3.2 Prototypes

3.2.1_Introduction

We have discussed the properties of vague concepts, and found that they can be modeled within a semantic framework that captures the fact that a particular object might belong to a given concept to a particular **degree**.

Another set of findings that appears to be related to vagueness is that for many concepts we can distinguish between good and not-so-good instances. A number of quite stable experimental findings points towards that direction. The best instances are called the **prototypes** of a concept. Prototype theory was developed in particular by Eleanor Rosch and her collaborators in a series of very important experiments in the seventies. We will discuss some of her major findings and theoretical models that were developed to interpret them.

3.2.2_Evidence for Prototypes

Some pieces of evidence for prototypes are:

Subjects generally can answer the question whether a particular object is a good example of a category, and that there is high agreement between subjects (cf. Rosch (1975)). For example, when asked to rank given items in terms of how good examples of *furniture* they are, the following list emerges: *chair, sofa, couch, table, easy chair, dresser, rocking chair, coffee table, rocker, love seat, chest of drawers, desk, bed, bureau, davenport, end table* etc; ending in *stove, counter, clock, drapes, refrigerator, picture, closet, vase, ashtray, fan, telephone*. That is, people are aware of how good a category instance a particular object is.

Category statements of the form *An X is a Y* can be verified quicker if X is a good instance of a category Y (Rosch (1973)). This holds especially for children.

If asked to **enumerate** instances of a concept, subjects in general start with the better examples and name less good examples only later (Rosch 1973, Battig & Montague (1969)).

Prototypes show up in **priming effects**. If subjects have the task to decide whether two words X_1, X_2 are identical after they have been exposed to a priming word Z that denotes a superordinate concept, then they can decide quicker if X_1/X_2 is a good example of Z than if it is a less good example of Z, or if it is no example of Z at all. This suggests that mentioning Z activates the more prototypical instances of Z (Rosch 1975).

Interestingly, prototypicality effects also show up when the superconcept is not even made explicit. Rosch & Mervis (1975) asked subjects to name properties of a list of items that belong to different superconcepts. Pairs of items that are prototypical for a category share more common properties than items than pairs of items in which one or both members are less prototypical for a superconcept.

Another technique that does not ask for superconcepts is **multidimensional scaling**. In this technique, subjects are asked to rank pairs of items that are instances of a superconcept, and pairs of the superconcept and instances of the superconcept. For example, with birds *robin* and *sparrow* are generally ranked more similar to each other than the pair *robin* and *duck*. The similarities between words is then rendered by the distance of points in a space, typically in a two-dimensional plane or also in a higher-dimensional space. (There are various algorithms that find such spatial representations for any given ranking of pairs.) The prototypical members of a category appear relatively closely together, and closer to the point representing the superconcept than other points or point clusters. Cf. the following representation for the similarity of various birds. Notice that the horizontal dimension represents variations in size, and the vertical dimension represents variations in ferocity, two common hidden dimensions that subjects use in classifying animals.

(30)

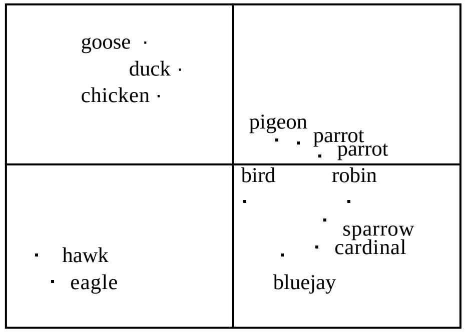

Short presentation: Discuss the first four experiments in Rosch & Mervis (1975).

3.2.3_What Prototypes tell us about Concepts

Prototype effects where initially used to show that the classical view of concepts as a set of defining criteria, like a set of features, could not be right. The argument goes as follows: If a concept were defined by a number of necessary and sufficient criteria, then subjects should not show any prototypes at all. If asked whether a given instance A belongs to a concept B, a speaker would check whether the criteria for B-membership apply to A. If yes, A is counted as an instance of B. As all the criteria of B-membership have to be checked, we should not expect that any instance is treated preferentially.

