Perceptual Meaning in TTR Judgement-based Semantics and Conceptual Spaces

Staffan Larsson

Dept. of Philosophy, Linguistics and Theory of Science
University of Gothenburg, Sweden

Abstract. We are developing a type-theoretical judgement-based semantics where notions such as perception, classification, judgement, learning and dialogue coordination play a central role. By bringing perception and semantic coordination into formal semantics, this theory can be seen as an attempt at unifying cognitive and formal approaches to meaning. The purpouse of this paper is to briefly compare judgement-based semantics to the theory of conceptual spaces. We argue that the former enables integration of perceptual aspects of meaning with those traditionally studied in formal semantics, and furthermore that it enables computational modeling and implementation of these aspects of meaning.

1 Introduction

How is linguistic meaning related to perception and the world? How do words acquire their meaning? These are two central questions for any theory of meaning in natural language (Miller and Johnson-Laird, 1976) . We are developing a type-theoretical *judgement-based semantics* where notions such as perception, classification, judgement, learning and dialogue coordination play a central role (Cooper, 2005; Cooper and Larsson, 2009; Larsson, 2011; Dobnik *et al.*, 2011; Cooper, 2012; Dobnik and Cooper, 2013; Cooper *et al.*, 2015; Larsson, 2013). Meaning is regarded as being acquired by an agent through its perception of, and interaction with, the world and other agents. This makes meaning agent-relative but essentially social (in the sense of being coordinated in interaction between individuals) and dynamic (in the sense of always being up for revision and negotiation as new perceptual and conversationally mediated information is encountered). By bringing perception and semantic coordination into formal semantics, this theory can be seen as an attempt at unifying cognitive and formal approaches to meaning. The purpouse of this paper is to briefly compare judgement-based semantics to the theory of conceptual spaces (Gärdenfors, 2000; Gärdenfors, 2004).

2 Judgement-based semantics and perceptual meaning

We will here take for granted that relating meaning to perception is of central importance when accounting for the meaning of linguistic expressions which refer to the physical world; for a motivation see Larsson (2013). Knowing the meaning of an expression is related to an agent's ability to identify perceived objects and situations which can be referred to by the expression. For example, knowing the meaning of "blue" is intimately connected with an agent's ability to correctly identify blue objects. Similarly, an agent's ability to assign a meaning to "a boy hugs a dog" is related to her ability to correctly classify perceived situations where a boy hugs a dog. An important difference to traditional possible worlds semantics (Montague, 1974) is that we focus on modelling perceptual mechanisms which start from perceptual raw data, rather than giving an abstract representation of them as functions from possible worlds and times to objects in a space of denotations as in the Montagovian approach. This allows us to provide a more concrete account of the nature of the link between language and the world.

To make the link between "low-level" perceptual data and "high-level" formal linguistic representations, we will use the notion of a (statistical) classifier, a computational device determining what class an item belongs to, based on various properties of the item. Crucially, the information fed to a classifier need not be encoded in some high-level representation language (such as logic or natural language). Instead, it may consist entirely of empirical raw data encoding "low-level" information about the item in question. The idea of using classifiers (or more specifically, connectionist models) to represent meanings was first put forward by Harnad (1990) as a way of addressing the "symbol grounding problem" in artificial intelligence, and is consistent with several theories of word meaning as grounded in sensory (or *embodied*) representations which have emerged during the last decade or so (Roy, 2005; Steels and Belpaeme, 2005).

To integrate classification of perceptual data with formal semantics, we are using TTR (Type Theory with Records), a framework developed with a view to giving an abstract formal account of natural language interpretation (Cooper, 2012), as our formalism and foundational semantic theory. TTR starts from the idea that information and meaning is founded on our ability to perceive and classify the world, i.e., to perceive objects and situations as being of types. In TTR, types are first-class objects, which allows perceptual classifier functions to be formalised and used in representing meanings of linguistic expressions together with the high-level conceptual aspects of meaning traditionally studied in formal semantics. Semantic phenomena which have been described using TTR include modelling of intensionality and mental attitudes (Cooper, 2005), dynamic generalised quantifiers (Cooper, 2004), co-predication and dot types in lexical innovation, frame semantics for temporal reasoning, reasoning in hypothetical contexts (Cooper, 2011), enthymematic reasoning (Breitholtz and Cooper, 2011), clarification requests (Cooper, 2010), negation (Cooper and Ginzburg, 2011), and information states in dialogue (Cooper, 1998; Ginzburg, 2012).

