Meanings and Prototypes

Studies in linguistic categorization

edited by

S. L. Tsohatzidis

London and New York

Meanings and Prototypes

Studies in linguistic categorization

edited by

S. L. Tsohatzidis

London and New York

First published 1990
by Routledge
11 New Fetter Lane, London EC4P 4EE

Simultaneously published in the USA and Canada
by Routledge
a division of Routledge, Chapman and Hall, Inc.
29 West 35th Street, New York, NY 10001

Typeset in 10/12 pt Times Roman by Linotron 202 at Columns of Reading
Printed in Great Britain by T. J. Press, Padstow, Cornwall

British Library Cataloguing in Publication Data
Meanings and prototypes : studies in linguistic categorization.
1. Linguistics
I. Tsohatzidis, S. L.
410

ISBN 0–415–03612–7

Library of Congress Cataloging in Publication Data
Meanings and prototypes : studies in linguistic categorization /
[edited by] Savas Tsohatzidis.
p. cm.
Includes bibliographical references.
ISBN 0–415–03612–7
1. Categorization (Linguistics) 2. Semantics. I. Tsohatzidis,
Savas
P128.C37M4 1990
401'.43—dc20 89–10910
CIP

Contents

Introduction

S. L. Tsohatzidis

There are fewer distinctions in any natural language than there are distinct things in the universe; if, therefore, the languages people speak are ways of representing the universe their speakers live in, it is not unreasonable to suppose that one central function of the various kinds of element constituting a natural language is to allow the much more varied kinds of thing populating the universe to be categorized in specific ways.

A prototype approach to linguistic categories is a particular way of trying to answer the question as to how this categorization proceeds. It involves two central claims. First, that, for the most part, linguistic categorization exploits principles that are not specific to language but characterize most, if not all, processes of cognition. Second, that one of the basic principles according to which cognitive and (under their influence) linguistic categories are organized is the prototype principle – that is to say, a principle whereby elements are assigned to a category not because they exemplify properties that are absolutely required of each one of its members, but because they exhibit to a greater or lesser extent (or are simply *expected* to exhibit to a greater or lesser extent) certain types of similarity with a particular category member that has been (naturally or culturally) established as the *best example* (or prototype) of its kind.

There are two fairly simple reasons why the prototype approach should be of interest to students of language. The first is that there is a significant amount of *prima facie* evidence suggesting that it is correct. The second is that if it *is* correct, then the view of linguistic categorization promoted by the dominant linguistic and philosophical schools of thought does not seem to be correct, since, on that view, linguistic categories are certainly *not* mere reflexes of general cognitive strategies (if they are cognition dependent at all), and they certainly *are* founded on principles which ensure 'yes or no' decisions on most (if not all) questions of category membership.

Now, prototype theory is fairly young by any standards, it is therefore not the case either that all the evidence that its practitioners could

legitimately hope to produce has been produced, or that all the conclusions they have derived from the evidence that they have already produced are the best conclusions one could possibly derive. So, the best way of serving the interest the theory naturally arouses would consist in trying, on the one hand, to enrich its empirical base, and, on the other hand, to clarify its conceptual foundations. These are precisely the areas where this volume intends to make original contributions: the first two parts contain chapters where various linguistic phenomena are analysed in ways that make essential use of the notion of prototypicality or of closely related notions. The last two parts contain chapters where the notion of prototypicality or closely related notions become themselves the object, rather than the instrument, of inquiry, and provide the opportunity for detailed statements of a variety of methodological attitudes towards several aspects of linguistic description.

The division between Part One and Part Two, as well as that between Part Three and Part Four, are much less sharp, and will be made clearer in the course of this introduction. Thematically, these divisions tend to reflect the fact that, within linguistics, the prototype approach was first of all presented as a novel way of dealing with matters of word meaning, and was only later extended to additional levels of linguistic representation. Accordingly, the papers in Part One present some new results of adopting the prototype approach in areas where questions of word meaning figure centrally, if not exclusively, whereas those of Part Two extend the same approach to areas where matters of word meaning arise incidentally, if at all. Similarly, the primarily methodological chapters in Part Three evaluate prototype theory specifically as a theory of word meaning, whereas those of Part Four assess it in ways that are both more indirect and more liberal as to what they take its potential range of application to be.

In what follows I will try to sketch some of the connections between the twenty-six chapters, trusting that the reader will not, at this stage, object to the amount of oversimplification that such an attempt may involve.

Part One

Most of the best known results on prototypically organized word meanings have been obtained, indirectly, from psychological studies of categorization of various kinds of concrete objects. Although the choice of such objects was probably necessary given the psychologists' experimental concerns, they are evidently not sufficient for supporting general statements on the nature of word meaning. For one thing, one might wish to know to what extent can prototype categories be relied upon even in the analysis of terms with ostensibly 'concrete' referents. For another thing – and most importantly – one might wish to know whether

prototypes show any promise of being involved in the analysis of terms whose referential targets are certainly not 'concrete' in the above sense. The first chapter of Part One offers a comprehensive answer to the former question, while the next four provide important elements of an answer to the latter.

In 'A survey of category types in natural language', Cecil H. Brown proposes to show that the development of names for concrete objects in natural languages follows eight routes, which correspond to the logically possible combinations of positive or negative valuations that a set of objects may receive with respect to three properties: artificiality of its members, configurational clues ensuring the identification of its members, and prototypicality of some of its members relative to others. The least that this proposal entails, then, is that the absence of prototypicality is just as much structurally important as its presence when a systematic description of concrete object categorization through language is sought. Brown, however, is interested in much more than the merely taxonomic significance of the three properties. He advances and defends various hypotheses as to why some of their logically possible combinations are more frequently instantiated in natural languages than others, why some of them are more likely to be superseded in the course of linguistic evolution than others, and why some of them are characteristically unstable in a way that others are not. As one might expect, each one of these hypotheses throws new light on questions regarding the nature of prototype categories. In this sense, the picture that emerges from Brown's survey is one in which such categories are perhaps less extensively involved in the development of concrete object naming than has been supposed, but where it is much clearer how and why they are involved, when they are.

