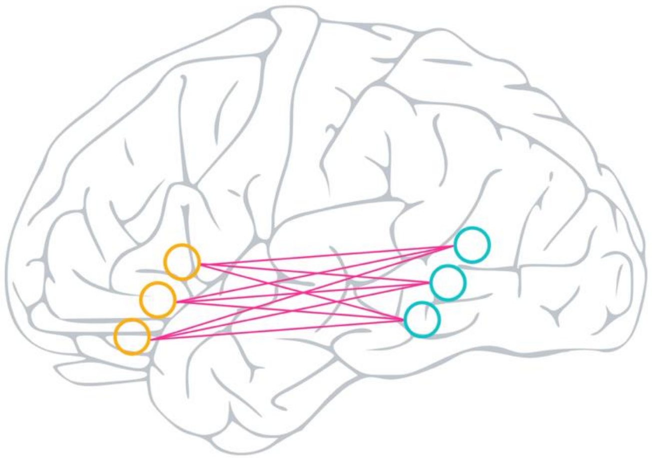

Two Models of Meaning: Revisiting the Principle of Compositionality from the Neurocognition of Language

Noortje Venhuizen
Harm Brouwer

Tilburg University

Journées LIFT2
16-17 oct. 2025, Paris

From words to utterances

How does the meaning of complex utterances derive from the meaning of the individual words that constitute these utterances?

Human language use is productive

- We produce/understand infinitely many unseen complex expressions

... and systematic

- The ability to produce/understand new expressions is intrinsically connected to the ability to produce/understand others

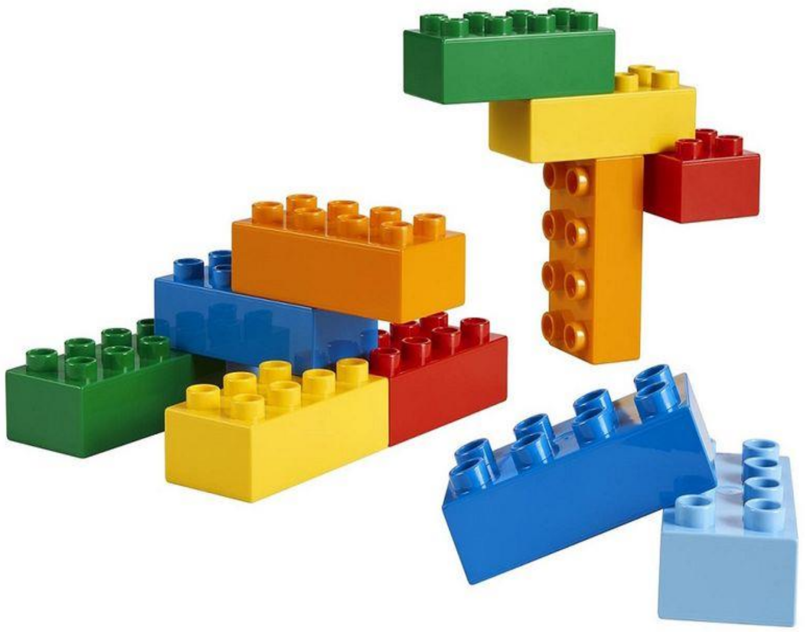

The Principle of Compositionality *aka Frege's principle*

"The meaning of a complex expression is a function of the meanings of its parts and of the syntactic rules by which they are combined"

- Barbara Partee, 1993

> Assumes close relationship between Lexical Semantics (LS) and sentence/utterance-level Compositional Semantics (CS)

> **But:** LS and CS have developed into surprisingly separate fields

Lexical Semantics: the meaning of words

- **Core observation:** Conceptual knowledge associated with individual words is both gradient and structured
- Semantic similarity: Concepts are related to each other to different degrees
- **Representational currency:** Semantic features (McRae et al., 2005)
  - Categories vs. local features vs. spaces

