

Concepts in Human Adults

James A. Hampton

City University London

In: *The Making of Human Concepts*

Editors. Denis Mareschal, Paul Quinn and Stephen Lea.

Oxford University Press. 2008.

Address:

Psychology Department

City University

Northampton Square, London EC1V OHB, UK

+44 20 70408520

hampton@city.ac.uk

www.staff.city.ac.uk/hampton

Draft: 15 January 2008

7413 words not including references.

8719 words including references

Physics tells us that the world that we live in is “really” composed of vast numbers of tiny particles, defined and individuated by their mass, charge, velocity and position. That is all there is, the rest being empty space. In contrast the world that we interact with in our every day lives is composed of rooms with windows, people and cats, and plates filled with milk and cornflakes. It is a world of individual objects and stuffs, most of which can be effortlessly identified and labelled as being of a particular kind. Like an illustration in a child’s reader, we could look around our environment and attach labels to all of the objects and kinds of things around us. We can also label events that occur across different time scales – breakfast, driving to the shops, going to college – as well as states, qualities, activities and the rest. These labels represent the repertoire of verbal concepts that we have for interpreting and interacting with the world around us. As far as we know language is not a necessary element for humans to possess this ability to conceptualise the world. Nonetheless language is clearly integral to human conceptual knowledge, to the extent that one cannot study one without the other. Language also provides a huge boost to the cognitive capacities of the human species in relation to other primates (see Waxman and other chapters in this volume). This chapter will focus on these verbal concepts, and primarily on the nature of the concepts underlying “kind” terms such as *cat*, *car*, or *breakfast*. For discussion of other kinds of conceptual knowledge see Murphy (this volume).

To situate the discussion in the light of earlier chapters, the conceptual repertoire of an average adult consists of a large amount of highly over-learned knowledge. Unlike the concepts acquired in the concept learning experiments described in Chapters 3, 4 and 5 with humans, or indeed Chapters 7, 8 and 9 with other species, an adult will have had daily exposure to numerous exemplars of familiar categories (breakfasts, cars, beds) for many years. We therefore need to keep an open mind about how such concepts will be represented. The models that best explain performance in concept learning experiments may have only limited relevance to how adult verbal concepts are either learned, represented or utilised. It is probable that most of our more

advanced conceptual knowledge is acquired verbally without any direct experiential input. My knowledge of malaria for example is entirely based on verbal accounts of the disease's causes, symptoms and transmission mechanism. Most people's concepts of a wide range of things is likely to be limited to what they have been told, what they have seen on TV or what they have read in books. Humans have the ability to be told the answer without needing to work things out for themselves. And often (although not always) that will be sufficient to generate an immediate change in behaviour. If I am told by a reliable source that a particular headache pill increases my risk of serious illness, then I will at once stop taking it and look for an alternative. I don't need successive trials with corrective feedback in order to change my behaviour.¹

The strong interdependence of language and concepts is central to understanding a vital distinction in discussing and exploring human adult concepts. The language faculty is a set of skills and knowledge contained in the head of a speaker, but a "Language" is also an abstract cultural artifact with an independent existence of its own. The same is true of the concepts that constitute the meaning of the terms in any language. When we speak of a concept such as *cat* there are two very different ways in which we can understand the notion of "concept". Taking a psychological stance, we consider the concept to be information stored in the brains of all people who can be reasonably be said to understand what a cat is. Rey (1983) called this the "conception" of a cat. We could imagine the research program for studying this concept as follows: take a random sample of competent adult speakers of English. Ask them if they know what a cat is, and if so to tell you all that they know about them. The set of information that is common to most of the respondents will correspond to the concept of cat as it is understood by this group – the "folk" concept of cat (see Chapter 17). There is obviously room for individual

¹ On the other hand health messages about fatty foods or tobacco may have less immediate effect!

variation in concepts studied this way. Some individuals may have “eccentric” views that disagree with beliefs commonly held by others. One might also differentiate the concept of *cat* of pet owners from those of farmers, an urban versus a rural concept, and so forth. Medin, Lynch and Coley (1997) provided an excellent example of this type of research when they contrasted the concept of “tree” as possessed by groups of taxonomists, landscape designers and parks workers. Importantly, it is not possible for everyone to be wrong about what a cat is. Anything fitting the descriptive information in the right kind of way has to be a cat, and nothing else can be.

On the other hand, we could turn to a dictionary and an encyclopaedia in order to learn what the term “cat” refers to. There we would learn another, overlapping, set of information including commonly known facts (about being mammals, hunters of mice etc) as well as less well known facts – for example about the evolution of the current race of domestic pets from desert cat ancestors in the Near East some 10,000 years ago. For this notion, it is at least conceivable that we could all be wrong about a concept. Scientific advance and careful systematic observation can redefine the categories of what “really” exists, so that, for example, whales are no longer classified as fish. The question is then which of these two approaches gives us the “real” concept of a cat?

