The interdisciplinary colloquium of the linguistic department, Tel Aviv University

Nouns, adjectives and *more*, Galit W. Sassoon, 25.01.2007

Obviously, nouns and adjectives are very different types of animals. The question is precisely what this obvious difference is that forms the cause for the linguistic contrasts between them.

1 VAGUE PREDICATES

The standard intensional semantics is a model of total information. Such a model determines for each individual and each property in a situation, context or world, whether it has that property or not. There is no third possibility – no gap containing individuals that one does not know if they have that property or not. But, when we look at the language, we see that some predicates do have a gap:

1) **Vague predicates:** *Tall, bald, large, hot, cool*
Have a denotation gap, $[P]^?$: Some entities are neither in $[P]^+_c$ nor in $[P]^-_c$:

2) **Non-vague ('sharp'):** *Bird, apple*
No denotation gap, by and large everything is in $[P]^+_c$ or in $[P]^-_c$:

3) Contrast I: Adjectives tend to be gradable and nouns tend to be sharp. But:
Chair is a vague noun (Kamp & Partee 1995); *even (number)* is a non-vague adjective.

Thus, it has been argued that the semantic interpretation is relative to information states (or ***contexts***) c, in which ***predicate denotations are only partially known***. *Vagueness models* consist of many such partial contexts, and they represent information growth: The order in which entities are categorized under the predicates through contexts and their extensions. We start with a zero context, c_0, where denotations are empty, and from there on, each context is followed by contexts in which more entities are added to the denotations. In a total context t, every entity is either in the negative or in the positive denotation of each predicate (figure 1).

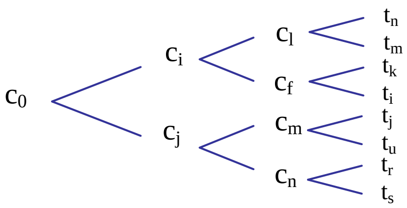

Figure 1: The contexts' structure in a standard vagueness model

For ex.: If in a context c the denotation of $[tall]^+$ consists of only very tall items, and the denotation of $[tall]^-$ consists of very short items, then c, we don't yet know if anything else, which is neither very tall nor very short, is tall or not.
Similarly, the denotation of *chair* in c may consist of only the prototypical type of chair, and the denotation of *non-chair* may consist only of things which are very clearly not chairs, say – the prototypical sofa. In such a context, we do not yet know if anything else (an armchair, a stool, a chair without a back, which is not used as a seat, not of the normal size, etc.), is a *chair* or not:
If $[chair]^+_c$ = {  } and $[chair]^-_c$ = {  }, then what is  ?

But each c is extended by a set of **total contexts (supervaluations)**, t – all the possibilities seen in c to specify the complete sets of *tall* and *non-tall* things, *chairs* and *non chairs* etc. A context c_1 is ***extended*** by another context c_2, $c_1 \leq c_2$, iff the positive and negative denotations of each predicate in c_1 are subsets of (or are extended by) the positive and negative denotations of that predicate in c_2 ($\forall P \in PRED$: $[P]^+_{c1} \subseteq [P]^+_{c2}$ and $[P]^-_{c1} \subseteq [P]^-_{c2}$).

2 GRADABLE PREDICATES

We can also distinguish gradable from non-gradable predicates:

4) **Gradable predicates** (*tall, bald, large, hot, cool*) can combine with comparatives *(more P; less P*) equatives (*equally P*), and superlatives (*the most P*)

5) **Non-gradable predicates** (*Bird, apple, chair, extinct, even (number)*) cannot occur (bare) in these structures (**more P,* less P, *as P as, *the most P*).

6) Contrast II: By and large, vague predicates (adjectives) are gradable. Nouns are not vague and hence not gradable.

Given the generalization in (6), gradability is often analyzed as vagueness dependent.
The total contexts are thought to represent different *standards* of precision (Lewis 1979). In some of them only very tall things are regarded as tall enough to be considered *tall*, in others more things are considered *tall*, etc.

It is generally, assumed (in for instance, Kamp 1975 and Kamp and Partee 1995) that we can make do with simplified vagueness models which contain but one partial context c (the ground model) and a set T_c of the total contexts t extending c. The intermediate steps between c and each t are thought to be unimportant (Figure 2).

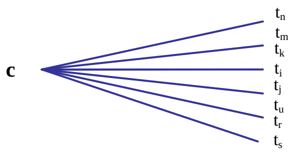

Standard height in t_n:	2 meter
Standard height in t_m:	1.95
Standard height in t_i:	1.9
Standard height in t_u:	1.85
…	…

Figure 2: The context structure in a simplified vagueness model Mc

Gradable predicate are standardly analyzed as mapping individuals to degrees (. For ex. *tall* maps its argument into a degree, $\deg_{(tall,x)}$, on the dimension *height*. Thus *tall* is associated with several elements (Kennedy 1999; Rotstein and Winter 2004):

7)

a. A set of *degrees*, S_{tall} (say – a set of number)

b. An *ordering* on this set $\leq_{tall}$, which states for each two degrees which one represents the larger degree under *tall*

c. A *unit of measurement* (say – centimeters)

d. A *dimension* which these degrees measure (*height*).

Sam is tall is considered true in a context c iff Sam's height, $\deg_{(tall,Sam)}$, reaches the standard for tallness in c.

8)

a. $[\ tall\ ]^{+}_{c} = \quad \{d \in D:\ \deg_{(tall,d,c)} \geq_{tall,c} \text{Standard}_{tall,c}\}$

b. $[Sam\ is\ tall\ ]_c = 1$ iff $\quad \deg_{(tall,[Sam]c)} \geq_{tall} \text{Standard}_{tall,c}$

If we do not know what the standard is, we can only consider as *tall* those entities which are tall in every total context above c. van Fraassen 1969 has defined ***super-truth*** as truth in every total context. Accordingly, *Sam is tall* is considered true (or super-true) in a partial context c iff Sam's degree of height, $deg_{(tall,Sam)}$, reaches the standard in every total context above c (9), that is, Sam's degree ought to exceed any degree which might yet be the standard of *tallness*. Thus, we can distinguish between the denotation $[tall]^{+}_{c}$ which consists of the things which are directly known to be *tall* in c, and the ***superdenotation***, $[tall]_c$ (in (10a)), which consists also of all the things which end up being tall in all the total possibilities – all the things which, given our knowledge, must be tall.

