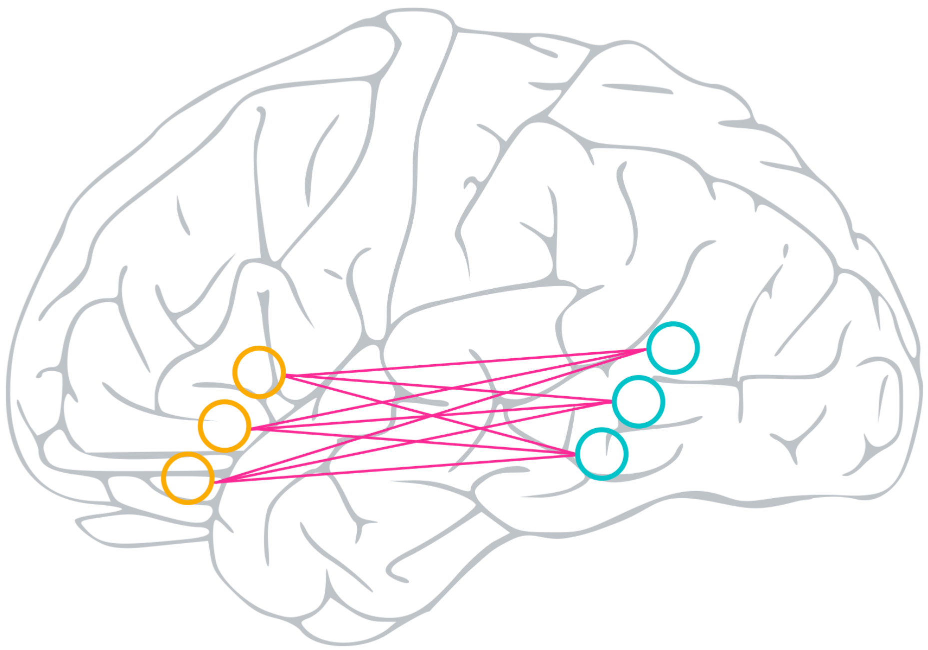

Neural Semantics

Bridging formal and probabilistic approaches to meaning

Semantic Theory — June 26th, 2018

Harm Brouwer

UNIVERSITÄT DES SAARLANDES

The Greatest Semanticist of them all …

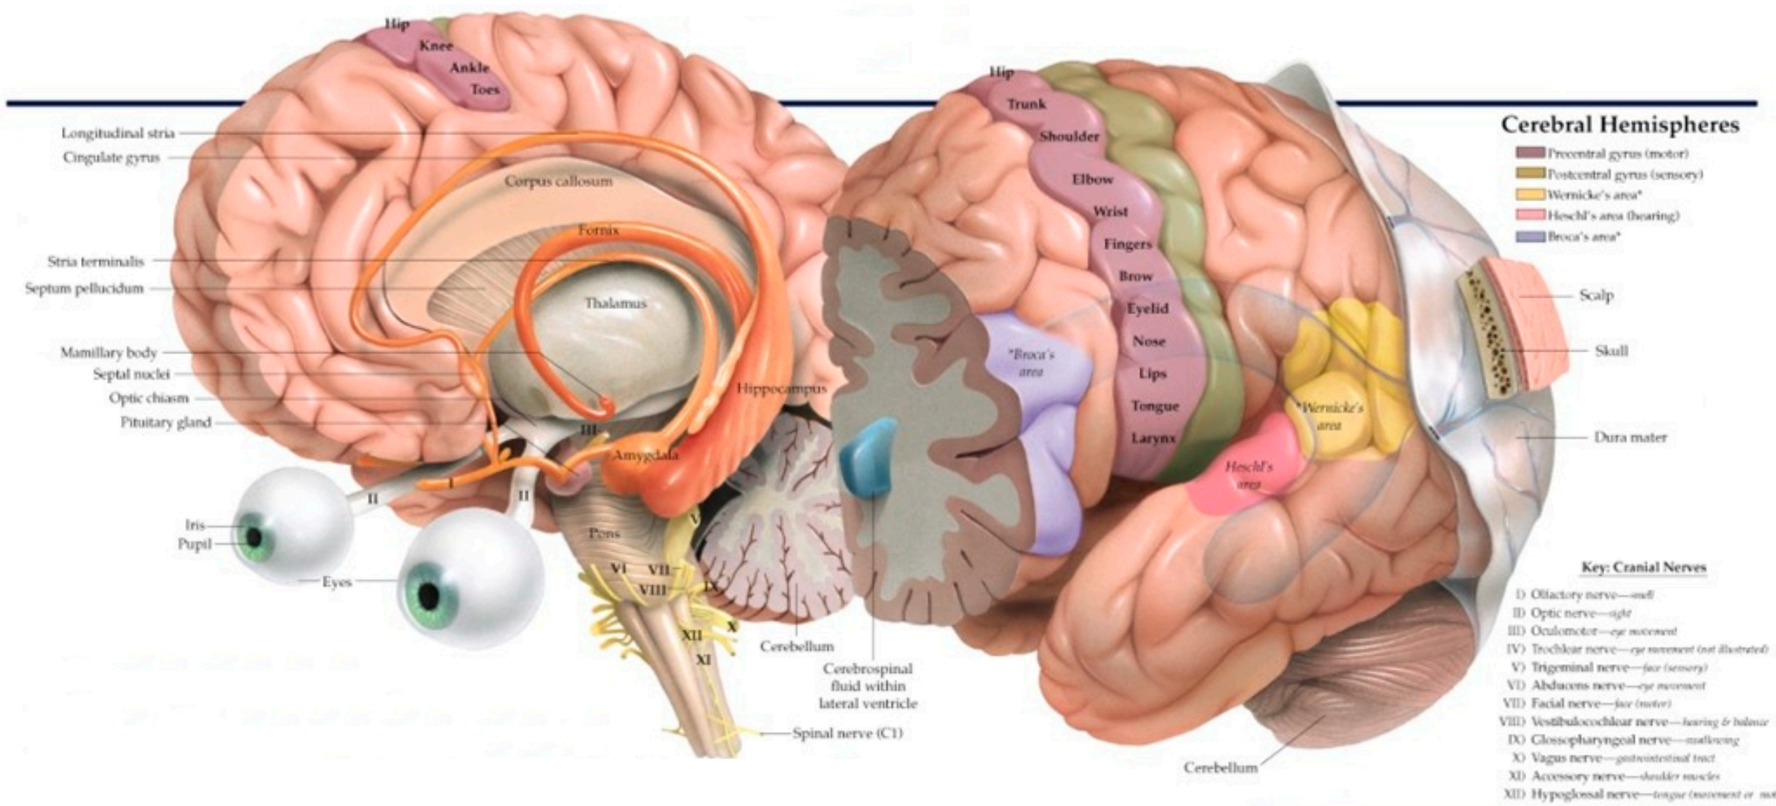

> Our **language comprehension system** is highly effective and accurate at attributing meaning to unfolding linguistic signal (~word-by-word)

>> This system's **representations** and **computational principles** are implemented in the **neural hardware** of the brain

>>> We should understand meaning construction and representation in terms of "brain-style computation" by identifying a **neural semantics**

Neural Semantics — Requirements

Neural Plausibility: assumed representations and computational principles should be implementable at the neural level [cf. Rumelhart, 1989]

Expressivity: representations should capture necessary dimensions of meaning, such as negation, quantification, and modality [cf. Frege, 1892]

Compositionality: the meaning of complex expressions should be derivable from the meaning of its parts [cf. Partee, 1984]

Gradedness: meaning representations are probabilistic, rather than discrete in nature [cf. Spivey, 2008]

Inferential: The derivation of utterance meaning entails (direct) inferences that go beyond literal propositional content [cf. Johnson-Laird, 1983]

Incrementality: As natural language unfolds over time, representations should allow for incremental construction [cf. Tanenhaus et al., 1995]