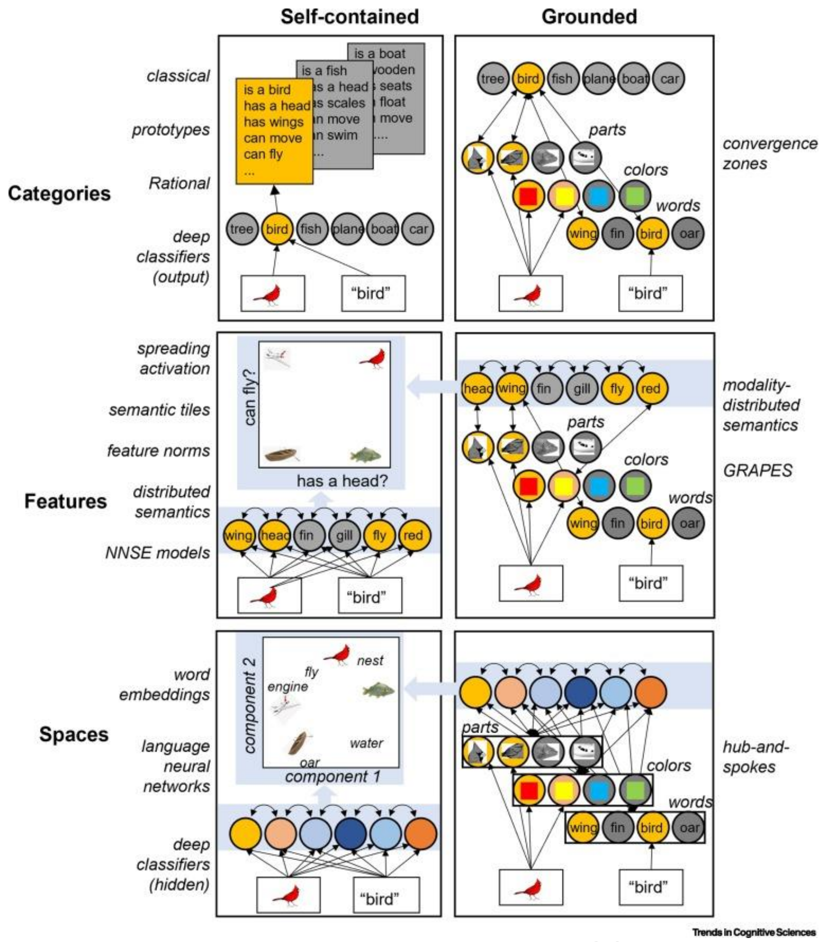

Frisby et al. (2023, *TiCS*)

Distributional hypothesis

> You shall know a word
> by the company it keeps!
>
> *- J. R. Firth (1957)*

- **Distributional semantics**: Lexical co-occurrence information across linguistic contexts (Latent Semantic Analysis, LSA, Landauer & Dumais, 1997; Hyperspace Analogue of language, HAL, Burgess, 1998; Dependency Vectors, DV, Padó & Lapata, 2007)
- **Word embeddings**: language model-derived vectors with abstract dimensions (word2vec, Mikolov et al., 2013a, 2013b; GloVe, Pennington et al., 2014; ELMo, Peters et al., 2018; BERT, Devlin et al., 2019; GPT, Radford et al., 2019)

Compositional approaches to LS

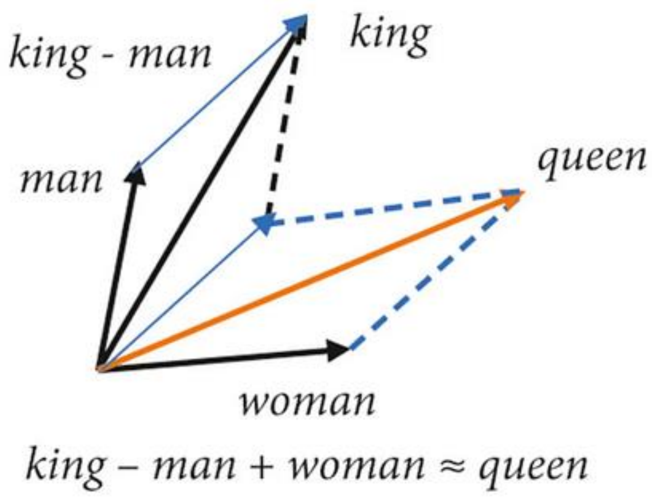

Vector operations as a proxy for semantic composition (e.g., Mitchell & Lapata, 2010)