One might think that, in moving away from names for concrete objects, one is diminishing one's chances for convincingly arguing that linguistically encoded meanings are, in any literal sense, organized prototypically. This, however, might well be a mistake. After all, saying that a concrete object is the 'best example' of its kind is an elliptical way of saying that humans tend (or, in some cases, decide) to regard it as the best example: in a world without perceiving minds, no object could possibly be a better, or a worse, example of anything than any other. If this is so, then it may also be that humans tend to conceptualize even the most abstract of ideas in ways that force them to regard some of their instances as best examples of their kinds, and to evaluate the rest on the basis of how well they resemble – or can be expected to resemble – the prototypical instances. Each of the next four chapters of Part One interprets this possibility as creating constraints on semantic representation, and the range of linguistic phenomena that they purport to be able to elucidate by thus interpreting it is quite remarkable.

In 'Possible verbs and the structure of events', William Croft argues that the familiar semantic trichotomy between inherently causative, inchoative, and stative verbs fails to account for the fact that every verb can systematically manifest each one of the supposed 'senses'. He then proposes that the non-rigid nature of the trichotomy should be viewed as a reflection of three types of perspectivization that are possible within a *single* cognitive model for events, which represents them as consisting of a cause, a process, and a state, and thereby prohibits the total exclusion of any one of these elements to the benefit of the others. He shows finally that this cognitive prototype is responsible for two types of interesting typological facts. First, that events conforming to it are given grammatical expression that is uniform across languages, whereas those that do not tend to be encoded idiosyncratically by each one of them. And second, that a verb whose intrinsic semantic value in a given language makes it especially suitable for the expression of a causative or of a stative event-view is morphosyntactically unmarked with regard to the expression of that view, whereas it becomes the domain of various marking processes when it expresses event-views that are less congruent with its intrinsic semantic value.

In 'Prototypical considerations on modal meanings', Steven Cushing argues that necessity and possibility modals in a natural language (in either their 'descriptive' or their 'prescriptive' readings) are in fact understood (and should be represented) as making an implicit appeal to prototypes of a higher order, in particular, to what speakers of that language understand as best theories of the (physical or moral) world. He then sets out to make the formal structure of that appeal explicit, and he thus arrives at a system of definitions that are significantly different from those obtainable from standard logical treatments of the modalities. He finally shows that these definitions make possible an orderly explanation of a variety of phenomena that were poorly understood or virtually unnoticed (for example, the difference between strong and weak modals of both the descriptive and the prescriptive varieties, the existence of non-accidental gaps in modal vocabularies, and the peculiarities of scope ambiguities in modal contexts).

In 'Belief ascription, metaphor, and intensional identification', Afzal Ballim, Yorick Wilks, and John Barnden describe an intelligent system which, in successfully ascribing beliefs to agents, uses an algorithm that treats the system's beliefs as prototypical – in other words, that takes the agents' beliefs to be identical with the system's own, unless there is evidence to the contrary. They then show that since, in successfully interpreting metaphors, the system can make use of precisely the same algorithm (which then amalgamates properties of the metaphor's 'vehicle' with those of the metaphor's 'tenor' unless there is evidence to the contrary), it could be plausibly maintained that metaphor interpretation

relies on the same fundamental process that is responsible for the ascription of beliefs and other propositional attitudes. They argue finally that belief ascription itself should in its turn be viewed as a phenomenon that is essentially metaphorical in nature (in the sense that it involves the treatment of an agent's mental states as a field for the metaphorical projection of other agents' mental states), and they conclude that, thus construed, the metaphoricity of belief casts serious doubt on some fundamental assumptions of formal semantic approaches to the analysis of propositional attitudes.

In 'Negated beliefs and non-monotonic reasoning', Ryszard Zuber examines the special behaviour that a wide variety of families of predicates (for example, factive, opaque, and emotive ones) are known to manifest with regard to negation, and seeks, on the one hand, a unified treatment of these peculiarities, and, on the other, an explanation of their existence. The former task he accomplishes by defining a notion of intensional negation that is noticeably different from those inherited from standard logical systems, and by characterizing each type of predicate in its terms. Concerning the latter task, he suggests that the explanation must be sought in the fact that, in their prototypical uses, all these predicates are associated with subjects denoting human beings, and that their special behaviour in negative contexts is a reflection of an implicit assumption to the effect that, because of what a prototypical human being is, certain forms of reasoning on its subject may be taken to be locally valid, although they are not of general validity.

What the above four contributions jointly suggest, then, is that, far from being relevant only to the analysis of processes of concrete object naming, prototype considerations are instrumental in characterizing far more abstract semantic domains, and that, in doing so, they provide original answers to questions that any of the currently available semantic theories would recognize as central (and that few of them could claim to have answered satisfactorily). The last three chapters in Part One deal with questions that are less central from the point of view just indicated (essentially because orthodox semantic theories do not seem particularly interested in systematically raising them) but are just as interesting in their own right. In 'Lexical hierarchies and Ojibwa noun derivation', Richard A. Rhodes shows how the idea that the senses of a morpheme may be organized in a way that is analogous to (and, at certain points, directly reflects) the categorization of physical entities around cognitive prototypes leads to a uniform account of the apparently unpredictable semantic contributions of a derivational suffix in an Alquonquian language. In 'Some English terms of insult involving sex organs', Keith Allan shows how some seemingly inexplicable constraints governing the interpretation of non-literal uses of certain vocabulary items can be satisfactorily explained when the use of such items is viewed as a

manifestation of conventionalized beliefs related to the prototypical referents of their literal counterparts. Finally, in 'The lexicographical treatment of prototypical polysemy', Dirk Geeraerts argues that if their ability to make sense of lexicographical practice is one condition of adequacy for semantic theories, then prototype views of word meaning meet that condition better than their classical alternatives, since it is the prototype rather than the classical view that can be shown to motivate the solutions to problems of complex categorization implicit in traditional lexicography.

