

Probabilistic compositional semantics, purely

Julian Grove and Jean-Philippe Bernardy

LENLS18, November 14, 2021

CLASP, University of Gothenburg

Motivation

Probabilistic semantics

In the last decade, lots of effort to connect formal semantics to mathematically explicit models of pragmatic reasoning...

Probabilistic semantics

In the last decade, lots of effort to connect formal semantics to mathematically explicit models of pragmatic reasoning...

- Rational Speech Act (RSA) models to formalize Gricean pragmatics (Goodman and Stuhlmüller, 2013; Lassiter and Goodman, 2013; Goodman and Frank, 2016; Lassiter and Goodman, 2017)

Probabilistic semantics

In the last decade, lots of effort to connect formal semantics to mathematically explicit models of pragmatic reasoning...

- Rational Speech Act (RSA) models to formalize Gricean pragmatics (Goodman and Stuhlmüller, 2013; Lassiter and Goodman, 2013; Goodman and Frank, 2016; Lassiter and Goodman, 2017)

Probabilistic semantics

In the last decade, lots of effort to connect formal semantics to mathematically explicit models of pragmatic reasoning...

- Rational Speech Act (RSA) models to formalize Gricean pragmatics (Goodman and Stuhlmüller, 2013; Lassiter and Goodman, 2013; Goodman and Frank, 2016; Lassiter and Goodman, 2017)

...generally, by dropping typed λ -calculus and encoding meanings in terms of probabilistic programming languages.

Probabilistic semantics

In the last decade, lots of effort to connect formal semantics to mathematically explicit models of pragmatic reasoning...

- Rational Speech Act (RSA) models to formalize Gricean pragmatics (Goodman and Stuhlmüller, 2013; Lassiter and Goodman, 2013; Goodman and Frank, 2016; Lassiter and Goodman, 2017)

...generally, by dropping typed λ -calculus and encoding meanings in terms of probabilistic programming languages.

- Church (Goodman et al., 2008)

Probabilistic semantics

In the last decade, lots of effort to connect formal semantics to mathematically explicit models of pragmatic reasoning...

- Rational Speech Act (RSA) models to formalize Gricean pragmatics (Goodman and Stuhlmüller, 2013; Lassiter and Goodman, 2013; Goodman and Frank, 2016; Lassiter and Goodman, 2017)

...generally, by dropping typed λ -calculus and encoding meanings in terms of probabilistic programming languages.

- Church (Goodman et al., 2008)

Probabilistic semantics

In the last decade, lots of effort to connect formal semantics to mathematically explicit models of pragmatic reasoning...

- Rational Speech Act (RSA) models to formalize Gricean pragmatics (Goodman and Stuhlmüller, 2013; Lassiter and Goodman, 2013; Goodman and Frank, 2016; Lassiter and Goodman, 2017)

...generally, by dropping typed λ -calculus and encoding meanings in terms of probabilistic programming languages.

- Church (Goodman et al., 2008)

Such programming languages are often *impure*: they allow for probabilistic effects, like sampling and marginalization, to occur at any point in a program.

Today's talk

Goal: show how a probabilistic semantics for natural language can be presented using only the simply typed λ -calculus (with products).

Today's talk

Goal: show how a probabilistic semantics for natural language can be presented using only the simply typed λ -calculus (with products).

- Achieve a seamless integration with approaches to meaning based on higher-order logic.

Today's talk

Goal: show how a probabilistic semantics for natural language can be presented using only the simply typed λ -calculus (with products).

- Achieve a seamless integration with approaches to meaning based on higher-order logic.

Today's talk

Goal: show how a probabilistic semantics for natural language can be presented using only the simply typed λ -calculus (with products).

- Achieve a seamless integration with approaches to meaning based on higher-order logic.

End up with a characterization of meanings as *probabilistic programs*, which are, nevertheless, pure (i.e., no real probabilistic effects).

Today's talk

Goal: show how a probabilistic semantics for natural language can be presented using only the simply typed λ -calculus (with products).

- Achieve a seamless integration with approaches to meaning based on higher-order logic.

End up with a characterization of meanings as *probabilistic programs*, which are, nevertheless, pure (i.e., no real probabilistic effects).