The cultural transmission of conceptual information through the use of language, both spoken and written, plays a key role in enabling us to use our concepts for communication. It also generates a sociological dynamic by which certain concepts are “owned” by particular groups in society, who construct and maintain the meaning of terms, often with perfectly explicit goals (as in the struggle against prejudice in areas such as gender, disability and mental health). It is therefore inaccurate to suppose that we learn our concepts simply through daily interaction with the world and by conversing informally with others. Consider that in most of the developed world a human will spend a minimum of 10 to 12 years in full-time education, and those who intend to pursue an

occupation that requires advanced levels of abstract thought will typically require 3 to 6 years more. Speaking as an educator, I feel it would be nice to think that at least some of our conceptual repertoire as adults was the result of this process.²

A brief consideration of history confirms the central importance of cultural transmission of concepts. Even a concept as simple as that of negative number was poorly understood until the end of the 18th Century (Boyer, 1968, p. 501) with some textbook writers in the mid-1700s still rejecting categorically the notion that two negative numbers could be multiplied together. The notion that zero is a number (rather than an absence of number) also took a surprisingly long time to evolve (Kaplan, 1999), just as did the physical notion that space could be “empty” (Aristotle had worked out that this was impossible, see Lafferty, 1993, p.91). Concepts in mathematics and science exemplify culturally transmitted concepts (see Carey, Chapter 14). Not only do school children around the world rapidly grasp concepts that eluded the most talented mathematicians of former times, but there is also an embarrassingly huge gulf between the conceptual grasp shown by the most able and least able in the same society, depending on their ability and opportunity to benefit from education.

The discussion so far has been to set the scene for our review of human adult concepts. As students of the human species we have two very different accounts to provide. One is to explain how the mind is able to develop, understand and competently use abstract relational systems such as mathematics, logic or scientific theories. Study of this question has been largely driven by theoretical and philosophical argument – what are the processes that drive scientific advance, and what sort of cognitive architecture would enable thinking of this kind to take place? There have

² Children’s concepts too may be heavily influenced by the topics that they are studying in class, as Lloyd, Braisby and Brace (2003) discovered.

also been fascinating studies of the processes of scientific reasoning and conceptual change in the lab (Dunbar, 1997). A quite different issue to explain is how the mind operates in everyday life, away from the rigors of abstract cogitation. To preview the rest of the chapter, it will be argued that most of the evidence about this second issue is that we depend very heavily on a similarity-based, flexible and relatively vague set of conceptual terms, often with little awareness of the deeper ontological questions that concern philosophers. People are aware of their conceptual cultural heritage, and are (in the right circumstances) willing to accept that their understanding of a concept may be incorrect if it doesn't correspond to how the "experts" in society understand the term. But for the most part they classify and label the world using inexact conceptual representations that they have picked up through being members of a community who share and fix the meanings through their daily usage of the conceptual terms.

The research to be reviewed should also be seen in its social context. It is almost exclusively research conducted with highly educated and conceptually able samples of university students, whose minds have had many years of training in "how to think". We should expect therefore that at least some of the results obtained may prove limited in generality (see e.g. Atran & Medin, 2008; Proctor & Keil, 2006; Ross & Tidwell, Chapter 17 this volume). Most importantly, in terms of context independent, abstract, logical thought, we should be cautious about taking the results of these experiments as representative of the conceptual competence of the average human adult.

Concepts in practice

How do we represent a concept in our minds? How can we determine the "content" of any particular concept and individuate it from others that we might possess? Some of our concepts have explicit definitions that we could learn and use: a fever is the raising of body temperature above its normal range due to internal processes, a prime number is a number divisible only by itself and 1, the crime of murder is the unlawful killing of a person with "malice aforethought". It

is in fact a major part of the job of legislators to find usable verbal definitions and of the courts to decide how to apply them in different circumstances. However most concepts in everyday use do not have definitions – a point now widely accepted (Fodor et al, 1980; Rosch & Mervis, 1975; Smith & Medin, 1981). Furthermore even when the definition may be known, people’s use of the term in everyday thought is likely to extend its application beyond its strict sense. Saturday night fever, or a game played at fever-pitch don’t explicitly require a raised body temperature. Murder in a naïve juror’s mind may correspond to a broader concept than the concept laid down in law – something more akin to the technical notion of homicide.

In the absence of definitions, it appears that we are forced to accept that “concepts in the mind” have more fluid, vague content. Four phenomena in particular led theorists to propose a prototype representation for concepts (Hampton, 1979, 1995, 2006; Rosch & Mervis, 1975). These are:

- a) **Vagueness.** Concepts refer to categories of things, but the borders of the categories are often poorly delineated, with people both disagreeing and even being inconsistent about what a concept refers (Hampton, 1998; McCloskey & Glucksberg, 1978).
- b) **Typicality.** Within a semantic category like *bird* or *car* there are some items that people agree are highly typical or representative of the category, while other equally familiar items may be atypical. Nightingales are typical birds (though rarely encountered) whereas ostriches are atypical (Rosch, 1975).
- c) **Opaque definitions.** People find it hard to give an explicit account of why objects should fall in a particular category. As remarked above, few of our concepts have explicit definitions (Hampton, 1979; McNamara & Sternberg, 1978).
- d) **Genericity.** Not being able to provide definitions does not mean that people do not know what is relevant to category membership. They will readily list aspects of a creature that are relevant to it being a bird, or aspects of a vehicle that make it a car. However these aspects are often only true of a majority of the category – they are generically but not

universally true (Hampton, 2008).³

Early versions of prototype theory (Hampton, 1979; Rosch, 1975) proposed that concepts capture similarity clusters of the kind studied in category learning experiments (See Chapter 3).

Effectively it was assumed that a prototype was constituted by a set of inter-correlated features each of which was predictive of category membership. Both membership and typicality in a category were determined by the weighted sum of features possessed by a candidate item.

Vagueness resulted from noise in the calculation of feature weights relative to some vaguely drawn criterion for membership. Typicality was simply a function of how many features an item possessed. Nightingales have more of the properties characteristic of most birds than do ostriches.

