Sense as sampling propensity

Michael Goodale
Institut Jean Nicod, Départment des études cognitives
École normale supérieure and Université Paris-Cité

November 17, 2023

1 The data

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

Drinks and beverages

(1) a. Drinks are consumed in bars.
b. Beverages are consumed in fast-food restaurants.

(2) a. ? Beverages are consumed in bars.
b. ? Drinks are consumed in fast-food restaurants.

(3) a. The French love food.
b. ? The French love comestibles.

It seems that someone can believe that every beverage is a drink and vice versa, yet (1) are much more felicitous than (2).

Odder still, this can also apply with names which refer to the same individual, even in transparent contexts (Saul 1997).

Superman and Clark Kent

(4) a. Clark Kent went into the phone-booth and Superman came out.
b. Superman went into the phone-booth and Clark Kent came out.

(5) a. Superman is successful with women
b. # Clark Kent is successful with women

There are many other cases where the extension of the referring terms might be identical, but with very different senses, connotations or stereotypes, such as:

- Slurs (i.e. many theories of slurs argue they are co-extensional)
- 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")

2 Sampling propensity

- Sampling propensity (Icard 2016) links probabilistic formalisms with the basic generative capacity of mind. The idea is that we sample or generate things according to some schema which can be modeled probabilistically.[1]
- Probabilities are not represented explicitly in the mind.
- Some things have high sampling propensity, while others have a low sampling propensity. When I think of a chair à propos of nothing, I am generating a chair according to some schema. Some chairs come to mind easily (a four-legged wooden chair), and others which do not (a 500 foot chair made entirely out of hamburgers).
- **Sampling propensity are not beliefs about frequencies.** A thought of an archaeological dig might draw to mind hoards like Sutton Hoo despite their extreme rarity.
- Sampling propensities are a kind of hyperintensionality; traditional intensional functions have no way of resolving these kinds of distinctions (as they only determine the extension in each world and so cannot distinguish things like "drink" and "beverage").

Lexical entries

Since any common noun has a sampling propensity, we need to enrich our lexical entries for at least categories. Any common noun, P is defined by the tuple: $\langle \mathbf{P}, \mathtt{P} \rangle$ where:

i) $\mathbf{P}$ is a standard $\langle e, t \rangle$ predicate.

ii) $\mathtt{P}$ is a generative procedure that samples from the extension defined by $\mathbf{P}$.

3 Quasi-quantification with μ

Definition. $\mu_{x \sim \langle \mathbf{P}, \mathtt{P} \rangle} \mathbf{Q}(x)$ is the expected value of $Q(x)$ where x is a sample from the probability distribution characterising the sampling propensity of P.

$$\begin{aligned} [\![\text{Lions have manes}]\!] &= \mu_{x \sim \langle \mathbf{Lion}, \mathtt{Lion} \rangle} hasMane(x) \\ &= \sum_{x \in \mathbf{Lion}} Pr_{\mathtt{Lion}}(x) \cdot hasMane(x) \end{aligned}$$

[1] This formalism is agnostic as to how the generative process works (and whether it is even random) but there are interesting possibilities using either causal models (Gopnik et al. 2004) or methods using manifolds (Goodale 2022)

"Lions have manes" is true to an extent proportional to our propensity to sample maned lions (where $Pr_{\texttt{Lion}}(x)$ is the probability of sampling x from `Lion`). This produces a **fuzzy logic**, albeit one that does not suffer from classical complaints about fuzzy logic (Kamp 1975) as standard predicates remain bivalent (provided that we managed scope correctly). Even if $\mu_{x\sim\langle\mathbf{P},\texttt{P}\rangle}\mathbf{Q}(x)$ has a truth value of 0.5, $\mu_{x\sim\langle\mathbf{P},\mathbf{P}\rangle}(\mathbf{Q}(x)\wedge\neg\mathbf{Q}(x))$ will necessarily have a truth value of 0.

Beverages and drinks revisited

Despite "drink" and "beverage" being co-extensional in every world, (2) shows that they are not freely substitutable. The difference between the two concepts is their sampling propensity: $\texttt{Drink} \neq \texttt{Beverage}$.

So, while `Beverage` will sample typical beverages (e.g. Coca-Cola), `Drink` will sample *typical drinks* (e.g. beer or wine) producing the difference between (1) and (2).

4 Names have the same structure as common nouns

- Individual concepts like Superman or Clark Kent are not just simple atoms (e.g. $[\![\text{Superman}]\!] \neq s$) (Matushansky 2006). Rather, **individuals have the same representation as category concepts** like beverage.
- Superman has the lexical entry of $\langle\mathbf{x},\mathtt{y}\rangle$ and Clark Kent has the entry of $\langle\mathbf{x},\mathtt{z}\rangle$. They have the same predicate but different sampling propensities.
- When we sample from an individual concept, we generate possible candidates that could be that individual under our current knowledge.