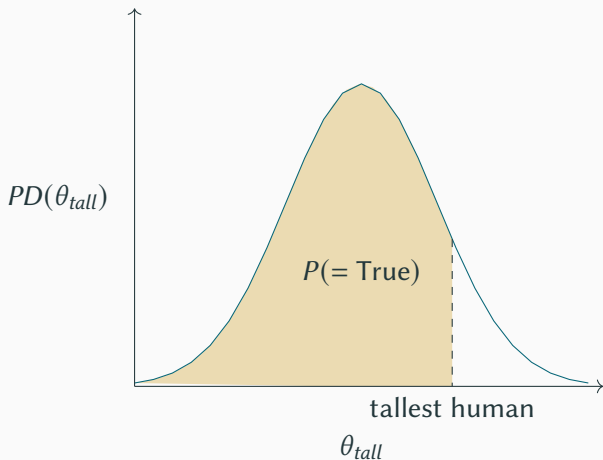
Such programs *describe* probability distributions over logical meanings.

Schematically...

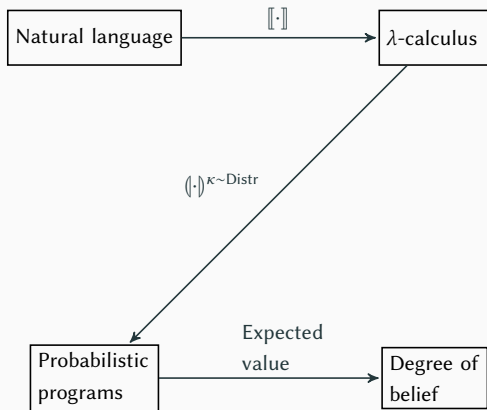
$$\llbracket \textit{someone is tall} \rrbracket = \exists x : \text{human}(x) \wedge \text{height}(x) \geq \theta_{\textit{tall}}$$

Schematically...

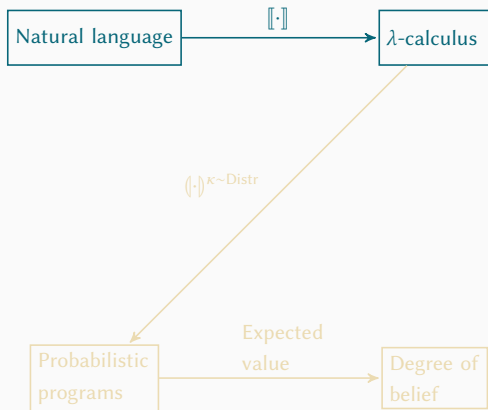
$$\llbracket \text{someone is tall} \rrbracket = \exists x : \text{human}(x) \wedge \text{height}(x) \geq \theta_{\text{tall}}$$



Our system



Up next



Formal semantics

Two strategies to formally interpret natural language, inherited from Montague:

Two strategies to formally interpret natural language, inherited from Montague:

- direct: right into set theory

Two strategies to formally interpret natural language, inherited from Montague:

- direct: right into set theory
 - denotations (entities, functions, etc.) are elements of sets

Two strategies to formally interpret natural language, inherited from Montague:

- direct: right into set theory
 - denotations (entities, functions, etc.) are elements of sets
- indirect: into a formal logic, e.g., the simply-typed λ -calculus/higher-order logic

Indirect interpretation

(1) Someone is tall.

Indirect interpretation

(1) Someone is tall.

$$\llbracket \textit{someone} \rrbracket = \lambda k. \exists x : \text{human}(x) \wedge k(x)$$

Indirect interpretation

(1) Someone is tall.

$$\llbracket \textit{someone} \rrbracket = \lambda k. \exists x : \text{human}(x) \wedge k(x)$$

$$\llbracket \textit{is} \rrbracket = \lambda x. x$$

Indirect interpretation

(1) Someone is tall.

$$\llbracket \text{someone} \rrbracket = \lambda k. \exists x : \text{human}(x) \wedge k(x)$$

$$\llbracket \text{is} \rrbracket = \lambda x. x$$

$$\llbracket \text{tall} \rrbracket = \lambda x. \text{height}(x) \geq \theta_{\text{tall}}$$

Indirect interpretation

(1) Someone is tall.

$$\llbracket \text{someone} \rrbracket = \lambda k. \exists x : \text{human}(x) \wedge k(x)$$

$$\llbracket \text{is} \rrbracket = \lambda x. x$$

$$\llbracket \text{tall} \rrbracket = \lambda x. \text{height}(x) \geq \theta_{\text{tall}}$$

Indirect interpretation

(1) Someone is tall.

