Sense as sampling propensity

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

November 17, 2023

Co-extensionality, generics and substitution

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

(1) a. Drinks are consumed in bars.

b. Beverages are consumed in fast-food restaurants.

Co-extensionality, generics and substitution

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

(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.

Co-extensionality, generics and substitution

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

(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).

Simple sentences and substitution

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

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

(5) a. Superman is succesful with women

Simple sentences and substitution

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

(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 succesful with women
b. # Clark Kent is succesful with women

Cases where extensionality doesn't seem sufficient

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)

Cases where extensionality doesn't seem sufficient

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)

Cases where extensionality doesn't seem sufficient

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”)

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 **without comitting to explicit probabilistic representations.**[1].

[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)

Sampling propensity applies to any potential individual and needn't have explicit probabilistical representations

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

Sampling propensities are part of the lexicon

Since any common noun has a sampling propensity, we need to enrich our lexical entries for at least categories (and potentially much more). Think of it as a *hyperintension*.

Lexical entries

Any common noun, P is defined by the tuple: $\langle \mathbf{P}, \mathtt{P} \rangle$ where:

i) **P** is a standard $\langle e, t \rangle$ predicate.

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

Quasi-quantification with μ

"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`).

Quasi-quantification with μ

"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 $\texttt{Lion}$).

Definition

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

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

Truth values are continuous

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).

$$\phi = 0.5$$
$$\neg\phi = 0.5$$
$$\phi \wedge \neg\phi = 0.5 \text{ ✗}$$

$$\mu_{x\sim\langle \mathbf{P},\mathrm{P}\rangle}\mathbf{Q}(x) = 0.5$$
$$\mu_{x\sim\langle \mathbf{P},\mathrm{P}\rangle}\neg\mathbf{Q}(x) = 0.5$$

Truth values are continuous

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).

$$\phi = 0.5 \qquad \mu_{x\sim\langle \mathbf{P},\mathrm{P}\rangle}\mathbf{Q}(x) = 0.5$$

$$\neg\phi = 0.5 \qquad \mu_{x\sim\langle \mathbf{P},\mathrm{P}\rangle}\neg\mathbf{Q}(x) = 0.5$$

$$\phi \wedge \neg\phi = 0.5 \text{ ✗} \qquad \mu_{x\sim\langle \mathbf{P},\mathrm{P}\rangle}(\mathbf{Q}(x) \wedge \neg\mathbf{Q}(x)) = 0 \text{ ✓}$$

Drinks and beverages revisted

Despite “drink” and “beverage” being co-extensional in every world, (2) shows that they are not freely substitutable because they have different sampling propensities.

$$\textbf{drink} = \textbf{beverage} \qquad \texttt{drink} \neq \texttt{beverage}$$

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).

Individual concepts are like categories

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.

$[\![\text{Superman}]\!] = \langle \mathbf{x}, \mathrm{y} \rangle$

$[\![\text{Clark Kent}]\!] = \langle \mathbf{x}, \mathrm{z} \rangle$

The “extension” of a individual concept are putative individuals

- The different atoms of an individual correspond to either different situations, times or some epistemic uncertainty.
- These possible individuals are like Stalnaker-like contexts.
- By learning information about an individual, we whittle down that individual’s extension just like we whittle down possible worlds with contexts.

Superman is a drink, Clark Kent is a beverage

An individual who believes Clark Kent is Superman will believe that **Superman** $=$ **ClarkKent**, but they could also have $\texttt{Superman} \neq \texttt{ClarkKent}$.

$$\textbf{Sample}(\texttt{Superman}) = \{x_1, \ldots\} \qquad \textbf{Sample}(\texttt{ClarkKent}) = \{x_2, \ldots\}$$

$$\begin{array}{llll} \textbf{Superman}(x_1) & \textbf{ClarkKent}(x_1) & [\![\text{succesful with women}]\!](x_1) & \neg[\![\text{works at the Daily Planet}]\!](x_1) \\ \textbf{Superman}(x_2) & \textbf{ClarkKent}(x_2) & \neg[\![\text{succesful with women}]\!](x_2) & [\![\text{works at the Daily Planet}]\!](x_2) \end{array}$$

`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.