Brouwer, Crocker & Venhuizen (2017)
In: *From Semantics to Dialectometry*

Today's lecture

I. A Framework for Neural Semantics

II. A Neural Model of Language Comprehension

III. Neural Semantics — A fertile approach?

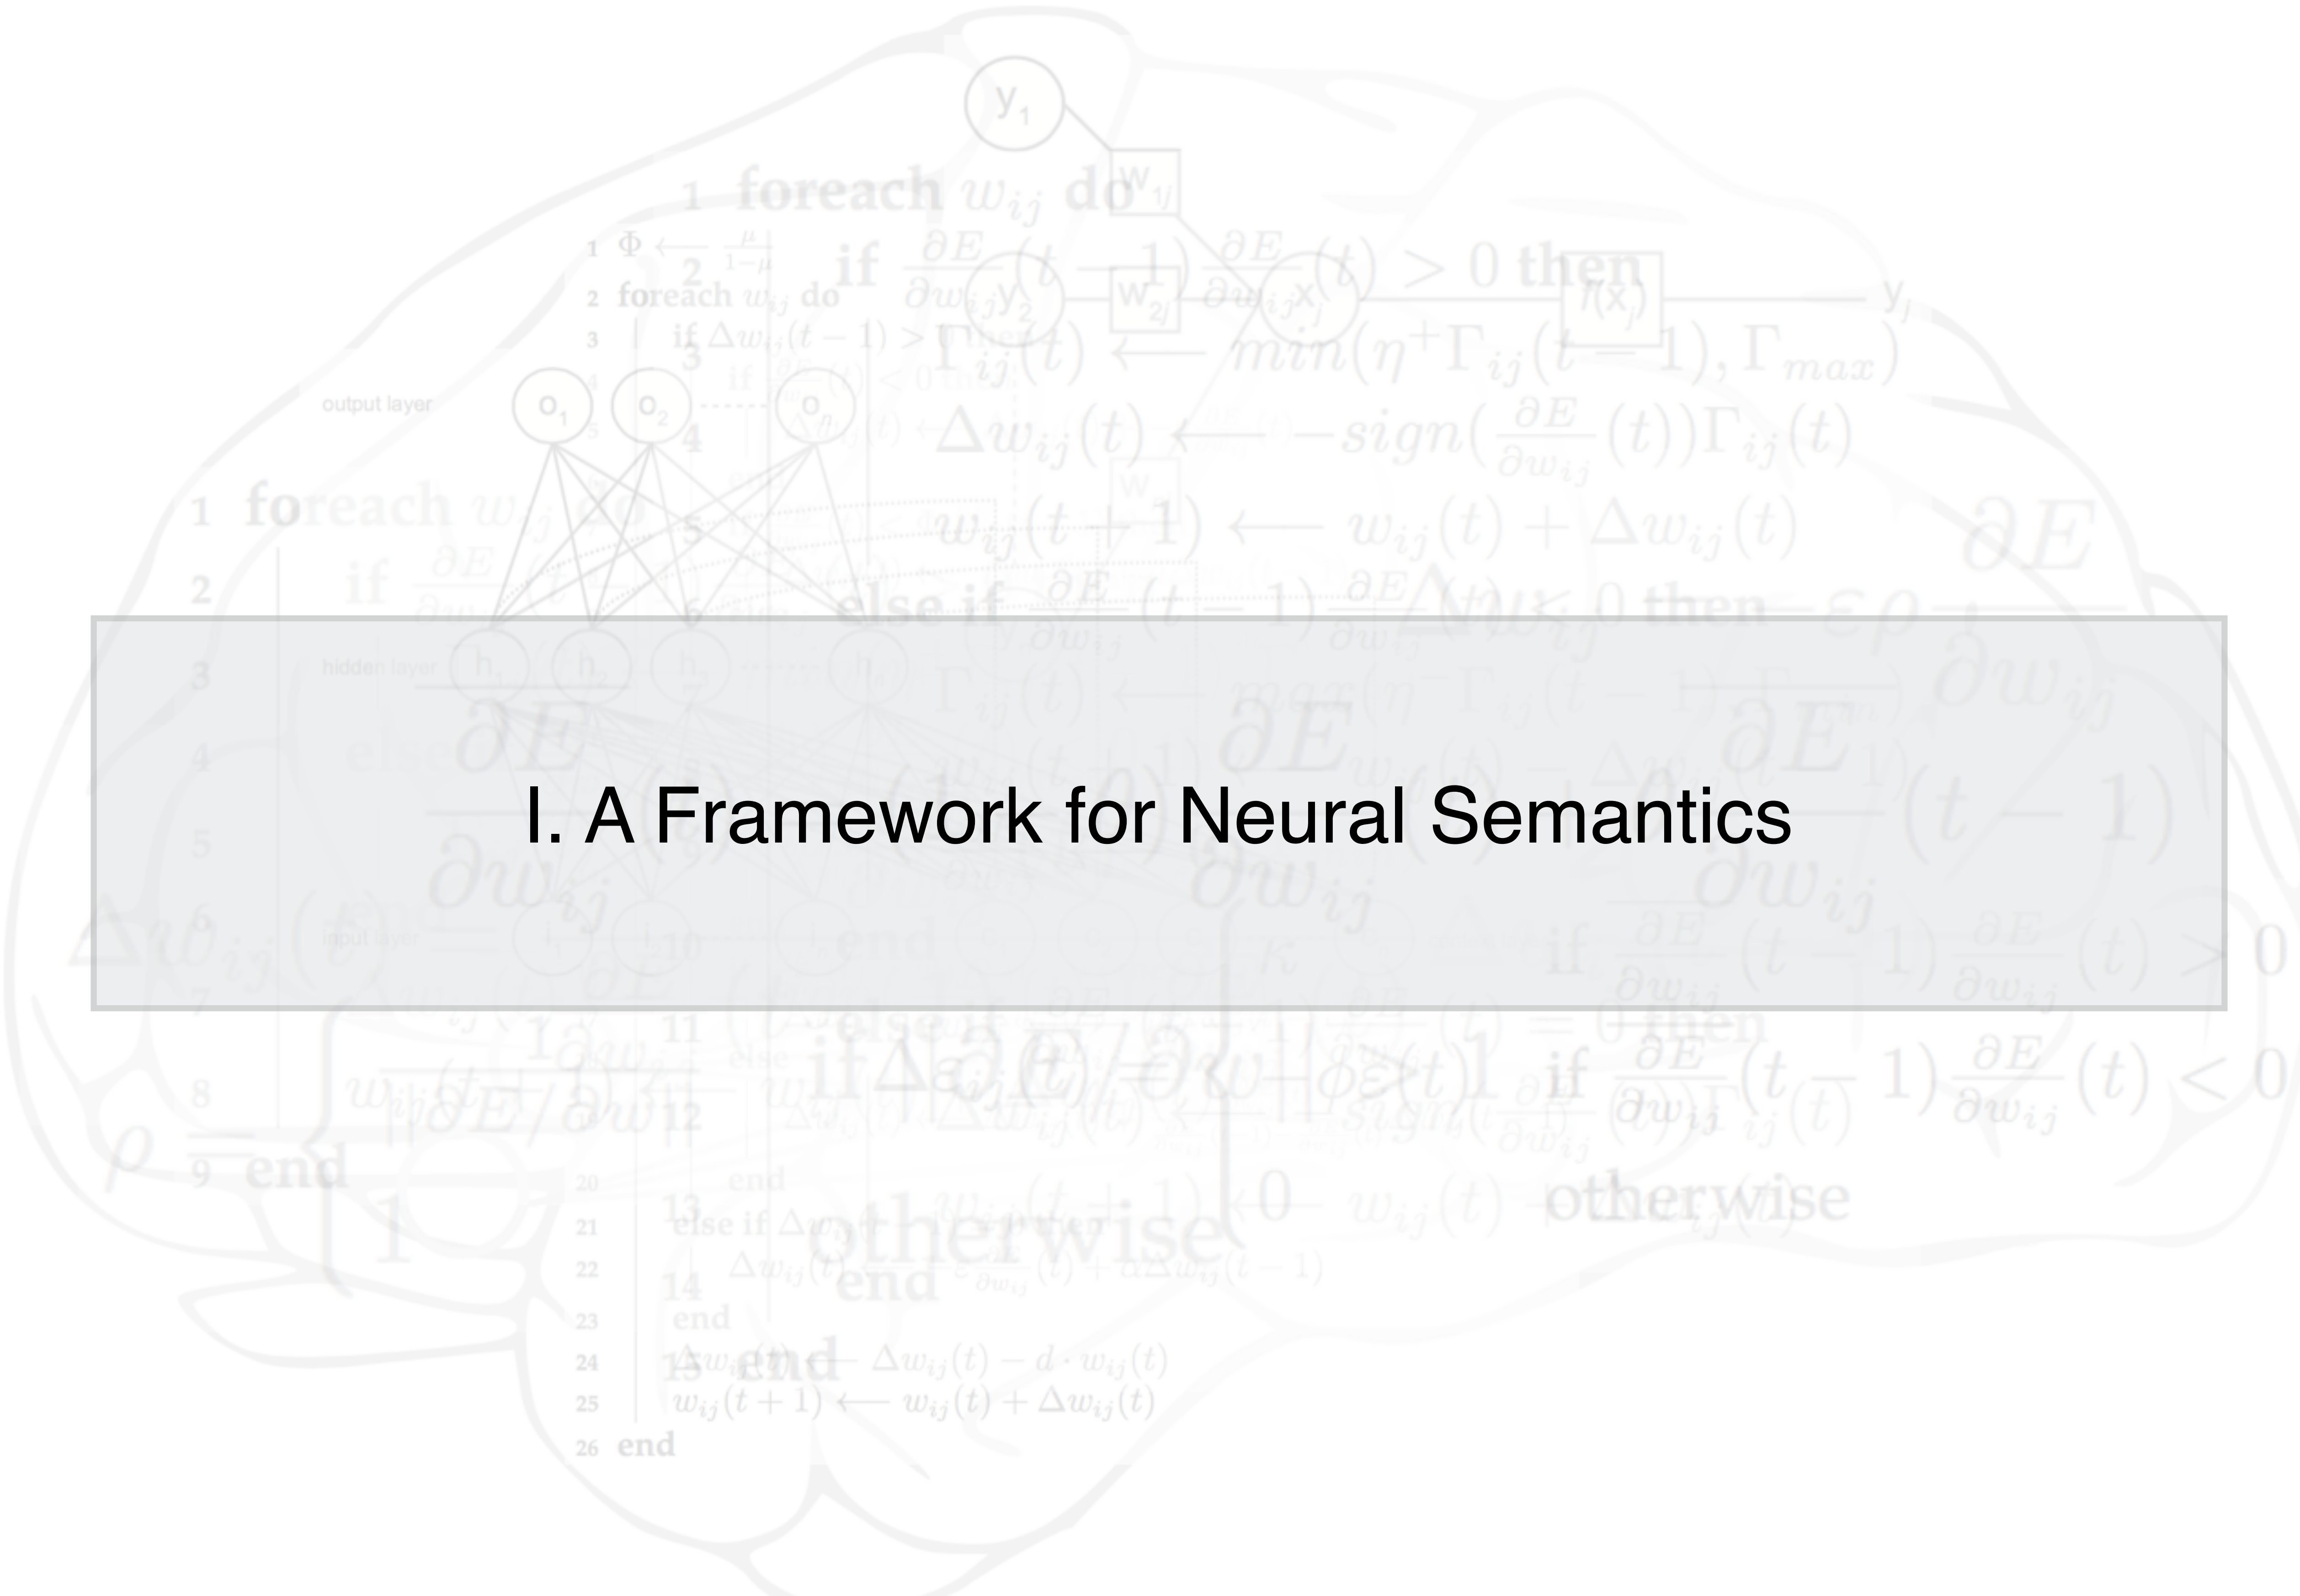

I. A Framework for Neural Semantics

Meaning in formal semantics

Meaning is defined as **truth-conditions** over logical models

$[\![\text{The Dutch team lost the race}]\!]^{M,g} = 1$

iff $[\![\text{lose\_race(dutch\_team)}]\!]^{M,g} = 1$

iff $V_M(\text{dutch\_team}) \in V_M(\text{lose\_race})$

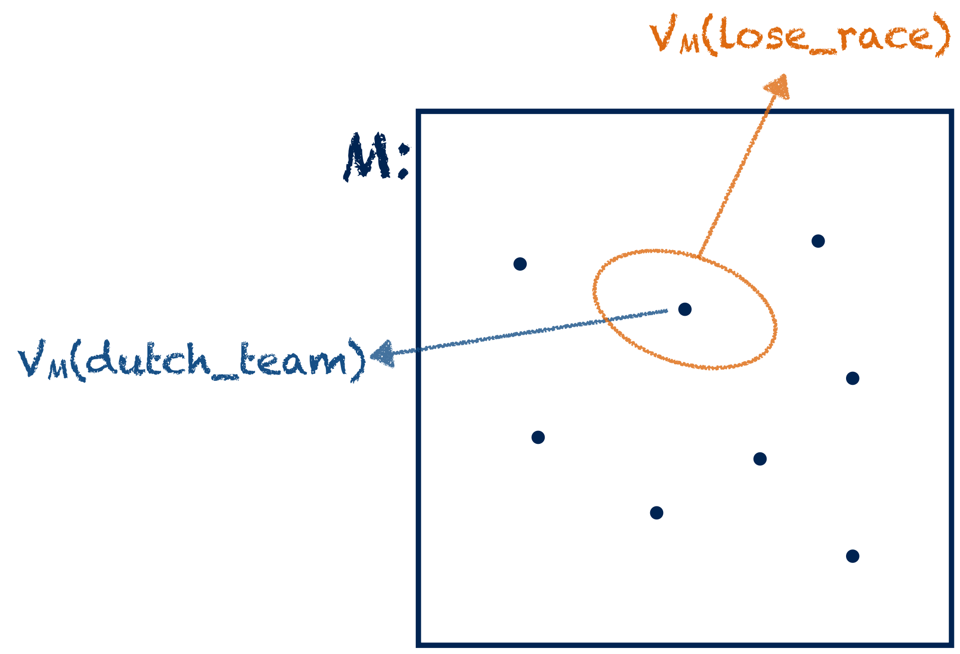

Intuition: Meaning is defined by the **situations** that satisfy an expression

Cognitively, truth-conditions derive from **experience** with the world

Idea: Model experience as **observations of states-of-affairs** in the world

Frege (1892)
Z. für Philosophie & phil. Kritik

Experiences as cues to meaning

Observations of SoAs can be captured by the set of logical models $\mathcal{M}$

M_1: enter(beth,restaurant), order(beth,dinner), eat(beth,dinner), enter(dave, cinema), order(dave,cola), ….

M_2: enter(beth,cinema), order(beth,popcorn), order(dave,cola), pay(dave), enter(thom, restaurant), leave(thom), ….

M_3: enter(beth,cinema), order(beth,popcorn), eat(beth,popcorn), order(dave,cola), pay(dave), order(thom, cola), ….

…

M_{150}: enter(beth,cinema), leave(beth), enter(dave,restaurant), pay(dave), enter(thom, cinema), order(thom, water), ….

> Each $M \in \mathcal{M}$ provides a **cue** toward the underlying truth-conditions

Brouwer, Crocker & Venhuizen (2017)
In: *From Semantics to Dialectometry*

Experiences as cues to meaning

Observations of SoAs can be captured by the set of logical models $\mathcal{M}$

M_1: enter(beth,restaurant), order(beth,dinner), eat(beth,dinner),
enter(dave, cinema), order(dave,cola), ….

M_2: enter(beth,cinema), order(beth,popcorn), order(dave,cola),
pay(dave), enter(thom, restaurant), leave(thom), ….

M_3: enter(beth,cinema), order(beth,popcorn), eat(beth,popcorn),
order(dave,cola), pay(dave), order(thom, cola), ….

…

M_{150}: enter(beth,cinema), leave(beth), enter(dave,restaurant),
pay(dave), enter(thom, cinema), order(thom, water), ….

> Each $M \in \mathcal{M}$ provides a **cue** toward the underlying truth-conditions

Brouwer, Crocker & Venhuizen (2017)
In: *From Semantics to Dialectometry*

Experiences as cues to meaning

Observations of SoAs can be captured by the set of logical models $\mathcal{M}$

M_1: enter(beth,restaurant), order(beth,dinner), eat(beth,dinner), enter(dave, cinema), order(dave,cola), ….

M_2: enter(beth,**cinema**), order(beth,popcorn), order(dave,cola), pay(dave), enter(thom, restaurant), leave(thom), ….

M_3: enter(beth,**cinema**), order(beth,popcorn), eat(beth,popcorn), order(dave,cola), pay(dave), order(thom, cola), ….

…

M_{150}: enter(beth,**cinema**), leave(beth), enter(dave,restaurant), pay(dave), enter(thom, cinema), order(thom, water), ….

> Each $M \in \mathcal{M}$ provides a **cue** toward the underlying truth-conditions

> Together, all $M \in \mathcal{M}$ reflect the world **truth-conditionally** and **probabilistically**

Brouwer, Crocker & Venhuizen (2017)
In: *From Semantics to Dialectometry*

Neural Semantics — Meaning

Formally, we can define the model space $\mathcal{S}_{\mathcal{M} \times \mathcal{P}}$ for a set of models $\mathcal{M}$ and propositions $\mathcal{P}$

Rows represent models that reflect observations of ‘states-of-affairs’ in the world

Columns represent meaning vectors of individual propositions

$$\mathbf{v}_i(p) = 1 \text{ iff } [\![p]\!]^{M_i} = 1$$

Meaning is defined in terms of **co-occurrence**

> Propositions with related meanings will be true in many of the same models

	p_1 = enter(beth,restaurant)	p_2 = ask_menu(beth)	p_3 = order(beth,cola)	....	p_n = leave(dave)
M_1	1	0	0		1
M_2	0	1	1		0
M_3	1	1	1	....	0
M_4	0	0	0		0
M_5	0	0	1		1
....					
M_m	0	1	1	....	1
		$\mathbf{v}(p_2)$			

Frank et al. (2003)
Cognitive Sci.

Brouwer, Crocker & Venhuizen (2017)
In: *From Semantics to Dialectometry*

Neural Semantics — Compositionality

Meaning vectors of **atomic situations** (propositions) are columns in $S_{\mathcal{M}\times\mathcal{P}}$

Meaning vectors of **complex situations** are derived compositionally:

$\mathbf{v}_i(\neg p) = 1$ *iff* $\mathbf{v}_i(p) = 0$ for $1 \leq i \leq m$

$\mathbf{v}_i(p \wedge q) = 1$ *iff* $\mathbf{v}_i(p) = 1$ and $\mathbf{v}_i(q) = 1$, for $1 \leq i \leq m$

> which gives functional completeness [e.g., $\mathbf{v}_i(p \vee q) = \mathbf{v}_i(\neg(\neg p \wedge \neg q))$]

Quantification is defined over the combined universe of $\mathcal{M}$: $U_{\mathcal{M}}=\{u_1,\ldots, u_k\}$

$\mathbf{v}_i(\forall x\phi) = 1$ iff $\mathbf{v}_i(\phi[x\backslash u_1] \wedge \ldots \wedge \phi[x\backslash u_k]) = 1$ for $1 \leq i \leq m$

$\mathbf{v}_i(\exists x\phi) = 1$ iff $\mathbf{v}_i(\phi[x\backslash u_1] \vee \ldots \vee \phi[x\backslash u_k]) = 1$ for $1 \leq i \leq m$

Frank et al. (2003)
Cognitive Sci.

Brouwer, Crocker & Venhuizen (2017)
In: *From Semantics to Dialectometry*

Neural Semantics — Probability

Meaning vectors inherently encode **(co-)occurrence probabilities** of situations—which may be atomic (~propositional) or complex (~conjunctive)

> The **prior probability** of situation **a**:

$$\mathbf{P(a)} = \sum_i (\mathbf{v}_i(a)) / m$$

> The **conjunctive probability** of situations **a** and **b**:

$$\mathbf{P(a \wedge b)} = \sum_i (\mathbf{v}_i(a)\mathbf{v}_i(b)) / m$$

> The **conditional probability** of situation **a** given **b**:

$$\mathbf{P(a \mid b)} = P(a \wedge b) / P(b)$$

	p_1 = enter(beth,restaurant)	p_2 = ask_menu(beth)	p_3 = order(beth,cola)	....	p_n = leave(dave)
M_1	1	0	0		1
M_2	0	1	1		0
M_3	1	1	1		0
M_4	0	0	0	....	0
M_5	0	0	1		1
....					
M_m	0	1	1	....	1

Frank et al. (2003)
Cognitive Sci.