$$\llbracket \textit{someone} \rrbracket = \lambda k. \exists x : \text{human}(x) \wedge k(x)$$

$$\llbracket \textit{is} \rrbracket = \lambda x. x$$

$$\llbracket \textit{tall} \rrbracket = \lambda x. \text{height}(x) \geq \theta_{\textit{tall}}$$

Functional application and β -reduction:

Indirect interpretation

(1) Someone is tall.

$$\llbracket \text{someone} \rrbracket = \lambda k. \exists x : \text{human}(x) \wedge k(x)$$

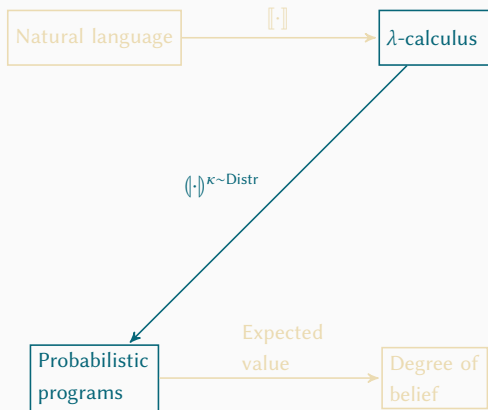
$$\llbracket \text{is} \rrbracket = \lambda x. x$$

$$\llbracket \text{tall} \rrbracket = \lambda x. \text{height}(x) \geq \theta_{\text{tall}}$$

Functional application and β -reduction:

- $\llbracket \text{someone} \rrbracket (\llbracket \text{is} \rrbracket (\llbracket \text{tall} \rrbracket)) \rightarrow_{\beta} \exists x : \text{human}(x) \wedge \text{height}(x) \geq \theta_{\text{tall}}$

Up next



The probabilistic interpretation

Let us assume that the non-logical constants of the logical language are finite in number and are ordered.

Let us assume that the non-logical constants of the logical language are finite in number and are ordered.

- (1) $\text{height} : e \rightarrow d_{tall}$ (2) $\text{human} : e \rightarrow t$
 (3) $(\geq) : r \rightarrow r \rightarrow t$ (4) $\theta_{tall} : d_{tall}$

Let us assume that the non-logical constants of the logical language are finite in number and are ordered.

- (1) $\text{height} : e \rightarrow d_{tall}$ (2) $\text{human} : e \rightarrow t$
 (3) $(\geq) : r \rightarrow r \rightarrow t$ (4) $\theta_{tall} : d_{tall}$
- A *context* (κ) is a tuple of type $\alpha_1 \times \dots \times \alpha_n$, where α_i is the type of the i^{th} constant.

Let us assume that the non-logical constants of the logical language are finite in number and are ordered.

- (1) $\text{height} : e \rightarrow d_{tall}$ (2) $\text{human} : e \rightarrow t$
 (3) $(\geq) : r \rightarrow r \rightarrow t$ (4) $\theta_{tall} : d_{tall}$
- A *context* (κ) is a tuple of type $\alpha_1 \times \dots \times \alpha_n$, where α_i is the type of the i^{th} constant.
- A context for this language would be of type $(e \rightarrow d_{tall}) \times (e \rightarrow t) \times (r \rightarrow r \rightarrow t) \times d_{tall}$.

A λ -homomorphism in a context

Given some context κ :

$$\langle c_i \rangle^\kappa = \kappa_i$$

(c_i is the i^{th} constant)

A λ -homomorphism in a context

Given some context κ :

$$\langle c_i \rangle^\kappa = \kappa_i$$

(c_i is the i^{th} constant)

$$\langle x \rangle^\kappa = x$$

(variables)

$$\langle \lambda x. M \rangle^\kappa = \lambda x. \langle M \rangle^\kappa$$

(abstractions)

$$\langle MN \rangle^\kappa = \langle M \rangle^\kappa \langle N \rangle^\kappa$$

(applications)

$$\langle \langle M, N \rangle \rangle^\kappa = \langle \langle M \rangle^\kappa, \langle N \rangle^\kappa \rangle$$

(pairing)

$$\langle M_i \rangle^\kappa = \langle M \rangle_i^\kappa$$

(projection)

Etc. (\diamond , logical constants)

Composing $(\cdot)^\kappa$ with $\llbracket \cdot \rrbracket$

If we compose the logical interpretation with the λ -homomorphism in some context κ :

Composing $(\cdot)^\kappa$ with $\llbracket \cdot \rrbracket$

If we compose the logical interpretation with the λ -homomorphism in some context κ :

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \kappa_2(x) \wedge \kappa_3(\kappa_1(x))(\kappa_4)$

Composing $(\cdot)^\kappa$ with $\llbracket \cdot \rrbracket$

If we compose the logical interpretation with the λ -homomorphism in some context κ :

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \kappa_2(x) \wedge \kappa_3(\kappa_1(x))(\kappa_4)$
- A truth value.