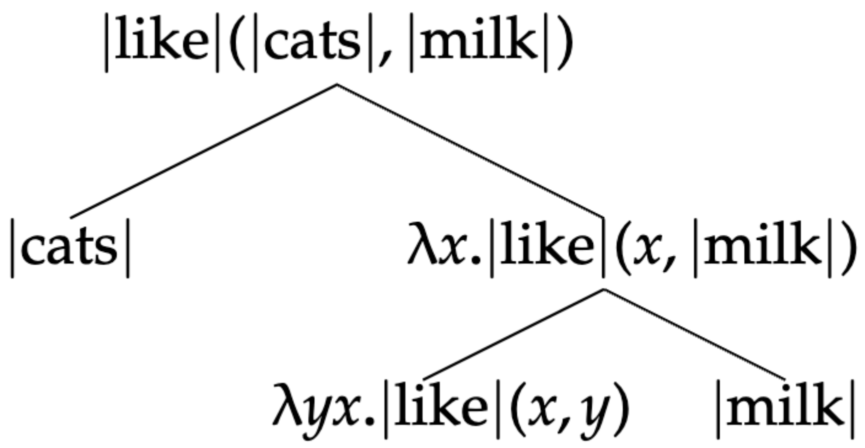

Combining LS representations into complex CS structures (e.g., Grefenstette & Sadrzadeh, 2015)

"Good at lexical semantics, bad at composition" (Pavlick, 2022)

Formal Semantics: sentence-level meaning

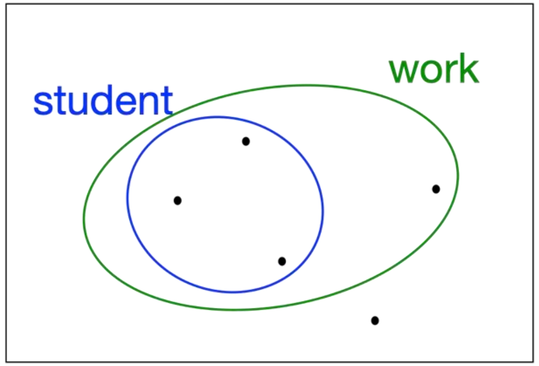

$[\![\forall x(\text{student}'(x) \rightarrow \text{work}'(x))]\!]^{M,g}$
$= 1$ *iff* $V_M(\text{student}') \subseteq V_M(\text{work}')$

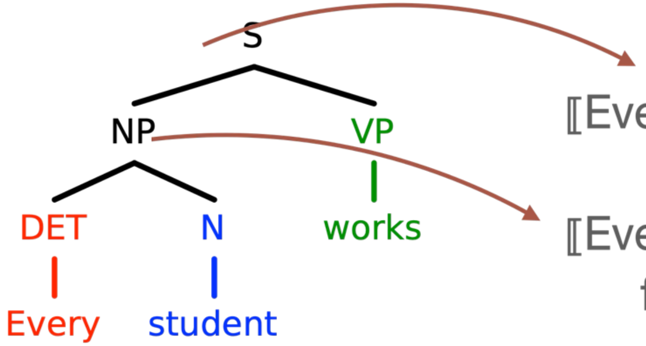

$[\![\text{Every student works}]\!]^{M,g} =$
$f_1([\![\text{Every student}]\!]^{M,g}, [\![\text{works}]\!]^{M,g})$

$[\![\text{Every student}]\!]^{M,g} =$
$f_2([\![\text{Every}]\!]^{M,g}, [\![\text{student}]\!]^{M,g})$

- Truth and entailment via model-theoretic interpretation
- Traditional formal semantics: composition as function application

> Incompatible with gradient notion of semantic similarity?

Distributional Formal Semantics (DFS)

$M_1 = \langle U_1, V_1 \rangle$

$p_1 \land \neg p_2 \land p_3 \land \ldots$

$M_2 = \langle U_2, V_2 \rangle$

$\neg p_1 \land p_2 \land p_3 \land \ldots$

$M_3 = \langle U_3, V_3 \rangle$

$\neg p_1 \land p_2 \land \neg p_3 \land \ldots$

…

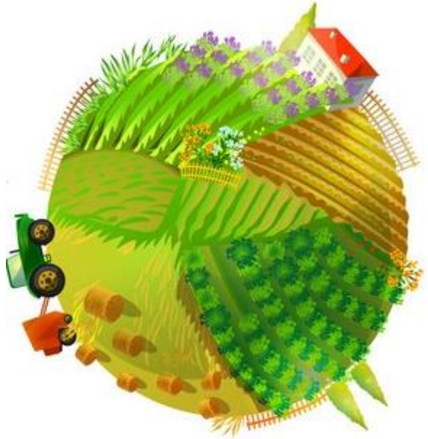

$M_n = \langle U_n, V_n \rangle$

$\neg p_1 \land \neg p_2 \land \neg p_3 \land \ldots$

- Individual models describe states-of-affairs over all propositions in $\mathbb{P}$
- The set of models $\mathbb{M}_{\mathbb{P}}$ defines a meaning space
- Propositional meaning defined by co-occurrence across models (~ world knowledge)