Brouwer, Crocker & Venhuizen (2017)
In: *From Semantics to Dialectometry*

Neural Semantics — Inference

How much is situation **a** understood to be the case from situation **b**?

$$\text{comprehension(a,b)} = \begin{cases} [P(a \mid b) - P(a)] / [1 - P(a)] & \textit{if } P(a \mid b) > P(a) \\ [P(a \mid b) - P(a)] / P(a) & \textit{otherwise} \end{cases}$$

P(a | b) > P(a) knowing **b** *increases* belief in **a**
➔ **a** is inferred to be the case from **b**

comprehension(a,b) = 1: knowing **b** took away all 'uncertainty' in **a** ➔ **b** ⊨ **a**

P(a | b) ≤ P(a) knowing **b** *decreases* belief in **a**
➔ **a** is inferred not to be the case from **b**

comprehension(a,b) = -1: knowing **b** took away all 'certainty' in **a** ➔ **b** ⊨ ¬**a**

Frank et al. (2003)
Cognitive Sci.

Brouwer, Crocker & Venhuizen (2017)
In: *From Semantics to Dialectometry*

World knowledge inferencing in $S_{\mathcal{M}\times\mathcal{P}}$

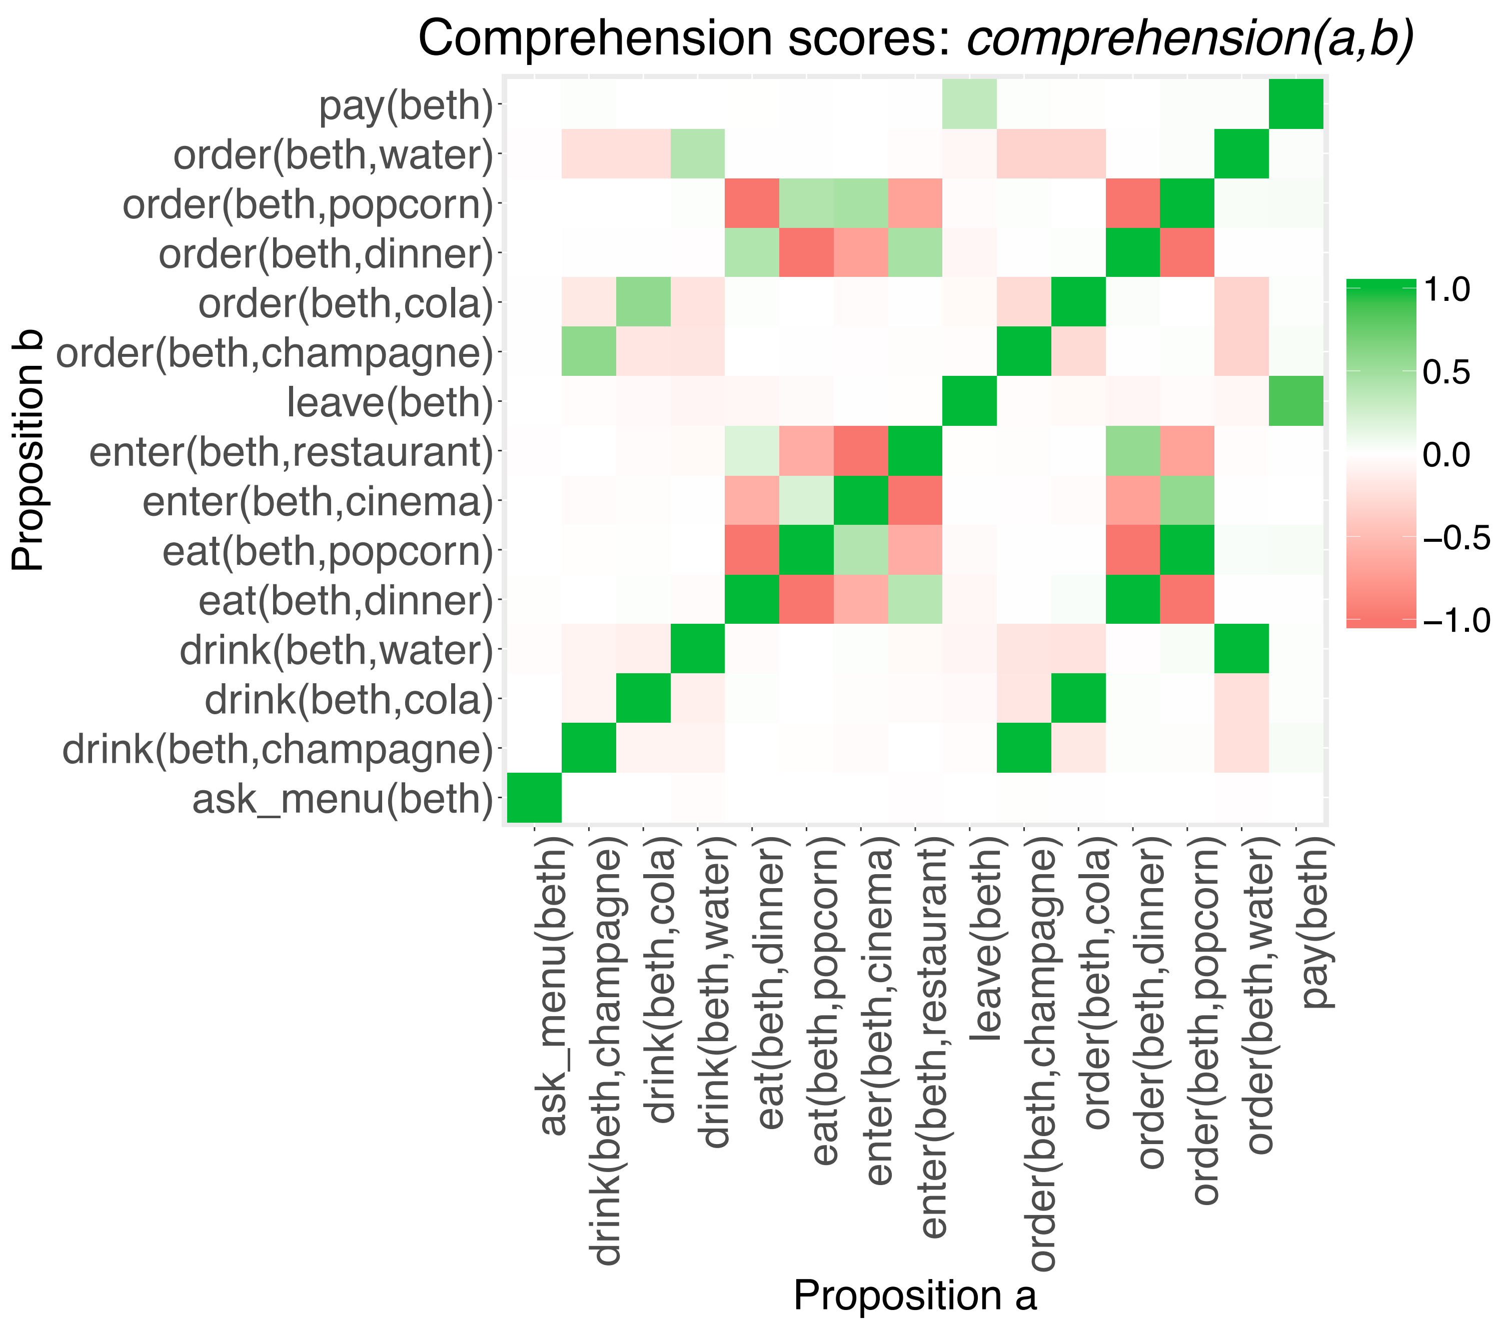

Next: Let's employ $S_{\mathcal{M}\times\mathcal{P}}$ in an **incremental comprehension model**