Composing $(\cdot)^\kappa$ with $\llbracket \cdot \rrbracket$

If we compose the logical interpretation with the λ -homomorphism in some context κ :

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \kappa_2(x) \wedge \kappa_3(\kappa_1(x))(\kappa_4)$
- A truth value.

Composing $(\cdot)^\kappa$ with $\llbracket \cdot \rrbracket$

If we compose the logical interpretation with the λ -homomorphism in some context κ :

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \kappa_2(x) \wedge \kappa_3(\kappa_1(x))(\kappa_4)$
- A truth value.

For example, say $\kappa = \langle \textit{height}, \textit{human}, (\geq), d \rangle$:

Composing $(\cdot)^\kappa$ with $\llbracket \cdot \rrbracket$

If we compose the logical interpretation with the λ -homomorphism in some context κ :

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \kappa_2(x) \wedge \kappa_3(\kappa_1(x))(\kappa_4)$
- A truth value.

For example, say $\kappa = \langle \textit{height}, \textit{human}, (\geq), d \rangle$:

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \textit{human}(x) \wedge \textit{height}(x) \geq d$

Composing $(\cdot)^\kappa$ with $\llbracket \cdot \rrbracket$

If we compose the logical interpretation with the λ -homomorphism in some context κ :

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \kappa_2(x) \wedge \kappa_3(\kappa_1(x))(\kappa_4)$
- A truth value.

For example, say $\kappa = \langle \textit{height}, \textit{human}, (\geq), d \rangle$:

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \textit{human}(x) \wedge \textit{height}(x) \geq d$

Composing $(\cdot)^\kappa$ with $\llbracket \cdot \rrbracket$

If we compose the logical interpretation with the λ -homomorphism in some context κ :

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \kappa_2(x) \wedge \kappa_3(\kappa_1(x))(\kappa_4)$
- A truth value.

For example, say $\kappa = \langle \textit{height}, \textit{human}, (\geq), d \rangle$:

- $(\llbracket \textit{someone is tall} \rrbracket)^\kappa = \exists x : \textit{human}(x) \wedge \textit{height}(x) \geq d$

Goal: **allow the context to be a random variable.**

Probabilistic programs

Definition

For any type α , a function of type $(\alpha \rightarrow r) \rightarrow r$ returns values of type α .

Definition

For any type α , a function of type $(\alpha \rightarrow r) \rightarrow r$ returns values of type α .

- Consumes a *projection function*: some f of type $\alpha \rightarrow r$.

Definition

For any type α , a function of type $(\alpha \rightarrow r) \rightarrow r$ returns values of type α .

- Consumes a *projection function*: some f of type $\alpha \rightarrow r$.
- Results in an r , by summing/integrating f over the possible values x of type α , weighting each $f(x)$ by the probability of x .

Definition

For any type α , a function of type $(\alpha \rightarrow r) \rightarrow r$ returns values of type α .

- Consumes a *projection function*: some f of type $\alpha \rightarrow r$.
- Results in an r , by summing/integrating f over the possible values x of type α , weighting each $f(x)$ by the probability of x .
- E.g., $\mathcal{N}(\mu, \sigma) : (d_{tall} \rightarrow r) \rightarrow r$

Definition

For any type α , a function of type $(\alpha \rightarrow r) \rightarrow r$ returns values of type α .

- Consumes a *projection function*: some f of type $\alpha \rightarrow r$.
- Results in an r , by summing/integrating f over the possible values x of type α , weighting each $f(x)$ by the probability of x .
- E.g., $\mathcal{N}(\mu, \sigma) : (d_{tall} \rightarrow r) \rightarrow r$
 - Represents a normal distribution with mean μ and standard deviation σ .

Definition

For any type α , a function of type $(\alpha \rightarrow r) \rightarrow r$ returns values of type α .