Venhuizen, Hendriks, Crocker, Brouwer (2019, *WoLLIC*; 2022, *Inf. Comp.*)

Propositional meaning in the meaning space

meaning vectors

formal models

	p_1	p_2	p_3	p_4	…
M_1	1	1	0	0	…
M_2	1	0	0	1	…
M_3	0	1	0	1	…
M_4	1	1	1	1	…
M_5	0	1	0	0	…
…	…	…			

$[\![p_j]\!]^{\mathbb{M}} := \mathbf{v}(p_j)$

where: $\mathbf{v}_i(p_j) = 1$ iff $M_i \vDash p_j$

- **Incremental inference-based probabilistic sampling:** Based on a set of propositions $\mathbb{P}$, we sample a set of models $\mathbb{M}_{\mathbb{P}}$—taking into account hard and probabilistic world knowledge constraints
- **Co-occurrence defines meaning:** Propositions with related meanings are true in many of the same models, resulting in similar meaning vectors

Formal properties of meaning vectors

Meaning vectors are compositional

- Negation: $\vec{v}_i(\neg p) = 1 - \vec{v}_i(p)$
- Conjunction: $\vec{v}_i(p \wedge q) = \vec{v}_i(p)\vec{v}_i(q)$

Meaning vectors are probabilistic

- Prior probability : $P(a) = \frac{1}{|\mathcal{M}|}\sum_i \vec{v}_i(a)$
- Conjunction probability: $P(a \wedge b) = \frac{1}{|\mathcal{M}|}\sum_i \vec{v}_i(a)\vec{v}_i(b)$
- Conditional probability: $P(a|b) = \frac{P(a \wedge b)}{P(b)}$

	p_1	p_2	p_3	p_4	
M_1	1	1	0	0	…
M_2	1	0	0	1	…
M_3	0	1	0	1	…
M_4	1	1	1	1	…
	0	1	0	0	…
	…	…	…	…	…

Probabilistic inference in the meaning space

$-1 \leq \textit{inference}(a,b) \leq 1$

- P(a|b) > P(a): Positive inference

$$\frac{P(a|b) - P(a)}{1 - P(a)}$$

- P(a|b) ≤ P(a): Negative inference

$$\frac{P(a|b) - P(a)}{P(a)}$$

Example: meaning space with
$|\mathbb{P}| = 51$ and $|\mathbb{M}_{\mathbb{P}}| = 150$