Brouwer, Crocker & Venhuizen (2017)
In: *From Semantics to Dialectometry*

II. A Neural Model of Language Comprehension

A Neural Model of Language Comprehension

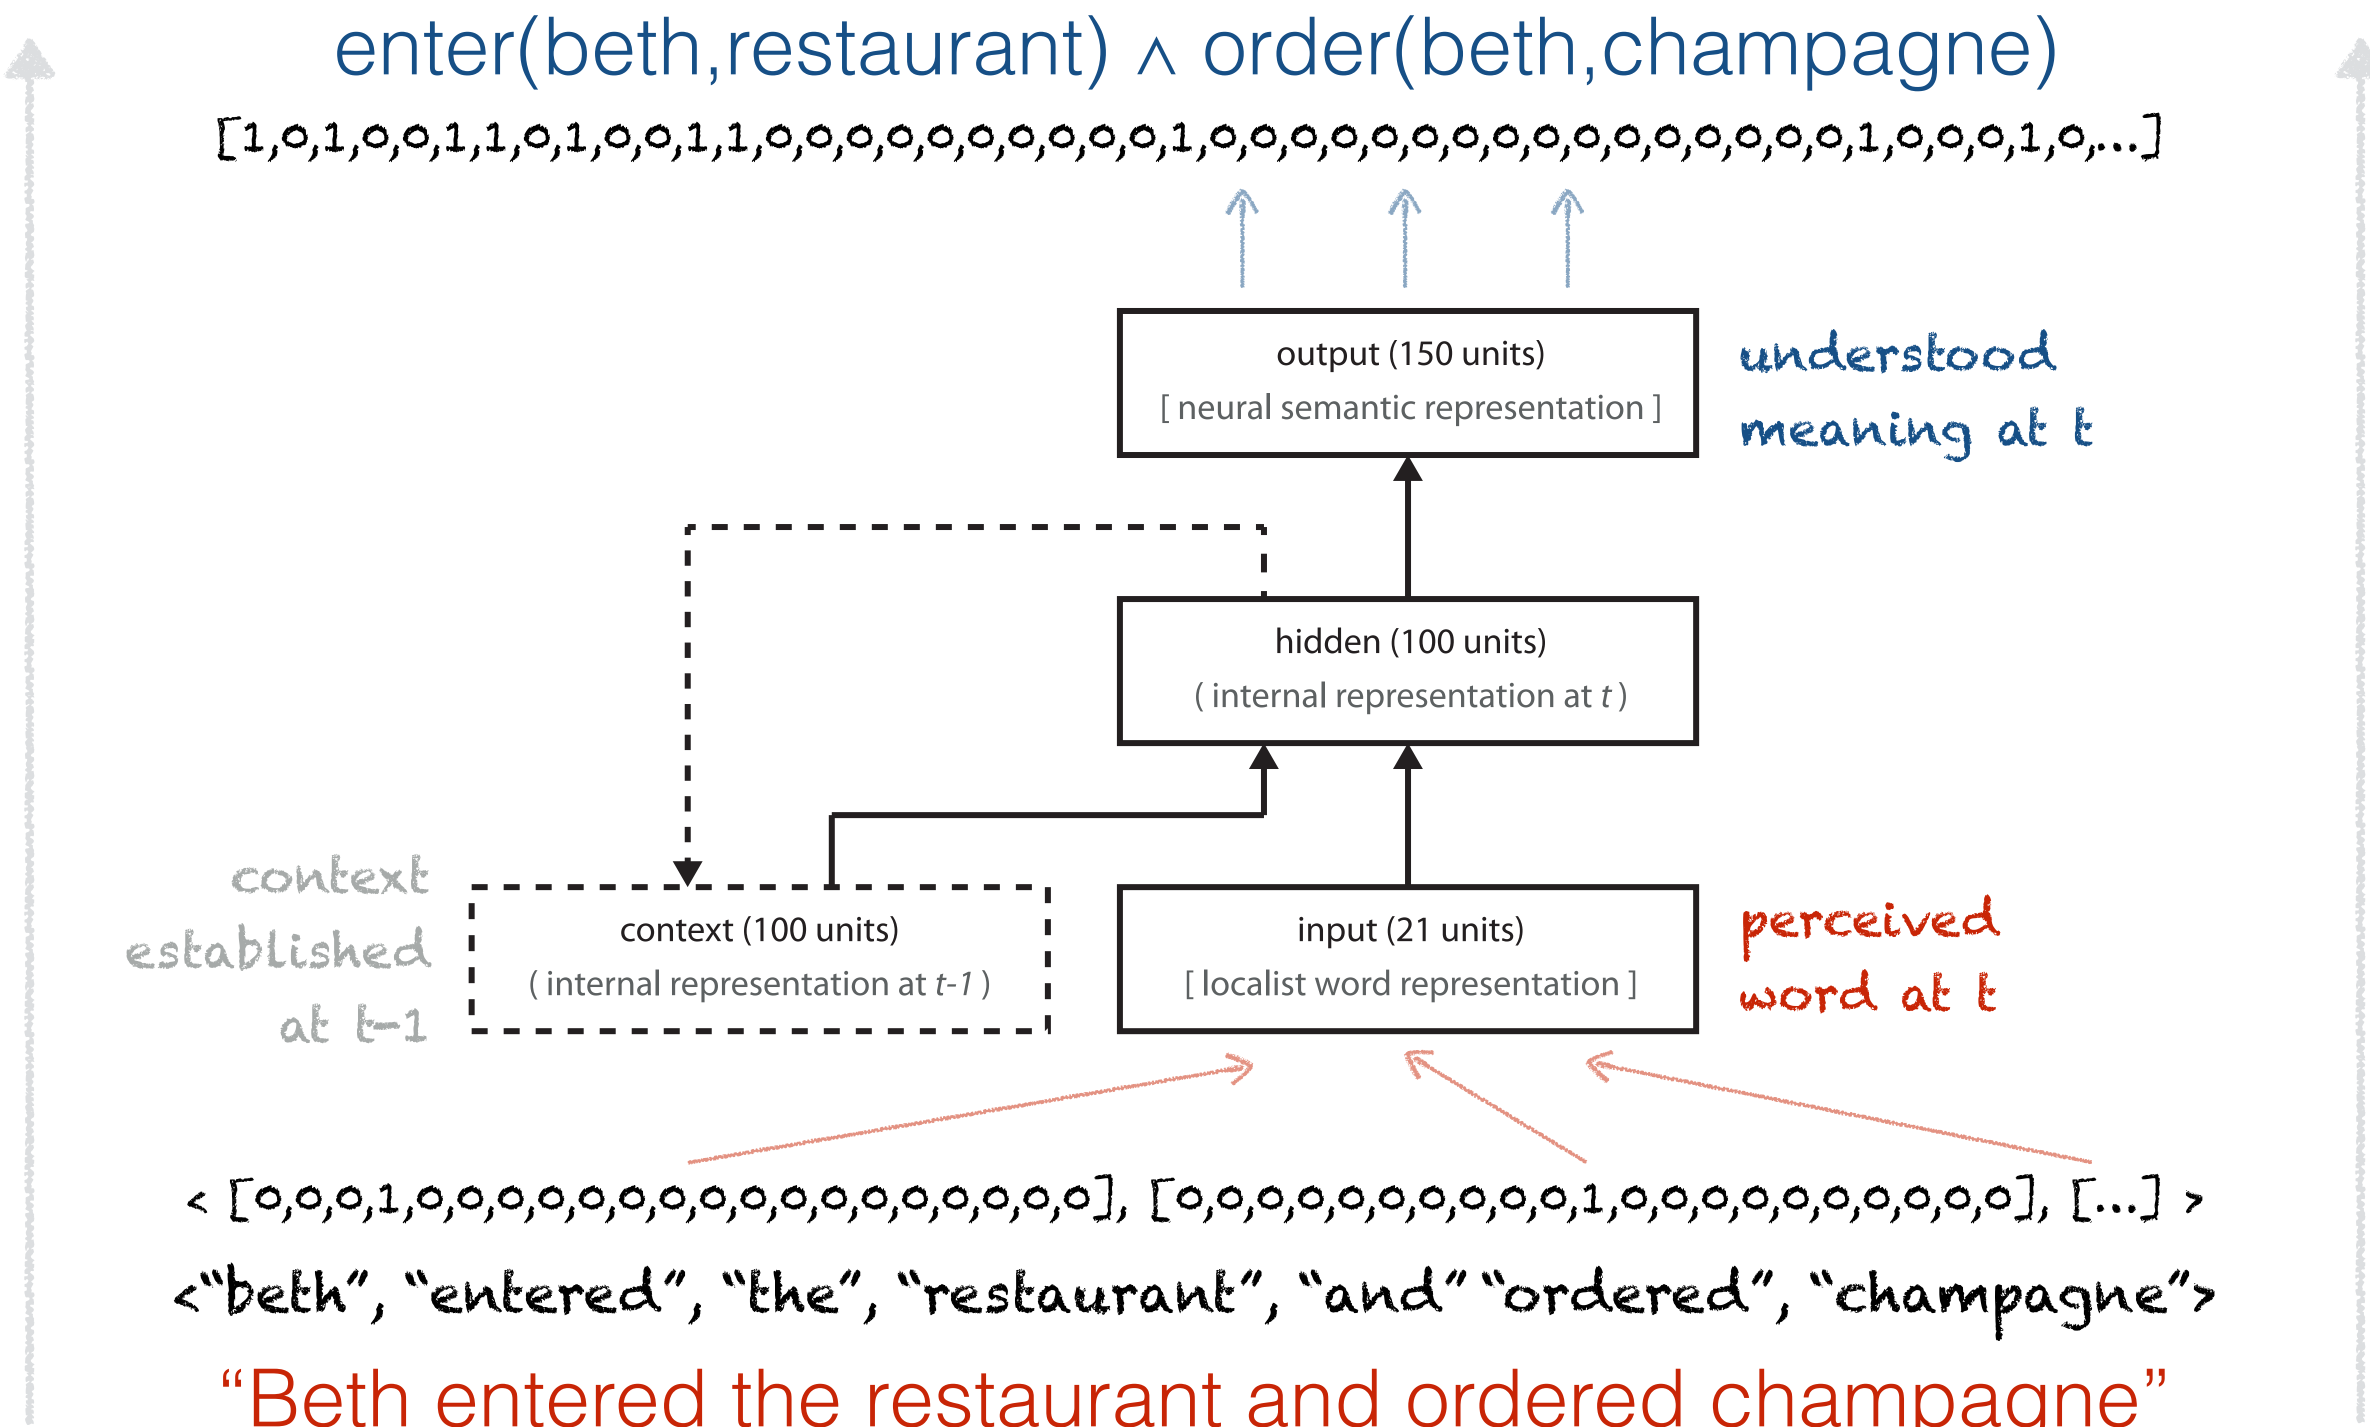

Elman (1990)
Cognitive Sci.

Venhuizen, Crocker & Brouwer (2018)
Discourse Process.

Comprehension is meaning-space navigation

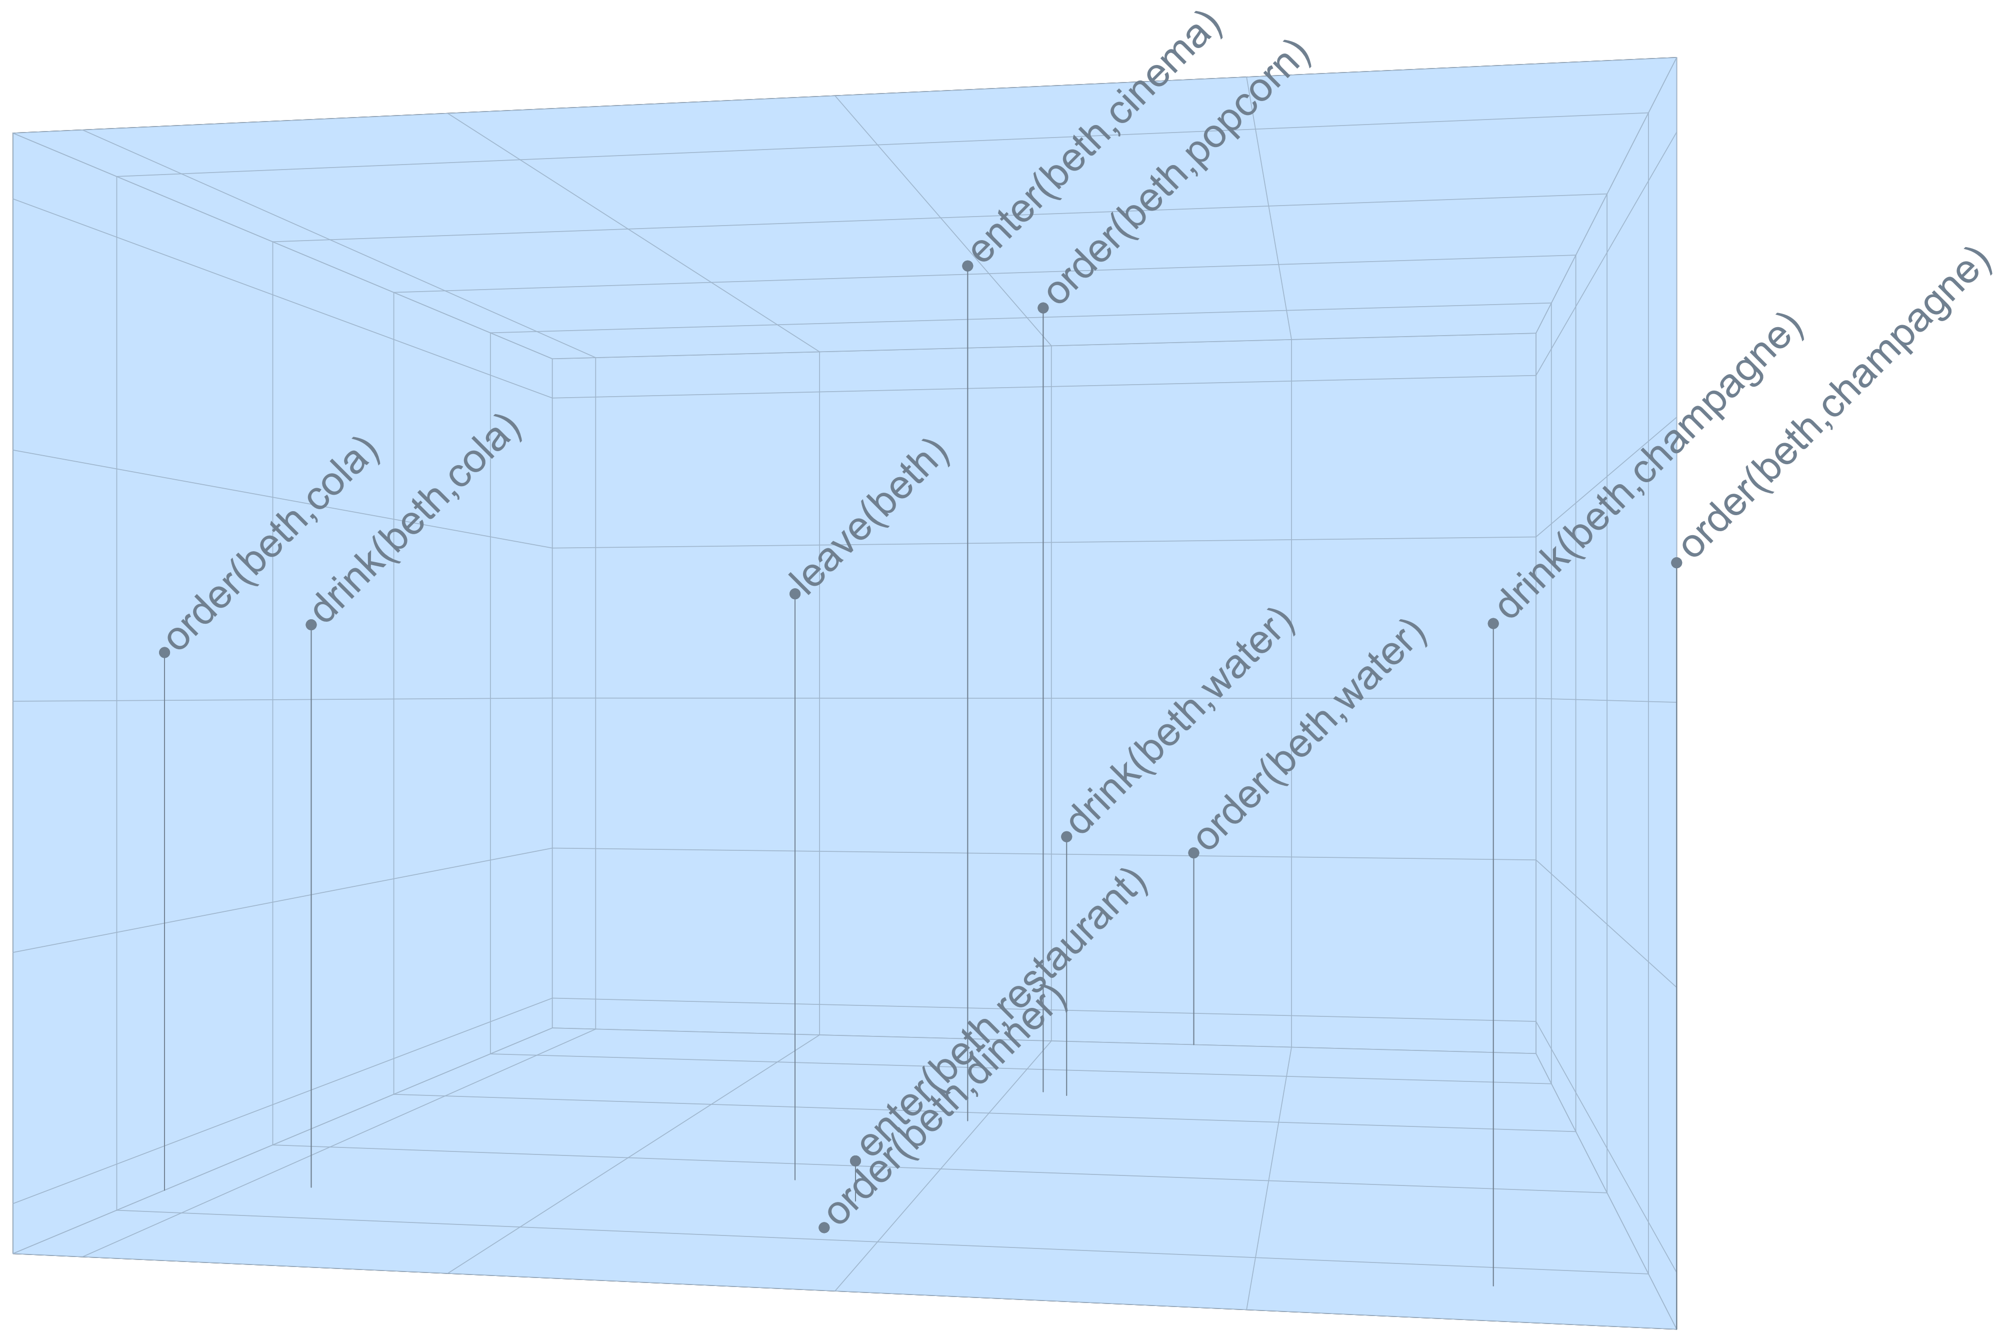

Multi-dimensional scaling: 150D ↦ 3D

Venhuizen, Crocker & Brouwer (2018)
Discourse Process.

