Psychology of Language and Communication 1997, Vol. 1, No. 2

NICK BRAISBY
London Guildhall University
BRADLEY FRANKS
London School of Economics

WHAT DOES WORD USE TELL US ABOUT CONCEPTUAL CONTENT?*

We suggest that different orientations on concepts have arisen from an emphasis on their different functions. An emphasis on the classification function of concepts leads to the theory-laden view, while an emphasis on their role in language understanding points in a different direction. Both adhere to the assumption that word use directly indicates conceptual content. We argue that Donnellan's (1966) distinction between different kinds of word use (the referential/attributive distinction) undermines this assumption, not only for definite and indefinite descriptions, but also for general terms. Indeed, the distinction raises a number of difficulties for the interpretation of extant empirical results (e.g., forced-choice categorisation, typicality ratings, attribute listings, and others), since designs have not attempted to control for the possibility of referential uses. We build on the view that the referential/attributive distinction can be traced to a speaker's distinct referring intentions in the perspectival view of concepts that we present. We briefly present a comparison between perspectives and common-sense theories, before presenting a range of experimental refinements that the perspectival view suggests may be required in order to tap conceptual content. The perspectival view also offers the potential for addressing some of the outstanding theoretical difficulties faced by the different orientations on concepts, including holism, concept combination and conceptual coherence.

Introduction

We argue that theoretical and methodological problems in the study of concepts may be traced to a failure to take into account the pragmatics of word use. The paper is structured as follows. We first sketch some difficulties with extant orientations on concepts, traceable to their emphasis on different roles of concepts in cognition. Despite their emphasising different facets of the exercise of concepts, these orientations nonetheless have a common assumption – that word use directly indicates conceptual con-

* Order of authors is alphabetical. We are grateful to Larry Barsalou, Robin Cooper, Dedre Gentner, Maciej Haman, James Hampton, Iven van Mechelen, Helen Moss, and the participants at the Workshop on Permanence and Change in Conceptual Knowledge (Kazimierz, Poland, November 1995) for discussions that have helped to give shape to the ideas presented here.

Correspondence should be addressed to: Nick Braisby, Department of Psychology, London Guildhall University, Old Castle Street, London, E1 7NT, UK. Tel: +44 171 320 1086. Fax: +44 171 320 1117. E-mail: braisby@unixa.lgu.ac.uk.

tent. This assumption is challenged by the referential/attributive distinction (RAD) made between different kinds of word use by Donnellan (1966), a distinction that we believe can be generalised to those kinds of referring expressions that have been the focus of empirical and theoretical approaches to concepts. We then demonstrate how the received interpretations of many empirical techniques (e.g., typicality ratings, attribute listing, forced choice categorisation) may be compromised by their lack of attention to the RAD in subjects' responses. Building on the view that RAD derives from different referring intentions underlying word use, we develop a framework for expressing the relationship between word use and conceptual content which revolves around perspectives. After sketching this framework, we finally outline some ways in which experimental investigations into concepts might be refined so as to be more sensitive to pragmatics.

2. Orientations on concepts and (some of) their difficulties

Research into concepts has spawned a range of theories with differing emphases and motivations; however, most theories appear to fall into one of two orientations, depending on their focus on different functions of concepts in cognition. One orientation takes it that concepts represent knowledge of categories, and as a result are inextricably connected with our common-sense theories about a category's domain. This view also stresses the role of concepts in classification, and attempts to respect findings concerning the context-dependence of classification by, for example, appeal to different theories underlying different acts of classification, giving rise to the possibility that concepts themselves may be unstable. The second orientation takes it that concepts represent the meanings of words, and so their concatenation provides for the interpretation of utterances and sentences. This view appears to prompt a characterisation of concepts as relatively stable, since the assumption of iterable content may be required to explain both the relative speed of language processing and the productivity of meaning.