The lack of explicit definitions and the ready production of generic properties were directly explained in terms of people representing the concept by a list of features each of which could be generally, but not necessarily, true of the concept. (The model does not require that *no* features are necessarily or universally true of a concept, only that such features on their own do not serve to define the concept. In Hampton, 1995, I also pointed out that by a judicious selection of feature weights a well-defined concept may also be represented as a prototype, so the fact that humans were defined by Plato as featherless bipeds does not imply that *human* may not be a prototype

³ An early treatise on logic made the distinction of metaphysically universal truths such as “*every human being is living*” and morally universal truths such as St. Paul’s rather un-Christian remark in Titus 1:12 that “*Cretans are always liars, evil brutes, and gluttons*”, or Horace’s equally pithy “*All singers have this fault: if asked to sing among their friends they are never so inclined; if unasked, they never stop.*” The authors commented that for moral truths “it is enough if things are usually this way”, Arnauld & Nicole, 1662/1996, p 114.

concept for Plato.⁴)

Subsequent research has led to the rejection of this relatively simple notion of a prototype. One suggestion has been to supplement the notion by proposing that the set of features are not a simple list, but are structured into a knowledge schema or frame (Barsalou & Hale, 1993; Hampton, 2006). The schema represents additional information about the relations (causal and other) between the features. For example a car has a motor, requires fuel and is self-propelled. These three features are of course intimately related, and their relationship is part of our concept of *car*. If the motor ran on clockwork and drove the windshield wipers only, the fuel was for heating the passenger compartment and the car was pulled along by sails it would be a highly atypical car, in spite of matching all three features.

Such schemas however may be quite patchy in terms of how much of the “true” cultural concept of *car* is represented. As Keil and his colleagues have shown in their studies of the Illusion of Explanatory Depth, people are often unaware of how little actual knowledge of causal links in complex artifacts they possess (Keil, 2003). People are probably quite good at conditional reasoning based on such schemas (if there is no fuel in the car it won’t go anywhere), but are usually very bad at being able to say just how the presence of the fuel is actually translated into the forward motion of the car. Schemas can still be thought of as prototypes for two important reasons: one, they provide abstract representations of the typical features of a class while losing details of individual members, and two, they represent the typical centre of a class but are vague about the borderline. They do not precisely specify how close the match to a schema needs to be in order for an individual (or subclass) to count as falling under that concept. Cars may be

⁴ Rumour has it that when Diogenes brought a plucked chicken to Plato in his academy, the latter was forced to amend his definition to “featherless biped *with broad nails*”.

powered by electricity or have three wheels with the steering at the rear, but yet still be cars.

In Hampton (1998) I investigated the degree to which the probability of a borderline item being categorized in a given category was simply a function of its typicality for that category, as would be predicted by prototype theory. I found that for biological categories (but not for others), categorization could deviate systematically from typicality where a creature was known to be of one kind but looked like another. The creatures in question were mainly bats and penguins, whales and dolphins. For example a whale was rated to be quite typical of a fish, but was unlikely to be categorized as such. Dissociations of categorization and typicality in verbal concepts have been reported elsewhere. Rips (1989) presented a number of studies in which the two variables were dissociated. Typicality is usually judged on the basis of surface appearance. The notion of being a typical example seems to require that an item look right. On the other hand categorization requires deeper hidden parts to be correct. Thus a bird that changes its appearance into an insect through toxic contamination was judged in Rips's study to be more typical of an insect, yet more likely to be a bird. Such results suggest that similarity to a category schema could be calculated in different ways depending on the task context (although Rips used his result to motivate the more radical claim that categorization is not based on similarity to a schema at all).

Hampton, Estes and Simmons (2007) replicated Rips's metamorphosis study and found that there were in fact quite large individual differences amongst college students in how they responded. When asked simultaneously to judge both typicality and category membership for each scenario, (as in Rips's study) a majority (17) judged that the creature was still of the same kind after the change, but a sizeable minority (12) made the reverse judgment. In a second experiment, where students in the relevant condition only judged category membership and not typicality, only 8 judged that the creature's kind had not changed, whereas 19 now thought it had. A third experiment confirmed that most students had the intuition that in fact the creature was no longer a

bird (as in Rips, 1989, the experiments all actually used multiple examples of the transformations and incorporated other conditions and controls). Hampton et al. argued that students were making this choice on the basis that the transformation, even though externally caused, was so great that it must have affected the inner workings of the creature to the point where it was now more of an insect than a bird.

The prototype theory has also been challenged by growing evidence for the role of exemplar knowledge in conceptual tasks (see Murphy, 2002, Ch.4; Smits et al., 2002; Storms, de Boeck & Ruts, 2000; 2001) which would suggest that the process of abstraction of conceptual prototypes is incomplete. For example Smits et al. (2002) considered how people would decide whether to categorize a number of unfamiliar exotic edible plant products as fruits or vegetables.

Classification was best modelled with an exemplar-based model that used similarity to individual known fruits and vegetables rather than similarity to an abstraction of each category to determine classification. Similarly Storms et al. (2000) found that typicality in categories is usually better captured by an exemplar model (summed distance in similarity space from stored exemplars) than by a prototype model (distance from a central average point of the category in similarity space).

