How to marry a star: Probabilistic constraints for meaning in context

Katrin Erk
Department of Linguistics
University of Texas at Austin
katrin.erk@utexas.edu

Aurélie Herbelot
Center for Mind/Brain Sciences
University of Trento
aurelie.herbelot@unitn.it

Most words can take on different meanings based on their context. Some influences come from *local context*, like selectional preferences, e.g. the agent of a sleeping event is generally an animate being. But *global context* also plays a role. (1) is a contrast pair with different senses of the word *ball* (sports equipment vs dancing event). Arguably, the sense of the predicate *run* is the same in (1-a) and (1-b), so the difference in the senses of *ball* must come from something other than direct semantic neighbors. We can characterize this influence as global topical context brought about by the presence of *athlete* in the first sentence, and *violinist* in the second.

(1) a. The athlete ran to the ball.
b. The violinist ran to the ball.

There is even a whole genre of jokes relying on a *competition* of local and global topical constraints: the pun. In sentence (2), the pun rests on two senses of the word *star*, which can be paraphrased as 'well-known person' and 'sun'.

(2) The astronomer married the star.

It is interesting that this sentence should even work as a pun: The predicate that applies to *star* (*marry*) clearly selects for a person as its theme. So if the influence of local context were to apply strictly before global context, *marry* should immediately disambiguate *star* towards the 'person' sense as soon as they combine. But the 'sun' sense is clearly present.

As Del Pinal (2018) points out, much of the flexibility of word meaning can be explained through conceptual knowledge associated with lexical items, rather than general pragmatic reasoning. We make a similar argument for topical context which, as we saw in (1), can be rooted in the conceptual correlate of lexical items (e.g. *athlete* invokes sport

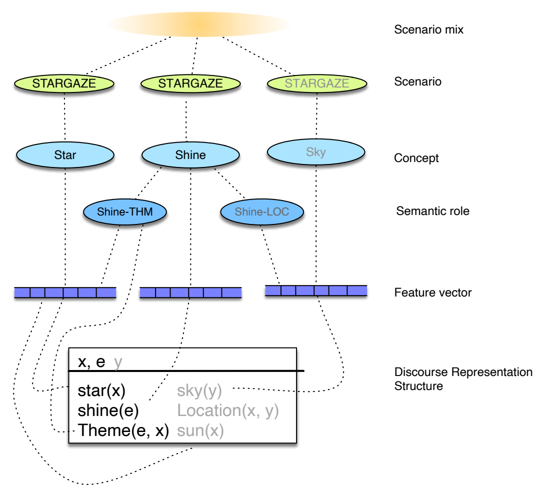

Figure 1: Illustration of the generative model using the sentence "a star shines".

equipment rather than high-society events).

We model both local and global context as a system of interacting, probabilistic constraints that let the listener imagine the scene described by a given utterance. So we obtain a probabilistic generative model that describes utterance understanding as a process of generating a description of a situation, subject to both local and global constraints associated with the lexical items. This process can be viewed as a formalization of Fillmore's "semantics of understanding" or 'U-semantics' (Fillmore, 1985), the aim of which is to give *"an account of the ability of a native speaker to 'envision' the 'world' of the text under an interpretation of its elements"* (p.235). The idea is that the listener uses the frames that are 'evoked' by the words in the utterance to *"[construct] an interpretation of the whole"* (p. 233), a full description of a scene, including elements that may not be explicit in the original sentence.

Proceedings of the Society for Computation in Linguistics (SCiL) 2021, pages 451-453.
Held on-line February 14-19, 2021

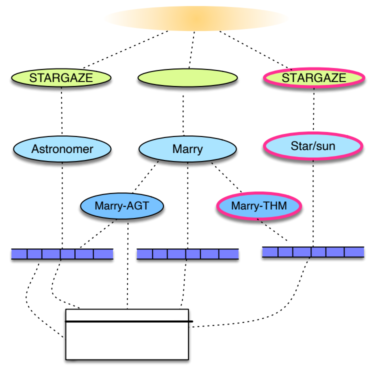

Figure 2: Competing constraints (outlined in red) in "The astronomer married the star."

Our model encompasses both a conceptual and a referential representation of meaning, where the generative process is such that the conceptual representation generates the conditions of a Discourse Representation Structure (DRS, Kamp and Reyle (1993)). Fig. 1 provides an illustration of the generative process using the sentence "A star shines". At the bottom of the figure is the DRS; everything above it constitutes the conceptual representation of the utterance. The referents and conditions of the DRS in black represent the original utterance, the material in grey is added as the listener 'fleshes out' the utterance, adding e.g. a sky as the *Location* of the star.

At the conceptual level, each referent in the DRS is associated with a concept (light blue). Selectional constraints (darker blue) and concepts jointly generate finer-grained feature vectors to account for meaning modulation, following e.g. Asher (2011) and McNally and Boleda (2017). The feature vectors then generate the DRS conditions for that referent. Each concept is further associated with a scenario (light green), where formally a scenario is simply a distribution over concepts. The sole global constraint is the distribution over scenarios (scenario mix, yellow). Through a soft constraint that prefers sparse scenario mixes, with only few different scenarios, we immediately obtain representations that tend to be topically coherent.

We formalize the process as a **situation description system (SDS)**. An SDS generates a representation of an utterance not as single situation description, as shown in Fig. 1, but a distribution over situation descriptions. For instance, *the star shines* might evoke a situation description containing a bright celestial object with 0.9 probability and a situation description with a witty entertainer with 0.1 probability. The evoked concepts provide something akin to traditional 'sense disambiguation': generating a concept CELESTIAL OBJECT rather than ACTOR in a given situation description clearly selects a given sense of the word *star*. However, our own notion of a word's meaning is more fine-grained and can be expressed as a function of the entire network of concepts, scenarios and attributes that accounts for its presence in the utterance.