- Consumes a *projection function*: some f of type $\alpha \rightarrow r$.
- Results in an r , by summing/integrating f over the possible values x of type α , weighting each $f(x)$ by the probability of x .
- E.g., $\mathcal{N}(\mu, \sigma) : (d_{tall} \rightarrow r) \rightarrow r$
 - Represents a normal distribution with mean μ and standard deviation σ .
 - $\mathcal{N}(\mu, \sigma)(f) = \int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(\mu, \sigma)}(x) * f(x) dx$

Some nice things about probabilistic programs (pt. 1)

You can turn any value (of any type α) into a trivial probabilistic program that returns just that value:

Some nice things about probabilistic programs (pt. 1)

You can turn any value (of any type α) into a trivial probabilistic program that returns just that value:

$$\eta : \alpha \rightarrow (\alpha \rightarrow r) \rightarrow r \quad (\text{'return'})$$
$$\eta(a) = \lambda f.f(a)$$

Some nice things about probabilistic programs (pt. 1)

You can turn any value (of any type α) into a trivial probabilistic program that returns just that value:

$$\eta : \alpha \rightarrow (\alpha \rightarrow r) \rightarrow r \quad (\text{'return'})$$
$$\eta(a) = \lambda f.f(a)$$

E.g., $\eta(jp) = \lambda f.f(jp)$:

Some nice things about probabilistic programs (pt. 1)

You can turn any value (of any type α) into a trivial probabilistic program that returns just that value:

$$\eta : \alpha \rightarrow (\alpha \rightarrow r) \rightarrow r \quad (\text{'return'})$$
$$\eta(a) = \lambda f.f(a)$$

E.g., $\eta(jp) = \lambda f.f(jp)$:

- The probabilistic program that returns Jean-Philippe with a probability of 1.

Some nice things about probabilistic programs (pt. 2)

You can pass the value returned by a probabilistic program m to a function k from values to probabilistic programs, in order to make a bigger, sequenced probabilistic program.

Some nice things about probabilistic programs (pt. 2)

You can pass the value returned by a probabilistic program m to a function k from values to probabilistic programs, in order to make a bigger, sequenced probabilistic program.

$(\star) : ((\alpha \rightarrow r) \rightarrow r) \rightarrow$ ('bind')

$(\alpha \rightarrow (\beta \rightarrow r) \rightarrow r) \rightarrow$

$(\beta \rightarrow r) \rightarrow r$

$m \star k = \lambda f.m(\lambda x.k(x)(f))$

Some nice things about probabilistic programs (pt. 2)

You can pass the value returned by a probabilistic program m to a function k from values to probabilistic programs, in order to make a bigger, sequenced probabilistic program.

$(\star) : ((\alpha \rightarrow r) \rightarrow r) \rightarrow$ ('bind')

$(\alpha \rightarrow (\beta \rightarrow r) \rightarrow r) \rightarrow$

$(\beta \rightarrow r) \rightarrow r$

$m \star k = \lambda f.m(\lambda x.k(x)(f))$

“Run m , computing x . Then feed x to k .”

Building probabilistic programs

We may now build probabilistic programs that return *contexts*.

Building probabilistic programs

We may now build probabilistic programs that return *contexts*.

- If a context is of type $\alpha_1 \times \dots \times \alpha_n$, then we seek a probabilistic program K of type $(\alpha_1 \times \dots \times \alpha_n \rightarrow r) \rightarrow r$.

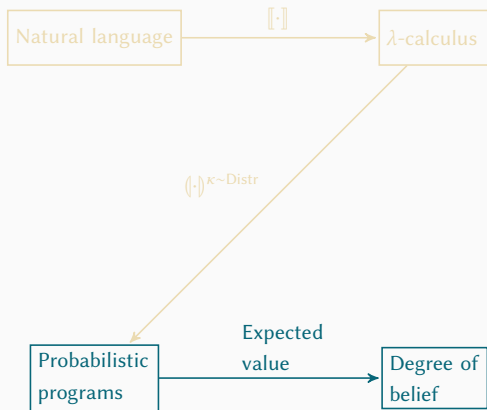
Building probabilistic programs

We may now build probabilistic programs that return *contexts*.