Not only is there an obvious tension between these two orientations, but each alone has independent difficulties. A primary motivation for the theory-laden view derives from the view that similarity – a foundation of earlier views of concepts – is a three-place relation rather than a two-place relation, holding between the two things to be judged for similarity and also the attributes or "respects" relative to which the judgement is made (Goodman, 1972; Medin, Goldstone, & Gentner, 1993). Common-sense theories are then taken to constrain the selection of those respects (Murphy & Medin, 1985). Two principal problems arise. First, common-sense theories are not the kinds of structure that enter readily into combinations, hence rendering problematic an account of complex concepts. Second, if concepts are inextricably linked with common-sense theories, which are also interconnected, this invites holism (Fodor, 1994). As such, there are no obvious constraints on conceptual content, and so no obvious constraints on the respects that enter into similarity judgements.

The language-based view is also confronted by difficulties, inasmuch as its assumption of stable conceptual content renders problematic providing accounts of the apparent dependence of classification on, for example, linguistic context (Roth & Shoben, 1983), domain (Gelman & Coley, 1991) and social contexts (Barsalou & Sewell, 1984).

3. The Relationship between Word Use and Concepts

In spite of the differences between these two orientations on concepts, they are both committed to the assumption that word use directly reveals the nature of conceptual content, and it is on this assumption that most experimental investigations of concepts are predicated[1]. As we shall see, however, this assumption is directly challenged by discussions of the RAD, and this necessitates a reinterpretation of many empirical results. Our viewpoint is that the two orientations can be at least partially reconciled once the relation between word use and conceptual content is recognised to be conditioned by pragmatic factors.

Donnellan (1966) first drew the RAD with regard to the use of definite descriptions (descriptions of the form 'the so-and-so...'). He claimed that these could be used in two distinct ways: attributively, to refer to whatever satisfies the description; and referentially, to 'pick out' some referent or referents in mind. He gives the example of someone who has usurped the throne nonetheless being referred to as 'the king' even though all concerned know the individual concerned not to be the true king. It is thus of the nature of referential uses that they can involve reference to things outside of a term's true denotation, even when the speaker and/or hearer know or believe these things to lie outside of the denotation. What referential uses demonstrate, then, is that the use of a word need not indicate that word's denotation, and hence word use need not indicate stable beliefs concerning category membership, i.e., conceptual content. Ludlow & Neale (1991) generalised the RAD to indefinite descriptions, claiming that these may be used either quantificationally (akin to the attributive) or specifically (akin to the referential). A specific use then may involve reference to a specific thing that the speaker has in mind and wishes to 'pick out,' even though the thing may not be physically present in the current context. Critically, the thing so referred to need not lie inside the term's denotation, and may be known or believed by speaker and hearer to lie outside of the denotation. We claim that a similar generalisation of the RAD holds for general terms (i.e., those that refer to categories), such as 'lions get fed at tea-time', including their use in predicate position, such as 'prawns are a type of fish'. The RAD here reflects the use of the general term to refer to either whatever kind constitutes the term's denotation (an 'attributive' use), or to some other kind or sub-kind that the speaker has in mind, and which may well be disjoint from the denotation. This generalisation receives substantiation from the empirical results that suggest that people may deploy two modes of categorisation – one based on denotation-determining attributes and one based on appearance attributes (Smith & Sloman, 1994).

Key to an understanding of the RAD and our analysis of concepts presented later is Kripke's (1977) suggestion that the RAD is essentially a pragmatic distinction, between different kinds of referring intention. Kripke distinguishes two kinds of intention: general intentions to refer to entities falling under a term's denotation, and specific intentions to refer on a particular occasion of use to some specific object. The referential, Kripke then claims, reflects a divergence between the specific and the general, so that the object picked out may not in fact fall under the denotation; the attributive reflects their convergence, so that the specific intention simply is to refer to whatever falls un-

[1] There is a caveat, of course, to this assumption, and this is that only conventional word use indicates conceptual content – metaphorical and analogical uses are at best only indirectly related to conceptual content.

der the denotation. The explication of the RAD in terms of the speaker's intentions renders the distinction as being pragmatic, one that cannot be captured solely in terms of the semantic properties of words and their referents (see also Bach, 1987; Recanati, 1987). Moreover, since Kripke claims the crucial distinction between the two types of use reflects divergence or convergence of intentions, one can think of the two types of use as reflecting opposing endpoints on a continuum. That is, uses may be more or less attributive or referential.