Comprehension is meaning-space navigation

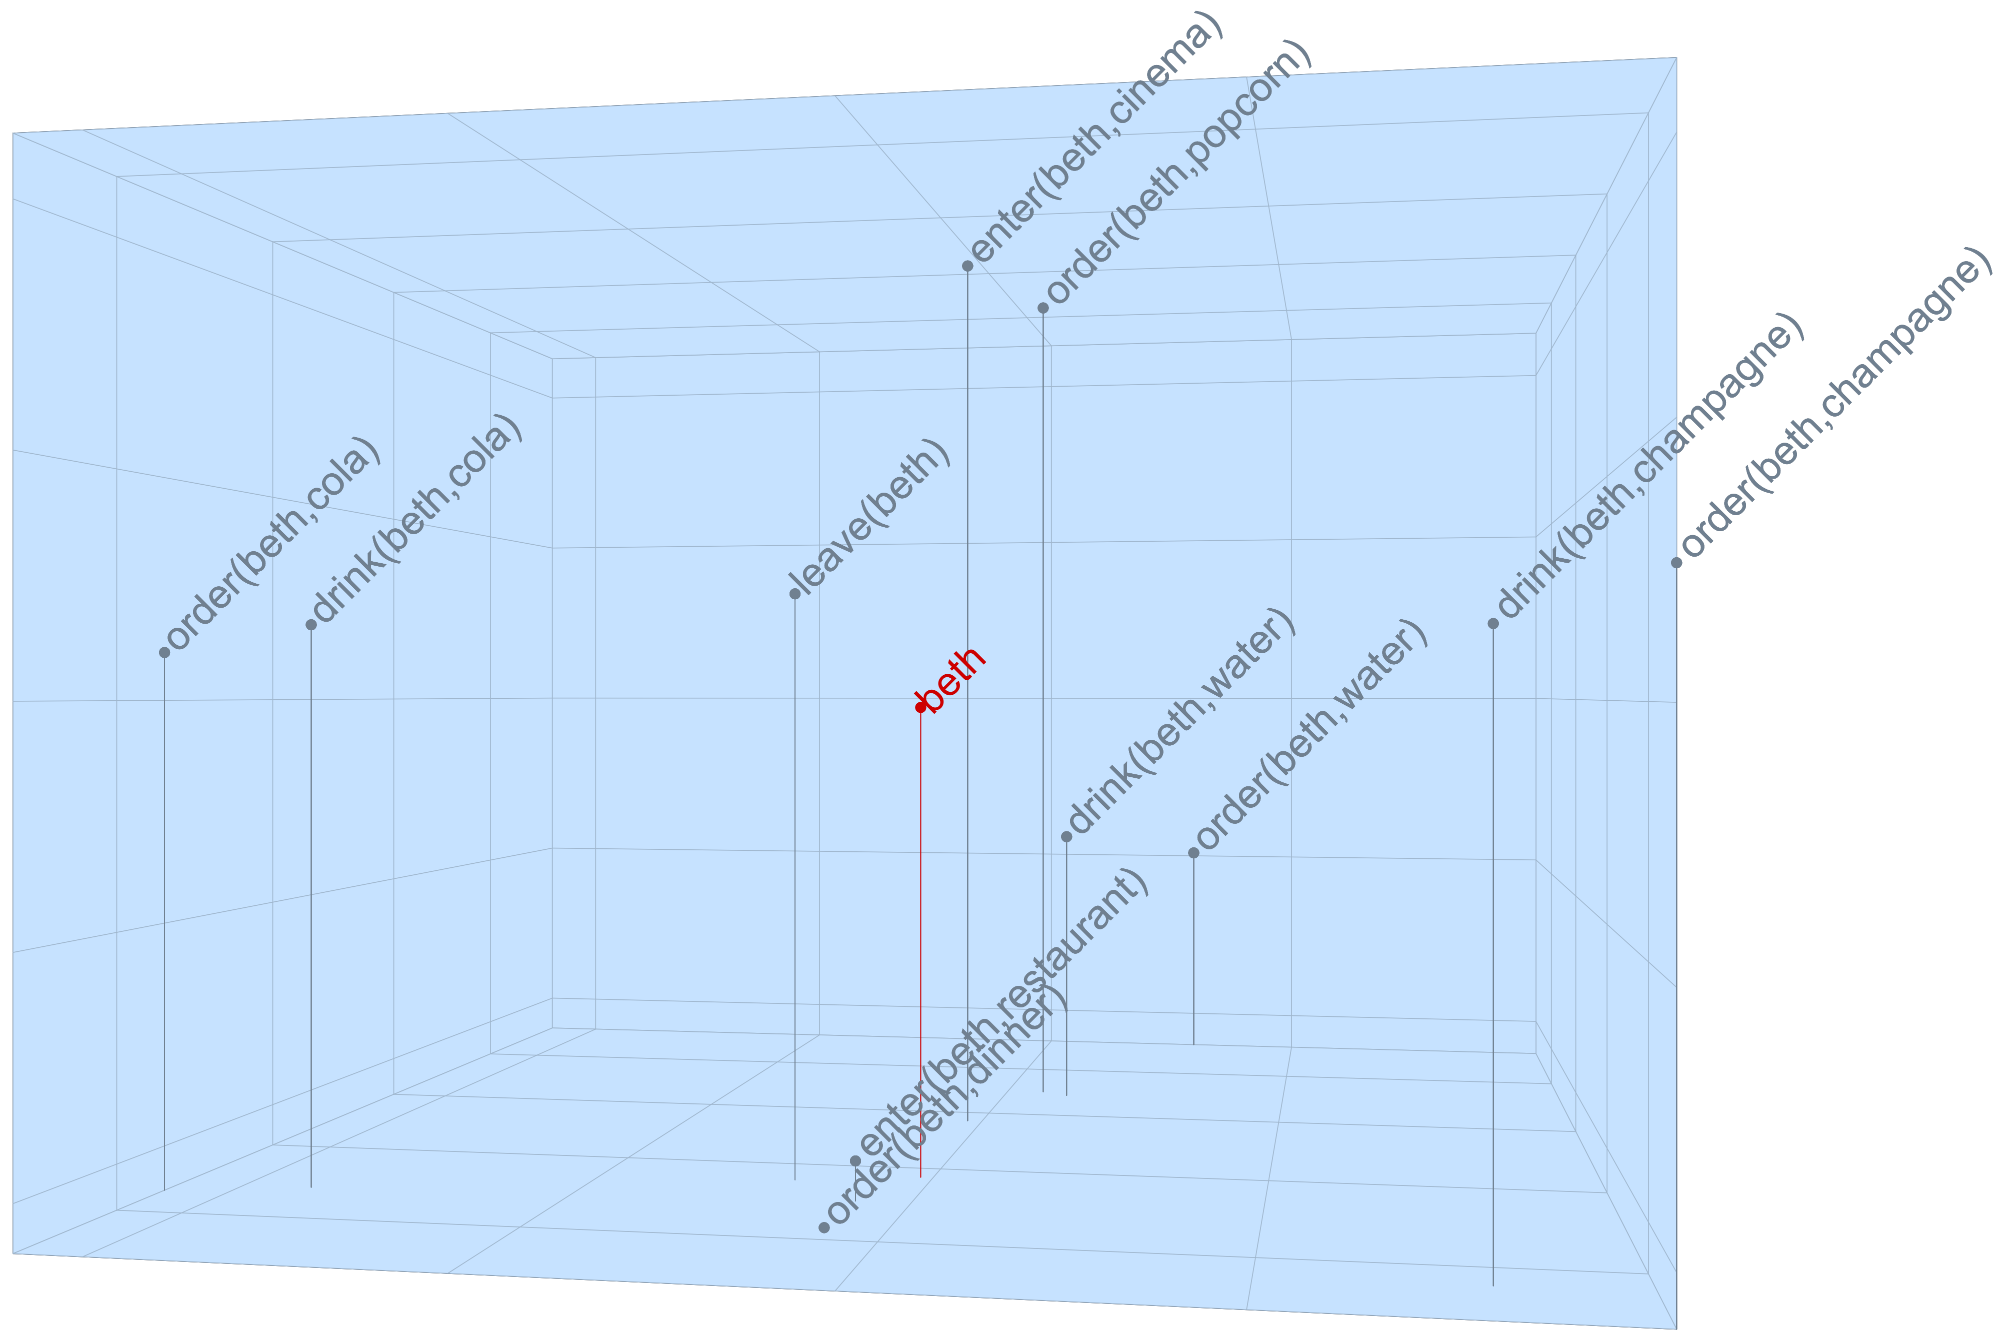

["beth"]

Venhuizen, Crocker & Brouwer (2018)
Discourse Process.

Comprehension is meaning-space navigation

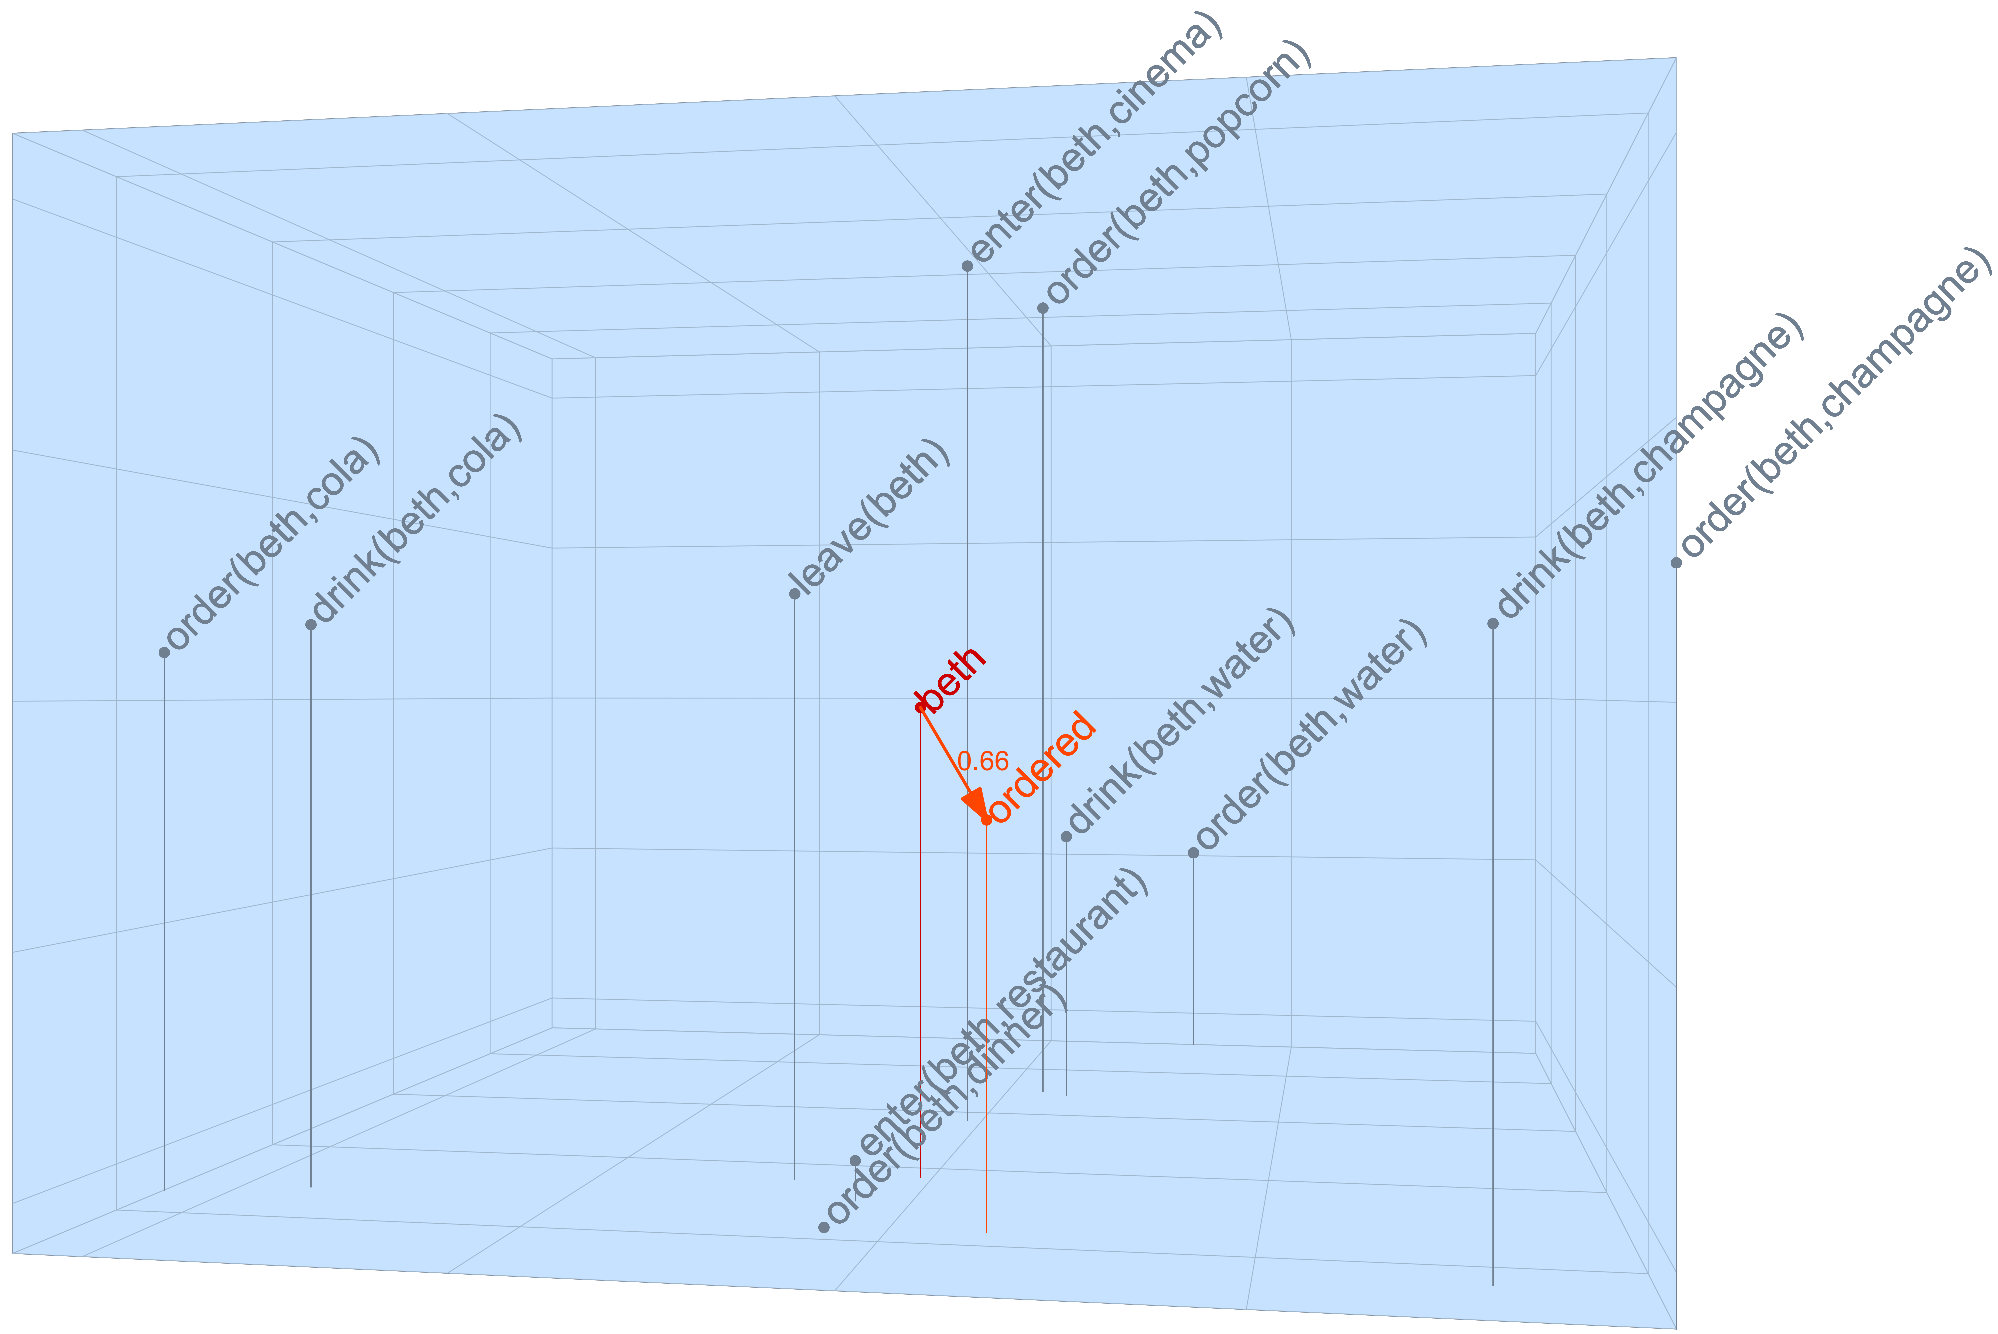

["beth", "ordered"]

(scalars ∝ distance)

Venhuizen, Crocker & Brouwer (2018)
Discourse Process.

