

Compositionality for Classifiers

Staffan Larsson

Centre for Linguistic Theory and Studies in Probability (CLASP)
Department of Philosophy, Linguistics and Theory of Science
Gothenburg University, Sweden
sl@ling.gu.se

1. Perceptual meanings as classifiers

In Larsson (2015), a formal semantics for low-level perceptual aspects of meaning is presented, tying these together with the logical-inferential aspects of meaning traditionally studied in formal semantics. Perceptual meaning is an important aspect of the meaning of linguistic expressions referring to physical objects (such as concrete nouns or noun phrases). Knowing the perceptual meaning of an expression allows an agent to identify perceived objects and situations falling under the meaning of the expression. For example, knowing the perceptual meaning of “blue” would allow an agent to correctly identify blue objects. Similarly, an agent which is able to compute the perceptual meaning of “a boy hugs a dog” will be able to correctly classify situations where a boy hugs a dog.

The key idea in Larsson (2015) is to model *perceptual meanings as classifiers of perceptual input*. Classifiers can of course be implemented in many ways. However, most if not all can be defined formally as mathematical functions. In the visual domain, the domain of a classifier function is numerical (e.g. real-valued, integer or binary) vectors and the range is a set of classes (or in the case of binary classifiers, equivalents of “yes” and “no”).

The crucial step in making use of classifiers in formal semantics is to regard them as (parts of) *representations of intensions* of linguistic expressions. Traditionally, the intension of an expression helps determine whether some item belongs to the extension of the expression. Here, this translates to using a classifier to help determine whether some perceptual data derived from some item can be used to classify that item as falling under the expression, i.e., to be included in its extension.

One feature of this model is that perceptual meanings are updated as a result of observing language use in interaction, which means that perceptual meaning is derived from visual features co-occurring with linguistic expressions, and thus essentially distributional in nature. Classifiers capture generalisations over the distributional data used for training them.

2. Type Theory with Records

The ability to update meanings in this way requires a framework where intensions are (1) represented independently of extensions, and (2) structured objects which can be modified as a result of learning. We use Type Theory with Records (TTR), a formal semantics framework which starts from the idea that information and meaning is founded on our ability to perceive and classify the world, i.e., to per-

ceive objects and situations as being of types. As an example of our approach, in Larsson (2015) we show how a simple classifier of spatial information based on the Perceptron can be cast in TTR.

3. Background

Perceptual aspects of meanings have been explored in previous research (Roy, 2005; Steels and Belpaeme, 2005; Kelleher et al., 2005; Skočaj et al., 2010). However, the connection to logical-inferential meaning and compositionality as traditionally studied in formal semantics has not been a focus of this body of work.

In computational linguistics, a well-known approach to dealing with low-level meaning aspects is the Vector Space Model (Turney and Pantel, 2010), where the context of linguistic expressions are encoded as vectors in a space, whose dimensions typically represent co-occurring linguistic expressions. Such models can also represent non-linguistic aspects of the context such as perceptually salient objects and relations. VSMs can also represent semantic gradience phenomena, and offer an account of learnability of meanings (van Eijck and Lappin, 2012). There has also been work on compositionality for VSMs (Mitchell and Lapata, 2008; Coecke et al., 2010; Grefenstette and Sadrzadeh, 2011). Although compositional VSMs can in principle represent non-linguistic aspects of context, it is not clear how the vectors resulting from compositional analysis are to be interpreted (van Eijck and Lappin, 2012), or what role they play in inferences of the kind typically studied in formal semantics.

More recently, there has been computational work which is more in line with the approach taken here, e.g. Kennington and Schlangen (2015).

4. Compositionality

We believe that the account of perceptual meaning presented in Larsson (2015) can be useful in connecting low-level (“subsymbolic”) distributional aspects of meaning explicitly to formal semantics in a detailed and integrated manner. Doing this would ideally enable making use of the rich body of work in formal semantics of natural language from the last 50 or so years in the development of a more complete account of meaning which would combine both high-level (logical-inferential) and low-level distributional and perceptual aspects. A crucial step is to demonstrate how the principle of compositionality, which is at the heart of formal semantics, can be applied also to perceptual aspects of meaning.