9) Supertruth: $[\textit{Sam is tall}\,]_c = 1$ iff $\forall t \in T, t \geq c$: $deg_{(tall,[Sam]t,t)} \geq_{tall,t} \delta_{tall,t}$

10) Superdenotations:

a. $[tall\,]_c = \cap\{\, [tall]^{+}_{t} \mid t \in T, t \geq c\}$

b. $[\neg tall\,]_c = \cap\{\, [tall]^{-}_{t} \mid t \in T,, t \geq c\}$

Thus, we can say that a comparative statement like *Dan is taller than Sam* is true in c, that is, the degree of *Dan* exceeds that *Sam*, iff *Dan* is tall relative to more standards, that is: *Dan is tall* is true in more total contexts above c, compared to Sam (Kamp 1975; Fine 1975). If Sam reaches a certain standard of tallness, Dan certainly reaches this standard, but not vice versa. Dan's height reaches certain standards which Sam's height does not reach.

11)

a. $[\textit{Dan is taller than Sam}]_c = 1$ iff $deg_{(tall,Dan,c)} >_{tall,c} deg_{(tall,Sam,c)}$ iff:
$\{t \in T \mid [\textit{Sam is Tall}]_t = 1\} \subseteq \{t \in T \mid [\textit{Dan is tall}]_t = 1\}$

b. $[\textit{a is more P than b}]^{+}_{c} = 1$ iff $\{t \in T \mid [P_{(a)}]_t = 1\} \subseteq \{t \in T \mid [P_{(b)}]_t = 1\}$

Problem 1:

The analysis of comparatives in (11) applies to gap members only. All the entities which are already known to be *tall* in c, are *tall* in all the total contexts extending c. Hence they are wrongly predicted to be all *equally tall*. But intuitively two tall individuals can stand in the relation *taller than* to each other.

12) Wrong Prediction of (11): $\forall d_1,d_2 \in [tall]_c$: d_1 and d_2 are *equally tall*
$\forall d_1,d_2 \in [\neg tall]_c$: d_1 and d_2 are *equally tall /short*

I show in my MA thesis, directly following Landman's 1991 analysis, that this problem is solved if we add back to the model ***the intermediate contexts*** (we go back to the standard vagueness model; figure 1). In such a model the gradable structure of *tall* reflects the order in which entities are added to the super-denotation (the order in which entities are learnt directly or by inference to be tall). Thus, two denotation members can stand in the relation *taller*. They do so iff one of them was added to the denotation in an earlier context.
So we replace (11) by (13) (Constrain L): d_1 *is more P* than d_2 in a context t iff:

Either the P-hood of d_1 is established before the P-hood of d_2 (i.e., in a context that precedes the context in which d_2 is added to the positive denotation),
Or the non-P-hood of d_2 is established before the non-P-hood of d_1 (i.e., in a context that precedes the context in which d_1 is added to the negative denotation).

13) Constraint L

$\forall t \in T$: $(\langle d_2,d_1\rangle \in [\leq P]^{+}_{t})$ In any total t, d_1 is *equally or more P* than d_2 iff:

$\forall c \leq t$: $(d_2 \in [P]_c \rightarrow d_1 \in [P]_c)$ & $(d_1 \in [\neg P]_c \rightarrow d_2 \in [\neg P]_c)$.

In any context c under t, if d_2 is P, d_1 is P, and if d_1 is $\neg$P, d_2 is $\neg$P.

For example, (14), *Dan is taller than Sam*, is true in t iff in every c leading to t:

If Sam is already considered *tall* (that is, Sam reaches the standard, be it what it may), then we can <u>infer</u> that Dan is *tall* (Dan definitely reaches the standard, given Dan's larger height). And if Dan is considered *not-tall* in c (that is, Dan does not reach the standard), then Sam is definitely *not tall* (Sam definitely does not reach the standard, given its lower height).

14) $[\textit{Dan is equally tall or taller than Sam}]_t = 1$ iff

$\forall c \leq t$: $([\textit{Sam}]_c \in [\textit{tall}]_c \rightarrow [\textit{Dan}]_c \in [\textit{tall}]_c)$ &

$([\textit{Dan}]_c \in [\neg\, \textit{tall}]_c \rightarrow [\textit{Sam}]_c \in [\neg \textit{tall}]_c)$.

15) $[\textit{Dan is equally tall or taller than Sam}]_c = 1$iff

$\forall t \geq c$: $[\textit{Dan is equally tall or taller than Sam}]_t = 1$ iff

$\forall c' \in C$, s.t. $c \leq c'$ or $c' \leq c$: $([\textit{Sam}]_c \in [\textit{tall}]_c \rightarrow [\textit{Dan}]_c \in [\textit{tall}]_c)$ &

$([\textit{Dan}]_c \in [\neg\, \textit{tall}]_c \rightarrow [\textit{Sam}]_c \in [\neg \textit{tall}]_c)$.

In every context c' under or above c, if d_2 is P, d_1 is P, and if d_1 is $\neg$P, d_2 is $\neg$P.

Thus, P's ordering in t is the order in which entities are learnt to be P or $\neg$P (whether directly or by inference) in the contexts under t.

<u>Problem 2:</u>

The noun *chair* is vague (things like stools may be regarded as neither chairs nor non-chairs) but it is not gradable (**more chair*). Vagueness turns adjectives, but not nouns, gradable. The nouns seem to be inherently non gradable. Hence, the standard linguistic theory dos not associate nouns with a gradable structure (a set of degrees, an ordering dimension etc.)