The different manifestations of the RAD converge on a single moral: word use is not universally indicative of conceptual content. In short, only when it is certain that a person is using a word or expression attributively, can that use support a direct inference to underlying conceptual content. If the use is a referential one, then there can be no guarantee that a speaker's word use illuminates his or her beliefs about the true nature of the object or category referred to.

This conclusion results in a number of concerns regarding existing experimental techniques. First, how can we be sure that the word use that has been taken to reveal conceptual content is in fact truly attributive, and not referential? Second, how can we be sure that subjects in experiments use and interpret words as referring to their denotations, and not as picking out other individuals or kinds 'in mind'? Third, is there any evidence that existing techniques tap into word use that is not attributive, but in fact referential?

4. The RAD and Empirical Findings on Concepts

An influential study, often cited as support for 'fuzzy' views of concepts and classification, is that of McCloskey & Glucksberg (1978). In a forced-choice categorisation task (in which subjects had to give a response of 'Yes' or 'No' to putative classifications), most subjects categorised squid and shrimp as fish, and almost 50% categorised lobster, octopus, and clam as fish; similarly, cucumber, eggplant, and squash were predominantly categorised as not being fruit. In contrast to the interpretation in terms of fuzziness offered by McCloskey & Glucksberg (which presupposes that all of the interpretations of general terms by subjects are attributive), we suggest that these findings are explicable on the view that subjects' interpretation of fish and fruit is referential, and not reflecting their stable beliefs concerning membership of these kinds. One interpretation may be that the kind of word use subjects employ depends on the audience that they take the experimenter to have in mind: for a scientifically minded audience, these uses would be inappropriate (a squash is a fruit), while for an audience whose categorisation serves a different purpose (e.g., a culinary one) they may be deemed appropriate.

Similar possibilities may hold for typicality ratings (where subjects are asked to rate on a scale how typical or representative putative category members are, e.g., Rosch, 1975; Armstrong, Gleitman & Gleitman, 1983). An intuitive ordering for typicality may be thought to arise from the following cross-classification of uses.

		Attributive	
		Yes	No
Referential	Yes	*High typicality*	*Intermediate typicality?*
	No	*Intermediate typicality?*	*Low typicality*

That is, if a categorisation statement ('An X is a Y') can be assented to on both attributive and referential interpretations of the terms X and Y, such a statement is likely to lead to a high typicality rating (of X in the category Y). And so on. Contrary to standard interpretations of intermediate typicality in terms of fuzziness of category membership, this view suggests a pragmatic explanation. Intermediate typicality ratings could potentially reflect one of two possibilities. Clear category members (satisfying attributive uses) may be rated intermediate if they fail to satisfy referential uses (e.g., a cucumber is a fruit, but a use of fruit may not enable a hearer to pick out a cucumber). Alternatively, a category non-member, not satisfying an attributive use, might nonetheless satisfy a referential use – perhaps where a category non-member nonetheless bears certain similarities, either perceptual or functional, to category members (e.g., a clam being classified as a fish). On this view, then, typicality ratings do not solely reflect attributive word use, hence category membership, hence conceptual content. Moreover this reinterpretation cannot be dismissed lightly, since there is no evidence that subjects were not employing referential interpretations of the general terms involved.

On similar grounds, it may be doubted that attribute listing experiments (where subjects are asked to list central attributes of categories, e.g., Rosch, 1975; Hampton, 1987) reveal conceptual content. Again, it is possible that subjects' interpretations of category terms and/or the attribute labels may not be solely attributive. A related possibility has already been raised by Tversky & Hemenway (1984). Their claim is that attributes are always generated with regard to an implicit contrast set, i.e., that attribute terms may be used simply to 'pick out' attributes from such a set, regardless of the accuracy of the term.

Given the potential for uncertainty concerning subjects' uses and interpretations, we see no reason to believe that the three techniques mentioned above are the only ones to be subject to the kinds of reinterpretation we have offered. In principle, any technique that employs word use as a means for uncovering or controlling for conceptual content, and which fails adequately to control for the pragmatic factors governing word use, will be subject to these kinds of concern.