Figure 1: Example of basic compositionality




Figure 2: Compositionality for degree modifier “far”

The key to dealing with compositionality for perceptual meaning, we believe, is to do compositionality not on the level of vectors, but on the level of classifiers. We thus take it that the perceptual meaning of a sentence can be composed from the perceptual meanings of the constituent expressions of the sentence.

5. Basic compositionality for classifiers

To demonstrate how compositionality for classifiers can be handled, we will deal with spatial location words (“right”, “upper”) and assume very basic grammar (allowing “upper left”, “lower right” etc). Furthermore, we assume sensors and classifiers corresponding to these words.

As a first proof of concept of compositionality in the framework put forward above, in Larsson (2015) we show how to compute the meaning of “upper right” from the meanings of “upper” and “right”. The compositional meaning of “upper right” is simply computed by merging the meanings of “upper” and “right” as in Figure 1.

6. Compositionality for degree modifiers

Of course we are not saying that all compositionality will be as trivial as in this example. For example, to correctly deal with a degree modifier such as “far right” one let the meaning of “far” modifying some parameter of the “right” classifier, rather than simply adding a further condition.

This assumes that the meaning of “far” specifies a function which is to be applied to the meaning of “right” (or whatever word appears adjacent to “far”). This would give a meaning of “far right” corresponding schematically to a classifier where the line dividing “right” from “not right” has been moved to the right (compared to the classifier for “right”), as shown in Figure 2.

7. Conclusion

We demonstrated two cases of compositionality for classifiers capturing distributional meanings. Note that in both cases, instead of trying to somehow compute compositionality on the level of vector representations, we instead computed it on the higher level of classifiers. In future work, we want to explore of this approach to compositionality can be applied to distributional meanings generally.

References

- Bob Coecke, Mehrnoosh Sadrzadeh, and Stephen Clark. 2010. Mathematical foundations for a compositional distributional model of meaning. *arXiv preprint arXiv:1003.4394*.
- Edward Grefenstette and Mehrnoosh Sadrzadeh. 2011. Experimental support for a categorical compositional distributional model of meaning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1394–1404. Association for Computational Linguistics.
- J Kelleher, F Costello, and J Vangenabith. 2005. Dynamically structuring, updating and interrelating representations of visual and linguistic discourse context. *Artificial Intelligence*, 167(1-2):62–102, September.
- Casey Kennington and David Schlagen. 2015. Simple learning and compositional application of perceptually grounded word meanings for incremental reference resolution. In *Proceedings of the Conference for the Association for Computational Linguistics (ACL)*, pages 292–301.
- Staffan Larsson. 2015. Formal semantics for perceptual classification. *Journal of Logic and Computation*, 25(2):335–369. Published online 2013-12-18.
- Jeff Mitchell and Mirella Lapata. 2008. Vector-based models of semantic composition. In *In Proceedings of ACL-08: HLT*, pages 236–244.
- Deb Roy. 2005. Grounding words in perception and action: computational insights. *Trends in Cognitive Sciences*, 9(8):389–396.
- Danijel Skočaj, Miroslav Janiček, Matej Kristan, Geert-Jan M. Kruijff, Aleš Leonardis, Pierre Lison, Alen Vrečko, and Michael Zillich. 2010. A basic cognitive system for interactive continuous learning of visual concepts. In *ICRA 2010 workshop ICAIR - Interactive Communication for Autonomous Intelligent Robots*, pages 30–36, Anchorage, AK, USA.
- Luc Steels and Tony Belpaeme. 2005. Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, 28(4):469–89, August. Target Paper, discussion 489–529.
- Peter D. Turney and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *J. Artif. Intell. Res. (JAIR)*, 37:141–188.
- Jan van Eijck and Shalom Lappin. 2012. Probabilistic semantics for natural language. *Logic and Interactive Rationality (LIRA)*, 2:17–35.