<u>Problem 2.1: Psychological adequacy</u>

Entity orderings in nouns:
The problem is that the last forty years of research in cognitive psychology have established beyond doubt that speakers consider certain entities as better examples of nouns than others (for instance, *robins* are often considered more typical *birds* than *ostriches*).
These ordering judgments are reflected in *online categorization time*: Most importantly, verification time for sentences like *a robin is a bird* is faster than for sentences like *an ostrich is a bird* (Rosch 1973; Armstrong et al 1983).

Ordering dimensions in nouns:
In addition, speakers associate nouns with ordering dimensions (features like *feathers, flying, nesting, singing* etc.)
The classical view considered these features definitional: Necessary and sufficient conditions for membership in the denotation. But Wittgenstein 1968 (1953) and Fodor et al 1980 have shown that this idea is rarely if ever met.
For example, it is already well known that counterexamples can be found to any definitional feature you would propose for natural categories like *games* or *bachelors*.
In addition, often, speakers are uncertain about the membership status of some entities and they vary their judgments in different times or contexts, refuting the assumption that there are clear-cut criteria. For example, tomatoes fall in between the categories *fruit* and *vegetables*.
While speakers rarely (only 3% of the time in average) change their minds about the category membership of clear instances, they do so much more often (above 20% of the times in average) with regard to the membership of borderline cases, like curtains for *furniture* or avocado for *vegetables* (Murphy 2002: 20).

Crucially, the features which people link with a category like *bird* are raising the typicality of entities in the category, that is, they *are* indeed ordering dimensions, which together help to measure the typicality (and membership likelihood) of entities in the category. Thus, the standard theory in cognitive psychology associates a concept like *bird* (or the word that denotes it) with a prototype, namely:

16) The basic prototype theory:

a. A set of ***dimensions***. The feature set of *bird*:
$\mathbf{F_{(bird,c)}}$ = {*feathers, flying, nesting, singing, small size ...*} $\subseteq$ PRED

b. Each dimension F has a weight W_F. For example, W_{flying} tells us how distinctive *flying* is of birds: How important *flying* is in discriminating birds from non-birds.

c. *The weighted mean hypothesis:*
For each entity d, its degree of typicality (or similarity to the prototype of P), $deg_{(d,P)}$, equals the weighted mean of d's degrees in the category features:
$\mathbf{deg_{(d,P)} = \Sigma(\{ W_F deg_{(d,F)} \mid F \in F^+_{(P)}\})}$

For instance, the typicality degree of a robin in *bird*, $deg_{(robin,\, bird)}$, is indicated by *the weighted-mean of its degrees in all the bird features:* How well it scores in *flies, sings, small* etc.
That is, the typicality of an entity in a category (e.g. *bird*) represents the extent to which it possesses the features that are distinctive of the category. The typical instances (*robins*) are more similar to the prototype, that is, they have more properties or they average better in the features, compared to the atypical instances (*ostriches*).

There are tight relations between the entity ordering and the denotation:
The basic cognitive theory argues that categorization is based on typicality (or similarity to the prototype). A certain degree of typicality functions as a standard, such that:

17) *The categorization criterion*: $[P] = \{d \in D \mid deg_{(d,P)} \geq standard_P\}$
An entity is classified in a category iff its typicality degree reaches the standard.

Indeed, there is abundant evidence showing that entities are positively classified iff their *average in the features* reaches criterion. Hampton 1998 analyzed the data about 500 items in 18 categories (McCloskey and Glucksberg 1978). He found a very strong coupling between their mean typicality ratings and the probability that they were categorized positively. There are also deviations, but they are highly systematic.
The deviations were shown to occur due to (i) *shift of weight*, in typicality judgments compared to membership judgments, towards non-definitional perceptual criteria (which increase the typicality of non-members) (ii) *unfamiliarity* (lack of knowledge about the features of members reduces their typicality), and (iii) the existence of *competing categories* (such as *kitchen utensil* and *furniture*, which reduces the likelihood of classification, but not the typicality of, say, a *refrigerator* in *furniture*).

It follows that the typicality ordering is determined by the order in which entities are added to the denotation (whether directly or by inference):
Dan is *more typical of a bird* than Sam iff Dan's bird-degree exceeds Sam's degree. That is iff: If Sam is already considered a *bird* in c (that is, Sam reaches the standard, be it what it may) then we can infer that *Dan* certainly does so, and hence it certainly counts as *a bird*. And if

Dan is known *not to be a bird* (not to reach the standard), we can infer that Sam, due to its lower typicality degree, does not reach the standard and is *not a bird*.
Thus, the gradable typicality structure of *bird* reflects the order in which entities are added to the super-denotation (the order in which entities are learnt directly or by inference, to be *birds*), as expected by constraint L.

Some very robust findings, the *order of learning effects*, form evidence for this generalization. Most importantly, typical instances are acquired earlier than atypical ones, by children and adults (Mervis and Rosch 1981; Rosch 1973; Anglin 1977; Murphy and Smith 1982). For example, birdhood is normally determined first for *robins* and *pigeons*, later on for *chickens* and *geese*, and last for *ostriches* and *penguins*. Similarly, non-birdhood is determined earlier for *cows* than for *bats* or *butterflies*:

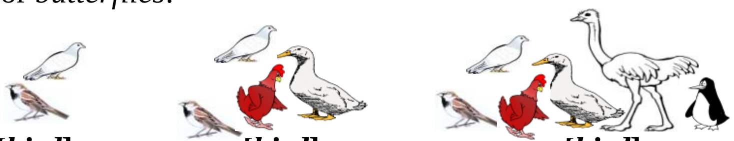

[*bird*]$\mathbf{c_0}$ **…** **[*bird*]**$\mathbf{c_i}$ **…** **[*bird*]**$\mathbf{c_n}$ **…** **[*bird*]**$\mathbf{t_s}$

Figure 3: A normal acquisition order for the category *bird* is highly indicative of the typicality structure

Second, language learners, learn faster if initial exposure is to typical category members (the crucial factor is not the amount of examples but their typicality; Mervis & Pani 1980).