An individual who believes Clark Kent is Superman will believe that $\mathbf{Superman} = \mathbf{ClarkKent}$, but they could also have $\texttt{Superman} \neq \texttt{ClarkKent}$. `Superman` samples atoms that wear red and blue leotards and are successful with women, whereas `ClarkKent` samples atoms which are mild-mannered and who work at The Daily Planet. In either case, the atoms sampled belong to *both* **ClarkKent** and **Superman**, they just have different sampling propensities such that (4) and (5) behave differently when we substitute them.

$$[\![\text{Superman is successful with women}]\!] = \mu_{x\sim\langle\mathbf{Superman},\texttt{Superman}\rangle}\,[\![\text{successful with women}]\!]\,(x)$$

$$[\![\text{Clark Kent is successful with women}]\!] = \mu_{x\sim\langle\mathbf{ClarkKent},\texttt{ClarkKent}\rangle}\,[\![\text{successful with women}]\!]\,(x)$$

$$\left[\!\!\left[\begin{array}{l}\text{Superman went into the phonebooth}\\ \text{and Clark Kent came out}\end{array}\right]\!\!\right] = \begin{array}{l}\left(\mu_{x\sim\langle\mathbf{Superman},\texttt{Superman}\rangle}\,[\![\text{went into the phonebooth}]\!]\,(x)\right)\\ \wedge\left(\mu_{x\sim\langle\mathbf{ClarkKent},\texttt{ClarkKent}\rangle}\,[\![\text{came out}]\!]\,(x)\right)\end{array}$$

We can think of the different atoms of an individual as corresponding to either different situations, times or some epistemic uncertainty. There is a clear analogy between the possible individuals in an individual concept and Stalnaker-like contexts. As we learn information about an individual, we whittle down the set of atoms that can be generated for that individual, just as how learning information whittles down the space of our possible worlds for the Stalnaker context.

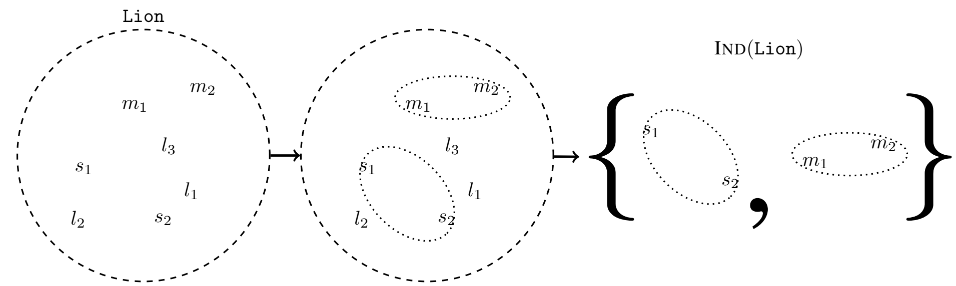

Figure 1: In a scenario where there are seven potential lions that can be sampled, but only two individual concepts that are lions, IND extracts the sampling propensities for those two lions while ignoring l_1, l_2, and l_3 which are potential lion atoms that aren't a member of any individual concept.

5 Quantification

Given individuals can be composed of multiple atoms, and our category concepts are composed of atoms, we need a way to pick out individual concepts from a category concept if we hope to quantify. We do this with a special operator: IND.

IND allows us to map from a category concept to the *set* of individual concepts whose extension are a subset of the category's extension.

Definition. Let $\langle\mathbf{P},\texttt{P}\rangle$ be a category concept and $\mathcal{I}$ be the set of individual concepts. IND is the function from a category concept to a set of individual concepts such that:

$$\textsc{Ind}(\langle\mathbf{P},\texttt{P}\rangle) = \{\langle\mathbf{Q},\texttt{Q}\rangle \mid \langle\mathbf{Q},\texttt{Q}\rangle \in \mathcal{I} \wedge \mathbf{Q} \subseteq \mathbf{P}\}$$

Using IND, traditional quantification becomes quantification over sets of individual concepts rather than sets of atoms.

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

$$\begin{aligned}[\![\text{Lions have manes}]\!] &= \mu_{x\sim\langle\mathbf{lion},\texttt{lion}\rangle} hasMane(x)\\ [\![\text{Every lion has a mane}]\!] &= [\![\text{Every}]\!](\textsc{Ind}(\langle\mathbf{lion},\texttt{lion}\rangle))(\lambda x.hasMane(x))\\ &= \begin{array}{l}\left(\lambda SP.\forall_{\langle\mathbf{L},\texttt{L}\rangle\in S}\left(\mu_{x\sim\langle\mathbf{L},\texttt{L}\rangle}P(x)\right)\right)\\ (\textsc{Ind}(\langle\mathbf{lion},\texttt{lion}\rangle))(\lambda x.hasMane(x))\end{array}\\ &= \forall_{\langle\mathbf{L},\texttt{L}\rangle\in\textsc{Ind}(\langle\mathbf{Lion},\texttt{Lion}\rangle)}\left(\mu_{x\sim\langle\mathbf{L},\texttt{L}\rangle} hasMane(x)\right)\end{aligned}$$

Where the type S is a set of individual concepts.