Comprehension is meaning-space navigation

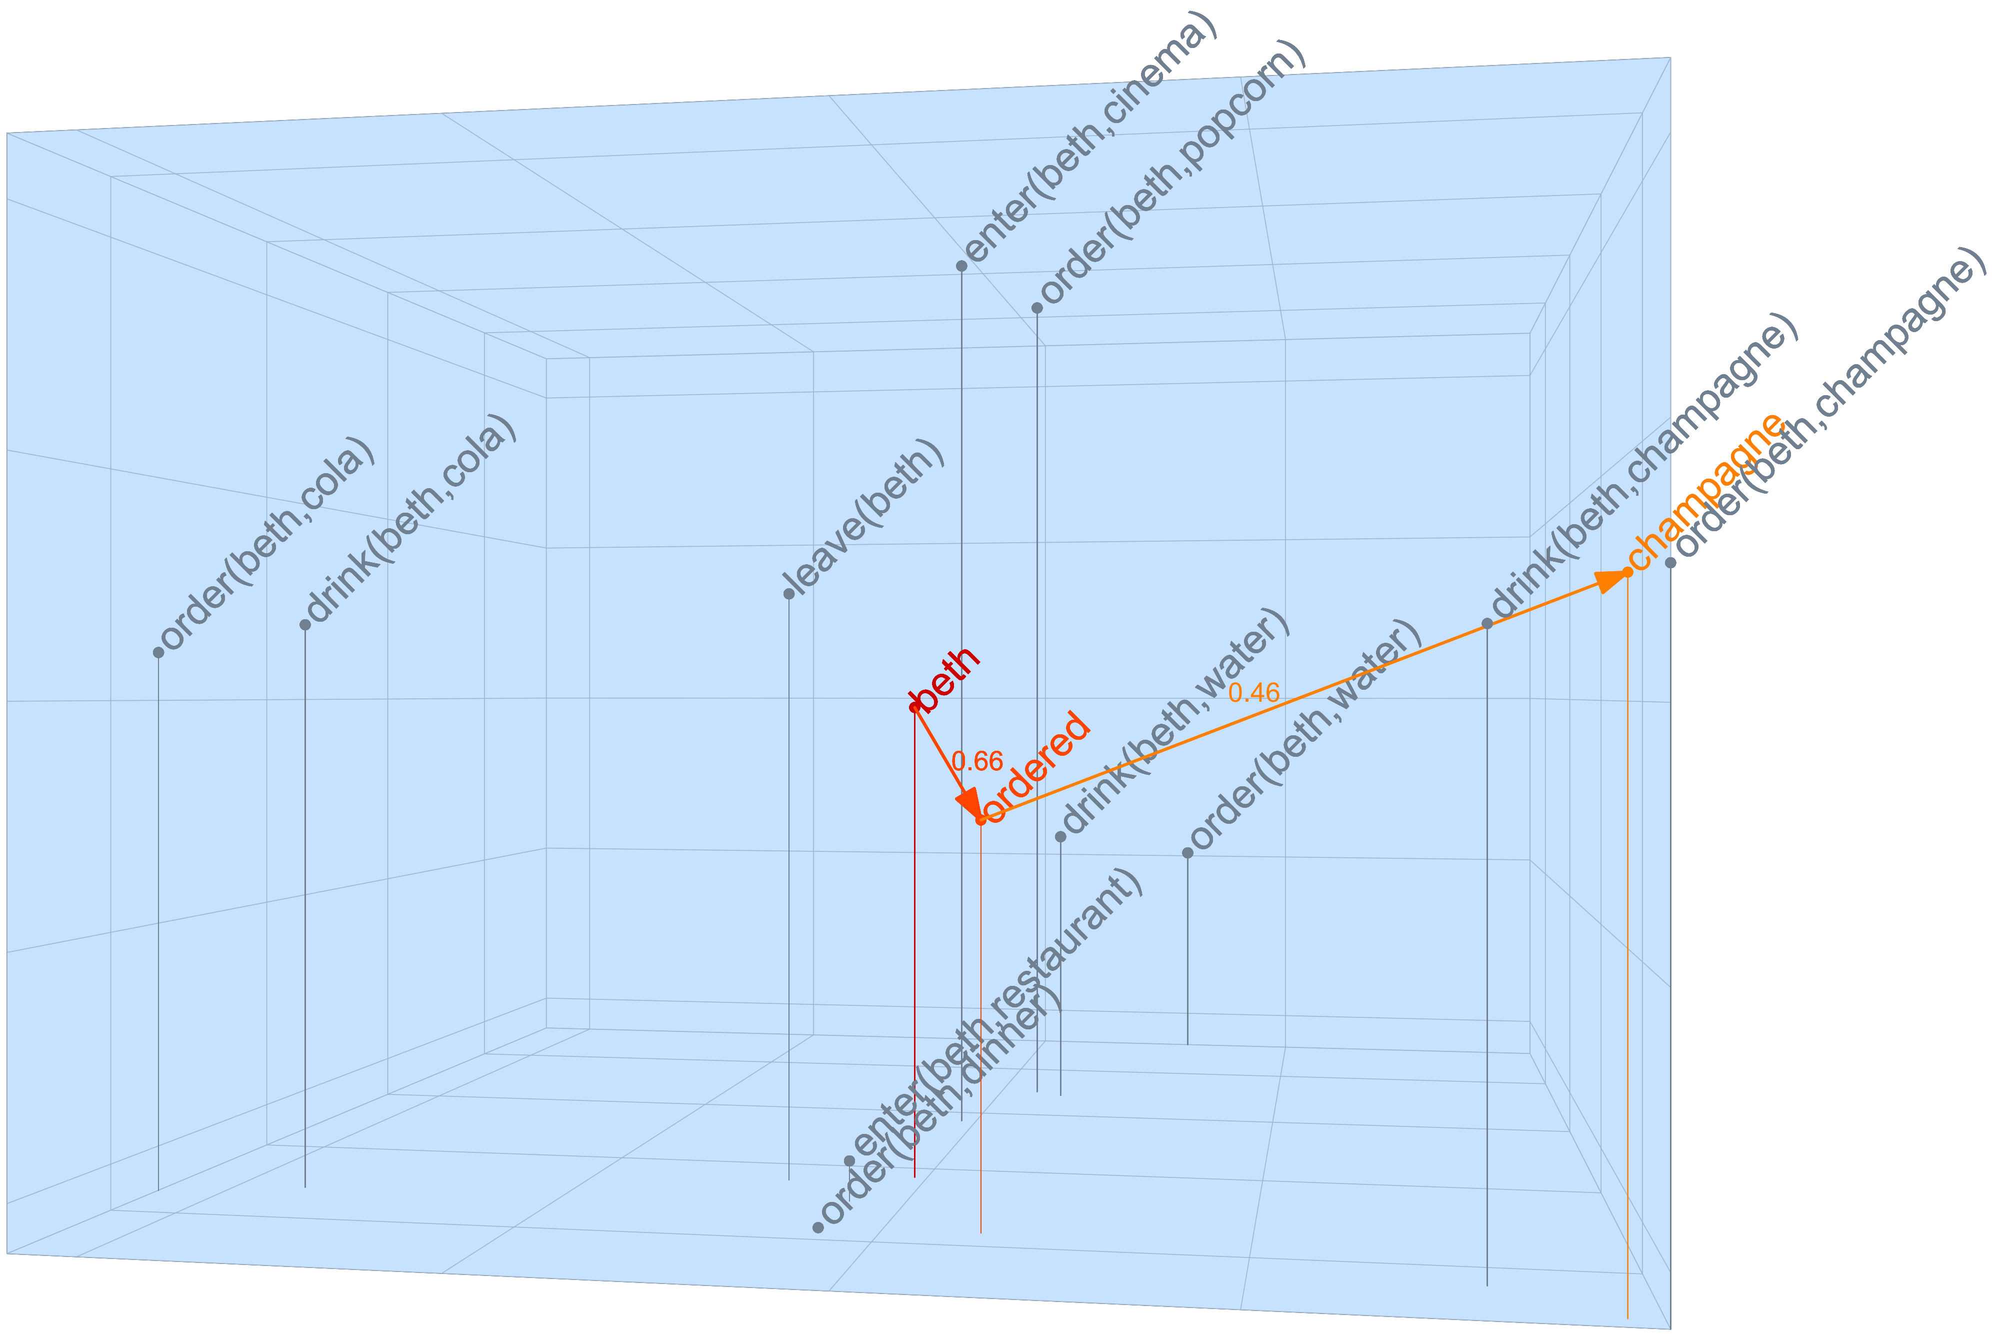

["beth", "ordered", "champagne"]

(scalars ∝ distance)

Venhuizen, Crocker & Brouwer (2018)
Discourse Process.

