

author

Michal Měchura

assignment title

Linguistic Categorization

degree

M.Phil. in Speech and Language Processing

module

LI 7870 Advanced Syntactic Theory

term

Hilary 2008

word count

4,091 words

Linguistic Categorization

1. Introduction

When a speaker produces a sentence in which it is difficult to determine whether a particular word is a noun or an adjective, what is the significance of this indeterminacy? One possible answer is that the speaker had only one particular category assignment in his or her mind, and the observed indeterminacy is an example of unintended ambiguity. Another possible answer is that the speaker intended the ambiguity, perhaps for rhetorical effect. A third possible answer is that there is no ambiguity at all, it is merely a case of the linguist's categories being inadequate. Otto Jespersen (1924, p. 183) expresses the opinion that word class ambiguity only exists when words are taken in isolation – while in actual usage, when context is available, word class assignment is obvious. Still, a lot of ink is spilled in the linguistic literature on questions such as “is the word *boy* in the phrase *a boy actor* a noun or an adjective?” It is straight question: is the phrase *a boy actor* ambiguous? If we answer yes, we will have to conclude that humans routinely produce ambiguous language. For reasons of parsimony, we should assume that humans ordinarily prefer not to speak and write ambiguously, and we should investigate the option that our categories of noun and adjective are inadequate. That would suggest that there is a word class that both *boy* and, say, *famous*, are members of – it is the class of words which can modify head nouns in a noun phrase.

Where does that leave the original notions of noun and adjective? Is this “third” class separate from them? Or is it the result of intersection between them? Are the boundaries of those categories crisp or fuzzy? And what if there are no such categories at all, merely individual attributes that a word may or may not have, such as the ability to modify words that have the ability to be modified?

Clearly, the central question here is one of linguistic categorization. How are categories such as word classes defined, and how is their membership decided when one encounters a word in a sentence? Such questions appear similar to questions of

categorization in general: how are categories such as DOG and CAT set up, and how do we decide whether an animal we are looking at right now is a dog and not a cat? Explaining categorization is more difficult than may seem at first, regardless of whether one is dealing with categories in linguistics or outside it.

The following section will review some observations and opinions on the categorization of entities such as physical objects, emotions, events, and so on – altogether entities other than linguistic objects like words and sentences. The motivation is that the treatment will provide useful insights into the categorization of linguistic objects, a subject we will return to in section 3.

2. Categorization

Categorization is the process of assigning individuals to categories. When you come across an entity in the real world, be it something concrete such as a physical object you are looking at right now or something abstract such as an emotion you are feeling at the moment, you instinctively assign it to some category or other: you conclude that the object you are looking at is a bird, or that the emotion you are feeling is fear. Humans do this because it helps them think about the entity and to communicate the entity to other humans by means of language. The process of categorization is possible because, presumably, humans carry in their heads a large ontology of categories. Importantly, the ontology is set up in such a way that practically every individual entity a human comes across during the course of his or her life can be assigned to some category or other.

There are three important questions to ask about categorization. Firstly, how are categories defined, what form do they take in the human mind and how can they be modelled outside it, for example in a computational grammar? Secondly, how do categories develop in the mind? Are they induced by the developing mind, are they taught explicitly, or are they acquired by means of some combination of both? And thirdly, what is the actual process that humans follow when assigning an individual to a category? These questions are important in philosophy and

psychology as well as linguistics. The following is an overview of some approaches to answer them.

2.1. *Ontologies*

An ontology, in this context, is the complete set of categories one has available in his or her mind. Human ontologies are large and the categories in them appear to be organized in networks of relationships such as subsumption (ANIMAL subsumes BIRD and DOG), partonymy (DOOR, WINDOW and ROOF are parts of HOUSE), telic relationships (CAR and BUS are both for TRANSPORT), and so on (Weissenhofer 1995, ch. 1.5). Some sections of these networks appear to be organized in top-down hierarchies that resemble artificially designed taxonomies, such as the familiar Linnean taxonomy of biological categories. Within such hierarchies, a basic level can often be found which contains the most readily recognized categories. For example, when looking at a coffee table, you would probably first categorize it as a TABLE, not as a COFFEE TABLE (which would be too specific for ordinary purposes) or as FURNITURE (which would be too generic). Similarly, when looking at a house, you would probably conclude that you are looking at a HOUSE, and not at an unstructured collection of WALLS, DOORS and WINDOWS. Therefore, the level in the hierarchy that concepts like TABLE and HOUSE occupy, can be characterized as the basic level (Rosch 1978, pp. 93–97 and Weissenhofer 1995, ch. 3.2.2).