Simple sentences revisted

Our same analysis of drink versus beverage is able to explain Superman and Clark Kent's simple sentences.

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

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

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

Quantification

$$[\![\text{Lions have manes}]\!] = \mu_{x\sim\langle\textbf{lion},\texttt{lion}\rangle} hasMane(x)$$

$$[\![\text{All lions have manes}]\!] = \forall_{x\sim\langle\textbf{lion},\texttt{lion}\rangle} hasMane(x))$$

$$[\![\text{Every lion has a mane}]\!] = \forall_{\langle\textbf{L},\texttt{L}\rangle\in\text{IND}(\langle\textbf{lion},\texttt{lion}\rangle)} \left(\mu_{x\sim\langle\textbf{L},\texttt{L}\rangle} hasMane(x)\right)$$

How do we quantify?

Individuals are now partially sets, same as categories.

$$[\![\text{Simba}]\!] = \langle \{s_1, s_2, \ldots\}, \texttt{Simba} \rangle \qquad [\![\text{Mufasa}]\!] = \langle \{m_1, m_2, \ldots\}, \texttt{Mufasa} \rangle$$

$$[\![\text{lion}]\!] = \langle \{s_1, s_2, m_1, m_2, l_1, l_2, l_3 \ldots\}, \texttt{lion} \rangle$$

How can we quantify over this when both s_1 and s_2 are in **lions**?

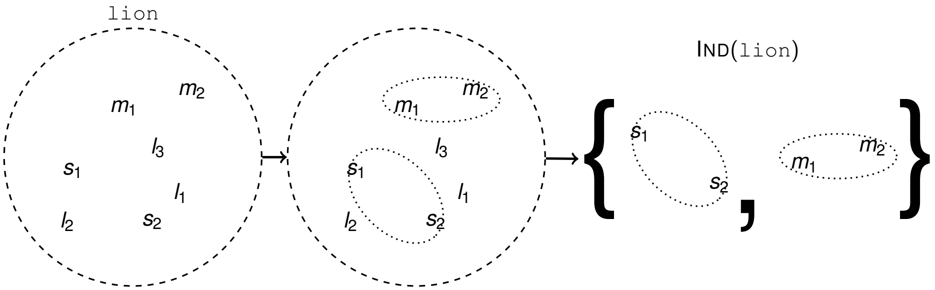

In a scenario where there are seven potential lions that can be sampled, but only two individual concepts that are lions, IND extracts the sampling propensities for those two lions while ignoring l_1, l_2, and l_3 which are potential lion atoms that aren't a member of any individual concept.

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:

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

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

Contrasting generics with universal quantification

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

Where the type S is a set of individual sampling propensities

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.

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

Where C is the type of a category concept.

Performance

To get performance, we can 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.

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

Performance

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

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

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

Performance and non-partitative all

"All" samples atoms from a predicate, "every" takes individuals from a predicate and then samples atoms from those individuals.

$$[\![\text{Lions have manes}]\!] = \mu_{x\sim\langle \textbf{lion},\texttt{lion}\rangle} hasMane(x)$$
$$[\![\text{All lions have manes}]\!] = \forall_{x\sim\langle \textbf{lion},\texttt{lion}\rangle} hasMane(x))$$
$$[\![\text{Every lion has a mane}]\!] = \forall_{\langle \textbf{L},\texttt{L}\rangle\in\text{IND}(\langle \textbf{lion},\texttt{lion}\rangle)} \left(\mu_{x\sim\langle \textbf{L},\texttt{L}\rangle} hasMane(x)\right)$$

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".

Thank you!

Towards concept composition

To compose a concept with another, we make a new concept which has a sampling propensity that is renormalised over its initial sampling propensity using a new operator, $\ominus$.

Definition

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

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

And with the extension of:

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

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