Third, competent speakers produce, or recall, typical category members before atypical ones (Rosch 1973; Batting & Montague 1969), they remember best the typical (or early acquired) instances (Heit 1997) and their features affect future learning, or remembrance, of new entities (Rips 1975; Osherson et al 1990). For example, when speakers are initially exposed to *joggers* that *wear expensive running shoes*, they frequently falsely recollect joggers that do not wear expensive shoes as *non-joggers* or as *joggers that do wear expensive shoes*. New facts are corrected so as to match earlier ones.

In sum, nouns behave very much like our semantics for gradable predicates expects.

The typicality (or graded concept structure) effects in nouns, are robust and pervasive. A prominent psychologist (Murphy 2002) has written that it would be very surprising to find a cognitive task that typicality does not affect. Thus, by assuming that nouns are non gradable, in order to account for their infelicity in the comparative, linguists pay a heavy price in terms of the dissociation between the basic semantics they assume for nouns and many other things that we know about them.

<u>Problem 2.2: Linguistic adequacy</u>

Some phenomena which remain unexplained if nouns are assumed to be non-gradable are completely linguistic (having to do with Gradability, Compositionality, Genericity etc.)

First, for example, we see in (18a) that it is sufficient to add the particle *of* to the comparative morpheme (as in *more of a bird*) and the interpretation of the noun *bird* turns gradable.
Second, it is very easy to turn a noun gradable by turning it into an adjective, either by modifying it with *typical* (18a) or simply by the adjectival morpheme 'y' (as in *birdy* or 18b).
This facts are hard to explain if nouns are merely (linguistically) non gradable.

18)

a. *A robin is more (typical) of* ***a bird*** *than an ostrich*
b. *The noun 'activity' is "nounier" / less "nouny" than the noun 'bird'*

Third, the typicality effects characterize very complex predicates. For example in (19) we see that a gradable structure pops up in a very complex noun phrase:

19) *…pretty much typical of a non-fan, non-entertainment, smart up-market British paper*

So typicality is highly productive. Within a context, we can produce typicality orderings for novel complex-concepts on the fly. It seems that *some* generative system plays a role here.
Forth, some adjectives are *multidimensional* (*healthy*, *intelligent*, *talented*, *good* etc.) For instance, the adjective *healthy* can be measured by dimensions such as *blood pressure*, *pulse* and *fever*. Now, the features of multidimensional (but not of one-dimensional) adjectives can be quantified over. This is demonstrated by the contrast between (20a) and (20b) (Bartsch 1986; Landman 1989). (20c) shows that quantification over dimensions is impossible in nouns.

20)
- a. *Maria is healthy in every respect / generally healthy / healthy wrt blood pressure*
- b. *? Maria is tall in every respect / ? Maria is generally tall / ? tall wrt height*
- c. *# Tweety is a bird in every respect /# generally a bird/ # a bird wrt flying*

But like the felicity of nouns in the comparative, also quantification over the dimensions becomes possible, if the noun is slightly modified (21). Again, this fact is hard to explain if nouns are not associated with ordering dimensions.

21) *Tweety is a typical bird in every respect / generally typical of a…/ typical of…wrt flying*

3 CAN WE CAPTURE THESE FACTS WHILE MAINTAINING THE ASSUMPTION THAT NOUNS ARE NEITHER VAGUE NOR GRADABLE (IN THE USUAL SENSE)?

This is precisely what Kamp and Partee's influential 1995 *supermodel* theory has attempted to do. Given its central status, I will now show that it has failed. Namely, if we are up to a correct analysis, we should give up the assumptions that nouns are non-vague and non-gradable.

22) In a supermodel (Kamp & Partee 1995) each predicate P and entity d is associated with:
- a. *A membership-degree*, $\mathbf{deg_{(d,P)}}$: The proportional size or measure m of the set of total contexts in which d is P: $deg_{(d,P)} = m(\{t\in T: d\in[P]^{+}_{t}\})$

  (m is a measure function from sets of total models to real numbers between 0 and 1:

  $m(T) = 1$; $m(\{\}) = 0$; $\forall T1,T2$, s.t. $T1\subset T2$: $m(T1)< m(T2)$ etc. [p. 153])

  For example, given the model in figure 2 (which is repeated below in figure 4), the membership degree of the prototypical chair in *chair* is 1 (because it is a chair in all total contexts); the degree of the blue sofa is 0 (because it is a chair in no total context); the degree of the strange armchair is 1/3 (because it is a chair in a third of the total contexts) etc.
- b. P has a prototype **p** – the best possible instance (examples are given in figure 5)
- c. $\mathbf{t\text{-}deg_{(d,P)}}$ is d's *typicality-degree* in P: d's distance from P's prototype

Figure 4: A Supermodel:

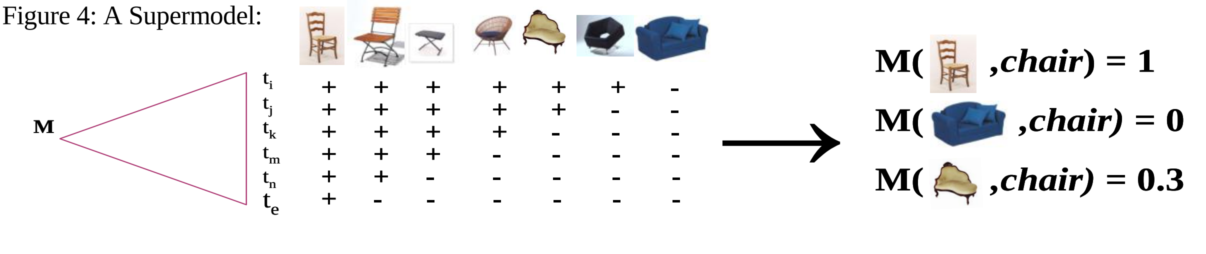

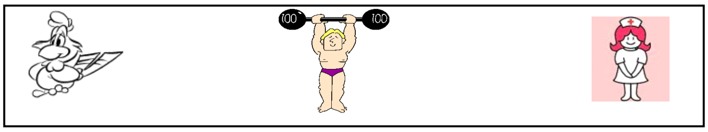

Figure 5: Examples of possible entity prototypes for *bird*, *male* and *nurse*

Kamp and Partee distinguish between different predicate types in the following ways:

23) Predicate types (Kamp and Partee 1995):

a. **+/– Vague**:
The denotations of non-vague predicates like *bird*, unlike those of vague predicates like *chair*, are total already in M. That is, everything is either a *bird* or a *non-bird*. There is no gap: $[bird]^{+}_{M} \cup [bird]^{-}_{M} = D$.

b. **+/– Prototype**:
Predicates like *tall* or *odd number*, unlike *bird*, *grandmother*, *red*, etc., have no prototype (because there is no maximal *tallness* or *oddness*).

c. **+/– Typicality-is-coupled-with-membership, t-deg ≅ deg**:
In predicates like *chair* the values of the typicality degree function, $\text{t-deg}_{(d,P)}$, are given by the values of the membership function: t-deg ≅ deg. Thus, the more typical entities fall under $[chair]^{+}$ in more of the total contexts in T. But in sharp nouns, like *bird*, typicality and membership are assumed to be dissociated: **t-deg ≠ deg**. (The original term: +/–*the-prototype-affects-the-denotation*).

	–Prototype	**+Prototype**	
		(deg ≠ **t-deg)**	(deg = **t-deg)**
+Vague	*tall, wide, heavy, not red*	*adolescent, tall tree*	*red, chair, shy*
–Vague	*Even, odd, inanimate, non bird*	*bird, grandmother*	∅

Table 1: Predicate types in Kamp and Partee's 1995 analysis

As for complex predicates, experimental findings reveal what is usually called the *conjunction effect* and *fallacy*. These were demonstrated by intuitive judgments such as (24a-b).

24)

a. The *conjunction effect*: Brown apples are *more typical* in *brown apple* than in *apple*
$\text{deg}_{(d,\, brown\ apple)} > \text{t-deg}_{(d,\, apple)}$

b. The *conjunction fallacy*: Brown apples are *more likely brown apples* than *apples*
$likely_{(d,\, brown\ apple)} > likely_{(d,\, apple)}$

Thus, Kamp and Partee replace deg in modified nouns like *brown apple* by *the modified membership function*, $\mathbf{deg}_{(d, brown\,/apple)}$ in (25). This function is basically given by d's degree in *brown*. The set of brown degrees which are assigned to apples, are linearly transformed (stretched), so as to range from 0 to 1. The result stands for the apple degrees in *brown apple*:

25) *The modified membership function for modified nouns:*
deg(d,*brown /apple*) = (deg(d,*brown*) – a) / (b – a)
(where a and b be the minimal and maximal *brown* degrees in $[\text{apple}]^{+}_{M}$).

For example, a brown apple ba may have degree 0.9 in *brown*; the minimal *brown* degree existing among the apples may be 0, because some apples are not *brown* at all; the maximal *brown* degree existing among the apples may be 0.95, assuming that no apple is maximally brown. If so $\text{deg}_{(ba,brown\,/apple)} = (0.9 - 0) / (0.95 - 0) = 0.974$. It exceeds $\text{deg}_{(ba,brown)}$, as desired.

Main Problems**:**

1. The membership degree functions fail to produce intermediate typicality degrees for denotation members. Positive denotation members always receive degree 1, and negative denotation members always receive degree 0. Now, in sharp nouns like *bird* the denotations are assumed to be completely specified in M. This is the standard way to distinguish them from vague predicates. So their membership degrees cannot indicate typicality. How are the typicality degrees indicated? Kamp and Partee do not explicate. This is a significant disadvantage given that sharp predicates form the central examples of the prototype theory. We do not account for the typicality effect unless we give up on the assumption in (23a).

2. Wrong predictions about the typicality judgments
Even if we give up the assumption in (23a), the degree functions would fail to predict the typicality judgments (*the sub-type effect*). For example, intuitively, ostriches are more typical in *ostrich* than in *bird*. But, the membership degree of entities in *ostrich* cannot exceed their degree in *bird*, because in any total context if an item is an *ostrich*, it is also a *bird*.
Also typicality in modified nouns (*the conjunction fallacies* and *effects*) is not dealt with correctly. For example, take *brown apples*. Kamp and Partee use a special function to order entities in *brown apple*. Essentially, it orders them only by their degree in *brown*. But this is not intuitive: Speakers consider bad examples of apples as bad examples of *brown apples* (Smith et al 1988).

3. Vagueness and context dependency in the typicality ordering
The degree functions are total, though knowledge about typicality is partial. If one bird sings and one flies, which one is more typical? We cannot tell: Both possibilities are imaginable. But Kamp and Partee's measure functions cannot represent this type of vagueness, or context dependency. They are fixed per supermodel (once and for all), not per total context.

4. It is hard to predict the prototypes of complex concepts from the prototypes of their constituents (Kamp and Partee 1995; Hampton 1997). Consider negations: What would the prototype of *non-bird* be: a dog, a day, a number? Similarly for conjunctions: What would the *male-nurse* prototype be, given that a typical *male-nurse* may be both an *atypical male* and an *atypical nurse* (ibid).

Figure 6: No prototypes for *non-bird*, *male nurse* and *tall*

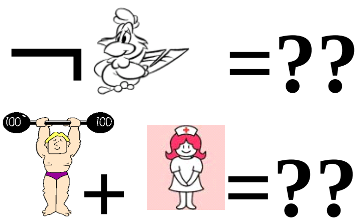

5. Obligatory modification by *typical of* in the comparative cannot be explained by the prototypes, because certain +Prototype predicates cannot occur in the comparative without modification by *typical of* or at least by *of* (e.g., *chair* and *bird*), while others can (e.g., *red*), and certain –Prototype predicates cannot do so (e.g., *non-bird*), while others can (e.g., *tall*).