Smith & Medin (1981) show that this is not necessarily right. First, we can plausibly assume that the defining features of lower-level concepts are supersets of the defining features of higher-level concepts, and that less typical instances have more features (we already mentioned this possibility), according to the following schematic model:

(31)	*Animal*	*Bird*	*Robin*	*Chicken*
	F_1	F_1, F_2	F_1,F_2,F_3	F_1,F_2,F_3,F_4

In this model, *robin* is a better instance for *bird* than *chicken* because it differs from *bird* in fewer features. Also, it is a better instance for *bird* than for *animal* because of the same reason. Tversky (1977) has suggested an explanation of similarity ratings in which (i) a concept A is more similar to a concept B if A shares more features with B, and (ii) a concept A is less similar to a concept B the more features A contains that are not in B.

Another problem posed by prototypicality effects to the classical view is that non-necessary features are critical. For example, flying birds are generally considered more typical, even though the ability to fly is not a necessary property of birds.

Furthermore, it is problematic that there are unclear cases of concept membership.

Smith & Medin 1981 discuss various modifications of the classical view. One is to take the features not as absolute, but as probabilistic. That is, we do not just assign features to concepts, but features plus the information about their probability and salience (the figures here represent the combined probability and salience of a feature).

(32)

Robin	*Chicken*	*Bird*	*Animal*
1.0 moves	1.0 moves	1.0 moves	1.0 moves
1.0 winged	1.0 winged	1.0 winged	
1.0 feathered	1.0 feathered	1.0 feathered	
1.0 flies	1.0 walks	0.8 flies	0.7 walks
0.9 sings		0.6 sings	
0.7 small	0.7 medium	0.5 small	0.5 large

This approach certainly can explain why there are unclear cases of being an instance of a concept and why non-necessary features may be crucial in determining membership. Also, prototypicality effects can be explained in a straightforward way: The prototypical instances retain the non-necessary features of a superconcept (e.g., fly and sing for birds), whereas less prototypical instances do not retain at least some of these features (e.g., chickens don't fly). Similarity judgements can be explained by the number and strength of shared features between two items. Also, similarity to other items has to be factored in: An instance A is less similar to a concept B, the more similar it is to a different concept C that contrasts with B.

We should assume that the non-necessary features of the superconcept are derived from the fact that many instances have those features. The probabilistic model supports Rosch & Mervis (1975) in their finding that prototypical members of a category share more features: These are the features on which the non-necessary features of the superconcept are based.

A problem of this approach is that the features are relatively unrestricted -- we have features like 'large', 'medium', 'small', that should be related. Smith & Medin 1981 suggest another model in which all instances of a concept are classified according to the same dimension. Example:

(33)

Robin	*Chicken*	*Bird*
1.0 animate	1.0 animate	1.0 animate
0.6 size	0.7 size	0.5 size
0.4 ferocity	0.4 ferocity	0.5 ferocity
1.0 flies	0.1 flies	1.0 flies

Problems of this approach are that it cannot capture relations between features (e.g., between the feature of being winged and the feature of being able to fly [an imperfect relation, to be sure]), or the featues for locomotion, 'walks' and 'flies'.

Smith & Medin 1981 also discuss a view that does not work with features at all, the **exemplar view**. The basic idea is that concepts are represented by their examplars, not by an abstract description in terms of features. More specifically, a concept is represented by all or the prominent exemplars encountered and classified so far. On classifying a new instance, speakers perform a global comparison of the new instance with the exemplars, using a general similarity relation. The idea is that, for example, a new instance of a robin will show greater similarity to all known instances of a robin than, say, to all instances of a bluejay. Prototypicality effects can be explained by assuming that not all exemplars are stored, but just "typical" exemplars, or that typical exemplars are used first for comparison with a new instance. This model is problematic as well for several reasons. In particular, it leaves unspecified which properties similarity judgements are based on, it does not mention possible relations between features (like being winged and being able to fly).

