The evolution of adjectival monotonicity

The evolution of adjectival monotonicity

A popular analysis sees adjectives as functions $\langle d, et\rangle$ from degrees to functions from individuals to truth values or, simplifying slightly, from degrees to the sets of individuals that have the relevant property to the specified degree. This accounts straightforwardly for the compositional semantics in so-called *measure* context, such as "Franka is 1.80m tall". However, so-called *bare context*, such as "Franka is tall", cause a type clash. To solve this problem, a silent morpheme *POS* of type $\langle\langle d, et\rangle, et\rangle$ that accepts adjectives as arguments is introduced resulting in the analysis "Franka is (POS) tall" [5]. POS is responsible for the semantic behaviour of gradable adjectives in bare contexts. POS is analysed as follows: $[\![\mathrm{POS}]\!] = \lambda G_{\langle\langle d,et\rangle,et\rangle}.\lambda x_e.\exists d[\mathbf{standard}(d)(\mathrm{G})(\mathbf{C}) \wedge \mathrm{G}(d)(x)]$, where $\mathbf{C}$ is a comparison population, G an adjective, and d a degree. Roughly, predicates of the form "is (POS) ADJ" are true for those individuals that have the property relative to ADJ to some degree d_i that verifies $\mathbf{standard}(d_i)(\mathrm{ADJ})(\mathbf{C})$.

The relation **standard** is true cross-linguistically (e.g. [2], [4], [7], [8], [9]) for any $\mathbf{C}$, G and d iff d exceeds a threshold on the scale associated with G that is computed relative to the comparison class $\mathbf{C}$ (we ignore issues relating to absolute and relative adjectives as well as antonymy in a way that is unsubstantial for the present discussion). For instance, "Maria is tall (for a Dutch person)" is true iff Maria has height to a degree that exceeds some standard calculated relative to the Dutch population. From this analysis follows that **standard** is monotonic in its degree argument, i.e. given a $\mathbf{C}$ and G, $\forall d_i \forall d_j[d_j \geq d_i \wedge \mathbf{standard}(d_i)(\mathbf{C})(\mathrm{G}) \rightarrow \mathbf{standard}(d_j)(\mathbf{C})(\mathrm{G})]$. This analysis predicts the validity of inferences such as "If Edinburgh is cold and Glasgow is colder than Edinburgh, then Glasgow is cold". Since the involved types do not constrain which sets of degrees verify **standard** given a $\mathbf{C}$ and G, the monotonicity of adjectival bare contexts with respect to degrees (henceforth simply "adjectival monotonicity") is an empirical fact about **standard** (see [1] for a familiar similar argument concerning quantifiers). The fact that monotonicity is not a semantic necessity raises the question of which pressures cause it to evolve. We present three computational models that built on top of each other to explain adjectival monotonicity.

The first model uses the Iterated Learning paradigm [6] to study how gradable adjectives evolve under a pressure coming from cognitive biases alone in a population of literal Bayesian agents. Our Iterated Learning model starts with a population of agents speaking random languages. Each language consists of three signals, each associated with a predicate modelling a possible meaning of the **standard** relation with respect to its degree argument. Some such predicates are monotonic, some are not. Each generation (except for the first) observes linguistic data produced by the previous generation and tries to infer the words' meanings. The learning happens via Bayesian update given the observed data. The agents' prior gives higher probability to meanings with smaller description length in a language we describe. Intuitively, this model studies the evolutionary trajectory adjectival meaning would take if all that people had to do was learn the meanings of adjectival bare contexts, but not use adjectives to communicate. Over many generations, under the proposed prior the meanings become monotonic but also degenerate, i.e. they cover either no degree or all the degrees. This is unlike the behaviour of natural language adjectives.

Intuitively, the reason why degenerate meanings would not evolve lies in the fact that they do not convey information about the world, and are thus useless for communication. We make this intuition formal in the second model by studying the effects of adding a pressure for communicative accuracy. The pressure for communicative accuracy is modelled as follows. In each generation except the first, the probability that agents in the following generation observe data produced by an agent a is proportional to the accuracy with which a communicates with its own cultural parent. Languages that are more successful for communication are learned more often by the following generation. The agents in this model are literal in the sense that they produce and understand language based solely on the meanings of words, and in particular do not perform inferences based on recursive mind-reading such as scalar implicatures. For literal agents, all the communicatively most efficient languages contain non-monotonic extensions. Therefore, in the second model over many generations languages with at least one non-monotonic meaning become widespread. This result is incompatible with the adjectival monotonicity observed in real languages.

In the third model, we implement more sophisticated agents using the Rational Speech Act (RSA) modelling paradigm [3]. Crucially, RSA agents are capable of calculating scalar implicatures with the help of a model of recursive mindreading. This reduces the loss in communicative accuracy caused by having monotonic languages. Therefore, pragmatic agents can accommodate their bias for simplicity while using language in a communicatively highly accurate way. Over many generations, agents in the third model evolve languages with non-degenerate monotonic meanings resembling real adjectives.

We conclude that the pervasiveness of monotonicity follows from a combination of human pragmatic skills and pressures coming from learnability and communication. The computational model we present thus offers an evolutionary explanation of adjectival monotonicity.

References

[1] Jon Barwise and Robin Cooper. "Generalized quantifiers and natural language". In: *Linguistics and philosophy* 4.2 (1981), pp. 159–219.

[2] Elizabeth Bogal-Allbritten. "Decomposing notions of adjectival transitivity in Navajo". In: *Natural Language Semantics* 21.3 (2013), pp. 277–314.

[3] Noah D. Goodman and Michael C. Frank. "Pragmatic Language Interpretation as Probabilistic Inference". In: *Trends in Cognitive Sciences* 20.11 (2016), pp. 818–829.

[4] Thomas Grano and Stuart Davis. "Universal markedness in gradable adjectives revisited: The morpho-semantics of the positive form in Arabic". In: *Natural Language & Linguistic Theory* 36.1 (2018), pp. 131–147.

[5] Christopher Kennedy and Louise McNally. "Scale structure, degree modification, and the semantics of gradable predicates". In: *Language* (2005), pp. 345–381.

[6] Simon Kirby, Tom Griffiths, and Kenny Smith. "Iterated learning and the evolution of language". In: *Current Opinion in Neurobiology* 28 (2014), pp. 108–114.

[7] Chen-Sheng Luther Liu. "The positive morpheme in Chinese and the adjectival structure". In: *Lingua* 120.4 (2010), pp. 1010–1056.

[8] Osamu Sawada and Thomas Grano. "Scale structure, coercion, and the interpretation of measure phrases in Japanese". In: *Natural Language Semantics* 19.2 (2011), pp. 191–226.

[9] Peter Svenonius and Christopher Kennedy. "Northern Norwegian degree questions and the syntax of measurement". In: *Phases of interpretation* 91 (2006), pp. 133–161.