Abstract: Probabilistic Prototypes and Compositionality

Introduction

The debate whether prototypes (Rosch and Mervis 1975; Rosch 1978) are compatible with compositionality has been raging for almost 40 years (some recent examples are (Hampton and Jönsson 2012; Gleitman et al. 2012)). In this paper, I will sketch an account of prototypes that is compositional, but I will argue that such entities make poor candidates for semantic (truth-conditional) representations. Instead, I suggest that compositional prototypes provide an excellent basis as pragmatic reasoning. In the final section of the paper, I survey what role there remains for a truth-conditional semantics, given a compositional model of pragmatic reasoning.

The compositionality problem: Prototype theorists tend to collect data on typicality via asking speakers to make graded judgments on instances of some type. However, although speakers might not rate goldfish very highly in terms of their typicality as fish or pets, goldfish may get rated as being highly typical of pet fish.

Probabilistic prototypes: One feature, in common for both pets and fish may be *species*. The low typicality of *goldfish* for *pet* could be explained by the relative infrequency of 'pet' being applied to goldfish (compared with dogs or cats). The low typicality of *goldfish* for *fish* could be explained by the relative infrequency of 'fish' being applied to goldfish (compared with cod or tuna). In learning words, speakers will generalise over the instances of uses of words that they experience. On a probabilistic learning model, the resulting representation can, in part, be captured with conditional probabilities. For example, $p(goldfish(x)|pet(x))$ will be comparatively lower than $p(dog(x)|pet(x))$. Furthermore, probabilistic prototypes can be shown to be compositional.

An Example: Pet Fish

Views based on conditional probabilities have also been criticised for failing to capture compositionality (Fodor and LePore 1991). If $p(goldfish(x)|pet(x)) =$ low and $p(golfish(x)|fish(x)) =$ low, why should $p(goldfish(x)|petfish(x)) =$ high?

To some extent Fodor and Lepore are right. Conditional probabilities taken in isolation are not compositional. The mistake is in thinking that we should take them in isolation. Put another way, individual conditional probability judgements may not compose, but (conditional) probability distributions do compose.

As a simplification, I will take a restricted range of species features: *dog*, *cat*, *goldfish*, *cod*, *tuna*. We also have three prototypes: *pet*, *fish*, *pet* ∧ *fish*. *Goldfish* will receive a lower value under *pet* than *dog* and *cat*, but *cod* and *tuna* will not get a 0 (cod and tuna could be pets). *Goldfish* will receive a lower value under *fish* than *cod* and *tuna*, but *cat* and *dog* will not get a 0.[1]:

species	*cat*	*dog*	*goldfish*	*cod*	*tuna*
$p(species(x)\|pet(x))$	0.40	0.40	0.18	0.01	0.01
$p(species(x)\|fish(x))$	0.01	0.01	0.18	0.40	0.40

If accurate, the above values could be used to explain high/low typicality ratings for species under *pet* and under *fish*. Yet, whilst typicality ratings (on some scale) needn't compose, these two distributions can be combined using standard Bayesian apparatus:

[1]This assumes that the 'cat' in 'catfish' and the 'dog' in 'dogfish' carry some form of compound information. However, this is not a vital assumption and could be dropped by those who deny that any fish could be rightly described using either 'dog' or 'cat'

$$p(species(x)|pet(x) \wedge fish(x)) = \frac{p(species(x)) \times p(pet(x) \wedge fish(x)|species(x))}{p(pet(x) \wedge fish(x))}$$

species	*cat*	*dog*	*goldfish*	*cod*	*tuna*
$p(species(x)\|pet(x) \wedge fish(x))$	0.095	0.095	0.62	0.095	0.095

Which gives the right result. Given a classification of being a pet fish (a pet and a fish), and a list of options of cats, dogs, goldfish, cod and tuna, one should have a far greater expectation that the object described will be a goldfish rather than any of the other possibilities. Probability distributions over sets of possible features are compositional, and could be taken to underpin strengths of typicality judgements.

Compatibility with Truth-Conditional Semantics

A striking feature of the above analysis of prototypes is that full-blown truth-conditions get no mention at all. Nothing in the probabilistic prototypes structure determines whether it would be true of a given cod or goldfish that it counts as a pet fish. However, the above account of probabilistic representations is far from being 'clearly internal' because the probabilities of features can be seen as grounded in semantic learning, namely in correlations between types of speech situations and types of states of affairs or events (see (Cooper et al. 2014) for an example of a comparable learning model). As such, there is no impediment to linking such structures with a compositional truth-conditional semantics.

However, if the kind of reasoning involved in composing prototypes such as *pet* and *fish* is, strictly speaking, pragmatic, the need for a denotationally truth-conditional semantics may be diminished. I assumed above that 'pet fish' should be interpreted as a conjunction of prototypes, but nothing in the combination of prototypes requires that this is so. Therefore the only role required for semantics may be to determine the logical relationships between parts of analysed sentences, thus making semantics impoverished with respect to being fully truth-conditional in a denotational sense. Put another way, given that one can provide a compositional probabilistic pragmatics, albeit one hung onto a semantic account of logical form, it remains unclear whether a philosophical notion of semantics that spells out full-blown truth-conditions for sentences is required.

References

Cooper, R., Dobnik, S., Lappin, S., Larsson, S., 2014. A probabilistic rich type theory for semantic interpretation. Proceedings of the EACL 2014 Workshop on Type Theory and Natural Language Semantics.

Fodor, J. A., LePore, E., 1991. "Why Meaning (Probably) Isn't Conceptual Role". *Mind and Language* 6 (4), 328–43.

Gleitman, L., Connolly, A., Armstrong, S. L., 2012. Can prototype representations support composition and decomposition? In: Werning, M., Hinzen, W., Machery, E. (Eds.), The Oxford Handbook of Compositionality. Oxford University Press, pp. 385–402.

Hampton, J., Jönsson, M., 2012. Typicality and compositionality: The logic of combining vague concepts. In: Werning, M., Hinzen, W., Machery, E. (Eds.), The Oxford Handbook of Compositionality. Oxford University Press, pp. 385–402.

Rosch, E., 1978. Principles of categorization. In: Cognition and Categorization. Lawrence Elbaum Associates, pp. 27–48.

Rosch, E., Mervis, C. B., 1975. Family resemblances: Studies in the internal structure of categories. Cognitive Psychology 7 (4), 573–605.