ENS DEC DÉPARTEMENT D'ÉTUDES COGNITIVES

PSL UNIVERSITÉ PARIS

Sense as sampling propensity

Michael Goodale

Institut Jean-Nicod, Départment des études cognitives
École normale supérieure

Are Superman and Clark Kent the same?

Saul (1997) argues that certain sentences show a harder version of Frege's (1892) puzzle or Kripke's (1979) puzzle. One can be fully aware that Superman *is* Clark Kent, and yet (1-3) have different meanings depending on whether one says Clark Kent or Superman.

Saul's simple sentences

(1) a. Clark Kent went into the phone booth and Superman came out.
  b. Superman went into the phone booth and Clark Kent came out.
(2) a. Clark Kent is a mild-mannered journalist.
  b. ? Superman is a mild-mannered journalist.
(3) a. Superman wears a red-and-blue leotard and saves the day.
  b. ? Clark Kent wears a red-and-blue leotard and saves the day.

Sampling propensity

The mind is a generative machine and is able to sample exemplars of concepts. Sampling propensity (Icard 2016) suggests we can model this process with probabilities without committing to the view that the process itself is probabilistic or that probabilities are represented in the mind.

Four principles of sampling propensities

- Probabilities need not be represented explicitly in the mind.
- The formalism uses probabilities, but the semantics cannot directly access probabilities.
- Some things have high sampling propensity, while others have a low sampling propensity (e.g. a wooden chair compared to a chair made of burgers).
- Sampling propensities are not beliefs about frequencies (i.e. one can be aware that most archaeological digs find only old pottery, and yet think of gold and treasure when they think of archaeological digs).

This formalism doesn't describe the sampling procedure itself and is thus agnostic about the structure of concepts, making it compatible with many other theories.

Quasi-quantification with μ

μ takes an extension, $\mathbf{p}$, and a sampling propensity, p, and samples from the extension using the propensity, returning the proportion of samples for which the predicate is true.

$$\begin{aligned}
[\![\text{Drinks are consumed in bars}]\!] &= \mu_{x\sim\langle\mathbf{drink},\mathsf{drink}\rangle}\,[\![\text{consumed in bars}]\!]\,(x)\\
&= \frac{1}{N}\sum_{i=1}^{N}[\![\text{consumed in bars}]\!]\,(\textsc{Sample}(\langle\mathbf{drink},\mathsf{drink}\rangle))
\end{aligned}$$

$$\begin{aligned}
\mu_{x\sim\langle\mathbf{P},\mathsf{P}\rangle}\varphi(x) &:= \lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\varphi(\textsc{Sample}(\langle\mathbf{p},\mathsf{p}\rangle)) && \text{(Performance)}\\
&:= \sum_{x\in\mathbf{p}}\varphi(x)\cdot P(x\sim\mathsf{p}) && \text{(Competence, by the law of large numbers)}
\end{aligned}$$

This is a *fuzzy logic*, but one that doesn't suffer from Kamp's (1975) classical complaints (and has nothing to do with analyses of vagueness). μ is omni-present in language and arises earlier in ontogeny (c.f. Leslie, 2008).

Drinks and beverages revisited

Despite the fact that every drink is a beverage and vice versa, μ can derive different truth values for (4a) and (5a) because they have different sampling propensities. For example, beverage may bring to mind Coca-Cola, whereas drink might bring to mind beer.
In other words, $[\![\text{drink}]\!] = \langle\mathbf{drink},\mathsf{drink}\rangle$, $[\![\text{beverage}]\!] = \langle\mathbf{beverage},\mathsf{beverage}\rangle$, and $\mathbf{drink} = \mathbf{beverage}$ but $\mathsf{drink} \neq \mathsf{beverage}$.

Generic universal quantification

Matthewson (2001) suggests that non-partitive uses of "all" and "most" contain an embedded bare plural on the basis of Salish data. **"All" has generic flavour because it operates over categories directly**, while "every" instead goes over individual concepts within a category.

$$\begin{aligned}
[\![\text{Lions have manes}]\!] &= \mu_{x\sim\langle\mathbf{lion},\mathsf{lion}\rangle}\,[\![\text{has a mane}]\!]\,(x)\\
[\![\text{All lions h X x}\langle\text{P,P}\rangle\text{ (x)ave manes}]\!] &= \forall_{x\sim\langle\mathbf{lion},\mathsf{lion}\rangle}[\![\text{has a mane}]\!]\,(x))\\
[\![\text{Every lion has a mane}]\!] &= \forall_{\langle\mathbf{L},\mathsf{L}\rangle\in\textsc{Ind}(\langle\mathbf{Lion},\mathsf{Lion}\rangle)}\left(\mu_{x\sim\langle\mathbf{L},\mathsf{L}\rangle}\,[\![\text{has a mane}]\!]\,(x)\right)
\end{aligned}$$

People often accept "all" statements if the corresponding bare plural is true (Leslie, Khemlani, and Glucksberg 2011). I predict this effect should decrease considerably if we use "every" rather than "all".

Are drinks and beverages the same?

Some common nouns that *seem* co-extensional in every world do not licence the same generic sentences.

Where co-extensional nouns differ.

(4) a. Drinks are consumed in bars.
  b. Beverages are consumed in fast-food restaurants.
(5) a. ? Beverages are consumed in bars.
  b. ? Drinks are consumed in fast-food restaurants.
(6) a. French people love food.
  b. ? French people love comestibles.

This is problematic for model-theoretic semantics where the intension of a common noun is a function from worlds to extensions; there is no way to distinguish *drink* from *beverage* nor *food* from *comestible*, despite their different connotations.