Part Two

The idea that grammatical constructions, *qua* grammatical constructions, carry a kind of meaning that is irreducible to the sum of the meanings of their constituents is not controversial. What has been, and still is, the subject of controversy is whether all the grammatically important properties of a construction can be ultimately explicated in semantic terms. While many influential grammatical theories have decided to proceed on the assumption that it is unlikely that this will turn out to be possible, the more ambitious project of trying to show that it may, after all, be possible has never failed to attract devotees. Prototype theory is currently giving new impetus to this project, for reasons that shouldn't be difficult to understand: if one can legitimately claim that some instantiations of a grammatical category are *better* instantiations than others, then one is implicitly claiming, first, that grammatical categories have a cognitively salient semantic basis (since it is only by reference to a basis of this sort that the relative representativity of their members could be realistically assessed) and second, that the claim that they have a semantic basis does not entail that they impose necessary and sufficient conditions for membership (and are therefore not open to some simplistic kinds of counterexample that have been raised against proposals for semantically based grammars in the past). Although neither of the above claims can be taken as conclusively established in all relevant respects, they seem to be well supported by much ongoing research, some aspects of which are represented in the first five chapters of Part Two of the volume.

In 'Settings, participants, and grammatical relations', Ronald W. Langacker outlines a grammatical theory where only cognitively motivated categories are recognized, and sets out to examine how basic grammatical relations could be best represented in its terms. He claims that such relations can be successfully accounted for by reference to a cognitive model representing the normal observation of a prototypical human action, and incorporating a fundamental distinction between the setting and the participants of an action scene. The subject- and object-

properties of constituents of various sentence types are then explicated as resulting from operations whose effect is to selectively accord linguistic representation to the various elements mentally instantiating the model, to vary the relative prominence with which the selected elements are encoded, and to respect or to reverse, in the course of the representation, the notional priorities implicit in the model underlying it.

In 'On the semantics of compounds and genitives in English', Paul Kay and Karl Zimmer note that genitive and nominal compound constructions in English are not always interchangeable, in spite of their fundamental structural similarity. They suggest that these differences should be semantically accounted for by associating the two constructions to prototype schemata which differ only in that the one representing genitives stipulates that their modifier nouns be individual terms, while the one representing compounds stipulates that their modifier nouns be class terms. They then show that observed deviations from these patterns are precisely the ones that one would expect, if the patterns were indeed prototypical: some proper nouns can exceptionally act as modifiers in compound constructions, but they are then interpreted as class terms; and some common nouns can exceptionally act as modifiers in genitive constructions, but they are common nouns with special conceptual ties with entities denoted by individual terms.

In 'A notional approach to the French verbal adjective', Roger McLure and Paul Reed show how a construction whose description seems to have been a consistent source of grammatical frustration – adjectival modification of nouns by present participles in French – can be satisfactorily analysed when it is viewed as a solution to the semantic problem of ascribing to an entity characteristic properties that cannot be regarded either as merely contingent or as strictly necessary. After eliminating a variety of possible alternative explanations of their distributional properties, they claim that it is precisely this intermediate conceptual region that French verbal adjectives prototypically grammaticalize, and they explain the different types of semantic effect that their permissible combinations with nouns may produce as different ways in which this prototypical meaning can, given a context, be exploited.

In 'Prototypical uses of grammatical resources in the expression of linguistic action', René Dirven draws attention to the striking variety of innovative syntactic frames within which a basic English speech activity verb may be used, and argues that these syntactic novelties have a conceptual basis (the same basis, in fact, that, in a morphologically richer language like German, would tend to activate equally diverse derivational processes): they are, he suggests, symbolic means for highlighting particular aspects of the folk model in terms of which the speech event denoted by the verb is understood – and, to this extent, they provide a basis for claiming that metaphorization is a phenomenon that is

manifestable not only on the lexical but also on the syntactic level.

Finally, in 'Towards a theory of syntactic prototypes', Margaret E. Winters identifies six features that could plausibly be regarded as jointly conferring prototypicality on grammatical constructions, argues that these features may themselves be organized prototypically (in the sense that some of them may be more prominent signals of conceptual centrality than others), and suggests that their relative prominence in any given language may itself be a function of diachronic pressures.

Varied as they obviously are, the applications of prototype notions to the analysis of lexical and grammatical meaning do not exhaust the range of linguistic phenomena in the description of which such notions might be fruitfully employed, any more than the use of the notion of prototype in psychological studies of concrete object categorization exhausts the range of psychological phenomena that could be analysed interestingly in its terms. Part Two of the volume concludes with three chapters, of which the first two exemplify realizations of such further possibilities in linguistics, and the last one in psychology. In 'Accent in prototypical *wh* questions', Dwight Bolinger argues that there are good reasons for claiming that one among the various stress patterns followed by *wh* interrogatives constitutes a prototype in terms of which the function of the others is understood, in much the same way in which instances of a conceptual category are said to constitute prototypes by reference to which the category status of less characteristic instances is determined. In 'Prototypical manners of linguistic action', Anne-Marie Diller argues that certain formal properties distinguishing performative from non-performative occurrences of speech act verbs can only be accounted by reference to conventionalized beliefs regarding the mental dispositions of prototypical performers of the speech acts that these verbs denote. Finally, in 'Where partonomies and taxonomies meet', Barbara Tversky reviews some recent psychological evidence which suggests that, just as categories are perceived as being organized around prototypical members, so individual category members are perceived as consisting of prototypical parts, and argues that this latter phenomenon opens an area of investigation that is not only interesting in itself, but adds a new dimension along which the analysis of the former could be further refined.