Interestingly, individual humans appear to have by and large identical sets of categories in their heads – that is part of the reason why we can understand each other when we communicate. There are minor individual differences, though. For example, not all people would agree that geographical territories like Wales, Scotland or Catalonia are countries. That suggests that people have slightly differing intuitions about what it means to be a country. Nonetheless, while we may disagree about some borderline cases, we all agree that clear-cut cases like France and Spain are countries. There tend to be cross-cultural and cross-linguistic differences between category systems as well: for example, in some cultures, rhubarb is

classified as a fruit while in others it is a vegetable. Some languages have lots of basic-level words for different kinds of bread while others recognize just *bread*.

2.2. Principles of categorization

Of all possible ontologies, why does the human mind have the particular ontology it has? We can follow Rosch (1978, pp. 92–93) to explain why some categories form in the mind while some other possible categories do not. First, the forming of categories is driven by a need for cognitive economy: we want to have the smallest possible number of categories in order to use up as little memory as possible, but at the same time we want to have the largest possible number of categories so that we can give attention to all relevant distinctions when categorising entities. If we only had a small number of very general, catch-all categories, we would lose sight of some important details in the real world: imagine not being able to distinguish between kinds of animals, such as between chickens and tigers. On the other hand, if we only had a large number of very detailed categories, we would lose sight of some important generalizations: imagine not being able to realize the commonalities between coffee tables and kitchen tables. The actual ontology we have is the result of a compromise between these two concerns.

That may help to explain why there are as many categories as there are but, it still does not explain why some categories are created while other possible ones are not. Rosch (1978, p. 92) observes that the world as we perceive it is not a random collection of equiprobably co-occurring attributes. It is a fact of life that, for example, the attributes “has feathers” and “has wings” occur together more frequently than in separation or in other combinations. Such clusters of attributes are more likely to become codified as categories in the human mind than other, uncommon combinations. We will return to the notion of categories as clusters of attributes later when we discuss the nature of syntactic categories.

2.3. *The “classical crisp categories”*

One property of categories that deserves scrutiny is their boundaries, and specifically, the question whether they are crisp or fuzzy. If one entity is categorized as a bird and another also as bird, can one of them be more “bird-like” than the other? Additionally, can something be a bird and, say, a mammal at the same time? There is a tradition in Western thought that likes to answer “no” to such questions and treat categories as crisp sets. A category is defined by listing any and all necessary and sufficient conditions for membership. Category assignment is constrained by the all-or-none principle (also known as the law of the excluded middle) which says that an entity is either in the set or outside it: there are no in-between cases. Additionally, all members of a category have the same status: no bird is more bird-like than another. This way of thinking about categories has its roots in the philosophy of Aristotle and has had a strong influence on modern Western thinking and society. Practically all categories created artificially are crisp. Things like country borders and voting age are crisp, for example. You are either in one country or another, but you cannot normally be in both. If you are younger than the legal voting age you cannot vote, even if you only fall short of the threshold by a single day.

The classical crisp approach runs into difficulties when applied to explain the meanings of words and the concepts they designate. It is rather difficult to craft full and exhaustive definitions for concepts such as CAT, HAPPINESS, VEGETABLE, FRUIT, COUNTRY OR NATIONALITY by listing necessary and sufficient conditions. Why is this difficult? The key is to notice that lexical concepts are, by and large, not artificially designed like country borders or voting age. Most people’s ideas of what a cat is and what happiness have formed spontaneously and independently of any conscious decisions – perhaps with the exception of some specialized terminology and foreign-language vocabulary which one may have acquired by definition. The fact is, we all know a cat when we see one but we don’t understand why – and it seems that the classical crisp approach is not providing any useful insights into the enigma.