3.2.4 Defining Features vs. Identification Features

Armstrong, Gleitman, & Gleitman (1983) replicate some of the findings of Rosch. They extend this type of research to include exact concepts like *odd number* or *triangle*. Their surprising finding is that even though subjects are fully aware of the fact that e.g. *odd number* is precisely defined, there are prototypicality concepts. For example, they find the following prototypicality ranking for odd and even numbers of two sets each, where subjects where asked how good an example for an odd of even number each number is.

(34) a. Even numbers: 4, 8, 10, 18, 34, 106
2, 6, 42, 1000, 34, 806

b. Odd numbers: 3, 7, 23, 57, 501, 447
7, 11, 13, 9, 57, 91

Factors that are relevant here are that high numbers are considered less typical for either concept, that factors of 10 are considered good examples for even numbers (cf. 1000), that numbers that contain more even digits are considered better examples for even numbers than numbers that also contain odd digits (cf. 42 vs. 34).

Armstrong, Gleitman & Gleitman argue that one has to distinguish between **defining features** (the definitional criteria of the classical analysis) and features that are used to **identify** whether an instance belongs to a concept. Even if the defining features are strict, the identificational features can lead to prototypicality effects.

> Short presentation: The experiments and discussion in Armstrong, Gleitman & Gleitman 1983.

3.3 Vagueness and Prototypes

3.3.1 The application of fuzzy sets to prototypicality effects

It seems quite natural to apply the tools developed for dealing with vagueness to prototypicality effects. If a concept A is represented by a fuzzy characteristic function c_A, then prototypical members should be mapped to a value close to 1, whereas less prototypical members should be mapped to a value substantial less than one, e.g.:

(35) a. If x is a robin: $c_{Bird}(x) = 1.0$
b. If x is a chicken: $c_{Bird}(x) = 0.8$
c. If x is a moa: $c_{Bird}(x) = 0.4$

This was suggested (and criticized) by Osherson & Smith (1981). More specifically, they suggested the following representation of concepts as a quadruple :

(36) Concepts: A, d, p, c ,
where A: a conceptual domain of objects (e.g., all birds)
d: a function from A×A to positive numbers (a distance metric) that expresses similarity
(e.g., if r is a robin, b a bluejay, o an ostrich: d(r,b) < d(r,o))
p   A the prototype of the concept, (e.g., a particular robin r),
and c is a function from A to [0,1], the characteristic function of the concept.

The distance function d should represent similarity between concepts. It has to satisfy certain requirements (a-c), and it has to stand in a monotone decreasing relation to the characteristic function (d):

(37) a. $d(x, x) = 0$
b. $d(x, y) = d(y, x)$
c. $d(x, y) + d(y, z) \geq d(x,z)$
d. $d(x,p) \leq d(y,p) \rightarrow c(y) \leq c(x)$.

This appears to be a fair implementation of prototypicality, in particular, of the exemplar view discussed by Smith & Medin 1981.

3.3.2_Problems with conceptual combination

Osherson & Smith point out problems of this representation when it comes to conceptual combination. The typical way of concept combination by intersection, union and negation in fuzzy set theory is as follows (recall that union and intersection are well motivated by general requirements; we could perhaps design different functions for negation):

(38) a. $c_{A \cap B} = \lambda x[\text{MIN}\{c_A(x), c_B(x)\}]$
b. $c_{A\ \ B} = \lambda x[\text{MAX}\{c_A(x), c_B(x)\}]$
c. $c_{\neg A} = \lambda x[1 - c_A(x)]$

They now discuss whether it is possible to derive composite concepts from their parts. For example (an example due to Fodor): How can we derive the concept of a *rich man's hous*, together with its prototype, given the concept of a *rich man* and its prototype, and the concept of a *house* and its prototype?

Osherson & Smith just consider simple adjective-noun combinations, and show that their represenation leads to problems. They consider the complex concept *striped apple* as an example. Consider the following objects:

(39) **a**: a regular apple with regular stripes;
b: a regular apple without stripes;

The meaning of *striped apple* can be given as follows, in terms of the characteristic functions of *apple* and *striped*:

(40) $c_{\text{striped apple}(}(x) = c_{\text{striped} \cap \text{apple}}(x) = \text{MIN}\{c_{\text{apple}}(x), c_{\text{striped}}(x)\}$

Now we clearly should have that **a** is a better exemplar of a striped apple than of an apple, as regular apples don't have stripes, which we can express as (a). But (a) contradicts (b) (which is implied by (40)), as this implies (c), a direct contradiction to (a).

