically. If you have the fancy to discuss fish, bubble gum, and bowling balls, you need only say so. What’s more, a suitably situated hearer will have no trouble recovering what you say, effortlessly tokening the concept FISH, BUBBLE GUM, AND BOWLING BALLS.

Speaker: We need to talk about a subject of grave political concern: fish, bubble gum, and bowling balls.

Hearer: I see.

Speaker: The problem with this country is its tendency to shamelessly isolate constituencies. We must act to promote a single, noble agenda. Fish, bubble gum, and bowling balls define this campaign. You are either for them or against them.

If this conversation is odd, it is because the oddness of the speaker’s views. Murphy and Medin may wish to insist that for the speaker to say something as wild as she does, she has to have a correspondingly wild “theory”, using the term loosely, as they do. For the moment, however, I want to concentrate on the hearer. The hearer has no trouble understanding the speaker. He may wonder why she would say such a thing; indeed, he may wonder what to make of her remarks. But he has no trouble recovering the literal meaning of her sentences. In the process, he easily tokens the (incoherent) concept FISH, BUBBLE GUM, AND BOWLING

BALLS. He needn't have anything remotely like a theory covering fish, bubble gum, and bowling balls. He needn't even have entertained this concept at some earlier time in his life. Yet he has no difficulty entertaining it when the situation demands, in this case when someone else has the corresponding category in mind and talks about it.

In short, incoherent concepts are readily available for use and, within certain limits, needn't impose any particular cognitive burden.⁴ This point shouldn't be unfamiliar. It is a consequence of our having a productive conceptual system that we are readily able to construct novel concepts. An examination of language comprehension manifests this aspect of the mind, but nothing essential turns on this one example. If someone had the will, she could easily entertain endlessly many new concepts, just for fun. Of course, this isn't to say that every concept is as easy to deploy as any other concept. Cognitive psychology offers many examples of concepts imposing varying cognitive demands upon a thinker. At the very least, there are such obvious facts as that the relative structural complexity of a pair of concepts will often determine which concept consumes more cognitive resources, such as memory.⁵ The present point is that arbitrariness in the extension of a concept needn't affect psychological processing. Witness the ease with which people token incoherent concepts in the course of language processing.

But I think we can say something stronger. In some cases, incoherent concepts *do* seem to offer greater challenges to a thinker than coherent concepts. However, these challenges aren't always particularly deep or mysterious. Don't forget: The coherence of a concept consists in something like the salience,⁶ or non-arbitrariness, of the collection to which it applies. A coherent concept just is one which applies to an ordinarily salient collection; an incoherent concept just is one which applies to an arbitrary collection. Often enough, the reason incoherent concepts impose greater cognitive challenges — to the extent that they do — is that they are concepts of arbitrary collections. In these cases, what more needs to be said? It is no surprise, for example, that a person's ability to recall a list is affected by whether she takes the members of the list to be related to one another. Twelve seemingly unrelated items will be much harder to recall than twelve items that fall naturally into three or four categories. In the latter case, people are able to chunk the items on the list so that for certain purposes it is as if they had to remember less. Yet all that's required for chunking is the recognition of an *organizing principle*. The principle could be as simple — and un-theory-like — as the fact that each member of the list is a fruit, a vegetable, or a dairy product.⁷

Perhaps Murphy and Medin mean to stress the issue of informativeness.⁸ Incoherent concepts, they might insist, are rarely informative and the informativeness of a concept increases with its coherence. For Murphy and Medin, a concept is informative to the extent that knowing that an object is in the extension of a concept provides one with relevant information about that object. For example, knowing that something is a cat presumably licenses the inference that it is an animal, that it has whiskers, that it has two eyes, and so on. That is, one can safely draw these

inferences just by knowing the object is a cat; one needn't have any substantial contact with the object, including perceptual contact. To repeat, Murphy and Medin might locate the coherence of a concept in its informativeness, forgetting about the issues related to efficiency and ease of learning and remembering.