Part Three

The undeniable heuristic value of the notion of prototypicality should not obscure the fact that its exact theoretical shape is less clear than one might have wished, especially when it is transferred from purely psychological to specifically linguistic domains of investigation. Since the first domain that has been affected by such a transfer is the domain of

lexical semantics, and since lexical semantics is a research area that is sustained by important theoretical traditions, one would expect lexical semanticists to be less than unreservedly prepared to embrace the new idiom and all its apparent consequences. The first three chapters of Part Three of the volume confirm this expectation, and they thus delineate one dimension along which more clarity could be systematically sought. Starting from independent considerations, these papers can be viewed as arguing for three main conclusions: that the range of semantic phenomena to which the notion of prototype could in principle be applied is more restricted than one tends to believe; that its successful application even in this properly delimited area cannot, at present, be taken to be unproblematic; and that even if it should turn out to be unproblematic it would not have the subversive effects that it is supposed to have on orthodox conceptions of word meaning.

In ' "Prototypes save": on the uses and abuses of the notion of prototype in linguistics and related fields', Anna Wierzbicka argues that many descriptions of word meanings that are directly inspired from prototype theory constitute in fact manifestations of either conceptual confusion or inadequate attention to linguistic facts. She then claims that reference to prototype representations is indeed necessary for the analysis of certain types of word meaning, but that it can be satisfactorily incorporated into standard forms of semantic description, without forcing them to abandon their claims to definitional adequacy. She concludes that the belief that such an incorporation could not be successfully implemented is an illusion, probably deriving from the mistaken assumption that definitionally adequate semantic descriptions should be entirely cast in non-mentalistic vocabulary.

In 'Prototype theory and its implications for lexical analysis', Adrienne Lehrer notes that there are aspects of word meaning which do seem to vindicate prototype theory, but which are neither unknown to nor inexplicable within fairly traditional theoretical frameworks. Turning then to certain hypotheses about word meaning which seem to follow specifically from prototype theory, she argues that, in some cases, they are, despite their interest, insufficiently precise to be tested, and, in some other cases, falsified by the relevant facts. She concludes by recommending that the semantic relevance of prototypes should not be taken to follow automatically from their psychological plausibility, and that the search for a specifically linguistic motivation of their occasional involvement in semantically sensitive areas would be well worth undertaking.

In 'Prototype theory and lexical semantics', D. A. Cruse draws attention to some important respects in which both the purely cognitive and the specifically linguistic interpretation of prototypicality is in need of serious reconsideration. Concerning the cognitive interpretation of the notion, he suggests that, among other things, it illegitimately conflates at

least three different respects in which a category member may be exemplary, it risks confusing two different conceptions of the opposition between gradable and non-gradable category membership, it underestimates the extent to which category boundaries may be sharp, and it overestimates the extent to which merely typical and properly prototypical category features can be strictly separated. Concerning the linguistic interpretation of the notion, he argues that, apart from inheriting most of the problems connected with the cognitive one, it pays insufficient attention to the fact that, alongside semantic properties that could plausibly be thought of as relating to prototypical conceptual representations, there are important classes of semantic properties for which no such relation could be postulated, since, despite appearances, these properties are radically word-specific (in the sense that they are properties of the words themselves and not of the concepts – prototypical or otherwise – that words may mediate).

The last two chapters of Part Three choose to concentrate not on general problems that a prototype approach to lexical semantics does or may have to face, but on some no less serious problems arising from specific analytical proposals that have been taken to be representative instances of the prototype approach at its best. In 'Representation, prototypes and centrality', Claude Vandeloise claims that a well-known analysis of the preposition *over* within a broadly prototype framework is in fact a good example of how some intuitively plausible notions may lead to wildly implausible theoretical conjectures when they are employed in a methodologically undisciplined way. And in 'A few untruths about "lie" ', I suggest that an equally well known analysis of the verb *lie*, in terms of a set of prototype features that are allegedly essential for characterizing both cases of clear applicability or inapplicability and cases of intermediate applicability of this term, rests in fact on highly questionable assumptions both with regard to what the clear cases are and with regard to what the proper explanation of the apparently intermediate cases should be.

Part Four

If the use of the notion of prototype in linguistics is indeed, as many of its advocates seem to believe, one among many signs of a paradigmatic shift that is currently under way in the study of language, then it may well be that attempts to emphasize the real or apparent shortcomings of prototype theory *vis-à-vis* more standard approaches to aspects of linguistic description miss (or, at least, misconstrue) the real issue. It would be much more appropriate, from that point of view, to emphasize instead the similarities between the prototype approach and certain other recent approaches which are just as sceptical as prototype theory has

become about standard views as to what a proper linguistic description should be; and, given this background of similarity, it would then be interesting to examine whether these approaches could be mutually reinforced in pursuing their partially overlapping goals. The five chapters of Part Four reflect very divergent research interests, but they are united in their refusal to take for granted some basic assumptions of linguistic analysis, as it is standardly practised; in doing so, they are led to implicitly or explicitly raise questions that have been at the centre of prototype research since its introduction in linguistics; and they thus offer some new perspectives within which the answers to those questions could be profitably sought.

In 'On "folk" and "scientific" linguistic beliefs', Roy Harris attacks a central thesis of modern linguistics which would seem to underlie a familiar kind of objection to prototype theory. The objection is, roughly, that, by taking speakers' untutored beliefs about the universe (including their linguistic universe) as a phenomenon that linguists not only should not disregard but should rather take as the basic force behind linguistic categorization, prototype theory encourages its practitioners to abandon the neutral stance that they should at all costs maintain towards their assigned objects of study. And the assumption behind the objection is, presumably, that there is a reliable basis for drawing a sharp distinction between 'folk' and 'scientific' linguistic beliefs, and for systematically preferring the latter when they appear to be in conflict with the former. Through a series of important arguments, however, Harris shows that the correctness, and, indeed, the coherence, of that assumption is highly questionable, and concludes that it is only by fully acknowledging (and by appropriately exploiting) its lay foundations that the study of language could adequately proceed. To the extent that prototype theory is one step in that direction, it would seem, then, to be reasonably strong in an area where it might have been thought to be particularly vulnerable.

