Concepts and Meaning:

Introduction to Special Issue on Conceptual Representation

James Hampton[1] and Helen Moss[2]

[1]City University, London, UK

[2]University of Cambridge, UK

Conceptual representation is arguably the most important cognitive function in humans. It stands at the centre of the information processing flow, with input from perceptual modules of differing kinds, and is centrally involved in memory, speech, planning, decision-making, actions, inductive inferences and much more besides.

It is therefore unsurprising that it is also a domain of interest to the full range of cognitive sciences – linguistics via lexical semantics, psychology through the use of concepts in thought and categorization and children's acquisition of concepts and word meanings, AI through the development of systems for knowledge representation, neuroscience through the recent development of interest in dissociations between knowledge domains and the role of perceptual and motor areas in concept representation, and finally philosophy which originally began the whole process of trying to find the basic building blocks of thought and knowledge. In this collection, we set our authors the task of attempting to draw together current thinking in their own field, and to lay out their views on the importance of their particular approach to concept representation. The result has been a series of papers that by and large have been broader and possibly more speculative than would normally appear in a top peer-reviewed journal such as *Language and Cognitive Processes*. We believe that our authors have responded with considerable courage to our call to consider a bigger picture, and we are delighted with the resulting articles.

In this introductory piece we aim to set the scene for the papers that appear later. It is necessarily a rather artificial task to take our eight papers and to cast them into a single framework. However, we were struck by the number of common issues and links that we found.

A major theme that emerges from the papers in this collection is the issue of the relation of thought to language – or concepts to meanings. **Wisniewski, Lamb and Middleton** present a review of a large number of studies investigating what initially appears to be a purely linguistic phenomenon – the distinction in many languages between count nouns such as *bicycle* and mass nouns such as *spinach*. The former can be counted (two bicycles), whereas the latter are modified by quantifiers suggesting they are undifferentiated masses (a lot of spinach). It has been argued that the distinction is largely a matter of convention, varying from language to language in an unpredictable way. For example, in English, *rain* is a mass noun but *shower* a count noun, *rice* a mass noun but *lentil* a count noun. In reviewing their evidence, Wisniewski et al. reveal some of the many subtle cognitive dimensions underlying the syntactic distinction. There are, to be sure, some specifically linguistic effects, often dependent on historical change. However, they show in

a series of experiments that, for example, people typically interact with the referents of mass (e.g. furniture) as opposed to count (e.g. vehicle) superordinate nouns in different ways, and that they consider object parts to be more important properties of count superordinates than of mass superordinates. Consequently Wisniewski et al. suggest that people conceptualise the two forms of superordinate concepts in different ways.

The syntactic forms of a language, then, can be shown to bear strong relations to conceptual distinctions in thought. The reverse conclusion is proposed by **Levy** in his review of Wittgenstein's Private Language Argument (PLA), and its implications for the philosophy and psychology of concepts. In an argument that is reminiscent of both functionalism and social constructionism, the PLA purports to show that the determination of the content of a person's internal representation of a concept is itself dependent on a publicly shared language. Levy presents different possible views of how we might have concepts that enable us to *think the same thoughts* as others – something which seems a desirable property for any theory of concepts. He describes Peacocke's theory of concept possession in which the individuation of concepts is achieved by determining the set of predispositions and abilities that a person must show in order for them to be deemed to possess the concept (Peacocke, 1992). To possess the concept of *bicycle* is to be able to do a number of things – for example, recognise one, know what it is for, how it works. Levy then provides an account of Fodor's theory of concepts as atomic symbols representing external properties (Fodor, 1988, 2000). He argues that both of these influential theories of concepts fail fully to account for the problems raised in Wittgenstein's PLA. He concludes by tentatively endorsing Dummett's Priority thesis that the "order of explanation" must take language first and concepts second. It follows that it will not be possible to provide an account of, say, concept acquisition, without first providing an account of language acquisition (Dummett, 1991). It is only through study of the usage of terms in a public language that we can have an independent way of fixing the contents of people's concepts.