In sum, the typicality effects are not accounted for by this theory.
I propose that nouns are gradable and multi-dimensional. Crucially, even non-vague denotations (like [*bird*]) are learnt gradually. This gradation forms an ordering, as constraint L predicts. The infelicity of nouns in the comparative is not due to lack of gradable meaning.

4 NOUNS VERSUS ADJECTIVES: PROPOSAL

In cognitive psychology today, one-dimensional adjectives like *tall* are treated as rule based, because categorization under them does not involves averaging over dimensions. Conversely, in nouns like *bird* or *house*, categorization *is* based on averaging ('similarity to the prototype'). This distinction is important because there is evidence that rule versus similarity based categorization tasks recruit different brain systems (Ashby and Spiering 2004; Photos 2004) and their acquisition course seem to be different (perhaps due to late maturation of the rule based brain system; Keil 1979; Zelazo et al 1996; Thomason 1994).
So the distinction between one dimensional and multi-dimensional categories seems to be a cognitively real distinction. Maybe it is this distinction that has been grammaticized into the two categories – nouns and adjectives:

26) Constraint F

a. All predicates P are associated with a feature set $F_{(P,c)} \subseteq PRED$

b. In ***one-dimensional adjectives*** (like *tall*) the feature set consists of one feature. The degree of each entity equals its degree in this feature. For example:
[*Dan is tall*]$_c$ = 1 iff $deg_{(Dan,height)} \geq standard_{tall}$.
$\forall d \in D$: $deg_{(d,P)} = deg_{(d,\sigma(F(P,c)))}$ & ($d \in [P]_c$ iff $deg_{(d,P)} \geq standard_{\sigma(F(P,c))}$)

c. In ***multidimensional adjectives*** (like *healthy*), in each context of use a *with respect to argument* (say – *wrt blood pressure*) selects one dimension, $wrt_{(P,c)}$, from $F_{(P,c)}$. The degree of each entity equals its degree in this dimension, and this dimension determines P's standard in that context of use. For example:
[*Dan is healthy wrt to blood pressure*]$_c$= 1 iff $deg_{(Dan,blood\ pressure)} \geq standard_{blood\text{-}pressure}$.
$\forall d \in D$: $deg_{(d,P)} = deg_{(d,wrt(P,c))}$ & ($d \in [P]_c$ iff $deg_{(d,P)} \geq standard_{\sigma(wrt(P,c))}$)

d. In ***nouns*** the degree of an entity equals the weighted mean of its degrees in the features. The feature weights w_f and averaging method Σ may vary between uses (Smith and Minda 1998). For example:
[*Tweety is a bird*]$_c$ = 1 iff $\Sigma(\{w_f deg_{(Tweety,F)} \mid F \in F_{(bird,c)}\}) \geq standard_{bird}$.
$\forall d \in D$: $deg_{(d,P)} = \Sigma(\{w_f deg_{(d,F)} \mid F \in F_{(P,c)}\})$ & ($d \in [P]_c$ iff $deg_{(d,P)} \geq standard_P$)

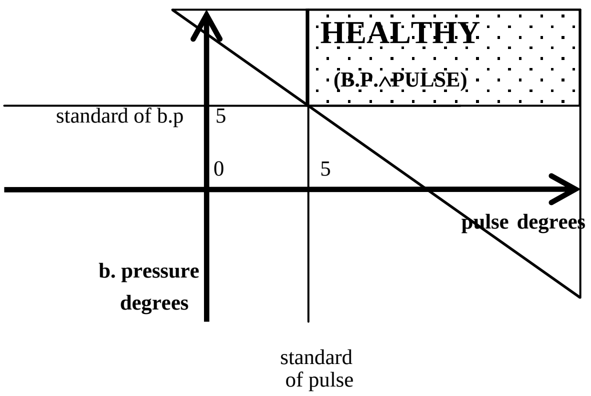

Figure 8: Multi-dimensional adjectives

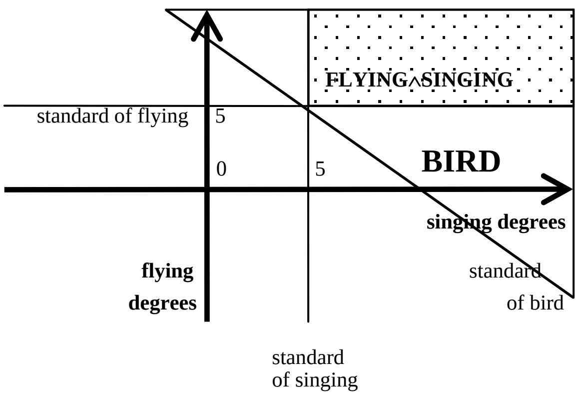

Figure 7: Nouns

1. Nouns:
The postulate in (26d) accounts for the fact that none of the features of nouns necessarily adds a categorization criterion. Even when typicality in *bird* is measured by typicality in *flying*, *singing*, and *nesting*, these features are not regarded as necessary for bird-hood (Wittgenstein 1953). Rather, entities are regarded as birds iff their *average in the features* reaches criterion (Hampton 1979).

Figure 7 demonstrates a two dimensional case, but the argument applies to n dimensional cases and all types of averaging functions. The axes stand for the set of degrees in *flying* and in *singing*. The set of birds is not given by feature intersection (that would wrongly give us the square to the right), but by averaging. An averaging formula like the one in (27) gives us the points above the straight line which represents the standard of *bird* (the triangle).

27) $[bird] = \{d \in D \mid w_{flying} deg_{(d,flying.)} + w_{singing} deg_{(d,singing)} + \ldots > standard_{bird}\}$

2. <u>Multidimensional adjectives:</u>
The postulate in (26c) accounts for the fact that each feature of a multidimensional adjective does add a categorization criterion. For instance, if health is measured by levels of *blood pressure*, *pulse* and *fever*, then one is *healthy* iff one is within the norm in all these respects. So the bare predicate *healthy* is interpreted as a conjunction of three one-dimensional adjectives (28a), that is, by feature intersection (28b). The set of healthy entities is the square in figure 7.