In 'Gestures during discourse: the contextual structuring of thought', Nancy L. Dray and David McNeill outline a decidedly naturalistic approach to linguistic description which seems to have significant additions to suggest to prototype accounts of linguistic categories. The distinctive feature of that approach (which is exemplified by some insightful analyses of gestural activity during discourse) is its claim that the value of linguistic elements should be viewed as a result not only of conventionally determined but also of contextually arising oppositions. And the systematic study of these latter could help, according to Dray and McNeill, not only to explain some linguistic choices that seem to lie outside the predictive power of prototype theory, but also to simplify the accounts of certain other choices that prototype theorists have already given.

In 'Why words have to be vague', Roger McLure proposes a

reinterpretation of some prototype phenomena in the context of hermeneutic phenomenology, and claims that this reinterpretation makes possible a deeper understanding of these phenomena, in two ways. First, by permitting their dissociation from certain unselfconsciously solipsistic philosophical views in terms of which they have been understood. Second, by providing a framework within which the essential instability of linguistic categorization that these phenomena highlight can be seen as a presupposition of, rather than as an obstacle to, the possibility of linguistic communication. Once the full implications of this reinterpretation are drawn, McLure suggests, prototype theory will be recognized as constituting a challenge to accepted modes of linguistic theorizing that is far more serious than has been supposed, even by its supporters.

In 'Schemas, prototypes, and models: in search of the unity of the sign', John R. Taylor examines the relation between prototype accounts of linguistic categorization and certain recent alternative accounts where schematic representations far more abstract than those sanctioned by prototype theory are claimed to make possible a more comprehensive account of linguistic facts, while at the same time doing justice to their cognitive basis. He suggests that, as far as their descriptive capabilities are concerned, the schematic and the prototypical view of categorization cannot be regarded as real alternatives, since all the basic results obtainable through the one could, in more or less complex ways, be translated into the idiom of the other. He argues, however, that, from the point of view of their overall plausibility, it is the prototype rather than the schematic view that is to be preferred, since the prototype idiom accommodates more naturally a greater number of types of linguistic category than the schematic idiom does.

Finally, in 'Psychologistic semantics, robust vagueness, and the philosophy of language', Terence Horgan draws attention to some important wider implications that past research on prototype categorization might have, as well as to some more refined ways in which it could itself be conducted in the future. He first outlines certain basic limitations of both the realist and the anti-realist conceptions of meaning in contemporary philosophy of language, and argues that these limitations can be transcended within a theory of meaning where the notion of cognitive prototype would play a central role. He then notes that the psychological modelling of that notion thus far has not been entirely satisfactory, essentially because it proceeded through minor emendations to classically inspired models of categorization, which are inherently ill-adapted to the representation of vagueness (and, hence, of an important aspect of prototypicality). He finally argues that there are good reasons for expecting that the radical departure from classical conceptions of categorization that is characteristic of the emergent connectionist paradigm in cognitive science will provide the means of constructing

models of prototypicality that will be not only philosophically suggestive, but also psychologically adequate.

I hope that the preceding remarks have sufficiently clarified the organizing principle of this volume, namely, to provide a view of prototype research that is appropriately balanced, first by maintaining proper proportions between analytical proposals and critical reflections, and second by making room for a significant degree of variation both in the choice of analytical objects and in the choice of critical targets. I also hope that, having been sufficiently aroused by these preliminaries, the reader will now wish to be in personal contact with the arguments of the individual chapters. What remains for me to do is to express my gratitude to those who, apart from the contributors, have made this volume possible. Henrietta Mondri and John Taylor played an important role in its inception. Jonathan Price took an even more significant part in the process leading to its completion. And Clelia Kachrilas was my unfailing source of support from beginning to end. To all of them, my sincere thanks.

Part One

On the content of prototype categories: questions of word meaning

Chapter 1

A survey of category types in natural language

Cecil H. Brown

The treatment of two or more distinguishable entities as if they were the same creates a category (cf. Mervis and Rosch 1981: 89). People create categories by assigning the same name or label to different things. When speakers of a language are in general agreement with respect to the different entities to which a single term applies, the pertinent category is a component of natural language. This chapter surveys types of category lexically encoded in natural language. Specifically, it focuses on categories whose membership is restricted to concrete objects such as plants, animals, toys, weapons, and tools, as opposed to abstract things such as war, love, religion, poetry, knowledge, and lies.

The present work attempts to show that categories of natural language can be profitably analysed by relating them to a system of category types defined in terms of three factors. These are (1) artifactual versus non-artifactual reference (+ AR vs. − AR); (2) Gestalt versus non-Gestalt motivation (+ GM vs. − GM); and (3) prototype/extension versus non-prototype/extension (+ P/E vs. − P/E). For example, as explained in detail presently, categories such as *screwdriver*, *cup*, *pen*, *chair*, *rope*, *button*, and *train* all belong to a single type of natural language category since all are plus for artifactual reference (+ AR), plus for Gestalt motivation (+ GM), and plus for prototype/extension (+ P/E). On the other hand, categories such as *raccoon*, *robin*, *black walnut tree*, and *dandelion* belong to a different category type since all are minus for artifactual reference (− AR), plus for Gestalt motivation (+ GM), and minus for prototype/extension (− P/E). There are, then, eight category types defined by all logically possible combinations of variables of these factors:

Category Type 1: − AR + GM − P/E
Category Type 2: − AR + GM + P/E
Category Type 3: − AR − GM − P/E
Category Type 4: − AR − GM + P/E
Category Type 5: + AR + GM − P/E
Category Type 6: + AR + GM + P/E

Category Type 7: + AR − GM − P/E
Category Type 8: + AR − GM + P/E

Artifactual reference

Concrete objects grouped in categories of natural language are either manufactured by humans (artifacts) or are natural kinds (non-artifacts) such as plants, animals, body parts, clouds, mountains, and rocks. A possible absolute universal of language is that artifacts and natural kinds are *never* included in the same category (putting aside 'categories' generated by metaphorical equations such as dipstick = penis). Thus, there is a clear distinction between categories which involve reference to artifacts (+ AR) and those that entail reference to non-artifacts (− AR).