5. The Perspectival View

Perspectives act as bridges between stable common-sense theories and the content associated with words on particular occasions of use. That is, within a given perspective, only some of the content expressed in the common-sense theory is associated with word use. Lexical concepts are posited, and are taken to express default conceptual content (some regularly accessed and deployed subset of attributive content). Senses for a term on particular occasions of use are (possibly non-monotonically) derived from the lexical concept together with lexical concepts associated with common-sense theories (Braisby, Franks & Myers, 1992). The access of these additional lexical concepts is governed by the perspective which the speaker adopts in classification, and which the hearer also attempts to adopt in arriving at an appropriate interpretation of the speaker's utterance (Braisby, 1990; Franks & Braisby, 1990; .Franks, 1991, 1995).

Perspectives have three components. First, perspectives are determined in part by the kind of referring intention adopted by the speaker. Although there are presumably a wide variety of possible referring intentions that a speaker may adopt, the general and specific intentions that Kripke considers to underlie referential and attributive uses are two obvious candidates. Second, perspectives are determined in part by the grounds available to

the speaker and hearer. The grounds constitute the facts concerning the context of discourse, including the immediate physical context, which are available to the speaker and hearer to inform and structure their communication. In particular, the grounds must provide support for particular referring intentions, i.e., a situation must contain a referent for the speaker to entertain an intention to refer (to that referent). Further, since perspectives act as bridges between the lexical concept and the sense generated in particular contexts, then the intention to refer to a particular object with a given term requires that a sense be generated for the term which represents properties possessed by the object. Third, perspectives are partly determined by the shared beliefs and suppositions of the speaker and hearer. For example, the way in which a speaker may classify a clam as a fish or as a non-fish may depend on the extent to which the hearer can be assumed to have knowledge of the criteria for classifying fish in different ways, e.g., biological, functional, culinary, etc.[2]

The following classification specifies some points of difference between perspectives and common-sense theories.

Common-sense theories	Perspectives
1) General across situations	1) Specific to types of situations
2) Wholly internal representations	2) Partly external/situational
3) Holistic	3) Non-holistic (restricted)
4) Audience-independent	4) Audience-relative
5) Relative stability across contexts	5) Instability across types of situation
6) "Semantic" (representations of objects/properties of objects)	6) "Pragmatic" (representations of intentions and communicative goals)
7) Explanations-based (causal relations between properties of objects)	7) Justification-based (support relation between components of properties)

The functional relations between the three components of perspectives will differ as between speakers and hearers. Speakers are presumed to execute classification, hence selection of a lexical concept and, where relevant other lexical concepts, in an appropriate mechanism for generating a sense for a term so as to satisfy their referring intentions. Hence the referring intentions are seen as determining the lexical concepts to be selected and the sense generated. Constraining this generation process, and the lexical concepts selected, are the grounds and shared beliefs. These functional relations are depicted in Figure 1.

The functional relations between these components are necessarily partially inverted for hearers. The hearer has, in general, no privileged access to the speaker's referring intentions: rather, these must be inferred on the basis of grounds and shared beliefs. If the hearer can assume that the speaker and hearer share roughly the same lexical concept, and the

[2] There are many antecedents to the perspectival view and much research from which we have drawn inspiration. We would particularly like to mention the works of Frege (1892) and Herb Clark (e.g., Clark, 1983; Clark & Marshall, 1983).

Figure 1. Hypothesized functional relations between concepts and perspectives in language production (i.e., for the speaker/utterer).