But this response isn't entirely convincing. Among other things, the informativeness of a concept — any concept — depends upon what a person knows and how skilled she is at deploying this information at suitable times. The concept CAT is informative in most settings because people know all sorts of things about cats and have little difficulty accessing this information. In the same way, however, *prime facie* incoherent concepts are also informative, so long as one is privy to relevant facts about their members and is skilled at manipulating these facts. Consider, for example, the concept BOOKS ON SHELVES THAT ARE LABELED HB171. I'm not sure whether Murphy and Medin would classify this concept as an incoherent one, but on any graded scale it is certainly much less coherent than, say, the concept BOOKS. Yet knowing that something is in the extension of this concept may be highly informative, especially to someone who knows the Dewey Decimal System. If something is a book on a shelf that is labeled HB171, then it is likely to be a book about economics. The point is that the informativeness of a concept isn't necessarily intrinsic to the concept. It derives from the knowledge associated with the concept, sometimes exotic or idiosyncratic knowledge.

Murphy and Medin tend to use a loose characterization of mental theories (Murphy & Medin 1985, p. 290):

When we argue that concepts are organized by theories, we use *theory* to mean any of a host of mental “explanations”, rather than a complete, organized, scientific account. For example, causal knowledge certainly embodies a theory of certain phenomena; scripts may contain an implicit theory of the entailment relations between mundane events; knowledge of rules embodies a theory of the relations between rule constituents; and book-learned, scientific knowledge certainly contains theories. Although it may seem to be glorifying some of these cases to call them theories, the term connotes a complex set of relations between concepts, usually with a causal basis.

On this characterization, just about any inter-conceptual relation embodies a theory. Then any knowledge structure counts as a theory, and the informativeness of a concept will, by definition, depend upon the mental theories in which it plays a role. The conclusion to draw is that Murphy and Medin are stuck in a dilemma. Either (i) mental theories are to be construed so loosely that they can't seriously be taken to constrain theories of concepts or else (ii) we need substantive restrictions on what counts as a mental theory, in which case mental theories won't exclusively explain conceptual informativeness.

The first horn of the dilemma should be obvious. If any inter-conceptual relation embodies a theory, then all knowledge structures are theories, and the theories associated with a concept include all of the knowledge structures that involve that concept. For any given concept, the theoretical apparatus associated with it will be

enormous. Consider, for example, all the trivial things you know about rabbits — that they are colored, that they take up space, that they eat, and so on. It can't seriously be taken as a defect of earlier theories of concepts (e.g., the Classical Theory, Prototype Theory) that they don't capture *all* of these inter-conceptual relations. This would be tantamount to the claim that a theory of concepts is insufficient if it fails as a comprehensive theory of memory and reasoning.

The other horn of the dilemma is that for any substantive constraints limiting the knowledge structures that count as theories, the informativeness of a concept will often result from its association with knowledge structures that these constraints exclude. Consider the concept CAT. I suppose Murphy and Medin would want to say that one's theory of cats — insofar as it is isolated from one's considerable atheoretical knowledge of cats — is restricted to causal/explanatory beliefs about the perceptual properties of cats and their manifest behavior, beliefs that offer rudimentary accounts of why they have these properties and exhibit these behaviors (cf. Medin & Ortony 1989). So, for instance, one's theory of cats will include knowledge about the typical shape of cats, their typical outward appearance, the fact that cats are animals and, as a consequence, that their internal structure is especially relevant to their kind membership.

At the same time, much of what one knows about cats can't be part of this theory, so long as we are to assume that the theory-like knowledge structures are substantive and hence limited in scope. It is difficult to give any definitive examples, since Murphy and Medin are cagey about what they count as theories. But here is one that seems intuitive: Cats don't like water. It is unlikely that a substantive restriction on the theory-like knowledge of cats would include this belief as part of the theory of cats.

Now a concept is supposed to be informative when knowing that an object is in its application yields relevant information about that object. It must be recognized, however, that the relevant information about an object depends upon a person's interest in that object, and a person's interest in an object will vary considerably depending upon who the person is and what his circumstances are. In my case, the interest in cats is often to make them less of a nuisance — one cat in particular, named "Ray". Knowing that Ray is a cat is highly informative for me, since it brings to light a way to make him less of a nuisance. I know that cats don't like water, so I know Ray doesn't like water, so, if I want Ray to stop bothering me, I know that I need only spray a little water in his direction. In short, the informativeness of a concept will often not be explained by its role in a mental theory.