Gestalt motivation

Concrete objects rarely are morphologically continuous, i.e., typically there is no continuum of objects grading from one to another with respect to similarity. Rather, there is usually a great deal of distinctiveness, making for obvious breaks or gaps among things. Hunn (1977) calls such gaps, when they apply to biological entities (non-artifacts), 'discontinuities in nature'. Clearly, discontinuities perceived by humans are not restricted to natural kinds. Cups, mugs, and glasses are no more or no less discontinuities than are maples, oaks, and walnuts.

Hunn (1977: 41–75) focuses on psychological processes through which discontinuities are translated into natural language categories. He notes that discontinuities in nature are underlain by feature or attribute clustering. Bruner, Goodnow, and Austin (1956: 47) illustrate this by citing the example of birds in general, creatures possessing feathers, wings, a bill or beak, and characteristic legs. Any one of the latter features is highly predictive of the others. For example, if a creature possesses feathers, it will invariably also have wings, a bill or beak, and characteristic legs. Thus attributes of the discontinuity 'birds in general' cluster together, or in other words are highly correlated with one another.

Hunn, following Bruner *et al.* (1956: 47), proposes that the mutual predictability of clustering features can lead to an expectancy in the minds of humans that attributes involved will be found together. For example, through exposure to different kinds of bird, people build up in their minds an expectation that feathers, wings, and so on, go together. Such an expectation underlies the conceptual development of the configurational or Gestalt property of 'birdness'. When such a property develops, inclusion of any particular object in a labelled bird category is contingent upon whether or not the object demonstrates the single feature 'birdness'. As a result, clustering features pertaining to the bird discontinuity become psychologically subordinated to the single Gestalt property.

A Gestalt property arises through the *recoding* of features or attributes (Hunn 1977: 46). The concept of recoding, borrowed from information theory, involves the notions of 'chunks' and 'bits' of information. Data organized by a restricted number of immediate or simultaneous judgements constitute chunks (Miller 1967). The amount of information which each chunk contains is described as a number of bits of information. Recoding essentially consists in taking a great number of chunks, each of which contain but a few bits, and reorganizing them into fewer chunks with more bits per chunk. Miller (1967: 24) gives the following example:

> A man just beginning to learn radio-telegraphic code hears each dit and dah as a separate chunk. Soon he is able to organize these sounds into letters and then he can deal with the letters as chunks. Then the letters organize themselves as words, which are still larger chunks, and he begins to hear whole phrases . . . I am simply pointing to the obvious fact that the dits and dahs are organized by learning into patterns and that as these larger chunks emerge the amount of message that the operator can remember increases correspondingly. In the terms I am proposing to use, the operator learns to increase the bits per chunk.

Bruner, Goodnow, and Austin (1956: 46) illustrate the recoding of attributes into a single Gestalt property by using the following biological example (cf. Hunn 1977: 47):

> The student being introduced for the first time to microscopic techniques in a course in histology is told to look for the *corpus luteum* in a cross-sectional slide of rabbit ovary. He is told with respect to its defining attributes that it is yellowish, roundish, of a certain size relative to the field of the microscope, etc. He finds it. Next time he looks, he is still 'scanning the attributes'. But as he becomes accustomed to the procedure and to the kind of cellular structure involved, the *corpus luteum* begins to take on something classically referred to as a *Gestalt* or configurational quality. Phenomenologically, it seems that he no longer has to go through the slow business of checking size, shape, colour, texture, etc. Indeed, 'corpus luteumness' appears to become a property or attribute in its own right.

Hunn (1977) restricts his discussion to the development of Gestalt properties relating to biological categories. However, it is clear that attribute recoding and resulting Gestalten are not limited to natural kinds. For example, as discussed in the above quotation of Miller, letters of radio-telegraphic code are recoded into words. Each word, then, constitutes a single Gestalt property. Words expressed in radio-telegraphic code are, of course, human artifacts. Such words differ from artifacts such as cups, mugs, and glasses, since they are not concrete objects. If words expressed in code can possess Gestalt properties, it seems clear that so

can concrete objects manufactured by humans, so long as these objects fall into discrete discontinuities (a point discussed at length presently).

An important assumption of the present discussion is that Gestalt properties typically motivate categories which relate to discontinuities. This is not to propose that such categories are always motivated by Gestalt properties. As noted in the above quotation, a student learning to identify the *corpus luteum* in a rabbit ovary may begin to do so by 'scanning the attributes', so at first the category is defined in terms of several features rather than in terms of a single configurational property. However, objects pertaining to most categories of natural language relating to discontinuities, especially *folk* categories known to all or nearly all speakers of a language (as opposed to specialist categories such as *corpus luteum*) ordinarily do not require close scrutiny (for pertinent attributes) for the purpose of class inclusion. In addition, I do not mean to imply that for any one category relating to a discontinuity, a Gestalt property *alone* motivates the category (another point to be discussed at length presently).

Not all categories in natural language relate to discontinuities. This is particularly clear when abstractions such as lies (falsehoods) are considered. There are, of course, no perceptual things that belong to the category called *lie* and, consequently, no perceptual discontinuity with which it is connected. Lies, then, do not have in common a certain Gestalt property, rather they relate to what Lakoff (1987: 113) calls a 'propositional model': 'Propositional models specify elements, their properties, and the relations holding among them.' Coleman and Kay (1981: 28), for example, have proposed a propositional model relating to the category *lie* involving a speaker (S) who asserts some proposition (P) to an addressee (A):

(a) P is false.
(b) S believes P to be false.
(c) In uttering P, S intends to deceive A.

Thus, a lie is characterized by the properties (a) falsehood, which is (b) intentional, and (c) meant to deceive. None of these properties, of course, is a perceptual property of a thing.