A recent book by Rogers and McClelland (2004) provides a detailed model of how both prototype and exemplar effects could be generated through the same simple associative learning process. They adapted a model proposed by Rumelhart (1990) in which a PDP feed-forward network is trained with facts such as *a robin CAN sing* or *a canary IS yellow*. The subject noun (robin) is activated in one input array, and the relation (CAN) in a second input array, and then connections are adjusted by gradual incremental learning in order to activate the correct element (sing) in an output array. The model shows a range of interesting effects, including learning broad category distinctions before narrower ones (e.g. animals versus plants before birds versus fish), and retaining broad categories over more specific ones when the network is damaged. Both of

these effects correspond to well-known effects in the literature (Mandler & McDonough, 1993; Quinn & Johnson, 2000; Warrington, 1975). The model abstracts away from individual inputs by forming generalisations over similar inputs, so that its representation of a category resembles a prototype. However it also retains enough within-category structure to be able to account for exemplar effects.

A particularly interesting feature of the model is that it provides a possible account of “centrality” effects. Sloman, Love and Ahn (1998) described how the features of a concept differ in their centrality for the concept. Centrality can be defined and measured in different ways, but the basic notion is the degree to which the feature could be changed without otherwise affecting the rest of the concept. Washing machines are typically white, but they could be yellow without any need to revise any of their other features. They typically also contain a drum for the clothes, and it would not be easy to change this feature without at the same time making major adjustments to the rest of the concept schema. Centrality is therefore related to the number of causal dependencies that a feature enters into within the schema that represents the concept. Sloman and colleagues have shown that centrality of an attribute is correlated with a number of other behavioural measures such as the importance of the attribute for predicting category membership.

Simple prototype or exemplar models have no obvious way to derive feature centrality. The prototype model weights features by their covariation with category membership, but does not represent the covariation or dependency amongst features. The exemplar model simply situates exemplars in a similarity space and has no more to say about the features of a concept than that. Rogers and McClelland’s model however retains information about the complex covariation of features within a domain, so that it can determine which features are more central in terms of the total pattern of correlation that they have with the other features. It captures what makes a concept coherent, without the need for a higher level of causal understanding to be superimposed

on the representation.

Psychological essentialism and beliefs about concepts

Because most humans live in social groups, and concepts are a fundamental part of communication with other people, we have seen how learning models that simply relate the mind's contents to its learning history are going to be inadequate. I have already made the important point that many of our concepts are not "ours" to use or change as we would wish but belong to our culture. I could choose to start including whales and other cetaceans in my category of fish, but I would then stand accused of using the word fish in a non-standard way (but see Dupré, 1999). Note however that the lay population has not accepted all of science's terms. Botanically speaking many of the vegetables that we eat are fruits – a green bean forms from the ovary on the plant after the flower has been fertilised, and contains the seeds for the next generation, just as a plum or an orange does. Sometimes indeed even established sciences have terminological problems. Recently the world press reported difficulties with the definition of the term "planet" in astronomy. Originally a large body circling a sun, the problem arose when it turned out that Pluto (the furthest recognised planet in our solar system) was smaller and had a less circular orbit than two other objects that had not been awarded the status of planets. A concept that appeared to be clear and well understood turned out to be vague and arbitrary to an embarrassing degree. The world is not always as neat as we would wish it to be.

These problems aside, the average student participant in concept experiments (no doubt like other members of the adult population) is aware that there are experts who may know more about a concept than they do. Kalish (1995) posed the following kind of problem to his participants. Imagine that there is an animal in the zoo, and that George says it is a possum while John says it is not. Does one of them have to be right and the other wrong, or could it just be a matter of opinion? Kalish found a systematic tendency to say that one of the two must be right –

particularly for biological kinds. This tendency to assume that there is an externally determined “correct” answer to a question of categorization has been associated with the thesis of Psychological Essentialism (Medin & Ortony, 1989). The thesis holds that people believe that things are of a particular kind because of some hidden essence deep within them. We may not know what the essence actually is, in which case our representation of the concept may contain an empty “place-holder”. (Keil et al., in press, have shown that both adults and children know about where expertise on different domains is to be found.)

Kalish’s result was replicated in a study by Claire Simmons as part of her PhD thesis (Simmons & Hampton, 2006). She investigated three different manifestations of essentialist beliefs. In addition to the question of fact versus opinion, students were asked to judge whether a category had all-or-none membership, as opposed to admitting of partial membership, and whether or not they would defer to an expert’s opinion if it differed from their own. Across a range of categories, there was evidence for stronger essentialist beliefs about biological categories than about artifacts. Willingness to defer to experts was however very limited (see also Braisby, 2001, 2004; Proctor & Keil, 2006). Interestingly, in common with the Hampton, Estes and Simmons (2007) study, there were systematic individual differences. Thinking that membership should be all-or-none correlated positively across individuals with thinking that categorization was a matter of fact rather than opinion. Degree of deference did not correlate with either of the other measures however. There is also reason to suppose that deference may be more a question of social conformity than of essentialist beliefs. Braisby (2001) found that when asked about whether a genetically modified salmon would still be a salmon his respondents deferred not only to the opinions of biologists but also (to a lesser extent) to the opinions of shoppers.

Concepts and causality

A recurring theme in the concepts literature since Murphy & Medin (1985) is that concepts are

part of a broader causal-explanatory with which we understand the world. As Rogers and McClelland (2004) point out, a simple learning model (such as could easily be found in other species) can learn about the centrality of different properties of a conceptual domain, without invoking notions of cause or explanation. However we do frequently invoke conceptual information as explanations for events and properties that we observe in the world.