The way in which concepts and word meanings are acquired is the key question in the paper by **Diesendruck**. He reviews a range of literature on the way in which children learn to label, categorize and reason about the world. From his review he draws the conclusion that there may be strong domain differences in the way that one might apply a principle like the Language Priority thesis. He finds that whereas artefact kinds such as tools, furniture and vehicles are highly susceptible to linguistic and cultural factors in learning, natural kinds like animals are relatively immune to such effects. Diesendruck notes that this evidence is consistent with a modular account of domain-specific learning, in which a naïve essentialism is applied very early on to the acquisition of categories of animals. Children understand very quickly that animal kinds

are based on deep causal principles and not on superficial appearance. Interestingly, the categorization of people appears to fall somewhere between animals and artefacts in the scale of essentialist thinking.

The issue of differences in representation and processing across conceptual domains is clearly one that arises in a number of the contributions to this volume (see for example the cognitive neuropsychological papers described later). **Sloman and Malt** provide a challenging paper in which they argue that artefacts should not be considered as "kinds" in the way that natural kinds are. They consider (among other positions) arguments for the categorization of artefact kinds being based on a belief in a common essence – such as the function that the object was intended to have when it was created. They claim that the only non-circular way to test such a notion is in relation to the linguistically determined name categories – the classes of objects that get called by the same name. But the cross-linguistic diversity and the context-dependent nature of naming is not consistent with any direct or stable relation between classes defined by intended function and those defined by naming data. They draw the conclusion that artefacts do not come grouped into stable kinds. Their paper therefore sits well alongside Levy's in the sense that both raise important issues about the relation between concepts and language. Levy questions the feasibility of determining conceptual content in the absence of a prior analysis of language use, whereas Sloman and Malt argue that in the case of artefacts different tasks (of which naming is one) will lead to different ways of creating conceptual groupings.

Keil's paper is concerned with an important aspect of conceptual thinking – the role of causal notions in the contents of our concepts. It has become a commonplace criticism of the simple similarity-based views of concept structure proposed by prototype or exemplar theories that our concepts play a crucial role in our theoretical understanding and explanation of the world and how it works. This "theory theory" was introduced in a seminal paper by Murphy and Medin (1985), and has had a particularly strong following among developmentalists. Demonstrations of causal theory effects in young children have ironically sometimes given the impression that young children have far more sophisticated concepts than do adults. Keil starts to unravel this puzzle, and to flesh out the often vague proposals of the theory theory. He presents a set of experiments introducing the Illusion of Explanatory Depth (IOED), in which he shows that people confidently believe they know how things work, but when challenged are forced to acknowledge that their understanding is superficial and even incoherent. Our concepts cannot therefore be like the concepts in scientific theories in which the definition and role of the concepts is clearly specified in axiomatic fashion. Keil suggests that what remains of the theory theory notion is that we have deep knowledge of a higher but coarser level of causal information. For example, children quite quickly develop an understanding of what dimensions of a domain are

likely to prove important for classification. They learn, for example, that colour is important in differentiating plants, but not in differentiating cars. They also develop a sophisticated understanding of the way in which human knowledge is divided up into domains of expertise. Knowing whom to ask if you have a desire to know more about a concept is a key part of the representation of the conceptual domain itself. Once again, domains of knowledge are a central theme here, as different dimensions and different forms of expertise are relevant in each domain.

Several of the papers in this volume address the issue of the relationship between conceptual representations and sensory-motor systems. Concepts allow us to access knowledge about the physical attributes of objects in the world, including colour, shape, motion, sound and texture as well as the actions we associate with them. Nevertheless, there is a traditional view in the fields of cognitive psychology and neuropsychology, that conceptual representations are different in kind from those computed within the perceptual input systems and motor output systems that feed into and out of them; that concepts are amodal symbolic representations, abstracted from their modality-specific perceptual bases. This position is challenged by **Barsalou** who develops the thesis that there is a common representational format for conceptual and sensory-motor processing, as part of his *situated simulation account.* He argues that situated simulation is dynamic, context-dependent and goal-driven. Barsalou marshals an impressive array of data from behavioural experiments, lesion studies and functional neuro-imaging studies in support of his claim that perceptual simulations represent concepts. Barsalou and colleagues have also conducted several studies using property generation and verification paradigms, which suggest an ubiquitous role for perceptual variables (such as property shape and size) in conceptual processing. In one study, for example, they found that subjects were slower to verify the properties of objects that would normally be occluded from view (e.g. the roots of a lawn), but that this was not the case if the concept was presented in a linguistic context that revealed those properties (e.g. *rolled-up lawn*).