- If a context is of type $\alpha_1 \times \dots \times \alpha_n$, then we seek a probabilistic program K of type $(\alpha_1 \times \dots \times \alpha_n \rightarrow r) \rightarrow r$.
- Then, for a sentence ϕ in the logical language, we may do:

$$K \star \lambda\kappa.\eta((\phi)^{\kappa}) : (t \rightarrow r) \rightarrow r$$

Up next



Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

$$P : ((t \rightarrow r) \rightarrow r) \rightarrow r$$

Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

$$P : ((t \rightarrow r) \rightarrow r) \rightarrow r$$
$$P(p) = \frac{p(\mathbb{1})}{p(\lambda b.1)}$$

Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

$$P : ((t \rightarrow r) \rightarrow r) \rightarrow r$$
$$P(p) = \frac{p(\mathbb{1})}{p(\lambda b.1)}$$

Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

$$P : ((t \rightarrow r) \rightarrow r) \rightarrow r$$

$$P(p) = \frac{p(\mathbb{1})}{p(\lambda b.1)}$$

- $\mathbb{1} : t \rightarrow r$ is an indicator function:

Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

$$P : ((t \rightarrow r) \rightarrow r) \rightarrow r$$
$$P(p) = \frac{p(\mathbb{1})}{p(\lambda b.1)}$$

- $\mathbb{1} : t \rightarrow r$ is an indicator function:
 - $\mathbb{1}(\top) = 1$

Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

$$P : ((t \rightarrow r) \rightarrow r) \rightarrow r$$
$$P(p) = \frac{p(\mathbb{1})}{p(\lambda b.1)}$$

- $\mathbb{1} : t \rightarrow r$ is an indicator function:
 - $\mathbb{1}(\top) = 1$
 - $\mathbb{1}(\perp) = 0$

Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

$$P : ((t \rightarrow r) \rightarrow r) \rightarrow r$$
$$P(p) = \frac{p(\mathbb{1})}{p(\lambda b.1)}$$

- $\mathbb{1} : t \rightarrow r$ is an indicator function:
 - $\mathbb{1}(\top) = 1$
 - $\mathbb{1}(\perp) = 0$
- In the above, it picks out the mass assigned to \top .

Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

$$P : ((t \rightarrow r) \rightarrow r) \rightarrow r$$
$$P(p) = \frac{p(\mathbb{1})}{p(\lambda b.1)}$$

- $\mathbb{1} : t \rightarrow r$ is an indicator function:
 - $\mathbb{1}(\top) = 1$
 - $\mathbb{1}(\perp) = 0$
- In the above, it picks out the mass assigned to \top .
- $\lambda b.1$ picks out the total mass (assigned to \top and \perp).

Computing probabilities

Once we have a probabilistic program of type $(t \rightarrow r) \rightarrow r$, we may compute a probability from it:

$$P : ((t \rightarrow r) \rightarrow r) \rightarrow r$$
$$P(p) = \frac{p(\mathbb{1})}{p(\lambda b.1)}$$

- $\mathbb{1} : t \rightarrow r$ is an indicator function:
 - $\mathbb{1}(\top) = 1$
 - $\mathbb{1}(\perp) = 0$
- In the above, it picks out the mass assigned to \top .
- $\lambda b.1$ picks out the total mass (assigned to \top and \perp).
- So, $P(p)$ is the probability that p returns \top .

An example

(1) Someone is tall.

An example

(1) Someone is tall.

Say our constants are:

1. $\text{height} : e \rightarrow d_{\text{tall}}$
2. $\text{human} : e \rightarrow t$
3. $(\geq) : r \rightarrow r \rightarrow t$
4. $\theta_{\text{tall}} : d_{\text{tall}}$

An example

(1) Someone is tall.

Say our constants are:

1. $\text{height} : e \rightarrow d_{\text{tall}}$
2. $\text{human} : e \rightarrow t$
3. $(\geq) : r \rightarrow r \rightarrow t$
4. $\theta_{\text{tall}} : d_{\text{tall}}$

Define K as:

$$K = \mathcal{N}(72, 3) \star \lambda d. \eta(\text{height}, \text{human}, (\geq), d)$$

An example (cont'd)

$$K \star \lambda \kappa. \eta(\langle \exists x : \text{human}(x) \wedge \text{height}(x) \geq \theta_{\text{tall}} \rangle^\kappa)$$

An example (cont'd)

$$K \star \lambda \kappa. \eta((\exists x : \text{human}(x) \wedge \text{height}(x) \geq \theta_{\text{tall}}))^{\kappa})$$

⋮

$$= \lambda f. \mathcal{N}(72, 3)(\lambda d. f(\exists x : \text{human}(x) \wedge \text{height}(x) \geq d))$$

$$= \lambda f. \int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) * f(\exists x : \text{human}(x) \wedge \text{height}(x) \geq y) dy$$

An example (cont'd) (cont'd)

Computing a probability:

An example (cont'd) (cont'd)