Some categories of natural language encompass concrete objects but, none the less, are not underlain by discontinuities. Rather, they are motivated solely by propositional models. This is true of categories such as *toy* and *weapon*. For example, membership of the class *toy* belongs to a highly heterogeneous group of objects (artifacts) such as dolls, rattles, spinning tops, jump ropes, and toy soldiers, none of which bear much perceptual similarity to one another. Toys, then, do not relate to a single perceptual discontinuity and, consequently, are not underlain by a single Gestalt property which could be called 'toyness'. Mere observation of a

sample of different objects called *toy* cannot lead to a conception of 'toyness' in the way that observation of a sample of birds can lead to a conception of 'birdness'. This is so because toys have little in common other than that they are used by people to play with. Consequently, the category *toy* is defined solely in terms of a propositional model in which 'object to play with' is a necessary element. (Lakoff [1987: 51] would identify the latter attribute as an 'interactional property'. 'Interactional' refers to relationships between humans and things, e.g., 'a chair is a thing to be sat upon by humans', or 'a knife is used by humans for cutting'.)

To propose that some categories are motivated by Gestalt properties is not to suggest that propositional models do not pertain to them as well. For example, while the category *squirrel* is typically Gestalt motivated, a propositional model also pertains to it which takes account of facts such as squirrels (a) 'store things to eat'; (b) 'are light and quick in their movements'; (c) 'are afraid of people and run away from them'; (d) 'people think of them as nice and a little amusing' (Wierzbicka 1985a: 165–6). Also included in such a propositional model are properties that reflect perceptual attributes of squirrels: (e) 'they have a big bushy tail'; (f) 'they have pointed ears'; (g) 'they are furry' (ibid.); and (h) 'they are brown in colour'. It is assumed in the present discussion that propositional models pertain to *all* categories of natural language.

In summary, some categories of natural language are underlain by Gestalt properties (+ GM) and some categories are not (− GM). Those which are not are defined solely in terms of propositional models.

Prototype/extension

Kronenfeld (1988; with Armstrong and Wilmoth 1985) proposes a general theory of semantics of word meaning centred on the concepts of category prototype (cf. Anglin 1977; Berlin and Kay 1969; Carey 1982; Fillmore 1978; Kay and McDaniel 1978; Kempton 1978, 1981; Rosch 1975d, 1977; Rosch and Mervis 1975) and category extension (influenced strongly by Lounsbury 1964a, 1964b and 1965). A prototype is the best exemplar of a category or, in other words, is most representative of things included in a class (e.g., the reddest red or the most cup-like cup). The prototype, then, is the core of a category which is 'surrounded by' other members that are not as representative of that class (Dromi 1987: 52).

Kronenfeld proposes that a category typically develops its membership by expanding or extending its boundary from a prototype to entities (non-prototypes) which are not identical to the prototype but which are perceived as relating to it in some manner (e.g., through perceptual similarity). Thus, for example, it is possible that penguins and ostriches are included in the category *bird* because they are similar in appearance to prototypic birds, i.e., probably robins for many speakers of American

English: penguins and ostriches, like robins, have feathers, wings, and a beak, but they differ from robins since they do not fly. A 'bird' category may be further extended from the prototype through inclusion of bats, although such an inclusion would not be based primarily on perceptual similarity but rather on behavioural similarity: robins fly and so do bats, but bats do not have feathers and a beak (or, for that matter, birdlike wings).

The essential tenet of Kronenfeld's theory is that prototypes are *a priori* to their corresponding non-prototypes. While the evidence Kronenfeld musters to support this contention is impressive, one cannot at present say that his theory has been proved. However, I am convinced of its validity mainly because my own recent research in ethnobiological classification (Brown 1984, 1985b, 1986, 1987a) reveals patterns that fit Kronenfeld's scheme. Some of these are discussed presently.

Kronenfeld's theory basically addresses the question of category development. 'Category development' in this paper is understood in three ways. First, development may be taken to refer to how a category has arisen in a particular language. Second, it may be taken to refer to how a category is acquired by children learning language or by people learning a second language. Third, it may be taken to refer to processes leading to the inclusion in a category of newly encountered potential exemplars.

In terms of Kronenfeld's theory, a category develops in a language by expanding its boundaries from a prototype to related non-prototypic exemplars. Thus, a word which in the past was referentially restricted to a prototype acquires a more comprehensive referential application. With respect to child language acquisition or second language learning, Kronenfeld's theory implies that learners will first apply a category label only to its generally recognized prototype and only later to non-prototypic members. With respect to deciding if a newly encountered thing is to be included in a category, Kronenfeld's theory implies that an inclusive decision will be based on a comparison of a new thing to a generally recognized prototype. If it relates in some way to a prototype (e.g., is perceptually similar to a prototype), it may be included; if not, it will not be included.

Category extension from a prototype to non-prototypes can involve relationships motivated by: (1) similarity; (2) metaphor; or (3) metonymy (Brown 1979; Lakoff 1987: 113–14). With respect to similarity, non-prototypes may be related to a prototype (a) if they are perceptually similar to a prototype or (b) if their propositional models specify some property or properties which are also possessed by a prototype. Thus, for example, assuming the robin to be a prototypic *bird*, a penguin may be included in the class *bird* by right of being perceptually similar to a robin (robins and penguins share feathers, wings, etc.), cf. (a) above. With respect to property specifications of propositional models, bats may

be included in the category *bird* based on the acknowledgement that bats, like robins, 'fly and eat insects' (in addition, typical bats are relatively similar in size to robins).

A relationship based on metaphor also entails some similarity between things, but similarity of a minimal kind. For example, the word *bird* is used in British slang to refer to a young, pretty woman. It is possible that this metaphorical usage is based on some perceived or propositionally specified similarity between attractive young ladies and a prototypic bird. Whatever this similarity may be, it is not obvious, at least to this author. Metonymy involves contiguity associations such as part to whole. Thus, for example, the English word *tube* labels a (prototypic?) part of a television set, which has become expanded in reference to the whole appliance (at least in slang).