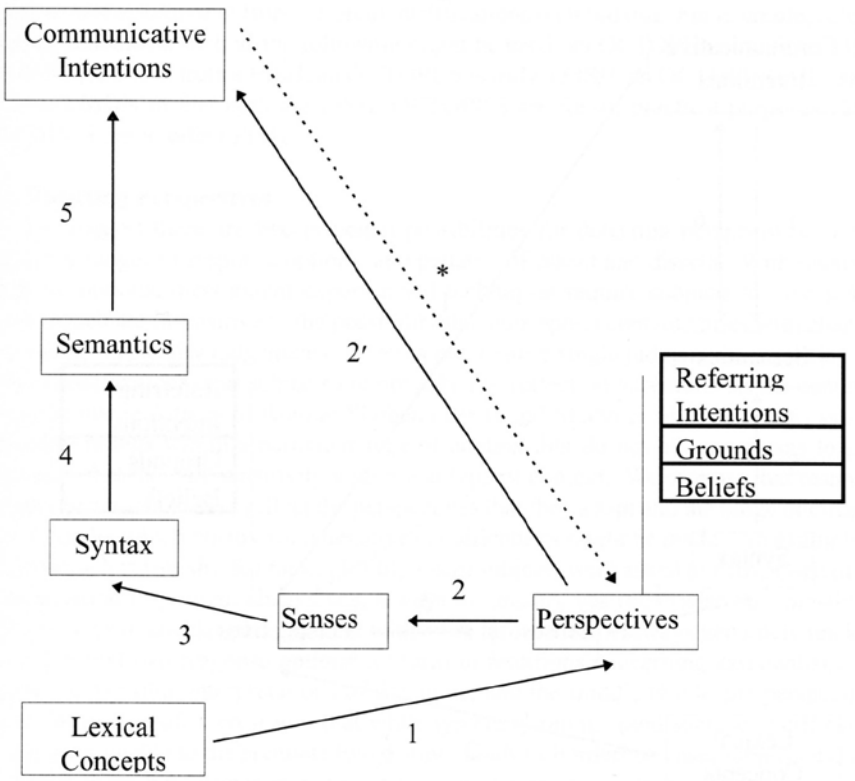

Note: 2 and 2¢refer to information hypothesized to flow concurrently. * refers to optional information flow which will occur if informational content resulting from 5 mismatches the perspective indicated by 2¢

content of this mismatches with the properties of what is the best candidate for the referent on the basis of grounds and shared beliefs, then the hearer may infer the speaker intends a referential use. Of course, the hearer may also adjust the assumption of shared social category membership (shared beliefs) which may then prompt the inference that the speaker and hearer do not share lexical concepts for the term, and perhaps that the speaker may have intended an attributive use on the basis of his or her lexical concept. In practice, the balance of evidence in favour of inferring one kind of intention or other depends upon the satisfaction of multiple constraints. These functional relations are depicted in Figure 2.

6. Addressing methodological problems concerning word use

There are a number of means by which we believe experimental attempts to uncover conceptual content may respect the methodological considerations concerning word use that we have outlined. Broadly, these fall into two types of refinement of experimental

Figure 2. Hypothesized functional relations between concepts and perspectives in language comprehension (i.e., for the hearer/audience)

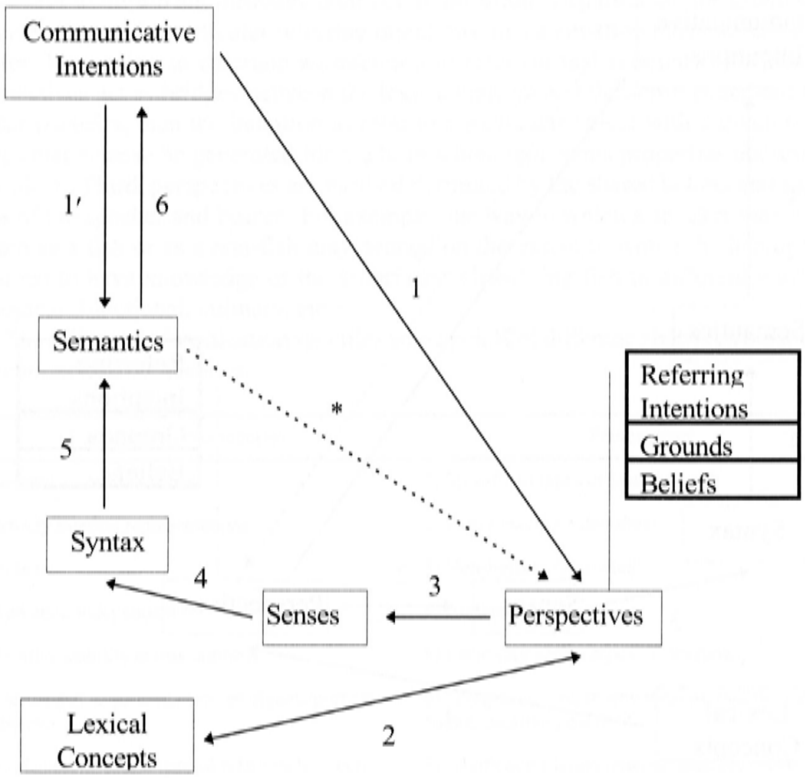

Note: 1 and 1¢refer to information hypothesized to flow concurrently. * refers to optional information flow which will occur if informational contents resulting from 5 and 1¢mismatch.

techniques: the first concerns attempts to manipulate perspectives explicitly; the second concerns attempts to detect the perspectives that subjects employ in deploying their concepts.