Comprehension is meaning-space navigation

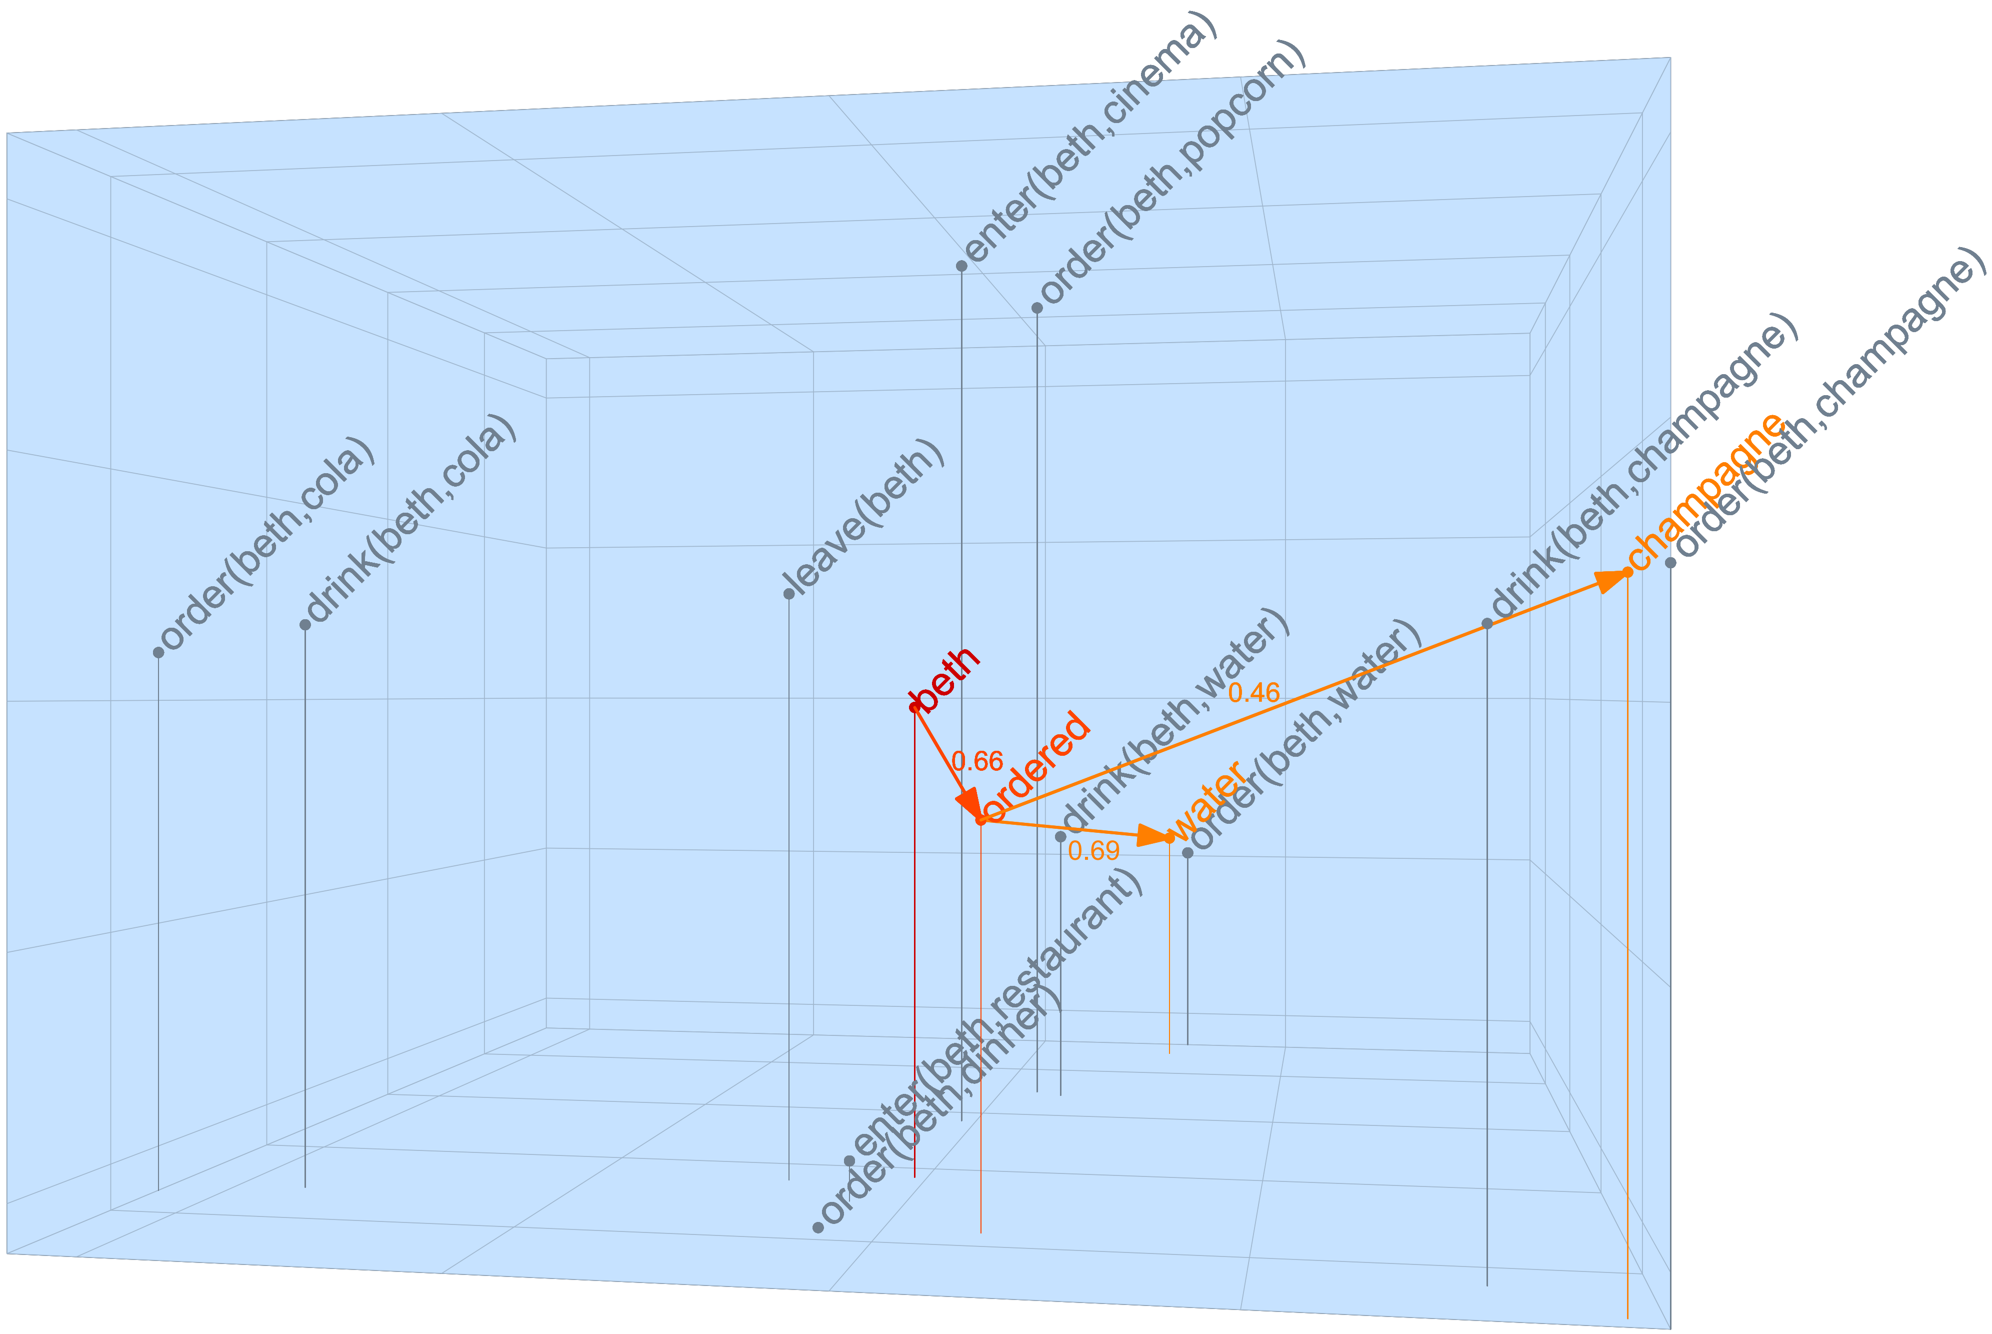

["beth", "ordered", "water"]

(scalars ∝ distance)

Venhuizen, Crocker & Brouwer (2018)
Discourse Process.

Comprehension is meaning-space navigation

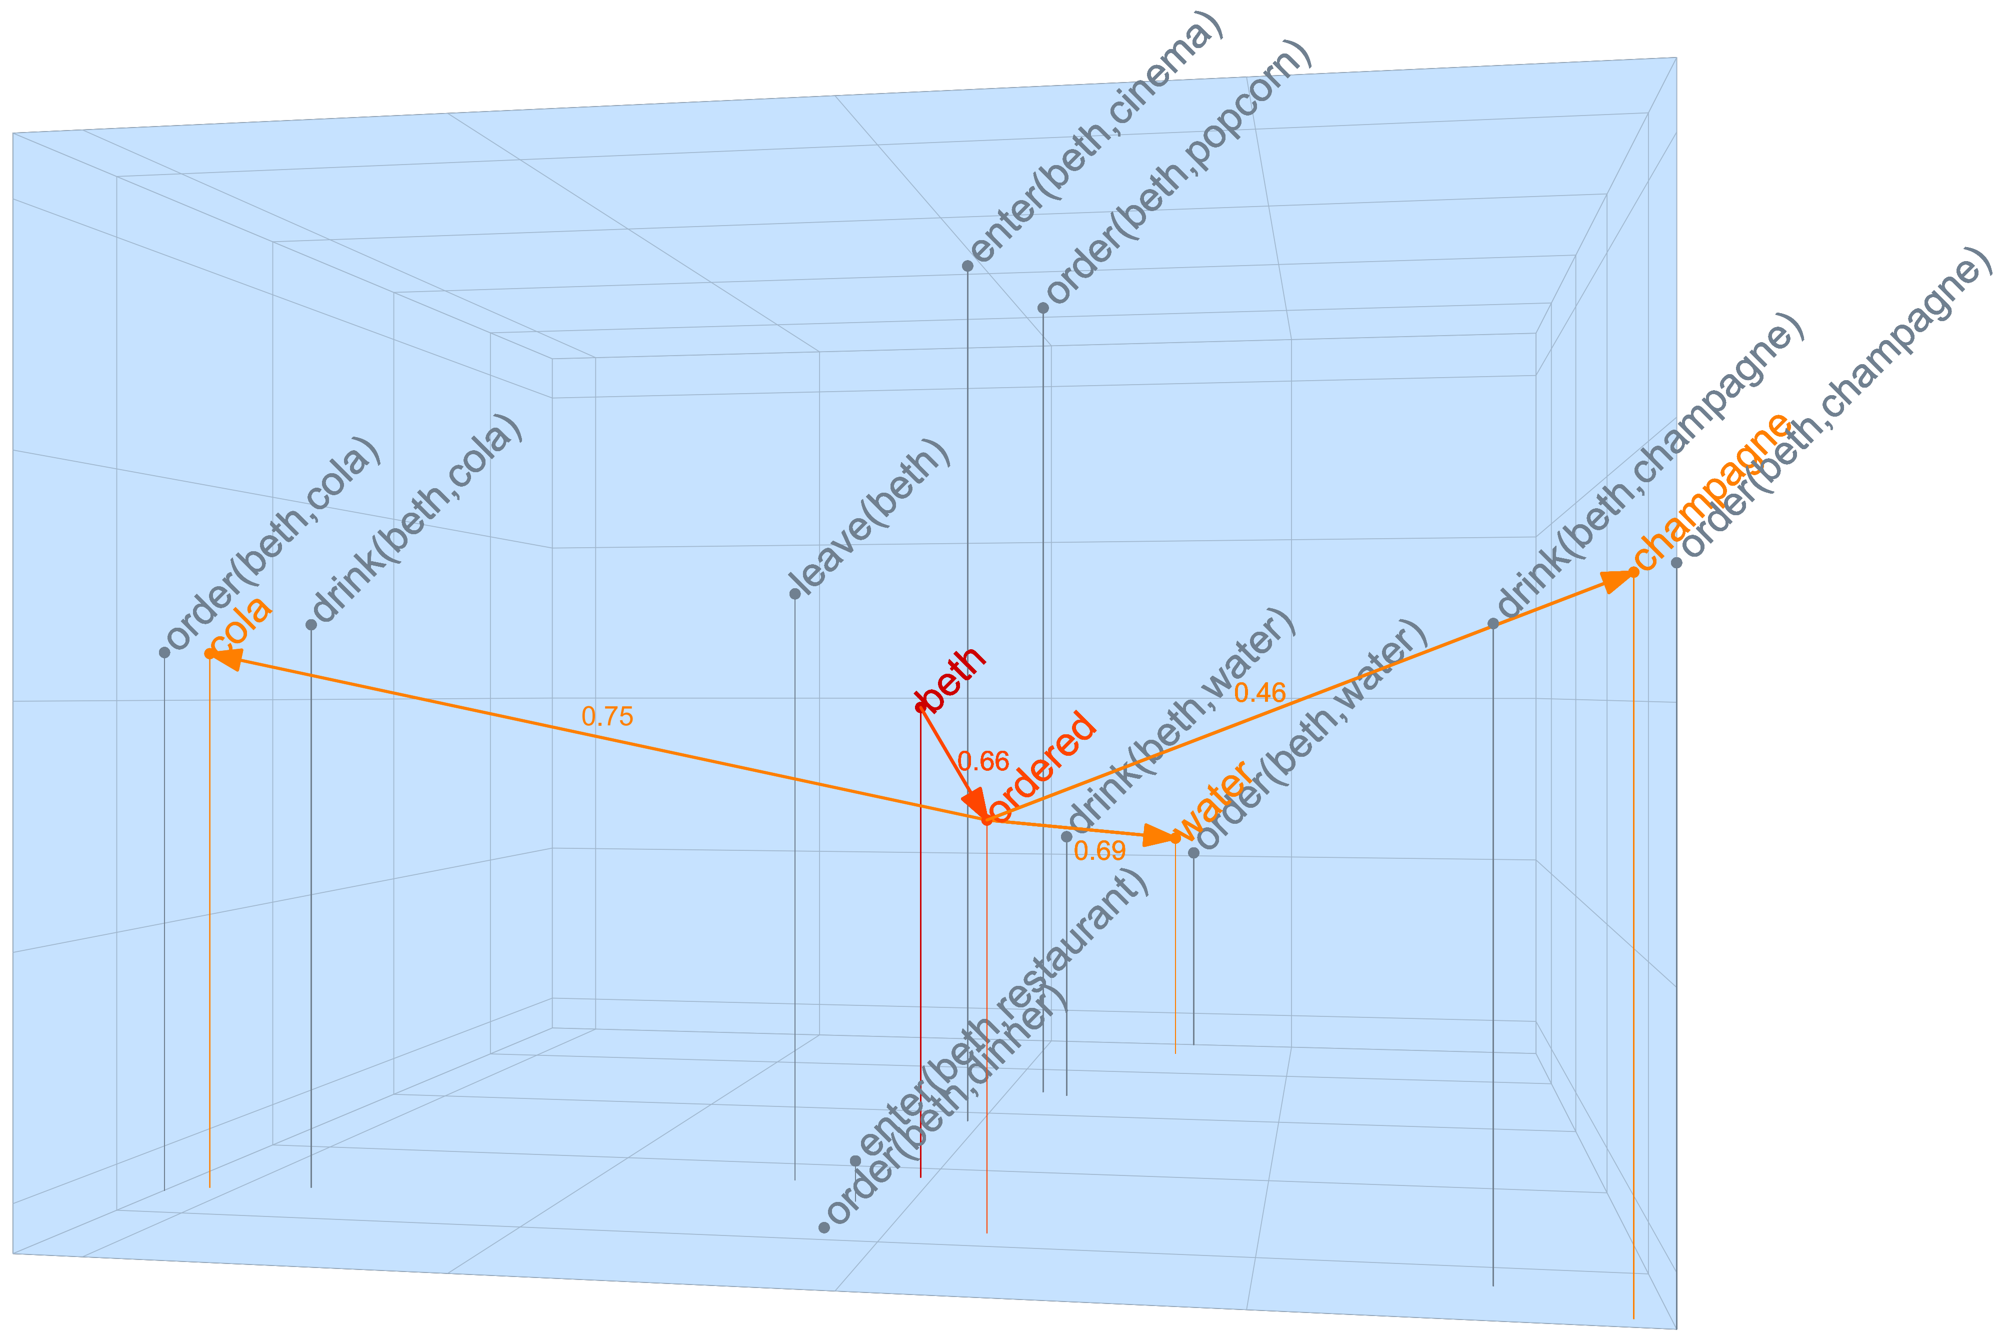

["beth", "ordered", "cola"]

(scalars ∝ distance)

Venhuizen, Crocker & Brouwer (2018)
Discourse Process.