Barsalou also interprets evidence from functional neuro-imaging studies of conceptual processing within the situated simulation framework. Many studies have found that areas of the brain adjacent to specific sensory-motor systems are differentially activated as a function of the conceptual category or type of property being processed; for example, visual areas are relatively more activated when people process objects for which the visual modality is important (e.g. animals), while certain motor areas are activated when people process motor-related objects such as tools. As Barsalou points out, the areas activated are not necessarily the primary sensory-motor areas that are involved in perception and action themselves, but rather adjacent areas, which leaves open a number of interpretations. Some studies have also interpreted domain effects in

terms of the nature of conceptual processing required rather than the sensory-motor contents of the representations *per se* (e.g. Tyler & Moss, 2001).

Studies of neuropsychological patients with impairments to the conceptual system provide further insights into the format and organisation of conceptual knowledge and its relation to sensory-motor systems. Barsalou suggests an integrative framework that reconciles apparently conflicting hypotheses about category and modality specific conceptual deficits. Essentially he suggests a series of convergence zones (c.f. Damasio & Damasio, 1994) which integrate increasingly widely distributed modality-specific representations located within the sensory-motor systems of the brain. The importance of neuropsychological data in theoretical development of these issues is reflected in the final two papers in this volume; **Rogers, Hodges, Lambon Ralph & Patterson** and **Saffran, Coslett, Martin & Boronat.** These two papers also address the issue of the relation between conceptual knowledge and sensory systems, and echo many of the points raised by Barsalou. First, Saffran *et al* put forward a view that has become quite widely accepted in the neuropsychological literature; that conceptual knowledge is distributed across a network of modality-specific attribute systems determined by the mode of input of acquisition. For example, there are held to be visual, auditory, verbal and kinaesthetic systems. This proposal has several elements in common with Barsalou's account, although it is framed within a more traditional semantic memory system, in which the attribute systems *may* be seen as amodal redescriptions of sensory-motor properties. Nevertheless, different parts of the conceptual representation are claimed to be preferentially accessed by different input modalities, and crucially the attribute systems are held to be localised in different regions of the brain and so may be selectively impaired by brain damage.

In support of their theoretical claims Saffran *et al* present data from a patient, BA who has a progressive disorder of conceptual knowledge, with a significantly greater difficulty accessing conceptual information from words than from pictures. The interpretation of BA's deficit highlights one way that Saffran *et al's* position diverges from that of Barsalou. Saffran *et al* propose that BA's impairment on semantic tasks with words is attributable to selective damage to a propositional/encyclopaedic attribute sub-system; words are held to initially access this type of conceptual knowledge (since they were the mode of input responsible for its acquisition), whereas pictures primarily access the visual store. Barsalou's situated simulation account does not incorporate the idea that verbal inputs preferentially access a propositional semantic system, since words are claimed to involve sensory-motor simulations in just the same way as pictures. Indeed, there would be no separate propositional sub-system.