A set of studies by Daniel Heussen for his PhD thesis (Heussen and Hampton, 2007) explored the ways in which people explain the properties of a concept. Heussen asked student participants to explain things such as why axes are dangerous. An initial qualitative analysis revealed that they frequently appealed to another property as an explanation. They said for example that axes are dangerous because they are sharp. Heussen then obtained ratings of how satisfying such explanations were for a sample of paired properties in the frame “[Concept X] has [property p] because it has [property q].” He predicted the variance in these ratings on the basis of different measures of the properties in question. Where applicable, these measures were taken in each direction (p to q and q to p). The measures were co-occurrence (“Of all man-made things that are dangerous, what percentage is also sharp?”), counterfactuals (“If axes were not dangerous, would they still be sharp?”), and the mutability of each individual property (“How difficult is it to imagine axes that are not dangerous/ axes that are not sharp?”). He also showed another group of participants a diagram listing the most salient properties of the concept, and asked them to draw arrows to indicate when one property “depended on another”. A measure of centrality was then taken in terms of the number and strength of links involving each target property.

In a subsequent regression analysis, we found that the primary predictor of the plausibility of an explanation with a β of .48 was the dependence of p on q – that is the degree to which being dangerous was judged to depend on being sharp in the network dependency diagram. In addition the two counterfactual measures predicted additional variance with β s of around .23 each – the

questions of whether if axes were not dangerous they would still be sharp, and if they were not sharp they would still be dangerous. Interestingly neither of the co-occurrence measures, reflecting intuitions about the statistical correlation between the properties predicted any variance in the plausibility of explanations⁵. Apparently the fact that people consider most sharp things to be dangerous did not enter into people's thinking when deciding that it makes sense to say that axes are dangerous because they are sharp. (For a detailed model of causal understanding see Sloman, 200x).

Concepts and truth

Another cognitive ability involving language and concepts is the ability to make true assertions (and the concomitant ability to lie). Concepts are the building blocks of propositions, and many propositions can be given "semantic value" in terms of being true or false in our world. If I declare that there is a hippopotamus in the room, then the truth of the statement will depend on their being a creature in the room and on that creature being a hippopotamus. Enabling words and sentences to refer to the world, that is "establishing reference" is a major part of the role of concepts in our thought. There is a large literature in the philosophy of mind and language, and in semantics itself, on how words and sentences establish reference and the conditions under which propositions of different kinds are true or false. The interface of this literature with the psychology of concepts has led to some interesting debates on the nature of concepts (e.g. Fodor, 1998).

Much of the debate turns on the question of whether concepts are constituted by the type of thing in the world that they refer to (externalism) or by the type of content that is represented in the mind (internalism). For most philosophers (e.g. Rey, 1983) the concept of *cat* is an abstract entity that refers to the class *Felis silvestris catus*. Most speakers of English will have the name of this

⁵ Reliability of the different rating measures used as predictors was comparable.

concept in their vocabulary, and will represent it in some way, which may be more or less correct. This externalist view would allow that we may all be wrong about the content of a concept. It could turn out that water is not after all H₂O, in which case we would all change our represented concept to match what we now know to be the case (see Putnam, 1975, and Kripke, 1980 for famous papers on this topic). The externalist view subscribes to the distinction that I made at the start of the chapter between cultural concepts and individual concepts. In addition, they would say that for many concepts (such as those referring to natural kinds in the world) the cultural concept itself has to “defer” to the real class of things in external reality.

The advantage of specifying conceptual content in terms of external criteria should be obvious for the treatment of truth and falsehood. If concepts can be individuated with external criteria, then the truth of propositions involving those concepts can be easily defined in terms of set-theoretic semantics. The statement “all cats purr” will be true just in the case that the class of “things that purr” is a superset of the class of “cats”. If on the other hand the notion of *purring* is something that depends on how an average educated speaker of English understands the meaning of the term, then it becomes problematic to separate out issues of truth from issues of meaning. Whether all cats purr or not will depend on how you understand what kind of creature may count as a cat, and what kind of sound should count as a purr. Truth becomes relative to an individual’s set of meanings, and can not be determined independently of individual contexts.

This relativity relates to another problem in philosophical logic, namely *vagueness* (Keefe & Smith, 1997). In deciding whether all cats purr, it is necessary to determine for every creature (a) whether or not it is a cat, and (b) whether or not it purrs. But many statements of this kind appear to rely on vague predicates. The classic examples relate to the so-called Sorites paradox. How does one determine the set of adult men for whom “X is bald” is true? If we start with those with no hair at all, then the predicate clearly applies. If we then consider someone with just a single

hair the statement is still clearly true. But if we then proceed adding one hair at a time, we are on the uncomfortable horns of a dilemma. Either we must say that there is a particular number of hairs at which “is bald” switches abruptly from true to false, or we must allow that adding just one hair can never change the truth of the statement in this way. The first appears counter-intuitive, since it would imply that one could have two individuals, one with just one hair less than the other, but we would call one bald and the other not bald. The second is even worse, since it implies that everyone is bald (or if one starts at the other end of the sequence, that no-one is bald.)

One way to handle vagueness is to stipulate that it doesn't exist. Thus, Frege (1952) declared that all meaningful propositions involving concepts had to be true or false, and any concept that admitted of partial truth was not a concept at all. Alternatively it has been suggested that natural language statements should allow of varied degrees of truth (so called “fuzzy truth”, Zadeh, 1965). If baldness is a matter of degree, then the truth of “X is bald” can be mapped onto the number of hairs through a continuous function (Hampton, 2007). There are however unresolved problems with this solution (Osherson & Smith, 1981, 1982, 1997; Kamp & Partee, 1995), in particular because once fuzziness is allowed into truth values, other desirable properties of truth functions tend to be lost. For example, Harold may be 50% bald, but “Harold is either bald or not bald” would appear to be necessarily true, while “Harold is both bald and not bald” should be necessarily false. No completely satisfactory way to derive this intuition has been proposed.