(41) a. $c_{\text{striped apple}}(\mathbf{a}) > c_{\text{apple}}(\mathbf{a})$
b. $c_{\text{striped apple}}(\mathbf{a}) = \text{MIN}\{c_{\text{apple}}(\mathbf{a}), c_{\text{striped}}(\mathbf{a})\}$
c. $c_{\text{striped apple}}(\mathbf{a}) \leq c_{\text{apple}}(\mathbf{a})$

Osherson & Smith also point out that we can form contradictory concepts, like *apple that is not an apple*, that end up classifying a borderline case of an apple (say, a 0.5 apple) as belonging to them to a certain degree (here, to the degree 0.5):

(42) a. $c_{\text{apple that is not an apple}}(x) = c_{\text{apple} \cap \neg \text{apple}}(x) = \text{MIN}\{c_{\text{apple}}(x), 1 - c_{\text{apple}}(x)\}$

We have encountered this problem in our discussion of multi-valued logic.

3.3.3_Supervaluations and prototypes

Kamp & Partee (1995) discuss the problems that Osherson & Smith encountered with the modelling of prototype theory by fuzzy sets. They show that using supervaluations can overcome some but not all the problems.

Kamp & Partee propose that interpretations are dependent on models M which distinguish between **positive** and **negative** extensions. Simple nouns are interpreted with respect to fuzzy characteristic functions c, for example:

(43) a. x $[\![\textit{apple}]\!]^{+}_{M}$ iff $0.95 \leq c_{\text{apple}}(x) \leq 1$
b. x $[\![\textit{apple}]\!]^{-}_{M}$ iff $0 \leq c_{\text{apple}}(x) \leq 0.05$

If c_{apple} maps x to values between 0.05 and 0.95 it is neither considered an apple nor not an apple. The specific numbers for what should count as an apple are a feature of the model M. A more liberal model M′, for example, would set the borderline for applehood at 0.8, or the borderline for non-applehood at 0.3, o both. Only those models M are admissible for which the following holds:

(44) If x $[\![\textit{apple}]\!]^{-}_{M}$ and y $[\![\textit{apple}]\!]^{+}_{M}$, then $c_{\text{apple}}(x) < c_{\text{apple}}(y)$

For any model M, a possible **completion** of M that is consistent with the fuzzy characteristic functions that are used in the model definition of M is the set of all functions M′ that satisfy the constraint (44). Let M* be the set of all completions of M; this is a **supervaluation** based on M. We now can definite the truth value of a sentence in terms of supervaluations:

(45) a. $[\![\Phi]\!]^{M*} = 1$, if $[\![\Phi]\!]^{M'} = 1$ for all completions M′   M*,
b. $[\![\Phi]\!]^{M*} = 0$, if $[\![\Phi]\!]^{M'} = 0$ for all completions M′   M*,
c. undefined otherwise (that is, if 0 in some and 1 in others)

This solves Osherson & Smith's puzzle of an apple that is not an apple. Let **c** be a borderline case of an apple, with $c_{\text{apple}}(\mathbf{c}) = 0.5$. We can represent the sentence that c is an apple by **c** $[\![\textit{apple}]\!]$, and *apple that is not an apple* by fuzzy set conjunction:

(46) a. x $[\![\textit{apple that is not an apple}]\!]^{+}_{M}$
if $0.95 \leq [\![\textit{apple}]\!]^{+}_{M} \cap [\![\textit{not apple}]\!]^{+}_{M}(x) \leq 1$
if $0.95 \leq \text{MIN}\{[\![\textit{apple}]\!]^{+}_{M}(x), [\![\textit{not apple}]\!]^{+}_{M}(x)\} \leq 1$
if $0.95 \leq c_{\text{apple}}(x) \leq 1$ and $0.95 \leq 1 - c_{\text{apple}}(x) \leq 1$
if $0.95 \leq c_{\text{apple}}(x) \leq 1$ and $0 \leq c_{\text{apple}}(x) \leq 0.05$
b. x $[\![\textit{apple that is not an apple}]\!]^{-}_{M}$
if $0 \leq c_{\text{apple}}(x) \leq 0.05$ and $0.95 \leq c_{\text{apple}}(x) \leq 1$