The framework also accounts for utterances where senses oscillate between two readings, as in *The astronomer married the star*, where the associated distribution over situation description would contain descriptions of *both* prominent readings. Figure 2 sketches one reading of the *astronomer* sentence, with the competing selectional constraint and scenario outlined in red. This reading would be obtained because of the sparseness preference for scenario mixes. Similarly, in cases where the sentence is ambiguous without being a pun(e.g. *I saw the star*), the probability distribution over scenarios will generate the relevant readings as separate situation descriptions. We can further imagine scenario probabilities to be influenced by previous discourse, leading to non-uniform priors.

Small-scale examples of the framework have been implemented to illustrate the behaviour of the system with respect to a) global constraints (when does *bat* evoke a gothic novel rather than a baseball game?); b) local constraints at the semantic role level (when hearing *the vampire eats*, which food will the listener most likely infer?); and c) modifier-head combination at the local feature level (what are the features of a *fanged bat*?). We finish our exposition by showing how the SDS can also remain agnostic about sense when presented with a pun.

The next task will be to scale our implementation to account for arbitrary English sentences. We would then like to evaluate the framework on its ability to simulate human behaviour on tasks like meaning similarity prediction and paraphrasing, as well as expectations on upcoming words.

References

Nicholas Asher. 2011. *Lexical Meaning in Context: A Web of Words*. Cambridge University Press, Cambridge.

Guillermo Del Pinal. 2018. Meaning, modulation, and context: A multidimensional semantics for truth-conditional pragmatics. *Linguistics and Philosophy*, 41(2):165–207.

Charles J. Fillmore. 1985. Frames and the semantics of understanding. *Quaderni di Semantica*, 6:222–254.

Noah D Goodman and Andreas Stuhlmüller. 2014. The Design and Implementation of Probabilistic Programming Languages. `http://dippl.org`. Accessed: 2020-6-19.

Hans Kamp and Uwe Reyle. 1993. *From discourse to logic*. Kluwer, Dordrecht.

Louise McNally and Gemma Boleda. 2017. Conceptual versus referential affordance in concept composition. In James Hampton and Yoad Winter, editors, *Compositionality and Concepts in Linguistics and Psychology*, volume 3. Springer.

A Appendix: Example output

We give here two illustrations of the behaviour of our system when given an utterance, showing how it implements particular characteristics of natural language comprehension, and how meaning contextualisation naturally derives from such characteristics. The system is implemented in the probabilistic programming language *WebPPL* (Goodman and Stuhlmüller, 2014).

To illustrate how a situation description system fills in the details of an utterance, we inspect the behaviour of the system when presented with the utterance *A vampire is eating*. We have set the verb *eat* to probabilistically take a patient with high probability, and a location with lower probability. We want to show how a listener fills in the details of the sentence using their prior world knowledge. We sample 2000 situation descriptions for the utterance in a system that has a single scenario containing the concepts VAMPIRE, EAT, BLOOD_ORANGE, BAT-ANIMAL, and CASTLE. We have set the patient role of EAT to strongly prefer food stuff (that is, the concept BLOOD_ORANGE has by far the highest probability), and the location to prefer buildings over other concrete objects. Despite not being explicitly realized in the utterance, the patient role is activated with probability 0.71 and the location role with probability 0.25. Table 1 shows the probabilities of situations descriptions with particular patients / locations. As we can see, our vampire's food is most likely to be oranges, and she is more likely to eat in a castle than located at an orange. Since we are implementing soft constraints, though, we do also retain small probabilities that she is eating another vampire, a castle or a bat.

Concept	p Patient	p Location
BLOOD-ORANGE	0.64	0.004
VAMPIRE	0.02	0.005
BAT-ANIMAL	0.03	0.004
CASTLE	0.02	0.24

Table 1: Probabilities of situation descriptions with particular patients / locations for the utterance *A vampire is eating*.

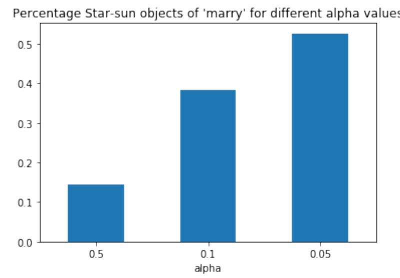

Figure 3: Percentage of STAR-SUN concepts for different values of the soft constraint that prefers few scenarios (the Dirichlet concentration parameter α.)

Secondly, we return to the pun example *The astronomer married the star*, which plays on conflicting constraints (one coming from the scenario level, one from the selectional preference of the verb *marry*). We now illustrate how our system retrieves both senses of *star* in different proportions. In particular, we show that we can control the interpreter's preference with respect to 'mixing scenarios' through a parameter of the probability distribution from which we draw a distribution over scenarios. This is the concentration parameter α of the Dirichlet distribution. By setting the parameter to prefer only having few scenarios in the mix, we introduce a competition of the scenario mix with the verb's selectional preference. We set the theme of MARRY to strongly prefer person fillers but also allow arbitrary object fillers with lower probability. Fig 3 shows the percentage of STAR-SUN concepts returned across 2000 situation descriptions, for different values of α. The α of 0.5 clearly prefers the *person* sense of STAR. As we decrease α and more strongly prefer to stay within a single scenario, however, more and more *sun* senses are introduced, resulting in the typical pun effect where the listener is left hanging between interpretations.