A possible solution was proposed by Kamp and Partee (1995). Without going into the technical detail, they proposed that there may be three types of truth value – true, false and undecided. Vague borderline cases would fall in the “undecided” category. To test this notion, I conducted some studies with my students Bayo Aina and Mathias Andersson to investigate whether people would actually recognise a category of “undecided” or borderline cases. The method we used was

to ask one group of people to categorize borderline cases (e.g. *is a tomato a fruit*) using just a Yes/No answer, and a second group to categorize them but with three response choices: Definitely Yes, Definitely No, and Not definitely sure. Our intuition was that if there were cases that a person knew were problematic in some way, then they would use this last category. Moreover, if asked to do the task again 2 weeks later, the same items would still be seen as problematic and would remain in the unsure category. In contrast we predicted that those participants who were forced to choose between yes and no would tend to choose a response in an arbitrary way and so would be more likely to give a different response when asked again at a later date. In the event, across two studies we found that the level of instability in categorization responses was identical in the two groups. People were no more reliable at selecting the items for which they were “definitely sure it’s in the category” versus the rest than they were at just categorizing the items with yes or no. The conclusion is that vagueness does not just apply to a certain set of items. The boundaries of the vague region are themselves just as vague as the category boundary.

Within the psychological literature on truth and vagueness, there is plenty of evidence that people fail to appreciate the set-theoretic approach to truth based on externally defined sets. Informal testimony from logic teachers suggests that the average person’s mind is unreceptive to logical theorems. Research into syllogistic reasoning confirms that people have a very hard time drawing inferences based on propositional form, and that most reasoning is performed through simulating situations in one’s mind, in effect constructing mental models (Johnson-Laird & Byrne, 1996). The reader may like to try the following puzzle taken from Lewis Carroll (1958, p.119):

No kitten, that loves fish, is unteachable.

No kitten without a tail will play with a gorilla.

Kittens with whiskers always love fish.

No teachable kitten has green eyes.

No kittens have tails unless they have whiskers.

The mind quickly fills with images of kittens, tails, whiskers, bowls of fish and lonely gorillas seeking a playmate. Logical form is buried beneath the conceptual content and discovering it takes considerable effort.

Verbal concepts also display a lack of respect for the constraints of logic. An early demonstration, following Randall (1976), was a study in which I showed that conceptual categorization is not necessarily transitive (Hampton, 1982). People may agree that X is a type of Y, and Y is a type of Z, but balk at the conclusion that X is also a type of Z. For example, car-seats are a type of chair, and chairs are a type of furniture, but car-seats are not furniture. The effect is not just due to fuzziness or partial category membership, since chairs are very typical kinds of furniture, mentioned by just about everyone when asked to list examples of the category. Rather, the effect demonstrates that people judge truth of semantic statements like these on the basis of the overlap of conceptual content, rather than on the basis of the inclusion of one extensional class within another.

Similar evidence of failure to respect the logic of extensions was found in a series of studies I conducted on how people categorize items in complex categories formed using logical connectives (see Hampton, 1997). When people judged what belongs in categories such as “*sports which are games*”, or “*fruits or vegetables*”, or “*dwelling that are not buildings*”, they frequently made judgements that contradicted the apparent logical relations involved of set intersection, set union and set complements. In Hampton (1987) I proposed a model for how the semantic content of two concepts could be integrated when forming a conjunctive concept such as *sports which are games*. Because it combines generic, default, features of the two concept prototypes, and then subjects them to a coherence check, the model actually predicts that people will not respect logical constraints in their judgments in just the way that was found.

A similar effect is a phenomenon discovered by Shafir, Smith & Osherson (1990) that I explored in a recent paper with Martin Jönsson (Jönsson & Hampton, 2006). People judged universally quantified sentences in which the subject noun was either unmodified or modified with an atypical adjective. Examples are:

- (1) All sofas have back-rests
- (2) All uncomfortable handmade sofas have back-rests.

There was an apparently irresistible tendency to consider the first sentence more likely to be true than the second. The reader should check their own intuitions before reading on. The problem with this intuition is that, of course, if (1) is true then (2) must necessarily follow. (A separate group of students confirmed that all uncomfortable handmade sofas were in fact sofas). Equally if (2) were false, then (1) must necessarily also be false since (2) provides a counterexample. We tried different variants of the task, including placing the sentences side by side and asking students to select which was more likely to be true, and replacing the quantifier with the phrases “Every single sofa..”, “All sofas always “ and “100% of sofas ..”. None of these manipulations removed the tendency to rate the unmodified sentence as more likely to be true. In our final experiment, we adopted a format in which each pair of statements were printed one above the other in pairs, and participants simply had to state whether each one was true or false. Filler sentences were included in which either (1) or (2) were clearly and legitimately more true. (For example it is quite correct to say that “all trucks are noisy” is more likely to be true than “all toy trucks are noisy”, since *toy trucks* may not be a subset of *trucks* in its default sense.) In this final experiment we managed to reduce the incidence of the fallacy to a more respectable level, although the trend was still in the same direction. We termed the effect the Inverse Conjunction Fallacy, in recognition of its close relation to Tversky and Kahneman’s (1983) more familiar Conjunction Fallacy, which is another example of people using conceptual similarity in a situation where extensional reasoning is called for.