Notice that we have neither **c** $[\![$*apple that is not an apple*$]\!]^{+}_{M}$ nor **c** $[\![$*apple that is not an apple*$]\!]^{-}_{M}$ as both areundefined — we have $c_{apple \cap \neg apple}(\mathbf{c}) = 0.5$. But for the supervaluation we have that:

(47) $[\![$*c is an apple that is not an apple*$]\!]^{M^*} = 0$.

The reason is by now familiar: There is no completion M′ of M that satisfies (44) for which **c** is both in the positive part and in the negative part of c_{apple}.

Kamp & Partee suggest the notion of a **presupermodel** which determines a set of possible completions of a model. A presupermodel is a pair M, c , where M is a partial model and c determines the fuzzy characteristic functions for predicates, e.g. c_{apple} is the fuzzy characteristic function for apples. A presupermodel M, c determines a set of possible completions M* of M as the set of those models M′ that complete M in accordance with c, that is, for which a rule like (44) holds for all predicates:

(48) M′ is a possible completion of M, c
iff for all predicates α, x $[\![\alpha]\!]^{-}_{M'}$ and y $[\![\alpha]\!]^{+}_{M'}$, then $c_{\alpha}(x) < c_{\alpha}(y)$

3.3.4 Striped apples, tall eight-year olds and all basketball players

The puzzle of striped apples posed by Osherson & Smith cannot be solved within supervaluation theory as it stands. Kamp & Partee argue that the meaning of an adjective is often readjusted within the context of the noun it applies to, or within the general context of the sentence. This readjustment is well-known for adjectives:

(49) a. a tall eight-year old / a tall basketball player

b. The kids in the kindergarten / The undergraduates built a tall snowman.

We concentrate here on the sensitivity with respect to the head noun (a). Kamp & Partee propose that the meaning of *tall eight-year old* should not be simply derived as the meaning of the (vague) *tall* and the (precise) *eight-year old*. Rather, we redefine (or **calibrate**) the interpretation of *tall* with respect to the meaning of the head noun.

The principles that determine this calibration are the following:

(50) a. Parallel structure effect: In a conjoined structure, each conjunct is interpreted with respect to the same common context.

b. Head primacy principle: In a modifier-head structure, the head is interpreted relative to the context of the whole constituent, and the modifier is interpreted relative to the
local context created from the former context by the intepretation of the head.

c. Non-vacuity principle: In any given context, try to intepret any predicate so that both its
positive and negative extension are non-empty.

For us, the latter part of (b) is important: In *tall eight-year old*, the modifier *tall* is interpreted relative to the local context created by the interpretation of *eight-year old*. Principle (c) then says that *tall* should be interpreted in such a way that it has a positive and a negative extension, that is in particular, that there should be some elements in the extension of *eight-year old* that are considered tall.

Kamp & Partee implement the idea of recalibration of an adjective meaning *A* in an adjective-noun construction *AN* as follows. Let us write *A/N* for the adjective *A* in the context of the head noun *N*. Assume a presupermodel M, c . For simplicity, assume the noun *N* is sharp, that is, for all x, $c_N(a)$ {0, 1}. Now we treat the best cases of *A* within $[\![N]\!]$ as definitely within the positive extension of *A/N*, and the worst cases of *A* as definitely in the negative extension of *A/N*, and the intermediate cases are adjusted proportionally. This is expressed by the following formula for the fuzzy characteristic function of *A/N*:

(51)
$$c_{A/N}(x) = \frac{c_A(x) - c^{-}_{A/N}}{c^{+}_{A/N} - c^{-}_{A/N}}, \quad \text{where} \begin{cases} c^{+}_{A/N} = \max\{c_A(y) \mid y \quad [\![N]\!]_M\} \\ c^{-}_{A/N} = \min\{c_A(y) \mid y \quad [\![N]\!]_M\} \end{cases}$$