Let me remind you of where we are. Murphy and Medin intend to exploit the theory analogy as a solution to the problem of conceptual coherence. We've been looking at their initial, but tentative, way of characterizing conceptual coherence. They say that concepts are coherent to the extent that they are, among other things, useful, efficient, easy to learn, and informative. It should be recalled, however, that these are tentative characteristics — indications of coherence, as it were — and that on one reading of Murphy and Medin, it is only the presence of conceptual

glue that's constitutive of conceptual coherence. We will examine the issue of conceptual glue in a moment. So far, what I've been doing is taking their initial characteristics of conceptual coherence as if they were essential to conceptual coherence and then asking whether conceptual coherence, defined by these characteristics, provides a motivation for the introduction of the theory analogy. I've been arguing that it does not.

In the end, the issue about learning, remembering, and efficiency and the issue about informativeness warrant the same response. The extent to which a concept is informative, easy to remember, etc., depends upon a person's knowledge involving the concept and other features of her cognitive economy, including her interests at the time. If all such knowledge is to count as a mental theory, then the theory analogy unpacks as a complete account of memory and reasoning, with no substantive constraints on what counts as a mental theory. On the other hand, if mental theories are to be constrained cognitive structures, extra-theoretical knowledge will account for conceptual coherence in an abundance of cases. Either way, the initial criteria for conceptual coherence pick out what looks like a heterogeneous set of psychological phenomena. Probably the best we can say by way of a general solution to the problem is that a concept's "coherence" is often increased when it figures in a salient knowledge structure — hardly an insight that should be credited to defenders of the theory analogy.

3. Conceptual Glue

Let us now turn to Murphy and Medin's more definitive characterization of conceptual coherence, the presence of conceptual glue. Again, conceptual glue is what "holds a concept together", and a concept, we are told, is coherent to the extent that it is "held together". Similarly, the members of some categories "seem to hang together", and a category is coherent to the extent that it "seems to hang together".

The slip between concepts and categories may be less innocent than we've been assuming so far. It is hard to tell, since Murphy and Medin fail to explicate the metaphor of conceptual glue. In my view, the metaphor suggests three questions that ought to be kept distinct, at least initially. I now propose to discuss these questions so that we can see what kinds of answers they motivate and how the theory analogy might be recruited.

3.1. THE ISSUE OF REPRESENTATION

Why does a primitive (i.e., unstructured) concept represent what it does? Concepts are representations, and, we may suppose, some concepts are composed of other concepts, ultimately of primitive concepts. Primitive concepts are ones that don't have other concepts as proper parts.⁹ Since their representational character is so important to concepts, a paramount concern for theories of concepts is how they represent what they do. In the first instance, this question concerns the primitive

concepts, since all other concepts are composed of these. So, supposing, for example, that the concept ANIMATE is primitive, what makes it that ANIMATE represents animates and not circles or tortoises or the color red?

Within the growing psychological literature endorsing the theory analogy, it is common to reject what are often called “feature-based” theories of concepts, of which the Classical Theory of Concepts and the Prototype Theory are supposed to be a paradigms. A natural (though not mandatory) way to construe features in this context is that they are the constituents of lexical concepts, that is, the concepts of which lexical concepts are composed. (A lexical concept is one which is typically expressed by a single morpheme in a natural language.) For instance, WINGS, FLIES, EATS, are commonly cited features for the concept BIRD. The idea is that BIRD is actually composed of these concepts, in the same way that the phrase “the white bird” is composed of the words “the”, “white”, and “bird”. When supporters of the theory analogy reject feature-based theories of concepts, they are often enough rejecting the view that lexical concepts have constituent structure and thereby endorsing the view that most lexical concepts are primitive. Then a theory of how primitive concepts represent what they do amounts to a theory of how lexical concepts represent what they do.

In this context, the theory analogy may be viewed as an explanation of the representational character of primitive concepts. It is clear that Susan Carey, another leading defender of the analogy, puts it to this purpose (Carey 1985, p. 198):

One solution to the problems of identifying the same concepts over successive conceptual systems and of individuating concepts is to analyze them relative to the theories in which they are embedded. Concepts must be identified by the roles they play in theories.