Comprehension is meaning-space navigation

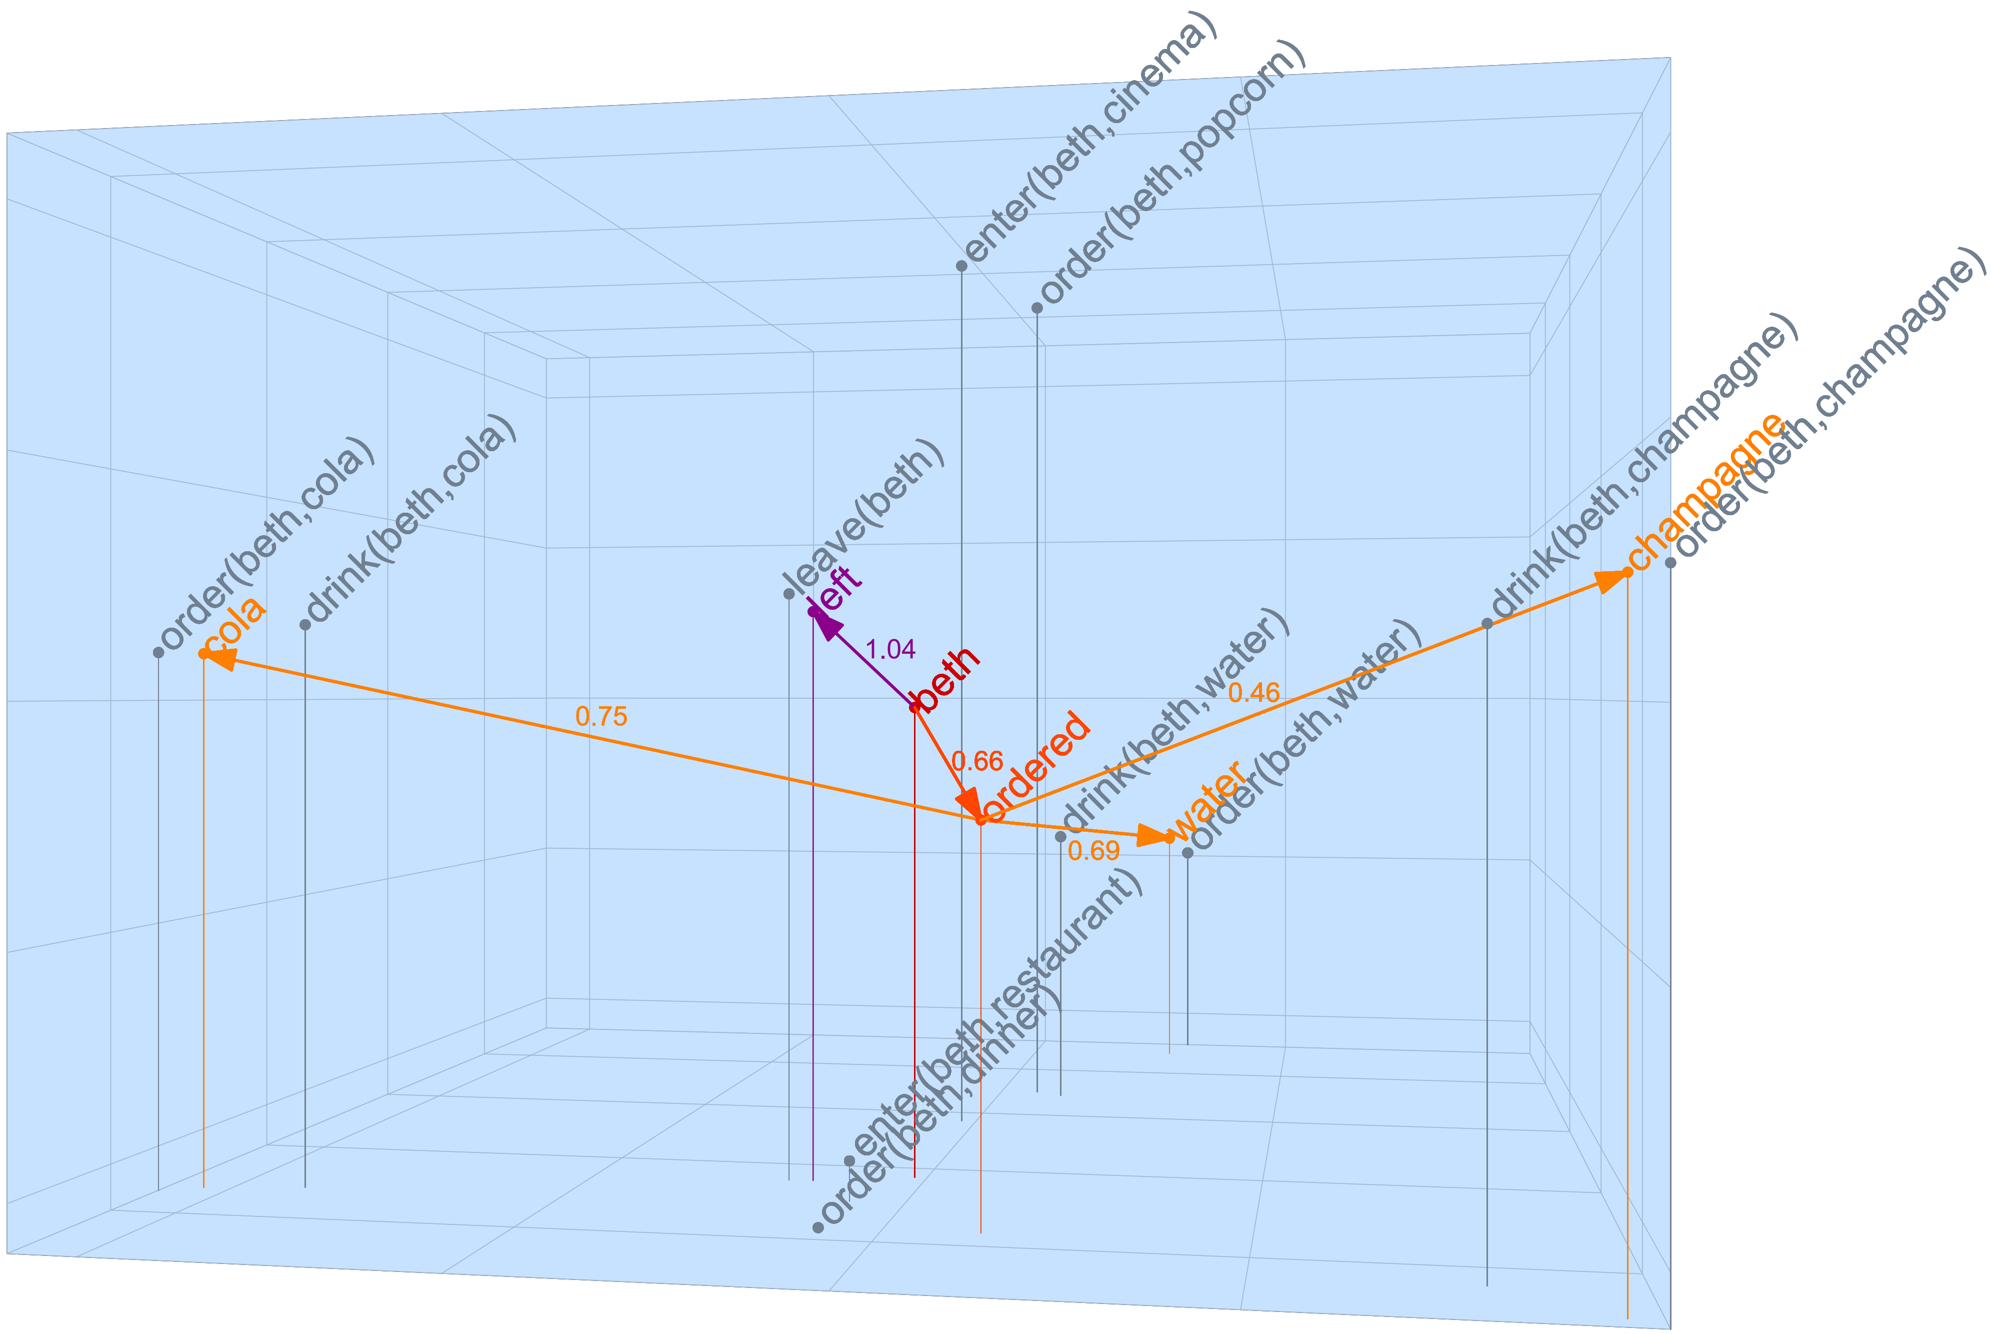

["beth", "left"]

(scalars ∝ distance)

Venhuizen, Crocker & Brouwer (2018)
Discourse Process.

What does the model ‘understand’?

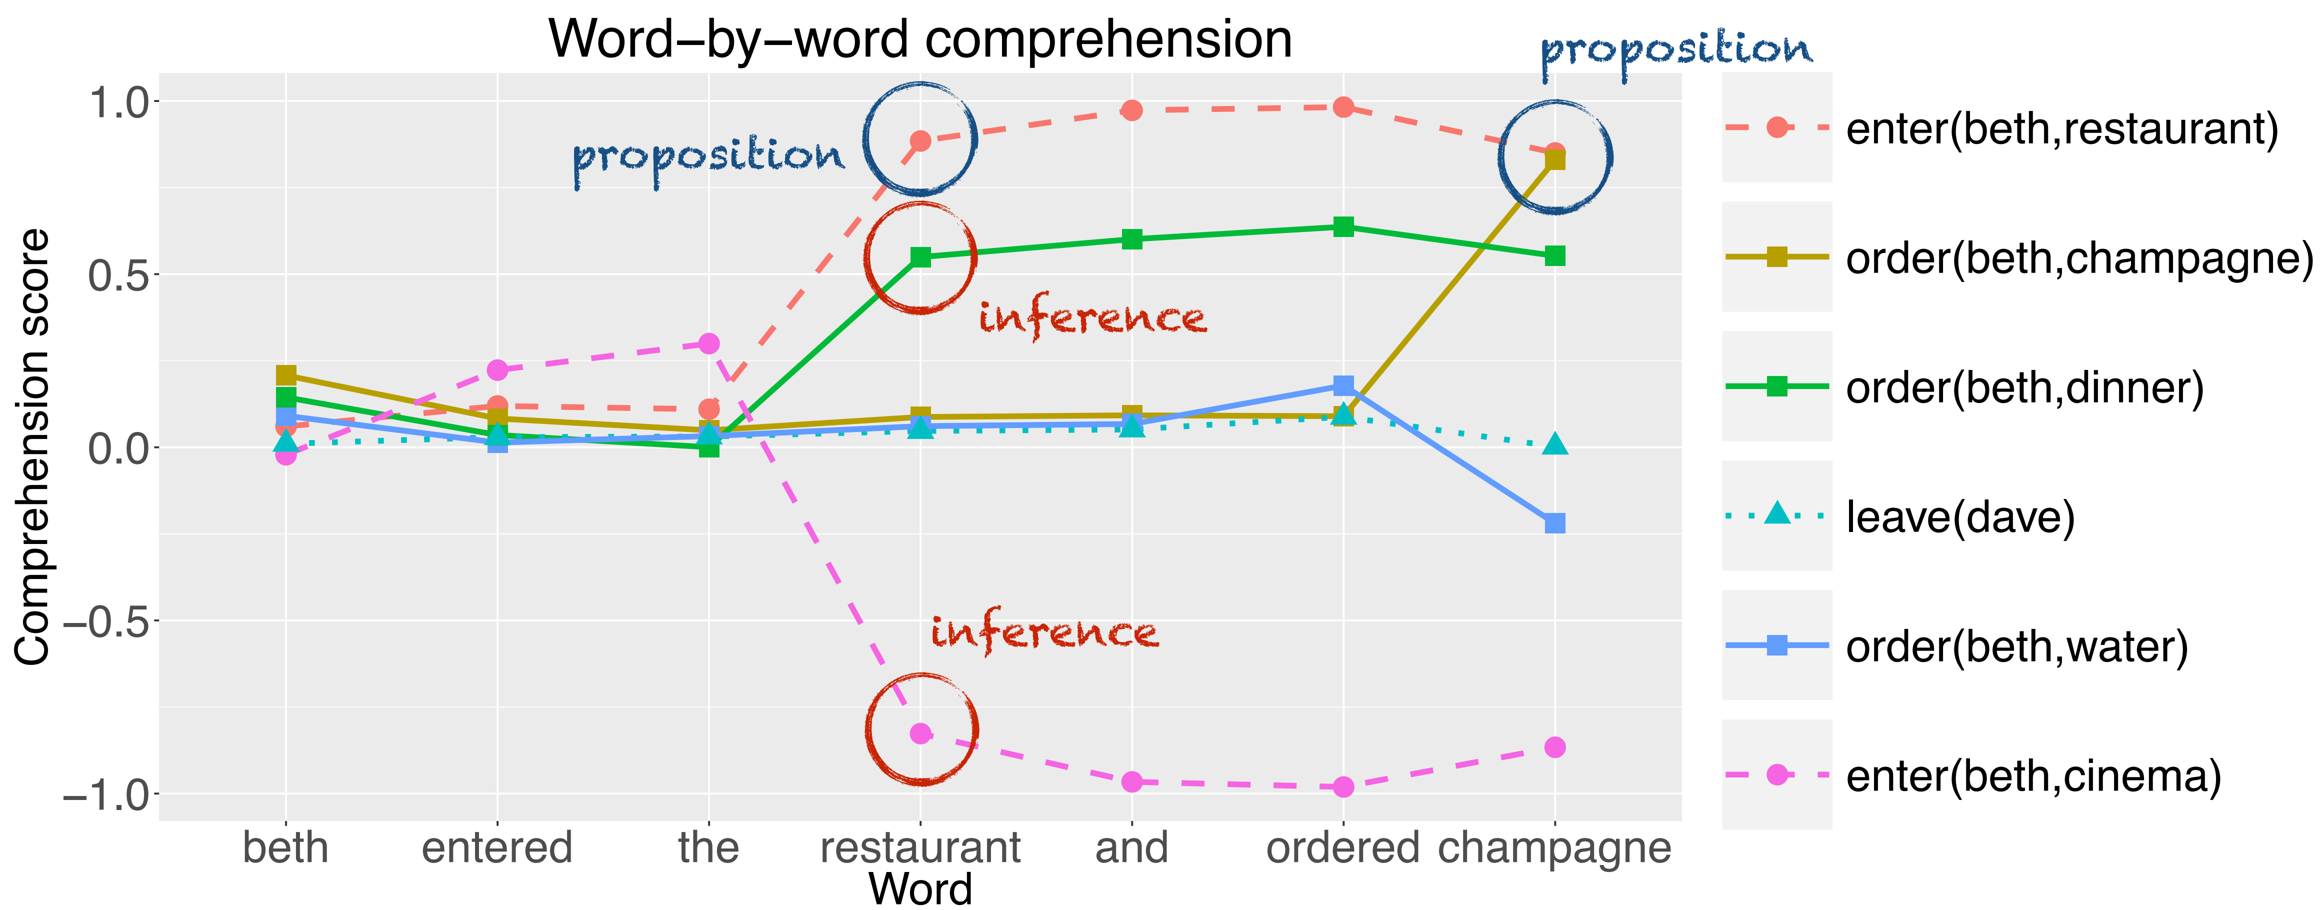

> Points in meaning-space capture meaning beyond literal propositional content; i.e., model engages in **direct knowledge-driven inferencing**

Venhuizen, Crocker & Brouwer (2018)
Discourse Process.

What does the model ‘understand’?