2.4. Gradience

Some of the problems with defining lexical concepts go away if we introduce the possibility of gradience: that is, the possibility for categories to not be “classical crisp”. It is useful to follow Aarts (2007) and to recognize two kinds of gradience:

- Subjective gradience, or gradience of membership within a category. Cases of typicality are an examples of subjective gradience: most people would agree that a robin is a more typical bird than a penguin. Thus, robin is an instance of the concept BIRD to a larger degree than a penguin is.
- Intersective gradience, or gradience between categories. If it is reasonable to believe that a bat is both a bird and a mammal, then that is an example of intersective gradience: a bat in that case would be an instance of both the concept BIRD and the concept MAMMAL.

The question now is, are lexical concepts crisp or gradient? If they are gradient, which of the two types of gradience do they possess? A demonstration of the difficulty to find crisp boundaries for lexical concepts is Ludwig Wittgenstein’s attempt to define the concept GAME (Wittgenstein 1953). He concludes that there is no property common to all instances of GAME, which we can read as implying that GAME is not a classical crisp category. Wierzbicka (1990), on the other hand, claims to have found such common properties, and says that the concept of GAME does have crisp boundaries. We can try to reinterpret the two positions in terms of the two types of gradience introduced above:

- For every individual event, is it possible to decide whether it is a game or something else (a fight, a military exercise, ...) but nothing in between? If not, then the concept GAME displays intersective gradience, and possibly other concepts do too.
- For all individual events that are games, are some more “game-like” than others? If so, then the concept GAME displays subjective gradience, and possibly other concepts do too.

Wittgenstein in effect claims that *GAME* does display intersective gradience and Wierzbicka claims that it does not. Neither says anything about subsective gradience. The question whether lexical concepts are gradient or not is unresolved. However, presuming that they *are* gradient (both intersectively and subsectively) has proved fruitful in providing powerful models of categorization.

2.5. *Prototype theory*

The idea that concepts are gradient has been developed further in prototype theory. In prototype theory, prototypes are the clearest cases of category membership (Rosch 1978, pp. 98–99). For example, the robin may well be a prototypical bird while a penguin is not. The prototype of a category can be an individual (such as one particular robin I remember having seen), a more specific category (such as the category of robins) or an abstraction, such as an abstracted amalgamate of robins, sparrows and blackbirds that only exists in the mind. In either case, a category is defined not in terms of its boundaries but in terms of its prototypes. Note that categories such defined effectively become fuzzy sets where membership is no longer a yes-no question, it is a question of degree. If we return to the idea of categories as clusters of attributes, then the more attributes an individual entity shares with the prototype(s) of a category, the more it is a member of that category. Prototype theory supports both subsective and intersective gradience:

- If an individual entity shares a criterial number of attributes with the prototypes of two categories, then it is a (peripheral) instance of both categories – this is an example of intersective gradience.
- If an individual entity shares a criterial number of attributes with the prototype(s) of only one category, but does not share all attributes with the prototype(s) of that category, then it is a peripheral instance of just that one category – this is an example of subsective gradience.

3. Linguistic categorization

We will now return to the problem of categorizing linguistic objects such as words, phrases and sentences and see if any of the ideas introduced in the previous section help in investigating the nature of linguistic categories. Most of the discussion will revolve around words and conventional word classes such as nouns and verbs, but the questions addressed and the answers suggested will be relevant to the categorization of larger units as well, such as phrases and clauses.

We must distinguish two sets of linguistic categories: those that are invoked in grammar books and other artefacts of linguists' work, and in opposition to those, those that presumably exist in the human mind. Linguistics has had occasion to coin terms such as noun and verb because it believes that some correlates of those concepts exist in the human mind and are invoked – unconsciously – during human language processing. What are such categories for? In the mind, they presumably serve as shortcuts to access information about the attributes of individual words – just like the concept *CAT* serves as a shortcut to access information about the attributes of cats in the real world. In linguistics, word classes are set up to describe the behaviour of words (syntactic and otherwise) at a level of useful generalization. Linguistic categorization is driven by the same requirement of cognitive economy as categorization in general: describing word behaviour individually for each word would lose sight of significant generalizations.