Here, $c^{+}_{A/N}$ stands for the highest value that an element of the extension of *N* achieves with respect to c, and $c^{-}_{A/N}$ for the lowest value. Assume that the highest value is 0.5 (that is, if *N* = *eight-year old* and *A* = *tall*, the tallest eight-year old is tall to degree 0.5), and the lowest value is 0 (that is, the least tall eight-year old is not tall at all). Then the formula gives us the following characteristic function for *tall/eight-year old* for the interpretation of *tall* in the context of *eight-year old*: Notice that the values of c_{tall} get multiplied by 2, guaranteeing that the tallest eight-year old is classified as a positive case of *tall*.

(52)
$$c_{tall/eight\text{-}year\ old}(x) = \frac{c_{tall}(x) - 0}{0.5 - 0} = 2\ c_{tall}(x)$$

Now consider the case of *tal basketball player*. Assume that the lowest value of c_{tall} for any baskedtball player is 0.8, and the highest value is 1. This gives us the following result, which guarantees that the smallest basketball player will get the value 0 (it's (0.8 — 0.8)/0.2, = 0), and the tallest basektball player will get the value 1 (it's (1—0.8)/0.2, = 1).

(53)
$$c_{tall/basketball\ player}(x) = \frac{c_{tall}(x) - c^{-}_{basketballplayer}}{c^{+}_{basketballplayer} - c^{-}_{basketballplayer}} = \frac{c_{tall}(x) -}{1 - 0.}$$

We can depict this recalibration graphically, as follows:

(54) Models ⟦*tall*⟧$^{+}$ ⟦*eight-year old*⟧, ordered after height

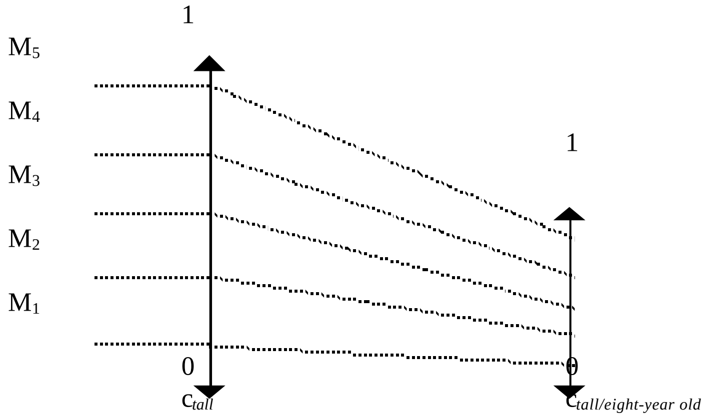

Notice that this is a solution to the *striped apple* puzzle: The problem was that the meaning rule presented in (40) is false. The characteristic function of *striped apple* is not just the minimum of stripedness and appleness; it must involve a recalibration of the adjective *striped* in the context created by *apple.* Now, striped apples are fairly unusual, but when we just consider the set of striped apples, we find that the better an apple classifies as having stripes, the better it is as an example of *striped apples*.

Notice that Kamp & Partee's rule does not take the possible vagueness of the head noun into account, which is perhaps the most severe criticism ofOsherson & Smith (1997). For example, of two striped apples, the one that is lessprototypical as an apple should also count as less prototoypical for a striped apple.

3.3.5 Stone Lions and Male Nurses

There are adjective/noun combinations that cannot be treated with the technique just outlined. Consider *stone lion*. Here we do not keep the meaning of *lion* constant and find out to which lions the adjective *stones* fits best. Rather, we have to revise the meaning of *lion* fist, to 'sculpture of a lion', and then apply the meaning of *stone*. After this shift, *stone lion* can be interpreted as regular predicate conjunction.

Kamp & Partee suggest that the intepretation of adjective/noun combinations discussed in the previous section is tried first. Only if there is no way to construe the modifier as applying to at least part of the extension of the head noun do we reinterpret the head noun.