28)

a. $[\textit{Dan is healthy}]_c = 1$ iff $\forall F \in F_{(healthy,c)}$: $[\textit{Dan is healthy wrt F}]_c = 1$

b. $[\textit{healthy}]_c = [\wedge F_{(healthy,c)}] = \{d \in D \mid deg_{(d,b.p.)} > standard_{b.p}\ \&\ deg_{(d,pulse)} > standard_{pulse} \ldots\}$

Thus, the denotation of multidimensional adjectives like *healthy* or conjunctions like *blood pressure and pulse* cannot be indicated by averaging on the features (or conjuncts): A formula like $\mathbf{w_{b.p} deg_{(d,b.p.)} + w_{pulse} deg_{(d,pulse)} > standard}$ wrongly selects the triangle in figure 8.

3. A <u>wrt argument</u> and hence <u>quantification over features</u> makes sense only when a predicate has several features which add categorization standards (only when, indeed, an entity can fall under it in one respect and not in another respect), namely, in multi-dimensional adjectives.
A determiner which quantifies over respects states how many of the features form categorization criteria in the context of use. So we can relax the meaning of *healthy* by accommodating an existential or a vague universal quantifier over the respects:

29)

a. $[\textit{Dan is healthy in every respect}]_c = 1$ iff $F_{(healthy,c)} \subseteq [\lambda F.\textit{Dan is healthy wrt F}]_c$
(iff *Dan is healthy* wrt every *health* feature).

b. $[\textit{Dan is healthy in some respect}]_c = 1$ iff $F_{(healthy,c)} \cap [\lambda F.\textit{Dan is healthy wrt F}]_c \neq \emptyset$
(iff *Dan is healthy* wrt some *health* feature).

c. $[\textit{Dan is generally healthy}]_c = 1$ iff, roughly, *Dan is healthy* wrt most of $F_{(healthy,c)}$.

But neither the one-dimensional adjective *tall*, nor the noun *bird*, have two different necessary criteria, so *wrt* modification and quantification over the features is infelicitous. Philosophers have argued that nouns do not have semantic necessary conditions for membership at all (Wittgenstein 1968 [1953]; Fodor et al 1980; Fodor 1998). Ordinary speakers, normally count an individual as a *bird* iff it *has bird genes*, or iff it *has a bird essence* (if they are a bit less educated). But crucially, speakers do not consider any other of the *bird* features to be a separate semantic necessary condition for birdhood.
Features which merely function as domain restrictors, but do not help to distinguish *birds* from *non birds*, are irrelevant. For instance, in most contexts *animal* restricts the set of birds but also the set of non-birds: We hardly ever call prime numbers *non-birds*.
But an expert on birds might in principle characterize birds by the possession of, say, 100 separate genes. Our proposal correctly predicts that in such a context, the expert might indeed describe new species which possess only 50% or 100% of the bird genes, as *birds in this respect but not in that respect*.

30) **+/–Q(uantification over features) & wrt-argument**:

a. **+Q/wrt** predicates (e.g., *healthy*, *talented*) have <u>several</u> necessary condition.

b. **–Q/wrt** predicates (e.g., *bird* or *tall*) have at most one necessary condition.

4. <u>The adjective *typical*</u> turns its argument gradable solely by adding a wrt role:

31)

a. $[$*Tweety is typical of a bird wrt flying*$]_c = 1$ iff $deg_{(Tweety,flying)} \geq standard_{flying}$.

b. $[$*Tweety is typical in every respect*$]_c = 1$ iff $F_{(bird,c)} \subseteq [\lambda F.$ *Tweety is typical wrt to F*$]_c$ (iff *Tweety is typical* wrt every bird feature).

This proposal predicts the intuition that *typical of P* has more categorization criteria than *P*, although it is hard to put a finger on the exact features which add criteria, if $wrt_{(P,c)}$ isn't specified and a determiner like *generally* is accommodated.

The same analysis for nouns with a morpheme like '*y*' ('*nouny*', '*primy*') predicts the ease with which a noun can turn adjectival.

5. <u>The comparative morpheme</u> denotes an operation on the dimension of its predicative arguments. It is undefined when there is no unique dimension.

Hence, it is clearly defined in one-dimensional adjectives. It is also defined for bare multi-dimensional adjectives, because they can be interpreted wrt the conjunction of features (which then forms a unique dimension).

But it is undefined in nouns, where the conjunction of features cannot form a dimension. The ordering given by it is different from the noun ordering. Birds which fall under all the features but have a relatively low average are less good in *bird*, but better in the conjunction, compared to birds which violate one feature but have a higher overall average.

Moreover, conjunctions of gradable predicates have at least 2 features. Hence, we predict that they would be non-gradable. And this prediction seems to be supported by the facts. For instance, I found (in 35 subjects) that, the preferred interpretation for *more bald and tall* is *balder and taller*. That is, *and* is not within the scope of *more*. So '*more*' does not take 2 features simultaneously.

32) <u>A possible semantic interpretation for *more* that would do the job:</u>

a. $[$*more P*$]_t = \{<d_1,d_2> \mid d_1$ and d_2 satisfy constraint b or c$\}$

b. $d_1 \in [\wedge F_{(P,t)}]_t$ but $d_2 \notin [\wedge F_{(P,t)}]_t$ (d_1 falls under all P's features but d_2 does not)

c. $\sum(\{w_F deg_{(d1,F)} \mid F \in F_{(P,t)}\} > \sum(\{w_F deg_{(d2,F)} \mid F \in F_{(P,t)}\}$ (if both fall under all P's features or both violate some feature, then d_1's average in P's features must exceed d_2's average).

In adjectives, the denotation is given by the intersection of the features ($[P]_t = [\wedge F_{(P,t)}]_t$). Hence, in adjectives (32b) gives the right result: Each member ends up more P than each non member (d_1 is more P than d_2 iff the degree of d_1 exceeds the degree of d_2 in P: $deg_{(d1,P)} > deg_{(d2,P)}$).