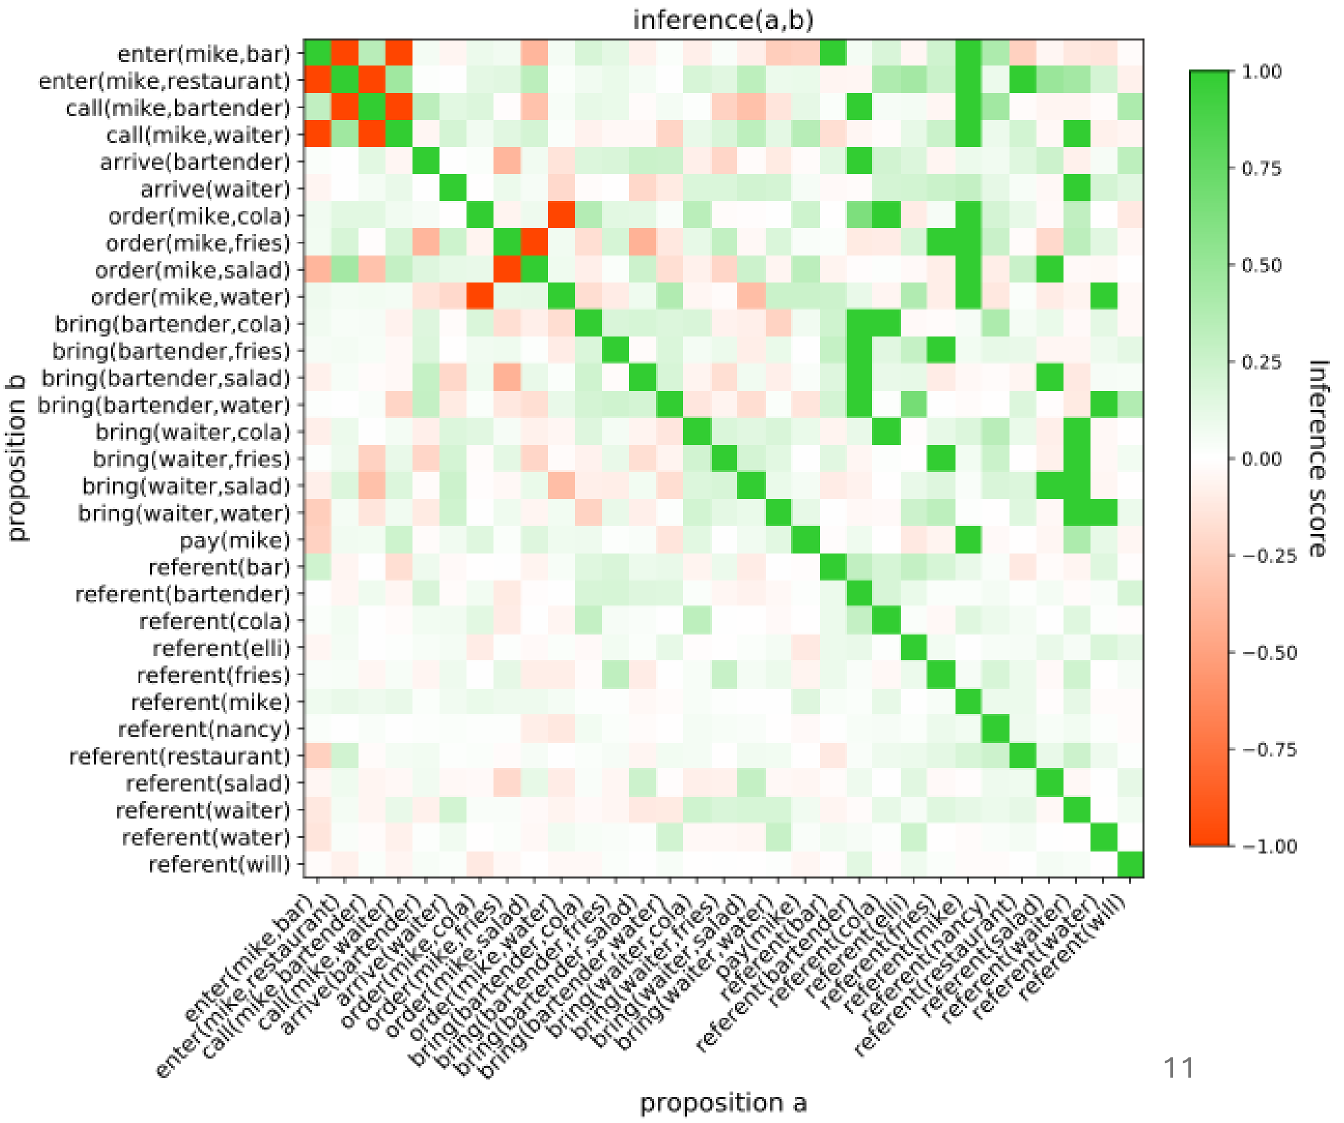

Interim: Two models of meaning?

LS representation

- Conceptual knowledge and structure
- Similarity driven by feature overlap

CS representation

- Discourse structure and inference
- Similarity driven by propositional co-occurrence

bar ~ restaurant

enter(bar) ~ call(bartender)

Q: How to compositionally combine LS into CS?

