

**Formal and distributional semantics model different notions of meaning**  
**Matthijs Westera & Gemma Boleda, Universitat Pompeu Fabra**

---

**1. Introduction.** In **formal semantics**, the meaning of an expression is typically captured in terms of truth conditions and reference to objects in (a formal model of) the external world. In **distributional semantics**, the meaning of an expression instead is an abstraction over its uses. Formal semantics is successful at characterizing *inferential aspects* of language (e.g., entailments, reference, scope readings) while falling short of capturing *conceptual aspects* (e.g., conceptual structure, similarity, association) – and vice versa for distributional semantics (Boleda and Herbelot 2016). Distributional semantics in its various guises has had a huge success in Cognitive Science and Computational Linguistics. In general, and because meaning comprises both inferential and conceptual aspects, formal and distributional semantics are regarded as two partial characterizations of a single notion of meaning, with complementary strengths and weaknesses (*ibid.*). By contrast, we propose to view them as **characterizing different notions of meaning**, as follows.

**2. Proposal** For reasons of scope we shall adopt without discussion the view that the meanings of sentences and their constituents are determined by **conventions of language use**: what an *expression* means, henceforth **expression meaning**, is constituted by what speakers of the relevant linguistic community tend to mean by it (e.g., Grice 1989). What a speaker means when using an expression on a particular occasion, henceforth **speaker meaning**, is constituted by the speaker’s communicative intentions on that occasion. (Other views on meaning may or may not require a different understanding of distributional and formal semantics than the one we will propose here.)

Speaker meaning, along with related notions such as information, entailment, belief, goal and intention, can be modeled by symbolic logic, in terms of truth conditions and reference to discrete objects in a formal model of the external world. This is of course an **abstraction** (cf. Marr 1982), but one which has proven its value over the ages, and we shall not discuss its usefulness further here.

Formal semantics effectively carries the abstraction over from speaker meaning to expression meaning, stipulating that not just the contents of speaker intentions but also bare words and sentences, irrespective of any particular usage event, can be ascribed well-defined, discrete logical meanings. This **extended abstraction**, while also very fruitful, has more serious shortcomings, and these have frequently been pointed out (the late Wittgenstein being an early example). Two issues stand out. First, the ambition to assign discrete formal semantic objects to expressions out of context is difficult to reconcile with the tremendous variation and gradience in the use of expressions, even by a single speaker, let alone in the relevant linguistic community (a notion that is itself intrinsically vague; cf. Chomsky 1986). Second, referents and truth conditions of expressions can often not be identified or even stated without reference to a pragmatic context or speaker intention (e.g. Recanati, 2004).

Distributional semantics, by contrast, does not rely on the extended abstraction. Instead, it embraces contextual variability, by assigning to each expression a compressed, graded representation of its distribution across contexts of use (typically a high-dimensional, numerical vector, induced from large amounts of naturally occurring data). Importantly, the notion of *context* in distributional semantics is not limited to textual context, but potentially includes all aspects of *extralinguistic* context, including aspects of speaker intentions such as reference to objects in a visual scene (Baroni 2016). Given the view we adopt here that expression meaning is a matter of conventions of use, this makes distributional semantics ideally suited for capturing expression meaning.

This raises the issue of what formal semantics effectively models (if not expression meaning). A large part of the answer, we propose, is that native speakers' judgments about sentences – the primary data against which formal semantic theories are empirically tested – in fact reflect their interpretations of these sentences *as uttered by some imagined speaker in a typical context* (e.g., Strawson 1950; Kadmon and Roberts 1986; Schwarz 1996; Bach 2002; Westera and Brasoveanu 2014). Put differently, what formal semantics actually models is closer to the notion of speaker meaning (see Borge 2009 for an independent argument towards the same conclusion, based on compositionality). Accordingly, we propose the following division of labor, which we detail in the talk:

While distributional semantics models expression meaning, formal semantics models (something closer to) speaker meaning.

**3. Outlook.** This view affords conceptual clarity, and it helps avoid certain common misunderstandings. First, it reveals the formal/distributional distinction to be independent of the symbolism/connectionism debate: the abstraction of formal logic can be dropped for expression meaning yet maintained for speaker meaning. It also reveals that, whereas distributional semantics is sometimes deemed more 'pragmatic' than formal semantics on grounds that it is usage-based (Kruszewski et al. 2016), the opposite is true in the sense that distributional semantics models conventions while formal semantics models intentions. We will clarify furthermore that the formal/distributional distinction is largely independent of the competence/performance and symbolic/sub-symbolic divides.

More practically, our proposal implies that we must not try to model formal notions like entailment with distributional semantics (something that has proven quite unsuccessful; Boleda and Herbelot 2016, a.o.), for this is not what it represents. It also means, in line with the work cited above, that reference, truth conditions, and entailments should be modeled as properties not of expressions as such, but of utterances (or expressions-as-uttered) by a speaker in a given context. Schwarz (1996) independently argues for this position as desirable to correctly interpret experimental results from psycholinguistics.

But our proposal's core contribution is that **it reframes the integration of distributional and formal semantics**. The two notions of meaning, hence the two types of semantics, have very different explanatory roles to play in a theory of language: Speakers rely on expression meaning as a tool for communicating speaker meaning, or, equivalently, they rely on the past usage of words in their linguistic community as a tool for communicating their intentions in the present (e.g., Recanati 2004, ch.8). Accordingly, integrating the two types of semantics is a matter of modeling this process, rather than, as previously suggested, combining their supposed complementary strengths into a single notion of meaning. We will argue that many concrete models in computational linguistics can be retroactively understood as already partially implementing our proposal.

---

**References:** Bach, K. (2002). Seemingly semantic intuitions. *Meaning and Truth*. Seven Bridges Press. • Baroni, M. (2016). Grounding distributional semantics in the visual world. *L.&L.Compass* 10(1). • Boleda, G. & Herbelot, A. (2016). Formal distributional semantics: Introduction to the special issue. *CompLing* 42(4). • Borge, S. (2009). Intentions and compositionality. *SATS*, 10(1). • Chomsky, N. (1986). *Knowledge of language*. Praeger. • Grice, H. (1989). *Studies in the Way of Words*. Harvard UP. • Kadmon, N. & Roberts, C. (1986). Prosody and scope. *CLS* 22. • Kruszewski, G., et al. (2016). There is no logical negation here... *CompLing* 42. • Marr, D. C. (1982). *Vision*. Freeman & Co. • Recanati, F. (2004). *Literal meaning*. Cambridge UP. • Schwarz, N. (1996). *Cognition and Communication*. Erlbaum. • Strawson, P. F. (1950). On referring. *Mind*, 59. • Westera, M. & Brasoveanu, A. (2014). Ignorance in context. *SALT* 24.

**Acknowledgments:** This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 715154). This paper reflects the authors’ view only, and the EU is not responsible for any use that may be made of the information it contains.