But in nouns, *more* fails to correctly represent their ordering, so it is incompatible with them. For example, by (32b), birds which fall under all the features but have a relatively low average, end up *more birds* than birds with a higher average degree, which violate one feature. That is, the former birds end up *more birds* than the latter, though their degree is lower ($deg_{(d1,bird)} < deg_{(d2,bird)}$).

6. <u>How does the average in the features affect the ordering in multi-dimensional adjectives</u>? Intuitively, we feel that it does. It is now obvious that the degree function in multi-dimensional adjectives, say – $deg_{(d,healthy)}$, should order the positive and the negative denotations separately, by say – the average in the features, and then it should transform the degrees, such that the degrees of non-members will never exceed the degrees of members:

33) <u>A possible implementation:</u> $deg_{(d,healthy)}$ =

a. If $d \in [P]_t$: $deg_{(d,healthy)} = (\sum(\{w_F deg_{(d1,F)} \mid F \in F_{(healthy,t)}\} - a_P) / (a_P - b_P)$
(where a_p and b_P are the minimal and maximal degrees in $[P]_t$).
The degrees of P instances are given by the weighted mean in the dimensions and they are normalized so as to range from 1 to 0.

b. If $d \in [\neg P]_t$: $deg_{(d,healthy)} = [(\sum(\{w_F deg_{(d1,F)} \mid F \in F_{(healthy,t)}\} - a_{\neg P}) / (a_{\neg P} - b_{\neg P})] - 1$
(where $a_{\neg p}$ and $b_{\neg P}$ are the minimal and maximal degrees in $[\neg P]_t$).
The degrees of non-P instances are given by the weighted mean in the dimensions and they are normalized so as to range from 0 to –1.

This directly accounts for the fact that in many adjectives non-members are not ordered at all. Part b in the definition of *deg* in (33) in these cases is simply missing.

34) *Dan and Sam are healthy. # Dan is more / less sick than Sam*
35) *Dan and Sam are intelligent. # Dan is more/ less stupid than Sam*

As for *typical*, this proposal correctly predicts that, say – non biological fathers, may be judged more typical fathers than biological fathers. None is, strictly speaking, *a typical father*, but among the things which are not *typical-fathers*, the real father may be much less typical.

7. <u>Typicality and complex phrases</u>
What do we do when we want to construct a typicality degree function and categorization algorithm for <u>modified nouns</u> like *male nurse, pet fish, pet bird* or *wooden spoon*?
One way to go is the 'nouny' way, namely, to move to a new space of features, such that the average in them will give us the denotation. This predicts two well known phenomena: First, the ***emergence of new features*** (for example, *talks* and *lives in a cage* for *pet bird*, and *big* for *wooden spoon*) and second, the ***failure of inheritance of features from the constituents to the conjunction*** (for example, *fish* typically *live in the open oceans* and *pets* are typically *warm and affectionate*, but *pet fish* are neither). We see that these effects are not despite, but *due*, to intersectivity! In addition, the new features can be deduced, in just the same way as they are deduced in basic lexical items (see 'bias-L' in 9 below).
Another way to go is the 'adjectival' way. Namely, to remain in the same space of features, and to use a split and normalized degree function such as the one in (33).
Finally, a third way to go violates intersectivity and creates new meanings: It uses averaging on the union of the constituents' features (Hampton 1997).

8. <u>Psychological adequacy</u>: <u>Economic memory representation</u>
We now predict (like the basic prototype theory) that knowledge of the *bird* features will trigger automatic categorization of new entities which are good enough in the features.
Indeed, Kiran and Thompson 2003 found that, **in aphasics and neural networks which were taught the category *features*,** training with atypical members (say, a chicken or a goose), results in spontaneous recovery of categorization of untrained more typical items (items which are better in the features, like ducks), but not of untrained less typical items (ostriches.)
Similarly, Reed and others show that in healthy subjects, **previously unavailable typical instances are frequently (falsely) assumed to be known** (Reed 1988; Mervis & Rosch 1981). Why? If less typical entities are denotation members, entities which are more typical (due to their high scores in the typicality features) should definitely be denotation members!
Thus, we have a reasonable story about the memory representation of predicate intensions (functions to potentially infinite sets). We only need to directly represent in memory a (small) finite sample of instances for each predicate. The membership of the other entities automatically follows, once they become accessible.

9. Psychological adequacy: Acquisition

We also have a reasonable story about the acquisition of the intensions. Constraints L and F make the following bootstrapping mechanism or bias very plausible:

36) Bias L:

In the lack of knowledge about the category features, the early acquired members are the basis for generalization:

a. First, the early acquired members are assumed to be best in the category features.

b. Second, the category features are assumed to be precisely those properties in which these members average better than other familiar objects (which were classified later, or were classified as non-members, and hence should be worse in the features).

c. Finally, classification of new entities is inferred from their average in these features.

For example, normally, robins are classified earlier under *bird* and hence are considered more typical birds than chicken. Consequently, features in which robins average better than chicken (like *small* and *flying*) are linked to *birds*: Categorization is based on average in them.
Indeed, it was found that, **unless the features are directly taught**, acquisition is delayed if early exposure is to atypical items (say ostriches for *bird*), or even to the whole category, but not to the typical items first! (Mervis & Pani 1980). This is unexpected if features are selected based on their frequency within and outside the category, as is usually assumed. It is expected if features are selected based on the early acquired entities together.
That first exposure to an atypical item slows down acquisition supports the proposal that such situations trigger wrong inferences (37a-b). For example, if my initial exposure to birds was through ostriches, I would think that the ostrich is a representative *bird* and that its features (*running*, *ostrich size* etc.) are typical of the category. The inferences would be canceled later on, when items that fall out of the category would be discovered to be averaging better in these features than category members. But this process would slow down acquisition. Interestingly, in certain children acquisition is completely blocked within the experiment time. They refuse to abandon inferences from bias-L. This gives further evidence that it is at work.

Conclusions

The new direction of explanation seems to be promising and fruitful, and it allows a considerable improvement in the psychological adequacy of the linguistic theory.

Thank you!

gala@post.tau.ac.il