Computing a probability:

$$\frac{\int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) * \mathbb{1}(\exists x : \text{human}(x) \wedge \text{height}(x) \geq y) dy}{\int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) dy}$$

An example (cont'd) (cont'd)

Computing a probability:

$$\frac{\int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) * \mathbb{1}(\exists x : \text{human}(x) \wedge \text{height}(x) \geq y) dy}{\int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) dy}$$
$$= \int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) * \mathbb{1}(\exists x : \text{human}(x) \wedge \text{height}(x) \geq y) dy$$

An example (cont'd) (cont'd)

Computing a probability:

$$\frac{\int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) * \mathbb{1}(\exists x : \text{human}(x) \wedge \text{height}(x) \geq y) dy}{\int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) dy}$$
$$= \int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) * \mathbb{1}(\exists x : \text{human}(x) \wedge \text{height}(x) \geq y) dy$$

An example (cont'd) (cont'd)

Computing a probability:

$$\frac{\int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) * \mathbb{1}(\exists x : \text{human}(x) \wedge \text{height}(x) \geq y) dy}{\int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) dy}$$
$$= \int_{-\infty}^{\infty} \text{PDF}_{\mathcal{N}(72,3)}(y) * \mathbb{1}(\exists x : \text{human}(x) \wedge \text{height}(x) \geq y) dy$$

...the mass of $\mathcal{N}(72, 3)$ less than or equal to the height of the tallest human

Bayesian inference (e.g., RSA)

RSA models: a popular application of probabilistic semantics.

RSA models: a popular application of probabilistic semantics.

The basic idea:

RSA models: a popular application of probabilistic semantics.

The basic idea:

- The RSA framework models a pragmatic listener, $L_1 \dots$

RSA models: a popular application of probabilistic semantics.

The basic idea:

- The RSA framework models a pragmatic listener, L ...
- ...who infers a distribution over meanings m from an utterance u , based on the probability that a pragmatic speaker, S , would make the utterance u to convey m .

RSA models: a popular application of probabilistic semantics.

The basic idea:

- The RSA framework models a pragmatic listener, L_1 ...
- ...who infers a distribution over meanings m from an utterance u , based on the probability that a pragmatic speaker, S_1 , would make the utterance u to convey m .
- Given a meaning m , the probability that S_1 would make the utterance u to convey m is related to the probability that a literal listener, L_0 , would infer m , given a literal interpretation of u .

Factoring by a weight / observing a premise

$$\mathit{factor} : r \rightarrow (\diamond \rightarrow r) \rightarrow r$$

$$\mathit{factor}(x)(f) = x * f(\diamond)$$

Factoring by a weight / observing a premise

$$\mathit{factor} : r \rightarrow (\diamond \rightarrow r) \rightarrow r$$

$$\mathit{factor}(x)(f) = x * f(\diamond)$$

$$\mathit{observe} : t \rightarrow (\diamond \rightarrow r) \rightarrow r$$

$$\mathit{observe}(\phi)(f) = \mathit{factor}(\mathbb{1}(\phi))(f)$$

$$= \mathbb{1}(\phi) * f(\diamond)$$

Say the type of the context is $\kappa = \alpha_1 \times \dots \times \alpha_n \dots$

Say the type of the context is $\kappa = \alpha_1 \times \dots \times \alpha_n \dots$

$$L_0 : u \rightarrow (\kappa \rightarrow r) \rightarrow r$$

Say the type of the context is $\kappa = \alpha_1 \times \dots \times \alpha_n \dots$

$$L_0 : u \rightarrow (\kappa \rightarrow r) \rightarrow r$$

$$L_0(u) = K \star \lambda \kappa. \text{observe}(\llbracket u \rrbracket^\kappa) \star \lambda \diamond. \eta(\kappa)$$

Say the type of the context is $\kappa = \alpha_1 \times \dots \times \alpha_n \dots$

$$L_0 : u \rightarrow (\kappa \rightarrow r) \rightarrow r$$

$$L_0(u) = K \star \lambda \kappa. \text{observe}(\llbracket u \rrbracket^\kappa) \star \lambda \diamond. \eta(\kappa)$$

$$S_1 : \kappa \rightarrow (u \rightarrow r) \rightarrow r$$

Say the type of the context is $\kappa = \alpha_1 \times \dots \times \alpha_n \dots$

$$L_0 : u \rightarrow (\kappa \rightarrow r) \rightarrow r$$

$$L_0(u) = K \star \lambda \kappa. \text{observe}(\llbracket u \rrbracket^\kappa) \star \lambda \diamond. \eta(\kappa)$$

$$S_1 : \kappa \rightarrow (u \rightarrow r) \rightarrow r$$

$$S_1(\kappa) = U \star \lambda u. \text{factor}(\text{PDF}_{L_0(u)}(\kappa)^\alpha) \star \lambda \diamond. \eta(u)$$