Carey’s strategy is to endorse a particular view of scientific theories and to give it a psychological twist. The view is that theoretical terms have their semantic properties in virtue of their roles within their respective theories (see, e.g., Kuhn 1962, Sellars 1956, Lewis 1970, 1972; Lewis 1972 is an especially good introduction to the literature). Carey takes this account of the semantics of theoretical terms and applies it to concepts as psychologists conceive of them, that is, to mental representation. The idea, in brief, is that mental representations have their semantic properties in virtue of their roles in proprietary cognitive structures that are theory-like, such as folk psychology or folk physics.

This view of the nature of mental representation has its difficulties, and I’ve discussed these elsewhere (Margolis 1995). Suffice it to say for now that it is untenable if the theory analogy is supposed to offer a general account of representation. The problem is that most concepts aren’t embedded in mental theories, if the notion of a mental theory is at all constrained. This is one of the reasons Susan Carey is explicit about restricting the use of the analogy to just what she calls “ontological concepts”. These are concepts that encode our most central conceptualizations, such as the difference between plants and animals or between animate and inanimate objects.

3.2. THE ISSUE OF CATEGORY UNITY

Why does a group of objects form a category? A theory of how it is that concepts represent what they do doesn't, at first, seem to have anything to say about what holds a concept together or even suggest a clear sense of togetherness that needs to be explained. This is one place where the distinction between categories and concepts might be to the point. Maybe it really is category unity with which Murphy and Medin are concerned. They are sensitive to the fact that any arbitrary collection of objects could be grouped together to form a category. What, then, makes one group of objects form a category and other groups mere collections of objects?

There are two answers to this question — one that is trivial and one that is substantive. In either case, it must be remembered that “category” is a technical term. For Murphy and Medin a category is any collection of objects that serves as the semantic value of a concept. Here is the trivial answer: A group of objects is a category just in case a concept applies to all and only members of that group. In other words, categoryhood is decided by the existence of a semantic relation of a certain sort holding between a representation and a group of objects. If the representation represents the group, then the group is a category. The issue of category unity, as we might call it, reduces to the issue of how concepts represent what they do. Notice that, in this case, we have a fairly clear sense of something being held together. The existence of a representation holds a category together in that it applies to all and only members of that category, however arbitrary the category may otherwise seem. The objects may have little more in common than their being within the application of a single concept. Then again they might have a lot in common. This brings us to the more substantive answer to question 3.2.

Many philosophers think that at least some human concepts reflect the nature of items in the world, rather than the other way around. So what makes a group of objects a category may, in the first instance, have quite a lot to do with extra-semantic properties of the world (see, e.g., Rey 1983 and Kornblith 1993, both of which borrow from Saul Kripke's and Hilary Putnam's seminal work in semantics). Natural kind concepts are the ones that tend to generate these intuitions. Consider, for example, the category *gold*. No doubt, many people have the concept GOLD, which applies to all and only samples of gold. But it is less clear that the category *gold* is circumscribed just by there being a concept that applies to such samples. Intuition suggests that it is something about the internal structure of these samples that makes them instances of gold and that chemistry, not psychology or semantics, tells us about what makes something gold. Again, we have a fairly clear sense of something being held together. The existence of extra-semantic facts holds a category together in that they determine which objects are members of a category. But if the problem of holding together a category is the problem of what determines category membership in this sense, then the problem may be of little concern to psychology.¹⁰ And if the problem is what determines category membership in the

other sense — the trivial sense — then it reduces to the problem of the nature of representation.

3.3. THE ISSUE OF CONCEPTUAL UNITY

What determines which rule of conceptual combination applies in a given case? Conceptual glue is supposed to hold something together. So far, the only sense that we've been able to make of the metaphor where it might be of psychological interest is as a problem about the nature of representation. Under this interpretation, it would be more apt to call the phenomenon "category glue". Concepts are the glue that hold categories together in that a concept applies to all and only members of its corresponding category. There is, however, a sense in which a concept might be held together which we should consider.

Complex concepts are composed of other concepts and ultimately of primitives. The constituents of a complex concept are literally its components, just as the constituents of an English phrase are literally its components. What this means isn't so clear, but a standard view is that it implies that one cannot token a complex concept without tokening its constituents; that is, if the complex concept is present, then its constituents must also be present. Suppose, for example, that the concept BLACK CAT is complex, composed of the concepts BLACK and CAT, as its expression in English suggests that it is. Then one cannot token the concept BLACK CAT without thereby tokening the concepts BLACK and CAT.