In the following discussion I focus only on associations based on perceptual and propositional similarity to the exclusion of those based on metaphor and metonymy. One reason for this is that I do not believe that unions of referents through metaphor and metonymy represent instances where united referents belong psychologically to the same category. For example, intuitively it does not seem appropriate to regard robins and attractive young women as members of a single class despite the fact that they are both called *bird*. Similarly, the union of a picture tube and the appliance of which it is a part does not constitute a category for me. In these cases use of a single term to denote two different things creates *polysemy*, not a category. Metaphor and metonymy, then, really relate to another area of linguistic enquiry which is, by the way, vast in its dimensions (cf. Lakoff 1987; Lakoff and Johnson 1980; Witkowski and Brown 1983, 1985; Witkowski, Brown, and Chase 1981; Brown 1979, 1983, 1985a, 1987a, b, 1989; Brown and Witkowski 1981, 1983). Thus, in this essay, the word *category* is restricted in use to labelled groupings of referents which are related only through perceptual and propositionally specified similarity.

Nowhere in his outline of prototype/extension theory does Kronenfeld discuss the essential nature of a prototype. Is a prototype simply a single exemplar? If a robin is the prototype of bird, do particular examples of robin constitute that prototype for different people? I think not. Rather, prototypes are themselves categories. Thus, to say that a robin is a prototypic bird is to propose that a class of similar creatures called *robin* is a prototype of bird.

Categories which constitute prototypes, like all categories, relate to propositional models and, in specific instances, may also relate to Gestalt properties. Above, in passing, I assume that a propositional model ordinarily specifies only the properties of a prototype of the class to which it relates rather than attributes of all possible members of the category. In making this assumption, I am following Wierzbicka (1985a). Thus, the

propositional model outlined for squirrel above lists those properties of squirrels that pertain to a prototypic squirrel. Most of the listed attributes also happen to pertain to squirrels in general, for example, squirrels 'store things to eat' and 'have a big bushy tail'. One of them, i.e., squirrels 'are brown in colour', does not pertain to squirrels in general, but rather only to my thoughts concerning the colour of a prototypic squirrel. Throughout this chapter, propositional models are to be understood as specifying properties of prototypes of categories rather than attributes of categories (although in some instances, such as *robin*, properties of prototypes and properties of categories turn out to be exactly the same).

Categories which serve as prototypes, then, are always underlain by propositional models which specify attributes of prototypes of those categories. In addition, categories which constitute prototypes, in some cases, may also be underlain by Gestalt properties. Thus, the prototypic bird, i.e., robin, relates to the Gestalt property of 'robinness' in addition to a propositional model which specifies the properties of a prototypic robin. An important implication of these observations is that prototypes themselves may have prototypes. An example of this is cited presently. Robin, however, is not such an example since the attributes of robin and the attributes specified in a propositional model relating to robin are the same.

While it is possible, it is highly unlikely that any significant portion of English-speaking people recognize a variety of robin among all birds called *robin* that they consider most robin-like. Little introspection is required to convince oneself that there is no prototypic robin and that this must be true for the vast majority of English-speaking people (specialists in ornithology possibly excepted). This is to suggest, then, that some labelled categories, such as *robin* and *raccoon* and others mentioned below, do not have prototypes at all, and, consequently, probably have not developed through prototype/extension.

In summary, some categories of natural language involve prototypes and extensions from prototypes resulting in inclusion of non-prototypes (+ P/E), and some categories do not have prototypes and probably have not developed through extension from prototypes (− P/E).

The survey

Category Type 1: − AR + GM − P/E

Above I propose that the category *robin* does not have a prototype and probably did not develop through prototype/extension. Consequently, it is assigned to a category type of natural language having the feature − P/E. In addition, *robin* is a biological category and, hence, shows the

feature − AR (non-artifactual). Finally, since *robin* relates to an unambiguous discontinuity in nature, it is almost certainly motivated by a Gestalt quality, i.e., 'robinness', and, hence, shows + GM. The features − AR, + GM and − P/E define a Type 1 category to which *raccoon*, *black walnut tree*, and *dandelion*, as well as many other categories of living things, in addition to *robin*, belong.

The reason why categories such as *robin*, *raccoon*, *black walnut tree*, and *dandelion* lack prototypes is threefold. First, exemplars of these categories in nature are very homogeneous. This point is underscored by the fact that these folk biological categories, with the exception of *dandelion*, bear a one-to-one relationship to scientific species for which biotaxonomists recognize no varieties. Since varieties of robin (the North American version, *Turdus migratorius*), raccoon, and black walnut are not identified by scientists, it should not be surprising to learn that ordinary folk do not recognize them either and, consequently, do not have a basis for distinguishing respective prototypes.

Second, relative lack of human interest in exemplars of such categories combines with class homogeneity to produce prototypeless groupings. For instance, even if scientific varieties of robins exist, humans may pay these creatures so little attention that varietal or other distinctions are not noticed. Consider *dandelion*. Most speakers of American English would be temporarily lost for words in response to the question 'What is the most dandelion-like dandelion you can think of?' An eventual response typically given might be 'All dandelions are alike, silly!' In fact, all things called *dandelion* by ordinary speakers of American English are not the same scientifically, since three species, one having two varieties, are recognized (at least in north-eastern and north-central North America, cf. Peterson and McKenny 1968: 170). That there is no prototypic dandelion for most people is linked to the fact that they do not distinguish differences among dandelions mainly because they have little positive interest in them and, consequently, pay them precious little attention.

Third, and, perhaps, most important is that these categories lack prototypes in part because prototypes have not been and are not now necessary to their development (as category development has been defined above). For example, dandelions are so similar to one another (homogeneous) that experiencing just a few exemplars should rather rapidly lead to the recoding of their attributes into the Gestalt property 'dandelionness'. This configurational quality, rather than a prototype, would underlie the diachronic development of the category in languages, the acquisition of the category by children and second language learners, and the decision of individuals to include in the class or exclude from the class potential exemplars.

When the membership of a category is especially homogeneous and people pay it little attention, prototypes are rarely, if ever, involved in category