6.1. Manipulating perspectives

We wish to raise two possibilities for manipulating perspectives: audience factors and experimental instructions. Audience factors concern the possibility that one experimentally controls for the kind of audience with regard to which a classification might be made. the supposition is that referential uses may be made in order to help the audience to 'pick out' a particular referent. For example, a biologist might refer to a whale as a large fish relative to a naive audience, perhaps for the purpose of 'picking out' an object in the context, but as a large mammal for a more sophisticated one. Indeed, the latter classification is more attributive in character. These uses may then in principle be teased apart by instructing subjects to classify objects relative to different audiences (e.g., describing objects to children, to experts, etc.).

Equally, experimental instructions may be manipulated in order to obtain uses that can be located precisely on the referential-attributive continuum. That is, subjects may be asked to assent to/dissent from different qualifications on word use. For example, categorisation statements such as the following might be used: an OCTOPUS is really a FISH, an OCTOPUS is in truth a FISH, an OCTOPUS is truly a FISH, an OCTOPUS is the same as a SALMON in that both are FISH, OCTOPUS are for all practical purposes FISH, OCTOPUS are in effect FISH.

6.2. Detecting Perspectives

We suggest there are two principal possibilities for detecting perspectives: use of expanded ranges of response options, and patterns of assent and dissent. With regard to response options, most extant experimental techniques require subjects to give judgements which are insensitive to the possibility that conceptual content varies with changes in context. Typicality judgements, for example, require single judgements regarding putative classifications, and subjects are not asked to reflect on variations across contexts. Even the manipulations of Roth & Shoben (1983) and Medin & Shoben (1988) ask for typicality ratings within a particular type of context, but do not allow subjects to give responses that indicate sensitivity within that type of context. We propose that response options be chosen so as to reflect the perspectives that they adopt and the range of circumstances under which positive and negative classifications might be made. An example of this approach is Braisby & Franks (1996), where subjects were asked to deliver one of six categorisation responses: always yes, always no, mostly yes (otherwise no), mostly no (otherwise yes), unclear (though clear with more information), always genuinely unclear. Thus, the first two response options conform to intuitions concerning essentialism, the last two to possible interpretations of fuzziness, and the middle two to the perspectival view. What was of interest was that while typicality ratings (and sentence verification times) were unable to differentiate two distinct kinds of borderline case, subjects did differentiate these borderlines in terms of the set of response options. Furthermore, the borderlines attracted many more perspectival response options than either the essentialist or fuzzy options. Though we are unable to report the details of this study, the moral we draw is that studies of concepts must pay considerable attention to the kind of responses that subjects are allowed to make.

The second possibility, patterns of assent and dissent, is inspired by the work of Bach (1987) on the RAD. Bach suggests that referential and attributive uses may be distinguished by patterns of assent and dissent to categorisation statements in the light of (possibly apparently contradictory) further information. For example, if classification of olive as a vegetable reflects a referential use (on the lines suggested earlier), then speakers, on being told that an olive is really a fruit, might argue on the following lines: "I was right to say that olives are vegetables in a functional sense, but wrong to imply that olives are vegetables in a biological sense". In contrast, if the categorisation reflected an (incorrect) attributive use, they might produce in response to being corrected: "I was wrong to say that olives are vegetables".