Say the type of the context is $\kappa = \alpha_1 \times \dots \times \alpha_n \dots$

$$L_0 : u \rightarrow (\kappa \rightarrow r) \rightarrow r$$

$$L_0(u) = K \star \lambda \kappa. \text{observe}(\llbracket u \rrbracket^\kappa) \star \lambda \diamond. \eta(\kappa)$$

$$S_1 : \kappa \rightarrow (u \rightarrow r) \rightarrow r$$

$$S_1(\kappa) = U \star \lambda u. \text{factor}(\text{PDF}_{L_0(u)}(\kappa)^\alpha) \star \lambda \diamond. \eta(u)$$

$$L_1 : u \rightarrow (\kappa \rightarrow r) \rightarrow r$$

Say the type of the context is $\kappa = \alpha_1 \times \dots \times \alpha_n \dots$

$$L_0 : u \rightarrow (\kappa \rightarrow r) \rightarrow r$$

$$L_0(u) = K \star \lambda \kappa. \text{observe}(\llbracket u \rrbracket^\kappa) \star \lambda \diamond. \eta(\kappa)$$

$$S_1 : \kappa \rightarrow (u \rightarrow r) \rightarrow r$$

$$S_1(\kappa) = U \star \lambda u. \text{factor}(\text{PDF}_{L_0(u)}(\kappa)^\alpha) \star \lambda \diamond. \eta(u)$$

$$L_1 : u \rightarrow (\kappa \rightarrow r) \rightarrow r$$

$$L_1(u) = K \star \lambda \kappa. \text{factor}(\text{PDF}_{S_1(\kappa)}(u)) \star \lambda \diamond. \eta(\kappa)$$

Conclusion

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus
- functional application

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus
- functional application

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus
- functional application

This semantics allows one to characterize:

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus
- functional application

This semantics allows one to characterize:

- probability distributions over possible denotations

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus
- functional application

This semantics allows one to characterize:

- probability distributions over possible denotations
- probabilities for formulae (given some distribution over contexts)

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus
- functional application

This semantics allows one to characterize:

- probability distributions over possible denotations
- probabilities for formulae (given some distribution over contexts)
- Bayesian update (or marginalization)

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus
- functional application

This semantics allows one to characterize:

- probability distributions over possible denotations
- probabilities for formulae (given some distribution over contexts)
- Bayesian update (or marginalization)
- RSA models (also, semantic learning)

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus
- functional application

This semantics allows one to characterize:

- probability distributions over possible denotations
- probabilities for formulae (given some distribution over contexts)
- Bayesian update (or marginalization)
- RSA models (also, semantic learning)

Summing up

It is possible to have probabilistic semantics for natural language that relies on the same machinery as used for logical interpretations:

- typed λ -calculus
- functional application

This semantics allows one to characterize:

- probability distributions over possible denotations
- probabilities for formulae (given some distribution over contexts)
- Bayesian update (or marginalization)
- RSA models (also, semantic learning)

...using the same logical language one uses to characterize linguistic meanings.

References

- Goodman, Noah D., and Michael C. Frank. 2016. Pragmatic Language Interpretation as Probabilistic Inference. *Trends in Cognitive Sciences* 20:818–829.
- Goodman, Noah D., Vikash K. Mansinghka, Daniel Roy, Keith Bonawitz, and Joshua B. Tenenbaum. 2008. Church: a language for generative models. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI'08*, 220–229. Arlington, Virginia, USA: AUAI Press.

- Goodman, Noah D., and Andreas Stuhlmüller. 2013. Knowledge and Implicature: Modeling Language Understanding as Social Cognition. *Topics in Cognitive Science* 5:173–184.
- Lassiter, Daniel, and Noah D. Goodman. 2013. Context, scale structure, and statistics in the interpretation of positive-form adjectives. *Semantics and Linguistic Theory* 23:587–610. Number: 0.
- Lassiter, Daniel, and Noah D. Goodman. 2017. Adjectival vagueness in a Bayesian model of interpretation. *Synthese* 194:3801–3836.