Larsson (2011) and Larsson (2013) show how a simple classifier of sensory information based on the perceptron can be cast in TTR, and how an agent can learn from interaction by training the classifier based on linguistic and perceptual input. Linguistic input, e.g. utterance of "That's red" (assuming a situation where two agents are inspecting a colour sample), is interpreted as a function from a situation where some object is in the (shared) focus of attention and there is a sensor reading (e.g. in the form of a real-numbered vector, but shown below as a colour patch) from a colour sensor, to a judgement that the situation is of a type where the object is judged to be red.

$$\begin{bmatrix} \text{foc-obj} & = \text{a} \\ \text{sensor}_{colour} & = {\color{red}\blacksquare} \end{bmatrix} : \text{red(a)}$$

To achieve this, the function contains a classifier which takes the (real-numbered vector corresponding to the perception of the) colour sample and produces a judgement whether the vector is within the borders of redness. For details on how classifiers are embedded into TTR functions, see Larsson (2013), which also includes a brief account of compositionality.

Categorical judgments of the kind exemplified above are of course not suited for accounting for vagueness and other gradient semantic phenomena. To remedy this, a probabilistic extension of TTR has been developed (Cooper *et al.*, 2014, 2015). For an account of vagueness in perception using probabilistic TTR, where judgements are associated with probabilities, see Fernández and Larsson (2014). Below is an example of a probabilistic judgement where a situation is judged, with a probability of 0.79, to be one where the object in the focus of attention is red.

$$\begin{bmatrix} \text{foc-obj} & = \text{a} \\ \text{sensor}_{colour} & = {\color{red}\blacksquare} \end{bmatrix} :_{0.79} \text{red(a)}$$

In probabilistic TTR, the result of an act of classification is represented as a probability distribution over type assignments, i.e., as a set of probabilistic judgements where the probabilities sum to 1.

3 Observations in conceptual spaces

There are interesting connections between the idea of using classifiers to model perceptual meaning, and the notion of conceptual spaces. Gärdenfors distinguishes three levels of modeling concepts and reasoning: symbolic, subconceptual, and conceptual.

Reasoning on the *symbolic level* is framed as operations on propositions expressed by symbolic structures (i.e., symbol manipulation according to explicit rules), and focuses on computing logical consequences (i.e. deductive reasoning). Gärdenfors also discusses connectionist models of meaning as an example of representations on a *subconceptual* level, where reasoning modelled by the activities of the artificial neurons. Concepts on the subconceptual are modelled "implicitly", in contrast to the more explicit *conceptual* level where concepts are modelled as geometrical structures (points, vectors and regions) in conceptual spaces. Conceptual reasoning is described in terms of distances in a space, and focuses on modeling reasoning about concepts, in particular inductive and nonmonotonic reasoning.

Corresponding to these three kinds of reasoning, Gärdenfors makes a distinction between three ways of describing an observation, which we will use to frame our discussion. On the symbolic level, observations are described in some specified language with a fixed set of primitive predicates. Denotations of predicates assumed to be known, and observational statements furnished to reasoner by incorrigible perceptual mechanisms. On the conceptual level, concepts characterised in terms of some underlying conceptual space, consisting of a number of "quality dimensions" (or *domains*). An observation on this level is an assignment to an object of a location in a conceptual space. For example, an observation that "x is red" is expressed by assigning x a point in colour space. Finally, on the subconceptual level an observation is regarded as something which is received by our sensory organs, or in general some kind of *receptors*, including e.g. measuring instruments.

4 Perceptual judgements and conceptual spaces

Below, we will explore the relation between conceptual spaces and classifiers as different (but to some extent complementary) ways of capturing perceptual meaning. Roughly, the correspondence is the following: classification events can be regarded as making a judgement as to whether an observation falls within that region in a conceptual space. Also, classifier learning can be regarding as defining areas (regions or volumes) within a conceptual space.

4.1 Sensor readings and the subsymbolic level

It appears fairly obvious that the sensor readings in TTR correspond to Gärdenfors' subsymbolic representations. Both represent low-level perceptual input using vectors, points or regions in vector spaces. In judgement-based semantics, sensor readings are the input to classifiers.

4.2 Types and the symbolic level

Gärdenfors' symbolic representations of observations appears to correspond to the types which result from judgements. However, in TTR we do not assume that there is a fixed set of primitive predicates. Instead, we are interested in modeling concept learning (Larsson, 2013). Perhaps even more importantly, we do not assume that the extensions (denotations) of predicates are known. Instead, we represent meanings of concrete expressions using classifiers which take some perceputal input and produce a judgement. The classifiers can be thought of as representing the *intensions* of linguistic expressions. Since these classifiers can be trained, they are also dynamic and learnable. Furthermore, TTR does not assume that perceptual mechanisms are always correct, nor that agents always agree on their perceptions of a situation or their judgements about the situation. Finally, a nice feature of TTR is not only properties but also relations are types which opens the way for compositional semantics involving predicates of arbitrary arity.

4.3 Classifiers, judgements and the conceptual level

According to Gärdenfors, an observation on the conceptual level is an assignment to an object of a location in a space. In TTR, this corresponds to an act of classification producing a (possibly probabilistic or graded) judgement concerning (the probability of) a situation being of a certain type, thus mediating between the sensor reading and the high-level "symbolic" types. In this way, classifiers connects subsymbolic observations and semantic concepts to "symbolic" reasoning.