At the same time, the presence of both BLACK and CAT doesn't suffice for the presence of the complex concept BLACK CAT. This should be obvious, since there must be a difference between the case where one is thinking about black cats and cases where one is just thinking about the color black and about cats, as when one entertains a list. So what makes it that BLACK and CAT form a complex concept in cases of the first sort and remain ununited in the second? Any answer to this question will appeal to the rules of conceptual combination: It is the application of a rule of combination that explains why two or more concepts form a single complex concept. But what determines which rules apply in a given case? Sometimes Murphy and Medin seem to think this is the problem that the theory analogy is supposed to address (Murphy & Medin 1985, p. 306):

B. Cohen and Murphy (1984) argued that it is impossible to explain how people form ... compound concepts using only *knowledge independent operations*. That is, they said that it is impossible to say in advance what a complex concept *XY* means knowing only the meaning of *X* and *Y*, but that extensive knowledge relating *X* and *Y* comes into play in order to arrive at just the right compound. In the context of our discussion, this point translates into the use of people's implicit theories and operations on concepts. For example, one's knowledge of the use of vehicles, their parts and what they do, and mishaps that happen to them can lead one to combine *engine* and *repair* to get "repair of an engine". One's knowledge about *experts* leads one to combine *expert* and *repair* differently.

The interpretation of a compound concept may be thought of as a hypothesis generated by background theories.

There is something deeply peculiar about the idea that one interprets one's compound concepts in light of background theories. It is not as if one entertains the concept ENGINE REPAIR and then wonders what concept it is, only to settle the matter by hypothesis formation and confirmation. That may be how it works in the case of linguistic exchanges — one tests hypotheses about the concepts one's interlocutor means to express with her words—but not in the case of thinking. Still, we can understand Murphy and Medin's point as follows: The rules that effect conceptual combination apply in context sensitive ways. In particular, of the many rules that can be used to form compound concepts, the right rule for a given case depends upon the concepts being combined. The mental theories associated with these concepts determine which is the correct rule. Since the application of a rule holds a complex concept together in that it unites its components, this is a way of construing conceptual glue. The rules of conceptual combination are the glue, and the theory analogy explains why one rule applies rather than another.

The example of how REPAIR combines with ENGINE and with EXPERT serves to illustrate how different rules will be used in similar-looking cases. EXPERT REPAIR and ENGINE REPAIR may look at first like the same type of construction. Both involve the occurrence of a concept that is typically expressed by a noun with the concept REPAIR, so one might think that their respective rules of combination are identical. However, they are not. We can tell this by noticing that “expert repair” intuitively means *repair done by an expert*, while “engine repair” intuitively means *repair of an engine*. If the corresponding conceptual constructions involved the same rule of composition then we would get one or other unintuitive meaning: “expert repair” would mean *repair of an expert* or “engine repair” would mean *repair done by an engine*. We need something to explain facts like these. Murphy and Medin think that it is our mental theories that determine the difference. “[O]ne's knowledge of the use of vehicles, their parts and what they do, and mishaps that happen to them can lead one to combine *engine* and *repair* to get ‘repair of an engine’. One's knowledge about experts leads one to combine *expert* and *repair* differently”.

Murphy and Medin are right that we will need an explanation of why one particular rule of combination applies rather than another, but it is doubtful that the theory analogy is of much use. Consider, for example, the knowledge that is necessary for effecting the combination of ENGINE and REPAIR. Most people are terribly ignorant about engines. Pace Murphy and Medin, most people know hardly anything about the parts of an engine, what they do, and the mishaps to which they are prone. I myself know almost nothing about engines; to me, an engine is a mystery. I can name a few crucial parts — the carburetor, the timing belt, the spark plugs — but I have no idea how these things work and, in most cases, what they do. Yet I am quite capable of effecting the conceptual combination. I know,

for example, that “engine repair” means *repair of an engine* and not *repair by an engine*. How do I know this?