In keeping with this orientation on subjects' classification in response to further information, and the resulting qualifications on that classification which they will accept, Braisby, Franks & Hampton (1996) examined the intuition underlying the claims of essentialism. Employing scenarios such as Putnam's (1975) putative discovery that all cats

are really robots controlled from Mars, they examined categorisation through assent to or dissent from positive and negative variants of three statement types: existential (cats exist/ cats do not exist), qualified (cats do exist, and people's beliefs concerning cats have changed/ there are no such things as cats, only robots controlled from Mars), membership (Tibby is a cat, though we were wrong about her being a mammal/Tibby is not a cat, though she is a robot controlled from Mars). Subjects' responses diverged from the predictions of essentialism, though only with regard to certain scenario types and certain statement types. An explanation for this is that some statements may be considered from different perspectives, and so need not universally indicate attributive interpretations (and hence provide evidence for or against essentialism or other theories of concepts). Importantly, subjects tended to provide classification judgements for a scenario that appeared contradictory (in effect, claiming that Tibby is a cat and Tibby is not a cat), and which support the perspectival view, on the view that one use of cat is more referential and the other more attributive.

7. Addressing theoretical difficulties in the study of concepts

Earlier we noted several theoretical difficulties in extant orientations on concepts. The first, holism, indicated that the restriction on attributes or 'respects' for comparison might not be forthcoming from common-sense theories. The perspectival approach suggests that appropriate restrictions on content for a task will not be provided by any a priori partitioning of conceptual content (along the lines, say, of the analytic/synthetic distinction). Rather, it is in the nature of perspectives to provide a distinction which depends on the referring intentions of the speaker (inferred by the hearer); depending on the different intentions, under different circumstances a different subset of the content of common-sense theories will be appropriate. In this way, referential intentions, in conjunction with grounds and mutual beliefs serve to constrain the nature and detail of the content of the sense for the referring expression. With regard to similarity, the perspectival view, restricting as it does the potential for holism that might otherwise issue from common-sense theories, also offers the prospect of restricting the 'respects' which are implicated in similarity judgements.

A related account may be offered of the balance between stability and flexibility of conceptual content. The flexibility of conceptual content, as expressed in the context-dependence of senses, need not indicate an instability in underlying representations. This is because perspectives act as the interface between the stable, enduring content of common-sense theories (as reflected, in part, in lexical concepts), and the context-dependence of senses. Moreover, changes to one concept need not entail wide scale revision of all of the content of all concepts and hence of all common-sense theories (as would be predicted by holism). Since the deployment of conceptual content is relativised to perspective, so changes to content will be restricted only to certain contents under perspectives appropriately related to the one in which the change originated.

The perspectival view also has purchase on the difficulties often encountered with regard to concept combination. Many instances of concept combination appear to involve the defeat of attributes inherent in the constituent concepts, and have been thought to require a non-compositional account. Within the perspectival view, some of these difficulties may be lessened by considering that the content of the constituent concepts varies according to perspective, such that within a perspective the content of the constituents may combine in a compositional way.

Lastly, perspectives offer an interesting contrast with the treatment of conceptual coherence in terms of common-sense theories. Such theories are deemed to provide explanatory relations between the attributes of concepts, and also between concepts, these relations providing for conceptual coherence. While we have said little concerning explanation in our exposition of the perspectival view, this is not because we consider explanatory relations unimportant. Indeed, we consider such relations to have a central role in explaining coherence, but adopt the view that everyday explanations are ultimately perspectival. Such a view owes much to the work of van Fraassen (1984), Achinstein (1984) and Putnam (1978), who consider that explanations are purpose- and audience-relative: what counts as an explanation depends upon why, by whom and for whom the explanation is sought. Perspectives can then be considered to provide partial explanations, for specific audiences, (Ruben, 1990) underpinning conceptual content and hence conceptual coherence.

8. Conclusions

Within the literature on concepts, approaches tend to subscribe to one of two orientations: one that emphasises the classificatory functions of concepts, and one that emphasises the role that concepts play in language understanding. While these orientations have led to views of concepts which often appear to contain elements of contradiction, we suggest that they may be reconciled via the central notion of perspectives. The reconciliation requires a reconsideration of experimental techniques that have been used to uncover conceptual content, and that future investigations of concepts explicitly attempt to detect and/or manipulate the perspectives relative to which classifications are made. Only then will inferences to conceptual content be sound.

References

Achinstein, P. (1984). A type of non-causal explanation. In P. French, T. Uehling jr. & H. Wettstein (eds.), *Causation and Causal Theories, Midwest Studies in Philosophy.* Minneapolis: Minnesota University Press.