Bipartite theory of concepts

A concept is a tuple, $\langle\mathbf{p},\mathsf{p}\rangle$ consisting of an extension, $\mathbf{p}$ and a sampling propensity, p.

$\mathbf{p}$

The extension, $\mathbf{p}$, is the set of e type objects corresponding to members of a concept or its extension (for now).

p

The sampling propensity, p is a function that samples from the extension, $\mathbf{p}$, which can be modeled probabilistically.

Superman and Clark Kent are the same as drinks and beverages

- Superman or Clark Kent are not just simple atoms of type e (e.g. $[\![\text{Superman}]\!] \neq s$). Rather, **individual concepts have the same representation as category concepts** like beverage.
- When we sample from an individual concept, we generate possible candidates that could be that individual under our current knowledge (our category extensions also will have potential members, e.g., unicorn's milk is a potential drink) .

$$[\![\text{Superman}]\!] = \langle\mathbf{S},\mathsf{S}\rangle \qquad [\![\text{Clark Kent}]\!] = \langle\mathbf{Ck},\mathsf{Ck}\rangle$$
$$\mathbf{S} = \mathbf{Ck} \qquad \mathsf{S} \neq \mathsf{Ck}$$

All Superman atoms are Clark Kent, and vice versa, but when one samples from S one gets a hero flying to save the day, while when one samples from Ck one gets a mild-mannered journalist.

$$\begin{aligned}
[\![\text{Superman saves the day}]\!] &= \mu_{x\sim\langle\mathbf{s},\mathsf{s}\rangle}\,[\![\text{saves the day}]\!]\,(x)\\
&= \sum_{x\in\mathbf{s}}[\![\text{saves the day}]\!]\,(x)\cdot P(x\sim\mathsf{s})\\
&= \sum_{x\in\mathbf{s}}[\![\text{saves the day}]\!]\,(x)\cdot P(x\sim\mathsf{s})
\end{aligned}$$

$$\begin{aligned}
[\![\text{Clark Kent saves the day}]\!] &= \mu_{x\sim\langle\mathbf{Ck},\mathsf{Ck}\rangle}\,[\![\text{saves the day}]\!]\,(x)\\
&= \sum_{x\in\mathbf{Ck}}[\![\text{saves the day}]\!]\,(x)\cdot P(x\sim\mathsf{Ck})\\
&= \sum_{x\in\mathbf{s}}[\![\text{saves the day}]\!]\,(x)\cdot P(x\sim\mathsf{Ck})
\end{aligned}$$

Traditional quantification

Since our extensions contain potential individuals, generalised quantifiers need a slight adjustment. Let $\langle\mathbf{p},\mathsf{p}\rangle$ be a category concept and $\mathcal{I}$ be the set of individual concepts.

$$\textsc{Ind}(\langle\mathbf{p},\mathsf{p}\rangle) := \{\langle\mathbf{q},\mathsf{q}\rangle \mid \langle\mathbf{q},\mathsf{q}\rangle\in\mathcal{I}\wedge\mathbf{q}\subseteq\mathbf{p}\}$$

With Ind, traditional quantification is over sets of individual concepts rather than sets of atoms.

$$[\![\text{Five lions have a mane}]\!] = [\![\text{Five}]\!](\textsc{Ind}(\langle\mathbf{lion},\mathsf{lion}\rangle))(\lambda x.\,[\![\text{has a mane}]\!]\,(x))$$

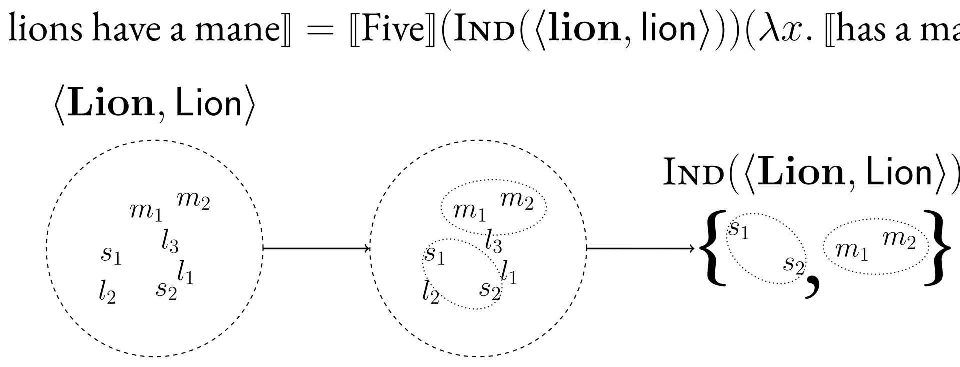

Ind extracts the sampling propensities for the two lion individual concepts, s and m, while ignoring l_1, l_2, and l_3 which are potential lion atoms that aren't a member of any individual concept.

Further applications

- Slurs (many theories of slurs argue they are co-extensional with neutral terms).
- Euphemisms (e.g. "urinate" and "piss" may draw to mind different scenarios).
- Circumlocutions (e.g. "kill" versus "cause to die", "people with disabilities" versus "the handicapped").

Handout

https://michaelgoodale.com/sense/

The research reported here was supported by Agence Nationale de la Recherche grant ANR-18-CE28-0008 (LANG-REASON, PI: Salvador Mascarenhas) and EURIPANR- 17-EURE-0012 (Investments d'Avenir Program) and by Ecole Doctorale Frontières de l'Innovation en Recherche et Education—Programme Bettencourt.

LaTeX TikZposter