The nature of the conceptual/perceptual interface is also the central issue addressed by **Rogers et al.** In a similar vein to Barsalou, Rogers *et al* distinguish between two broad views of the conceptual/perceptual relationship. One is a multi-modal distributed network, in which conceptual representations emerge from the associations among the representations that subserve perception, recognition and action across different modalities. This is contrasted with the "traditional" view, in which there is a central conceptual system, separate to "pre-semantic" modality-specific input systems, such as structural descriptions for object recognition, or lexical representations in the language system. In their exposition, Rogers *et al* refer to the former approach as a *process-based* account of conceptual information – semantic memory serves the function of associating sets of modality specific representations, while the latter account is described as a *content-based* approach, in that semantic memory is defined as having content or meaning, while non-semantic representations such as structural descriptions (i.e. stored representations of object shapes) do not. The content-based view of the conceptual system has been supported by reports of neuropsychological patients with apparent dissociations between impaired semantic memory and intact structural descriptions, as revealed in tasks such as object decision – where patients are asked to distinguish between line drawings of real versus chimeric objects or creatures. This has been interpreted as evidence for the functional (and perhaps neural) separation of conceptual and perceptual representations. Rogers *et al* present data from a set of novel object decision tasks, challenging this apparent dissociation. When materials are appropriately controlled for object typicality, patients with semantic deficits do show problems in object decision; they tend to accept as real those chimeric objects which have features typical of members of the category, and reject real objects with atypical features (e.g. they may reject a drawing of a camel as being a non-object, because it has an atypical feature - a hump-, but accept a chimeric picture of a camel with its hump removed, since this produces a more typical animal shape).

Thus there is considerable common ground across this group of three papers concerning the interface between conceptual and sensory-motor systems. All three challenge the view that there is a single static system of amodal conceptual representations. All three propose the distribution of conceptual knowledge across multiple sets of modality-specific attributes, which may be neurally as well as functionally distinct. And all three seem to share the assumption that concepts emerge from the interactive activation or simultaneous simulation of multiple modalities of information, perhaps implemented in terms of cross-modal convergence zones. As reviewed in each of these papers, many sources of psychological, neuropsychological and neuro-imaging data are consistent with these general claims.

However, there remain critical differences between the approaches taken in the three papers. For example, although Saffran *et al* and Rogers *et al* both claim that conceptual knowledge is distributed over multiple modality-specific attributes, rather than residing in a central amodal semantic system, they differ in the nature of the proposed relation between structural and semantic systems. Saffran *et al's* framework suggests a number of semantic subsystems, each storing semantic properties relevant to a specific modality (visual semantics, tactile semantics etc), with a separate level of pre-semantic perceptual systems corresponding to each modality (e.g. structural descriptions for visual objects). For Rogers et al, there is no such distinction between a semantic and structural level of representation; rather there is a unitary semantic system - better described as a process or function than a level of representation - which consists of the associations among the modality-specific structural or perceptual representations. Moreover, the Saffran *et al* framework is consistent with the view that the sensory-motor properties within each semantic subsystem are represented as amodal symbolic redescriptions, distinct from the modality-specific representational format of the input/output systems to which they are linked. In this respect it is very different to Barsalou's situated simulation model.

In summary, the eight papers in this volume address many central issues in the study of conceptual representations. We hope that the juxtaposition of these papers in the current volume - tackling related questions but from very different perspectives - will prompt debate as to whether it is the similarities or the differences among the approaches that are the more significant.

References

Damasio, A.R., & Damasio, H. (1994). Cortical systems for retreival of concrete knowledge: The convergence zone framework. In C. Koch & J.L. Davis (Eds.), Large-scale neuronal theories of the brain (pp. 61-74). Cambridge, MA: MIT Press.

Dummett, M. (1991). *Frege and other philosophers.* Oxford: Oxford University Press.

Fodor, J. (1988). *Concepts: where cognitive science went wrong.* Oxford: Clarendon Press

Fodor, J. (2000). Doing without what's within: Fiona Cowie's critique of nativism*. Mind,* 110, 99-148.

Murphy, G. L., & Medin, D.L. (1985). The role of theories in conceptual coherence. Psychological Review, 92, 289-316.

Peacocke, C. (1992*). A study of concepts.* Cambridge, MA: MIT Press.

Tyler, L.K., & Moss, H.E. (2001). Towards a distributed account of conceptual knowledge. *Trends in Cognitive Sciences, 5.6, 244-252.*