Most likely, conceptual combinations aren’t entirely knowledge-independent operations. But Murphy and Medin are definitely wrong about the richness of this knowledge. Their analysis of the ENGINE example suggests that the required knowledge for effecting the combination is considerable, that one needs a host of beliefs about the structure of engines, how they work, how their parts causally interact; in short, that one have a theory of engines. But surely all one need know is that engines are inanimate objects and that a repair is a result of an action. Since inanimate objects can’t be agents, “engine repair” can’t mean *repair done by an engine*. The rule that effects this combination is thereby excluded.¹¹

Their explicit discussion of this example suggests that Murphy and Medin wouldn’t want to insist that the knowledge that engines are inanimate embodies a theory. This knowledge hardly seems “extensive”, to use their term. What is more, if one were to insist that this knowledge does embody a theory, it would be hard to see what stops any inter-conceptual relation from embodying a theory. Then the significance of the idea that there are mental theories and that they are important to cognition would be lost.

4. Conclusion

For Murphy and Medin, the motivation for the theory analogy comes from the problem of conceptual coherence, but, as I’ve pointed out, there are several issues that are being conflated under one rubric. If we understand the problem of conceptual coherence as the need to explain the usefulness, informativeness, efficiency, etc., of certain concepts, then the problem isn’t sufficiently unified to warrant a unique substantive solution. Probably the best we can say is that is that a concept’s coherence is often increased when it figures in a salient knowledge structure. If, instead, we understand the problem of conceptual coherence as the need to explain the presence of conceptual glue, we are owed a more precise characterization of what conceptual glue is. On one interpretation conceptual glue concerns the issue of how representations represent what they do. But Murphy and Medin can’t have a general account of representation, since people don’t have mental theories for most concepts. On other interpretations, the theory analogy verges on being irrelevant or empty. It is irrelevant to the extent that the problem of conceptual glue concerns what makes something a member of such categories as *bird* or *tree* — natural kinds, whose membership conditions are to be given by scientific investigation outside of psychology. It is empty to the extent that the problem of conceptual glue concerns knowledge-sensitive conceptual combinations. The knowledge for combination is generally atheoretical.

These days people talk of there having been a momentous shift in psychological theorizing from feature-based theories of concepts to theory-based theories. This may be so, but the content of the new theory remains murky and its utility question-

able. Until these issues are clarified, the new consensus surrounding theory-based theories is, at the very least, premature and perhaps unwarranted.

Notes

¹The term “concept” is used a number of ways in philosophy and cognitive science. Throughout, I follow psychological usage, according to which concepts are mental representations. Mentioned concepts are indicated by capital letters; words, by quotation marks.

²See also Medin & Wattenmaker (1987), Medin (1989), and Murphy (1991), where Murphy and Medin aim to elaborate their views. My discussion sticks to their original article because I don’t think they’ve said anything elsewhere to illuminate their original position, though interested readers may wish to follow up on this matter. For other endorsements of the theory analogy, see Carey (1985, 1991); Keil (1989), and Gopnik & Meltzoff (1997).

³A referee for this journal has suggested that degrees of coherence appears to be a pragmatic question, depending upon the goal at hand. I agree and take up the issue below when examining the claim that the degree to which a concept is coherent may reflect the degree to which it is informative.

⁴I want to emphasize that I am only claiming that on a plausible interpretation of incoherence, incoherent concepts are easy to use in certain ways. I am *not* claiming that, in general, any concept is just as easy to use as any other concept. This point is explained in the text, but it’s worth emphasizing since a number of readers have failed to make this distinction.

⁵Perhaps a more interesting illustration is the difference between concepts that occur in different levels of a cognitive hierarchy, e.g., basic concepts as opposed to super- and sub-ordinate concepts (Rosch et al. 1976). In general, basic level concepts are easier to learn and to access in a variety of psychological contexts.

⁶Throughout, I use “salient” in a relatively intuitive sense. Rather than try to construct a rigorous definition, I think it is best to stick to examples. A referee for this journal has asked whether there isn’t a pragmatic dimension to salience, just as there may be a pragmatic dimension to conceptual coherence (see note 3). This seems right, since intuitively, whether a collection is salient depends upon one’s interests, goals, and other contextual factors. Again, I discuss some of these issues below, when I examine the idea that the coherence of a concept turns on how informative it is.

⁷This isn’t to say that the organizing principle couldn’t be a mental theory of sorts; however, I do think it severely undermines the scope and importance of the theory analogy. As Murphy and Medin put it, the theory analogy is supposed to offer a novel, perhaps unique, solution to the problem of coherence. But on the current reading of the problem, a simple, ordinary belief could generate coherence; theories needn’t come into it.