5.1 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. I differ here slightly and propose that the sampling propensity is accessed directly providing the generic flavour without using any genericity-specific quantifiers.

$$\begin{aligned}
[\![\text{All lions have manes}]\!] &= [\![\text{All}]\!]\,(\langle \mathbf{lion}, \texttt{lion}\rangle)(\lambda x.hasMane(x)) \\
&= (\lambda CP.\forall_{x\sim\langle \mathbf{C},\texttt{C}\rangle} P(x))(\langle \mathbf{lion}, \texttt{lion}\rangle)(\lambda x.hasMane(x)) \\
&= \forall_{x\sim\langle \mathbf{lion},\texttt{lion}\rangle} hasMane(x))
\end{aligned}$$

Where C is the type of a category concept.

5.2 Performance and μ

The original definition for μ and "all" captures competence but performance falls out straightforwardly. We simply estimate the expected value by sampling n times (where n is a contextually determined parameter that corresponds to something like effort) and take the average.

$$\begin{aligned}
[\![\text{Lions have manes}]\!] &= \mu_{x\sim\langle \mathbf{Lion},\texttt{Lion}\rangle} hasMane(x) \\
&= \frac{1}{n}\sum_{i=1}^{n}(\lambda x.hasMane(x))(\mathbf{sample}(\texttt{Lion}))
\end{aligned}$$

By the law of large numbers, as $n \to \infty$, the two definitions will almost surely converge to the same value.

$$\mu_{x\sim\langle \mathbf{P},\texttt{P}\rangle}\phi(x) = \sum_{x\in\mathbf{P}} Pr_{\texttt{P}} \cdot \phi(x) \qquad \text{(Competence)}$$

$$\mu_{x\sim\langle \mathbf{P},\texttt{P}\rangle}\phi(x) \approx \frac{1}{n}\sum_{i=1}^{n}(\lambda x.\phi(x))(\mathbf{sample}(\texttt{Lion})) \qquad \text{(Performance)}$$

Likewise, the competence of "all" requires the predicate to apply to every atom in a concept, whereas in performance, we will only sample a handful. "Every" works over individual concepts rather than bare atoms and so depends on what individual concepts you have, rather than the sampling propensity.

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

Some studies have shown people often accept "all" statements if the corresponding bare plural is true (Leslie, Khemlani, and Glucksberg 2011), and here I predict the effect should decrease considerably if we use "every".

References

[1] Michael Goodale. *Manifolds as conceptual representations in formal semantics*. Master's thesis, Ecole normale supérieure. June 2022. URL: https://ling.auf.net/lingbuzz/006797.

[2] Alison Gopnik et al. "A theory of causal learning in children: causal maps and Bayes nets". eng. In: *Psychological Review* 111.1 (Jan. 2004), pp. 3–32.

[3] Thomas Icard. "Subjective Probability as Sampling Propensity". en. In: *Review of Philosophy and Psychology* 7.4 (Dec. 2016), pp. 863–903.

[4] J. A. W. Kamp. "Two theories about adjectives". In: *Formal Semantics of Natural Language*. Ed. by Edward L. Keenan. Cambridge: Cambridge University Press, 1975, pp. 123–155.

[5] Sarah-Jane Leslie, Sangeet Khemlani, and Sam Glucksberg. "Do all ducks lay eggs? The generic overgeneralization effect". en. In: *Journal of Memory and Language* 65.1 (July 2011), pp. 15–31.

[6] Lisa Matthewson. "Quantification and the Nature of Crosslinguistic Variation". In: *Natural Language Semantics* 9.2 (2001). Publisher: Springer, pp. 145–189.

[7] Ora Matushansky. "Why Rose is the Rose: On the use of definite articles in proper names". en. In: *Empirical Issues in Formal Syntax and Semantics 6* (2006). Ed. by O Bonami and P Cabredo Hofherr, pp. 285–308.

[8] Jennifer M. Saul. "Substitution and Simple Sentences". en. In: *Analysis* 57.2 (1997), pp. 102–108.

A Towards concept composition

In order to handle to attributive uses of adjectives, we need a kind of concept composition. We use a special operator: $\ominus$ which restricts a category according to some open proposition. This creates a new sampling propensity which is renormalised over its initial sampling propensity.

Definition. $\ominus(\langle \mathbf{Q}, \texttt{Q}\rangle, \phi)$ is the concept with the sampling propensity defined by the probability distribution:

$$\frac{Pr_{\texttt{Q}}(x)}{\sum_{x\in\{x\mid x\in\mathbf{Q}\wedge\phi(x)=1\}} Pr_{\texttt{Q}}(x)}$$

And with the extension of:

$$\{x \mid x \in \mathbf{Q} \wedge \phi(x) = 1\}$$

Since the restricted region will differ from the rest of the extension, we will get a novel sampling propensity.

$$[\![\text{Male lions have manes}]\!] = \mu_{x\in\ominus(\langle \mathbf{Lions},\texttt{Lions}\rangle,\lambda x.male(x))} hasMane(x)$$

Acknowledgements

This work benefited from many discussions with Salvador Mascarenhas as well as some with Benjamin Spector and Janek Guerrini. This work was funded by Agence Nationale de la Recherche grant ANR-18-CE28-0008 (LANG-REASON, PI: Salvador Mascarenhas) and a PhD fellowship from the Learning Planet Institute and Université Paris-Cité.