In sum, all of these demonstrations suggest that people are committed to thinking conceptually in terms of internally constituted meanings rather than in terms of external classes. Psychological truth is a very different beast from logical or semantic truth, and they should not be confused.

Concepts as grounded in perception and action

Much of the preceding discussion of how people use their concepts is compatible with a theory of conceptual representation propounded by Barsalou (2003, 2008). Barsalou challenges the traditional assumption that conceptual information is represented in the mind as an *amodal symbol*. Instead he suggests that concepts are never completely divorced from the experiential component of their instantiations. Encounters with cats will leave context-dependent traces of multi-modal experiences, including affective responses, vision, sound and smell. These are used to construct a *simulator* – a brain mechanism that allows us to reconstruct and interrogate representations of cats under different conditions. If the simulator is unable to construct a cat that is not at the same time an animal, then the system can infer that all cats are animals. Space does not permit a review of the increasing evidence for this view, but much of the evidence comes from showing that perceptual properties of a concept continue to influence the processing of the concept even when it is presented in abstract symbolic form (e.g. as written text). It should be clear that this approach would lend itself very readily to explaining the flexibility and vagueness of adult human concepts.

Conclusions

Research on adult human concepts has expanded greatly in the last three decades, and in addition to psychology there is also important and relevant work being done in linguistics, semantics, artificial intelligence, philosophy and neuroscience. In fact the topic is a prototype of the interdisciplinary interests of cognitive science. The applications of the research also spill over

into a range of areas, including branding and marketing of consumer products and companies, social stereotyping, racism and prejudice, and public understanding of science. The challenge for the future is to encourage researchers to cross the disciplinary boundaries and to start to integrate the different theories and perspectives. We need to take our concept of *Concept* from its current prototype stage (fluid, vague, similarity-based, polymorphous) into an explicit scientifically respectable concept, playing its role in a causal theory of mental representation and cognition.

References (1305 words)

- Arnauld, A., & Nicole, P. (1662/1996) *Logic or the Art of Thinking*. Translated and edited by J. V. Buroker. Cambridge: Cambridge University Press.
- Atran, S., & Medin, D. L. (2008). *The Native Mind and the Cultural Construction of Nature*. Cambridge, MA: MIT Press.
- Barsalou, L.W. (2003). Situated simulation in the human conceptual system. In H. E. Moss and J. A. Hampton (Eds.) *Conceptual Representation* (pp. 513-562). Hove: Psychology Press.
- Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, 59, in press.
- Boyer, C. (1968). *A History of Mathematics*. New York: Wiley.
- Braisby, N. R. (2001). Deference in Categorisation: Evidence for Essentialism? In Moore, J. D. & Stenning, K. (Eds.), *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society* pp. 122-127. Mahwah, NJ.: Erlbaum.
- Braisby, N. R. (2004). Deference and Essentialism in the Categorization of Chemical Kinds. In R. Alterman & D. Kirsch (Eds.), *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society* pp. 174-179. Mahwah, NJ: Erlbaum.
- Carroll, L. (1958). *Symbolic Logic and the Game of Logic*. New York: Dover Publications.
- Dunbar, K. (1997). How scientists think: Online creativity and conceptual change in science. In T. B. Ward, S. M. Smith, & S. Vaid (Eds.) *Conceptual structures and processes: Emergence, discovery and Change* pp. 461-493. Washington DC: APA Press.

- Dupré, J. (1999). Are whales fish? In D. L. Medin and S. Atran, (Eds.) *Folkbiology*. pp. 461–476.
Cambridge, MA: MIT Press.
- Fodor, J. A. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, J. A., Garrett, M., Walker, E., and Parkes, C. (1980). Against definitions. *Cognition*, 8, 263-367.
- Frege, G. (1952). On sense and reference. In P. Geach & M. Black (Eds.), *Translations from the philosophical writings of Gottlob Frege*. Oxford: Basil Blackwell.
- Hampton, J. A. (1979). Polymorphous Concepts in Semantic Memory. *Journal of Verbal Learning and Verbal Behavior*, 18, 441-461.
- Hampton, J. A. (1982). A Demonstration of Intransitivity in Natural Categories. *Cognition*, 12, 151-164.
- Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, 15, 55-71.
- Hampton, J. A. (1995). Testing Prototype Theory of Concepts. *Journal of Memory and Language*, 34, 686-708.
- Hampton, J. A. (1997). Conceptual combination. In K. Lamberts & D. R. Shanks (Eds.), *Knowledge, Concepts and Categories* (pp. 135-162). Hove: Psychology Press.
- Hampton, J. A. (1998). Similarity-based categorization and fuzziness of natural categories. *Cognition*, 65, 137-165.
- Hampton, J. A. (2006). Concepts as Prototypes. In B. H. Ross (Ed.), *The Psychology of Learning and Motivation: Advances in Research and Theory*, Vol. 46 (pp. 79-113). Amsterdam: Elsevier.
- Hampton, J. A. (2007). Typicality, Graded Membership and Vagueness. *Cognitive Science*, 31, 355-383.
- Hampton, J. A. (2008). Stability in concepts and evaluating the truth of generic statements. In F.