Kamp & Partee are sceptical if it is possible to develop a theory of compositionality for prototypes. First, they argue (with Armstrong, Gleitman & Gleitman) that there is no systematic connection between membership in an extension and prototypicality. Furthermore, the factors that determine prototypes lack the degree of systematicity which a compositional account of prototype choice would presuppose.

They show this with the example of *male nurse*. The extension of this concept is clear (the constituents *male* and *nurse* are close-to-sharp concepts, and we have simple set intersection). However, the prototypes for *male nurse* are difficult, if downright impossible, to develop from the prototypes of *male* and *nurse*. They will depend for people on the examples of male nurses they have experienced (in real life, on TV, etc.).

3.4 Stereotypes

A discussion of prototypes should include mentioning the notion of **stereotypes**. These properties that speakers typically attribute to the members of a class and that can have a great impact for their construction of prototypes.

Stereotypes can be expressed by generic sentences. For example, having a mane is a stereotype for lions, we can express this by

(55) Lions have a mane.

As a matter of fact, only male lions have a mane, and so having a mane is considered as likely, for lions, as being male (actually, far less likely, as only adult male lions have a mane). However, *Lions are male* is definitely not considered to be a true generic statement.

Stereotypes can be factually false, e.g. the stereotype that snakes are slimy, and many stereotypes about particular social groups.

The concept of stereotype was introduced by Putnam (1975) into semantics. Putnam broke down the meaning of a lexical expression into several components, including its extension and some stereotypical properties. These properties are considered "core facts" (about the extension) which everyone speaking the language in question must know. Their truth-functional status is, however, indetermined; they might well be false for many or even all of the entities in question.

Armstrong, Sharon Lee, Lila R. Gleitman & Henry Gleitman. 1983. What some concepts might not be. *Cognition* 13, 263-308.
Battig, W.F. & W.E. Montague. 1969. Category norms for verbal items in 56 categories. *Journal of experimental psychology monograph 80 No. 3 Pt. 2*.
Fine, Kit. 1975. Vagueness, truth and logic. *Synthese* 30, 265-300.
Kamp, Hans. 1975. Two theories about adjectives. In Edward L. Keenan (ed.), *Formal semantics of natural languages* 123-155: Cambridge University Press.
Kamp, Hans & Barbara Partee. 1995. Prototype theory and compositionality. *Cognition* 57, 129-191.
Lakoff, George. 1970. A note on vagueness and ambiguity. *Linguistic Inquiry* 1, 357-359.
Lewis, David. 1970. General semantics. *Synthese* 22, 18-67.
Mehlberg, Henryk. 1958. *The reach of science*: University of Toronto Press.
Osherson, Daniel & Edward E. Smith. 1997. On typicality and vagueness. *Cognition* 64, 189-206.
Osherson, Daniel N. & Edward S. Smith. 1981. On the adequacy of prototype theory as a theory of concepts. *Cognition* 9, 35-58.
Pinkal, Manfred. 1985. *Logic and lexicon*. Dordrecht: Kluwer.
Putnam, Hilary. 1975. The Meaning of "Meaning". In K. Gunderson (ed.), *Language, Mind and Knowledge*. Minneapolis: University of Minnesota Press.
Quine, Willard V.O. 1960. *Word and object*. Cambridge, Mass.: MIT Press.
Rosch, Eleanor. 1973. On the internal structure of perceptual and semantic categories. In T. Moore (ed.), *Cognitive development and the acquisiton of language* 111-144. New York: Academic Press.
Rosch, Elaine. 1975. Cognitive respresentations of semantic categories. *Journal of Experimental Psychology* 104, 192-233.
Rosch, Eleanor & Carolyn S. Mervis. 1975. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology* 7, 573-605.
Smith, Edward E. & Douglas L. Medin. 1981. *Categories and concepts*. Cambridge, Mass.: Harvard University Press.
Tversky, Amos. 1977. Features of similarity. *Psychological Review* 84, 327-352.
van Fraassen, Bas. 1971. *Formal semantics and logic*. New York: Macmillan.
Williamson, Timothy. 1994. *Vagueness*. London: Routledge.
Zadeh, Lofti. 1965. Fuzzy sets. *Information and control* 8, 338-353.