Armstrong, S. L., Gleitman, L. R., & Gleitman, H. (1983). What some concepts might not be. *Cognition*, 13, 263-308.

Bach, K. (1987). *Thought and Reference*. Oxford: Clarendon Press.

Barsalou, L. W., & Sewell D. R. (1984). Constructing representations of categories from different points of view. Emory Cognition Project Report No. 2. Emory University, Atlanta, GA.

Braisby, N. R. (1990). Situating word meaning. In R. Cooper, K. Mukai, & J. Perry (Eds.), *Situation theory and its applications*, I. CSLI: Stanford.

Braisby, N. R., & Franks B. (1996). Why concepts only appear to be fuzzy. In submission.

Braisby, N. R., Franks, B., & Hampton, J. A. (1996). Essentialism, word use and concepts. *Cognition*, (59), 247–274..

Braisby, N. R., Franks, B., & Myers, T. F. (1992). Partiality and coherence in concept combination. In J. Ezquerro & J. M. Larrazabal (Eds.), *Cognition, semantics and philosophy*. Dordrecht: Kluwer.

Clark, H. H. (1983). Making sense of nonce sense. In G. B. F. d'Arcais & R. J. Jarvella (Eds.), *The process of language understandin*g. Chichester: John Wiley and Sons.

Clark, H. H., & Marshall, C. (1981). Definite reference and mutual knowledge. In A. Joshi, B. Webber, & I. Sag (Eds.), *Elements of discourse understanding*. Cambridge: Cambridge University Press.

Donnellan, K. S. (1966). Reference and definite descriptions. *The Philosophical Review*, 75, 281-304.

Fodor, J.A. (1994). Concepts: A potboiler. *Cognition*, 50, 95-114.

Franks, B. (1991). Sense generation and concept combination. In B. Franks (Ed.), *Word meaning and concepts*. ESPRIT BRA Deliverable R2.4. Centre for Cognitive Science: University of Edinburgh.

Franks, B. (1995). Sense generation: a "quasi-classical" approach to concepts and concept combination. *Cognitive Science*, 19(4), 441–506.

Franks, B., & Braisby, N. R. (1990). Sense generation or how to make a mental lexicon flexible. In *Proceedings of the 12th annual conference of the cognitive science society*. Hillsdale, NJ.: LEA.

Frege G. (1892). On Sense and Meaning. In Geach P. & Black M. (eds.), *Translations from the Philosophical Writings of Gottlob Frege*. Totowa, NJ.: Barnes & Noble Books.

Gelman, S., & Coley, J. (1991). Language and categorization: the acquisition of natural kind terms. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on language and thought*. Cambridge: Cambridge University Press.

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman, *Problems and Projects*. Indianapolis: Bobbs-Merrill.

Hampton, J.A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory and Cognition*, 15, 55-71.

Kripke, S.A. (1977). Speaker's reference and semantic reference. In French, P.A., Uehling, D., & Wettstein, H. (Eds.), *Contemporary Perspectives in the Philosophy of Language*. Minneapolis: University of Minnesota Press.

Ludlow, P., & Neale, S. (1991). Indefinite descriptions: in defense of Russell. *Linguistics and Philosopy*, 14, 171- 202.

Medin, D., Goldstone, R., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, 100 (2), 254-278.

Medin, D. & Shoben, E. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, 20, 207-238.

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

McCloskey, M. & Glucksberg, S. (1978). Natural categories: well-defined or fuzzy sets? *Memory and Cognition*, 6, 462-472.

Putnam, H. (1975). The meaning of "meaning". In *Mind, language and reality*, volume 2: *Philosophical papers*. Cambridge: CUP.

Putnam, H. (1978). *Meaning and the Moral Sciences*. London: Routledge.

Rosch, E.H. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104, 192-233.

Roth, E.M. & Shoben, E.J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, 15, 346-378.

Ruben, D.H. (1990). *Explaining Explanation*. London: Routledge.

Smith, E.E. & Sloman, S.A. (1994). Similarity- versus rule-based categorisation. *Memory & Cognition*, 22(4), 377-386.

Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, 113(2), 169-191.

van Fraassen, B. (1980). *The Scientific Image*. Oxford: Oxford University Press.