3.1. *How linguistic categories are defined*

A word class is typically explained in the literature as an aggregate of attributes. There are different kinds of attributes that go into explaining a word class, including the following (after Crystal 1967):

- **Morphological:** some words are (or could in principle be) classified based on their morphological make-up. In English, words that end in a particular morpheme, such as *-ism* or *-ology*, are nouns, word ending in *-ize* or *-ify* are verbs, and so on. Similarly, the inflectional endings a word is capable of taking

can betray the word's class: words ending in *-ly* are typically adverbs, and so on. The difficulty, of course, is that in most languages, such clues are only that: they are mere clues, and additional criteria need to be considered as well.

- **Semantic:** word classes are often outlined intuitively in terms of what they denote, such as the notorious definition of a noun as denoting a “person, place or thing”. There is, however, a considerable lack of correspondence between such notional categories and grammatical word classes: obviously, not all nouns are people, places or things. Still, there seems to be something intuitively satisfying about such “definitions”, and they constitute a good start for somebody to start thinking about word classes. Possibly, word classes (or, more precisely, their mental correlates) emerged during the evolution of language in correspondence to exactly such notional categories, but later diverged from them and acquired a life of their own through a process of “disassociation of variables” (Halliday *et al.* 2004, chapter 1.2).
- **Syntactic:** in many cases, the membership of a word in a word class can be at least partially determined by answering tests such as “can it appear as the subject of a sentence?”, “can it co-occur with a determiner?”, “can it appear in attributive position to a noun?” and so on. Notice that the tests make reference to other word classes as well as other kinds of linguistic categories, such as sentence functions (subject). This causes the definitions to be circular, which may or may not be a problem, depending on the purpose of the grammar and its intended use. But in effect, such criteria are abstractions from substitution tests in which we test which words can and cannot be inserted grammatically in frames such as “the ~ man” or “the man is ~”.

In theory, a word class should be definable by some combination of such attributes: morphological, semantic and syntactic. For example, the class of adjectives could, at least in principle, be defined by answering tests such as the following:

- Can the word inflect for degree?
- Does the word refer to a quality?

- Can the word co-occur with intensifiers such as *very*?
- Can it occur in the frame “a/the ~ N” where N is a noun?
- Can it appear in a predicative position after verbs like *be*, *seem* and *become*?

The key question is, how many and which attributes should be chosen, and what weights should be attached to them? Similarly to non-linguistic objects, words do not appear to come with random collections of equiprobably co-occurring attributes. Some clusters of attributes occur together more often than separately or in other combinations. Thus, for example, the ability to inflect for degree co-occurs with the ability to appear in a predicative position more often than, for instance, with the ability to inflect for tense.

3.2. Gradience in linguistic categories

It appears that linguistic categories display the same properties of gradience as non-linguistic ones. For example, when investigating the properties of words we commonly class as adjectives, we observe that:

- some can be graded but others cannot (*bigger* but **aliver*),
- some can modify a noun but others cannot (*a brave man* but **an afraid man*),
- and some can appear in predicative position but others cannot (*the failure is obvious* but **the failure is utter*).

In fact, there does not appear to be a single attribute that all adjectives share and no other word class possesses. The same difficulty has already appeared in Wittgenstein’s attempt to define the concept *GAME*. Such “defining” attributes are either extremely elusive, or do not exist at all. Word classes might then be defined in terms of prototypes: the central members of a word class would possess all relevant attributes while more peripheral members only some, thus displaying subsecutive gradience. Additionally, words that share the attributes of more than one class would, in effect, display intersective gradience.

3.3. Basic-level categorization

Word classes such as nouns and adjective seem to constitute the basic level of linguistic categorization: they are the first categories that spring to mind when the question of categorization arises. They are more basic than over-specific category assignment such as “manner adverbs” and “proper nouns”, and more basic than over-general assignments such as “closed class”, “grammatical word”, and so on. For some purposes, however, staying at the basic level is not granular enough, even if one adopts a model of gradience such as prototype theory. Declaring that a word, for example *afraid*, is a peripheral adjective, is not sufficient. If we wish to write a generative grammar, we need to know which attributes exactly the word has and which ones it does not have. Describing it just as an adjective is not granular enough.