- J. Pelletier (Ed.), *New Directions in Cognitive Science*. Oxford: Oxford University Press (in press).
- Hampton, J. A., Estes, Z., & Simmons, S. G. (2007). Metamorphosis: Essence, Appearance and Behavior in the Categorization of Natural Kinds. *Memory & Cognition*, *35*, 1785-1800.
- Heussen, D., & Hampton, J. A. (2007). 'Emeralds are expensive because they are rare'; plausibility of property explanations. In S. Vosniadou, D. Kayser, & A. Protopapas (Eds.), *Proceedings of Eurocogsci07. The European Cognitive Science Conference*, pp. 101-106. Hove: Erlbaum.
- Johnson-Laird, P. N., & Byrne, R. M. J. (1996). Mental models and syllogisms. *Behavioral and Brain Sciences*, *19*, 543-546.
- Jönsson, M. L., & Hampton, J. A. (2006). The Inverse Conjunction Fallacy. *Journal of Memory and Language*, *55*, 317-334.
- Kalish, C. W. (1995). Essentialism and graded membership in animal and artifact categories. *Memory & Cognition*, *23*, 335-353.
- Kamp, H., & Partee, B. (1995). Prototype theory and compositionality. *Cognition*, *57*, 129-191.
- Kaplan, R. (1999). *The Nothing that is; a Natural History of Zero*. London: Penguin Books.
- Keefe, R. & Smith, P. (1997). Theories of vagueness. In R. Keefe & P. Smith (Eds.), *Vagueness: a reader* (pp. 1-57). Cambridge: MIT Press.
- Keil, F. C., Stein, C., Webb, L., Billings, V. D., & Rozenblit, L. (in press). Discerning the Division of Cognitive Labor: An Emerging Understanding of How Knowledge is Clustered in Other Minds. *Cognitive Science*.
- Kripke, S. (1980). *Naming and Necessity*. Cambridge, MA: Harvard U.P.
- Lafferty, J.M. (1993). History of Vacuum Science: A Visual Aids Project. In P. A. Redhead (Ed.) *Vacuum Science and Technology; Pioneers of the 20th Century*, pp. 91-106. New York: Springer.
- Lloyde, J., Braisby, N. & Brace, N (2003). *Seeds of Doubt: Are Children Taught to be*

- Essentialists? In: F. Schmalhofer, R. M. Young & G. Katz (Eds.), *Proceedings of EuroCogSci03*, pp. 211-216. Mahwah, NJ: Erlbaum.
- Mandler, J. M., & McDonough, L. (1993). Concept formation in infancy. *Cognitive Development*, 8, 291-318.
- Medin, D. L., Lynch, E. B., & Coley, J. D. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology*, 32, 49-96.
- Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning* (pp. 179-195). Cambridge: Cambridge University Press.
- Murphy, G. L. (2002). *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.
- Osherson, D. N., & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 11, 35-58.
- Osherson, D. N., & Smith, E. E. (1982). Gradedness and conceptual conjunction. *Cognition*, 12, 299-318.
- Osherson, D. N., & Smith, E. E. (1997). On typicality and vagueness. *Cognition*, 64, 189-206.
- Proctor, C., & Keil, F. C. (2006). Differences in Deference: Essences and Insights. In R. Sun & N. Miyake (Eds.), *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (pp. 1974-1979). Mahwah, NJ: Lawrence Erlbaum.
- Putnam, H. (1975). The meaning of 'meaning'. In H. Putnam (Ed.), *Mind, Language, and Reality: Philosophical Papers, Volume 2* (pp. 215-271). Cambridge: Cambridge U.P.
- Quinn, P. C., & Johnson, M. H. (2000). Global-before-basic object categorization in connectionist networks and 2-month-old infants. *Infancy*, 1, 31-46.
- Randall, R. A. (1976). How tall is a taxonomic tree? Some evidence for dwarfism. *American Ethnologist*, 3, 543-553.

- Rey, G. (1983). Concepts and stereotypes. *Cognition*, 15, 237-262.
- Rips, L. J. (1989). Similarity, typicality and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning* (pp. 21-59). Cambridge: Cambridge University Press.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition*. Cambridge, MA: MIT Press.
- Rosch, E. R., & Mervis, C. B. (1975). Family resemblances: studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.
- Rumelhart, D. E. (1990). Brain style computation: Learning and generalization. In S. F. Zornetzer, J. L. Davis, and C. Lau (Eds.), *An Introduction to Neural and Electronic Networks*, pp. 405-420. San Diego, CA: Academic Press.
- Shaffir, E., Smith, E. E., & Osherson, D. N. (1990). Typicality and reasoning fallacies. *Memory & Cognition*, 18, 229-239.
- Simmons, C. L., & Hampton, J. A. (2006) Essentialist Beliefs about Basic and Superordinate Level Categories. Poster presented at the Annual Convention of the Psychonomic Society, Houston Texas, November.
- Sloman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, 22, 189-228.
- Smith, E. E., & Medin, D. L. (1981). *Categories and Concepts*. Cambridge, MA: Harvard University Press.
- Smits, T., Storms, G., Rosseel, Y., & De Boeck, P. (2002). Fruits and vegetables categorized: An application of the generalized context model. *Psychonomic Bulletin and Review*, 9, 836-844.
- Storms, G., De Boeck, P., & Ruts, W. (2000). Prototype and exemplar based information in natural language categories. *Journal of Memory and Language*, 42, 51-73.
- Storms, G., De Boeck, P., & Ruts, W. (2001). Categorization of unknown stimuli in well-known natural concepts: A case study. *Psychonomic Bulletin and Review*, 8, 377-384.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293-315.

Warrington, E. K. (1975). Selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology*, 27, 635-657.

Zadeh, L. (1965). Fuzzy sets. *Information and control*, 8, 338-353.