Induction is indeed closely related to concept formation, since our concepts are formed and learned by induction from observations (including observations in interaction). Gärdenfors' claim that "induction can be seen as establishing connections between various kinds of input" is echoed in TTR in that classifiers are trained by generalising over several instances, thereby connecting several instances. An advantage of classifiers is that they are straightforwardly implementable, and that classification on sensory input is a well-studied research area.

5 Conclusion

We have compared judgement-based semantics with conceptual spaces, and concluded that there are important similarities but also some differences. One aim of TTR judgement-based semantics is to formalise semantic classification and learning in detail, to enable integration of these aspects of meaning with those traditionally studied in formal semantics, and to enable computational modeling and implementation of these aspects of meaning. By using statistical classifiers, we connect to machine learning theory, giving access to a host of classification methods and associated learning algorithms.

Bibliography

Breitholtz, E. and Cooper, R. (2011). Enthymemes as rhetorical resources. In *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, pages 149–157, Los Angeles (USA).

Cooper, R. (1998). Information states, attitudes and dependent record types. In *ITALLC98*, pages 85–106.

Cooper, R. (2004). Dynamic generalised quantifiers and hypothetical contexts. In *Ursus Philosophicus, a festschrift for Björn Haglund*. Department of Philosophy, University of Gothenburg.

Cooper, R. (2005). Austinian truth, attitudes and type theory. *Research on Language and Computation*, **3**, 333–362.

Cooper, R. (2010). Generalized quantifiers and clarification content. In P. Łupkowski and M. Purver, editors, *Aspects of Semantics and Pragmatics of Dialogue. SemDial 2010, 14th Workshop on the Semantics and Pragmatics of Dialogue*, Poznań. Polish Society for Cognitive Science.

Cooper, R. (2011). Copredication, quantification and frames. In S. Pogodalla and J.-P. Prost, editors, *LACL*, volume 6736 of *Lecture Notes in Computer Science*, pages 64–79. Springer.

Cooper, R. (2012). Type theory and semantics in flux. In R. Kempson, N. Asher, and T. Fernando, editors, *Handbook of the Philosophy of Science*, volume 14: Philosophy of Linguistics. Elsevier BV. General editors: Dov M. Gabbay, Paul Thagard and John Woods.

Cooper, R. and Ginzburg, J. (2011). Negation in dialogue. In *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA).

Cooper, R. and Larsson, S. (2009). Compositional and ontological semantics in learning from corrective feedback and explicit definition. In J. Edlund, J. Gustafson, A. Hjalmarsson, and G. Skantze, editors, *Proceedings of DiaHolmia, 2009 Workshop on the Semantics and Pragmatics of Dialogue*.

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

Cooper, R., Dobnik, S., Lappin, S., and Staffan, L. (2015). Probabilistic type theory and natural language semantics. Under review.

Dobnik, S. and Cooper, R. (2013). Spatial descriptions in type theory with records. In *Proceedings of IWCS 2013 Workshop on Computational Models of Spatial Language Interpretation and Generation (CoSLI-3)*, pages 1–6, Potsdam, Germany. Association for Computational Linguistics.

Dobnik, S., Larsson, S., and Cooper, R. (2011). Toward perceptually grounded formal semantics. In *Proceedings of the Workshop on Integrating Language and Vision at NIPS 2011*, Sierra Nevada, Spain. Neural Information Processing Systems Foundation (NIPS).

Fernández, R. and Larsson, S. (2014). Vagueness and learning: A type-theoretic approach. In *Proceedings of the 3rd Joint Conference on Lexical and Computational Semantics (*SEM 2014)*.

Gärdenfors, P. (2000). *Conceptual spaces - the geometry of thought*. MIT Press.

Gärdenfors, P. (2004). Conceptual spaces as a framework for knowledge representation. *Mind and Matter*, **2**(2), 9–27.

Ginzburg, J. (2012). *The Interactive Stance*. Oxford University Press, New York.

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, **42**(1990), 335–346.

Larsson, S. (2011). The ttr perceptron: Dynamic perceptual meanings and semantic coordination. In *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA).

Larsson, S. (2013). Formal semantics for perceptual classification. *Journal of Logic and Computation*.

Miller, G. A. and Johnson-Laird, P. N. (1976). *Language and perception*. Belknap Press.

Montague, R. (1974). *Formal Philosophy: Selected Papers of Richard Montague*. Yale University Press, New Haven. ed. and with an introduction by Richmond H. Thomason.

Roy, D. (2005). Grounding words in perception and action: computational insights. *Trends in Cognitive Sciences*, **9**(8), 389–396.

Steels, L. and Belpaeme, T. (2005). Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, **28**(4), 469–89. Target Paper, discussion 489-529.