It seems from the account so far that a great deal is known about the nature of linguistic categories, especially word classes – even if there is not much agreement as to what the word classes actually are. Very little is understood about the nature of the attributes that are used to define them, however. For example, it is not certain whether the number of attributes needed to describe the behaviour of all adjectives is finite or infinite. If attributes have a correlate in the human mind, then it would be reasonable to assume that their number must be finite, as human memory has a finite capacity. Another unresolved question is whether attributes, like the categories they define, can be gradient. So far, it seems that a word either has an attribute or does not: an adjective either can appear in predicative position or cannot. Can there be in-between cases? Some evidence suggests that gradience could exist in attributes too. For example, the adjective *alive* is normally believed to be unable to appear in an attributive position before a noun: we cannot say **an alive man*. However, there are attested examples of exactly such usage, which, in their context, do not seem like performance errors. Does this mean that *alive* possesses the attribute only to a certain degree? Or is it simply an example of insufficient

analysis? Would we, if we tried hard enough, find a more accurate set of attributes that describe *alive*'s behaviour without recourse to gradience?

This investigation of linguistic categories started with the statement that categories are set up in order to describe the behaviour of words at a useful level of generalization. It appears that individual attributes – and not classical word classes like nouns and adjectives – are the right level of description if one wishes to describe the behaviour of individual words. It is obvious that if one stays at the level of nouns and adjectives, one loses sight of some important details. Declaring that *afraid* is an adjective does not say whether it can or cannot appear in predicative position. Declaring that *afraid* is a peripheral adjective does not answer the question either. Thus, the category of adjective is not fine-grained enough for many tasks in linguistics. Perhaps, individual attributes should be the proper categories descriptive linguistics investigate.

4. Conclusion

It appears that there are strong parallels between the categorization of non-linguistic objects such as animals and emotions, and linguistic objects such as words and phrases. They are both driven by a need for cognitive economy, they both appear to have a level of basic categorization, and they both show signs of gradience. The purpose of this essay has been to highlight that. Additionally, turning to linguistic categories, the essay has attempted to dissect the notion of a word class and to show that, at least for some purposes, the proper object of linguistic categorization should be the individual attributes that define word classes.

References

Aarts, B. (2007) *Syntactic Gradience: The Nature of Grammatical Indeterminacy* Oxford: Oxford University Press

- Crystal, D. (1967) 'English', in *Lingua* 17. Reprinted as 'English Word Classes' in B. Aarts; D. Denison; E. Keizer; G. Popova (2004) *Fuzzy Grammar: A Reader* Oxford: Oxford University Press
- Halliday, M. A. K.; Ch. Matthiessen (2004) *An introduction to functional grammar (3rd edition)* London: Arnold
- Jespersen, O. (1924) 'Parts of Speech' chapter 4 of *The Philosophy of Grammar* London: George Allen & Unwin. Reprinted in B. Aarts; D. Denison; E. Keizer; G. Popova (2004) *Fuzzy Grammar: A Reader* Oxford: Oxford University Press
- Rosch, E. (1978) 'Principles of Categorization' in E. Rosch; B. B. Lloyd (eds.) *Cognition and Categorization* Hillsdale, NJ: Lawrence Erlbaum Associates. Reprinted in B. Aarts; D. Denison; E. Keizer; G. Popova (2004) *Fuzzy Grammar: A Reader* Oxford: Oxford University Press
- Weissenhofer, P. (1995) *Conceptology in terminology theory, semantics and word formation* Vienna: TermNet
- Wierzbicka, A. (1990) ' "Prototypes save": on the uses and abuses of the notion of "prototype" in linguistics and related fields' in S. L. Tsohatzidis (ed.) *Meanings and Prototypes: Studies in Linguistic Categorization* London: Routledge. Reprinted in B. Aarts; D. Denison; E. Keizer; G. Popova (2004) *Fuzzy Grammar: A Reader* Oxford: Oxford University Press
- Wittgenstein, L. (1953) 'Family Resemblances' in *Philosophical Investigations* Oxford: Blackwell. Reprinted in B. Aarts; D. Denison; E. Keizer; G. Popova (2004) *Fuzzy Grammar: A Reader* Oxford: Oxford University Press