⁸It is difficult to tell whether for Murphy and Medin the usefulness of a concept is anything more than its informativeness. I’ll assume that they take these to be one and the same notion.

⁹A referee has asked whether I am assuming a direct connection between the stock of primitive concepts and the undefinable terms of natural language. This is a tricky question since there are a variety of notions that may be at play when one posits that a concept has structure (Laurence and Margolis, in press). For instance, the concept corresponding to an undefinable term may still have structure; it just can’t have definitional structure. For present purposes, think of natural language as a guide to which concepts are *not* primitive. If a concept is encoded by a definable term or a non-idiomatic phrase, then it is complex. If it is encoded by an undefinable term, then it may be primitive, but that depends upon what other types of conceptual structure there are besides definitional structure.

¹⁰I take it that it is an open question to what extent other concepts are like natural kind concepts. The present point is that natural kind concepts illustrate how questions about conceptual coherence and questions about category coherence are separable.

¹¹The idea that conceptual combination requires considerable knowledge is probably wrong-headed for independent reasons. The problem is that the more knowledge that is required for a combination,

the more likely that the conceptual system can't be compositional (Fodor and Lepore 1992, chapter 6).

References

- Carey, S. (1985), *Conceptual Change in Childhood*, Cambridge, MA: MIT Press.
- Carey, S. (1991), 'Knowledge Acquisition: Enrichment or Conceptual Change?' In S. Carey and R. Gelman, eds. *The Epigenesis of Mind: Essays on Biology and Cognition*, Hillsdale, NJ: Lawrence Erlbaum Associates.
- Cohen, B. and Murphy, G. (1984), 'Models of Concepts', *Cognitive Science* 8, pp. 27–58.
- Fodor, J. and Lepore, E. (1992), *Holism: A Shopper's Guide*, Cambridge, MA: Blackwell.
- Gopnik, A. and Meltzoff, A. N. (1997). *Words, Thoughts, and Theories*, Cambridge, MA: MIT Press.
- Keil, F. (1989), *Concepts, Kinds, and Cognitive Development*, Cambridge, MA: MIT Press.
- Kornblith, H. (1993), *Inductive Inference and Its Natural Ground: An Essay in Naturalistic Epistemology*, Cambridge, MA: MIT Press.
- Kuhn, T. (1962), *The Structure of Scientific Revolutions*, Chicago: Chicago University Press.
- Laurence, S. and Margolis, E. (in press), 'Concepts and Cognitive Science', In E. Margolis and S. Laurence, eds. *Concepts: Core Readings*, Cambridge, MA: MIT Press.
- Lewis, D. (1970), 'How to Define Theoretical Terms', *Journal of Philosophy* 67, pp. 427–446.
- Lewis, D. (1972), 'Psychophysical and Theoretical Identifications', *Australasian Journal of Philosophy* 50, pp. 249–258.
- Margolis, E. (1995), 'The Significance of the Theory Analogy in the Psychological Study of Concepts', *Mind & Language* 10, pp. 45–71.
- Medin, D. (1989), 'Concepts and Conceptual Structure', *American Psychologist* 44, pp. 1469–1481.
- Medin, D. and Ortony, A. (1989), 'Psychological Essentialism', In S. Vosniadou and A. Ortony eds., *Similarity and Analogical Reasoning*. New York: Cambridge University Press.
- Medin, D. and Wattenmaker, W. (1987), 'Category Cohesiveness, Theories, and Cognitive Archeology', In U. Neisser, ed. *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, New York: Cambridge University Press.
- Murphy, G. (1991), 'Meaning and Concepts', In P. Schwaneflugel, ed. *The Psychology of Word Meanings*, Hillsdale, NJ: Lawrence Erlbaum Associates.
- Murphy, G. & Medin, D. (1985), 'The Role of Theories in Conceptual Coherence', *Psychological Review* 92, pp. 289–316.
- Rey, G. (1983), 'Concepts and stereotypes.' *Cognition*, 15, pp. 237–262.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M. and Boyes-Braem, P. (1976), 'Basic Objects in Natural Categories', *Cognitive Psychology*, 8, pp. 382–439.
- Sellars, W. (1956), 'Empiricism and the Philosophy of Mind.' In H. Feigl and M. Scriven, eds. *Minnesota Studies in the Philosophy of Science, I*. Minneapolis: University of Minnesota Press.