Neural perspective: evidence from ERPs

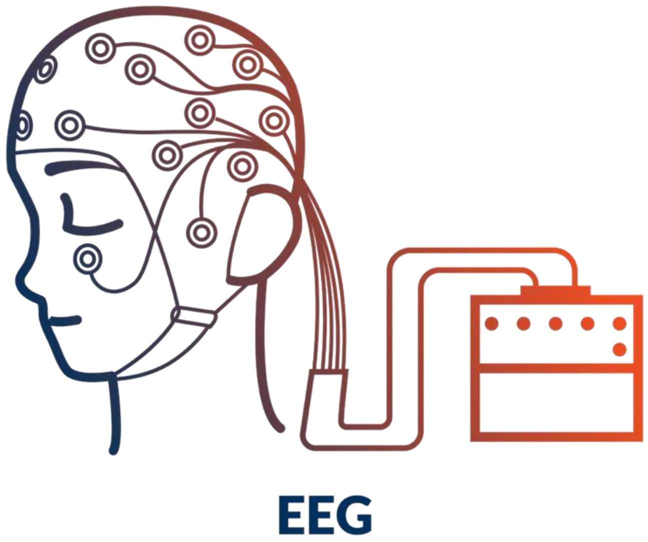

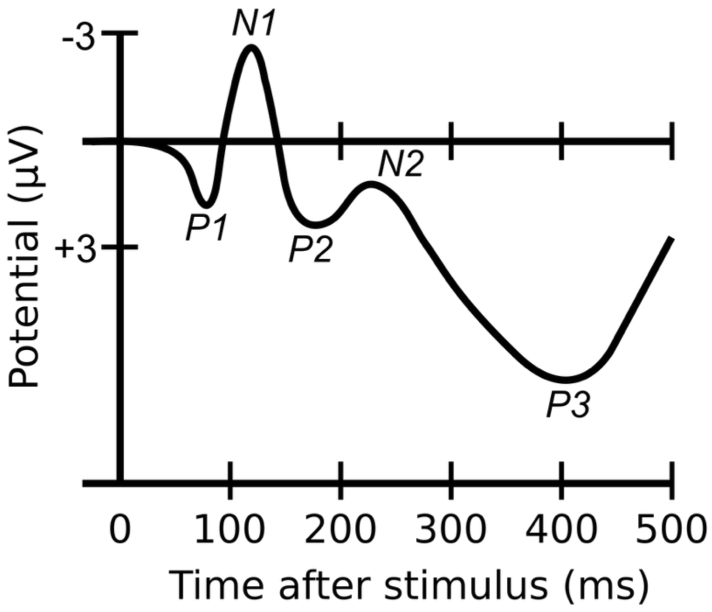

- **Event-Related Potentials (ERPs):** stimulus-locked EEG measurements
- **Components:** Systematic fluctuations in voltages, reflecting specific computational operations carried out in response to stimulus

Retrieval-Integration (RI) Theory

He spread his warm bread with [socks/butter]
[Kutas & Hillyard, 1980]

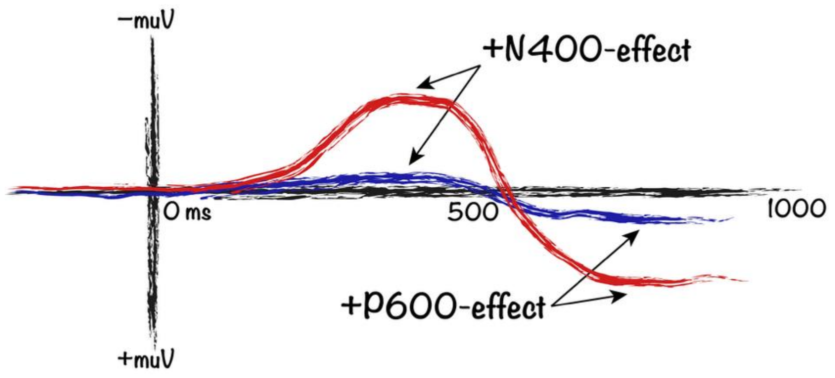

N400 ~ **retrieve**: (word form, utterance context) → LS representation

→ Retrieval of word meaning from long-term memory; facilitated if (part of) meaning is pre-activated due to lexical or contextual priming

P600 ~ **integrate**: (word meaning, utterance context) → CS representation

→ Word-by-word construction, reorganization, or updating of an utterance meaning representation—a mental model of the unfolding discourse

Brouwer et al. (2012, *Brain Res.*; 2017, *Cogn. Sci.)*

Neurocomputational implementation of RI

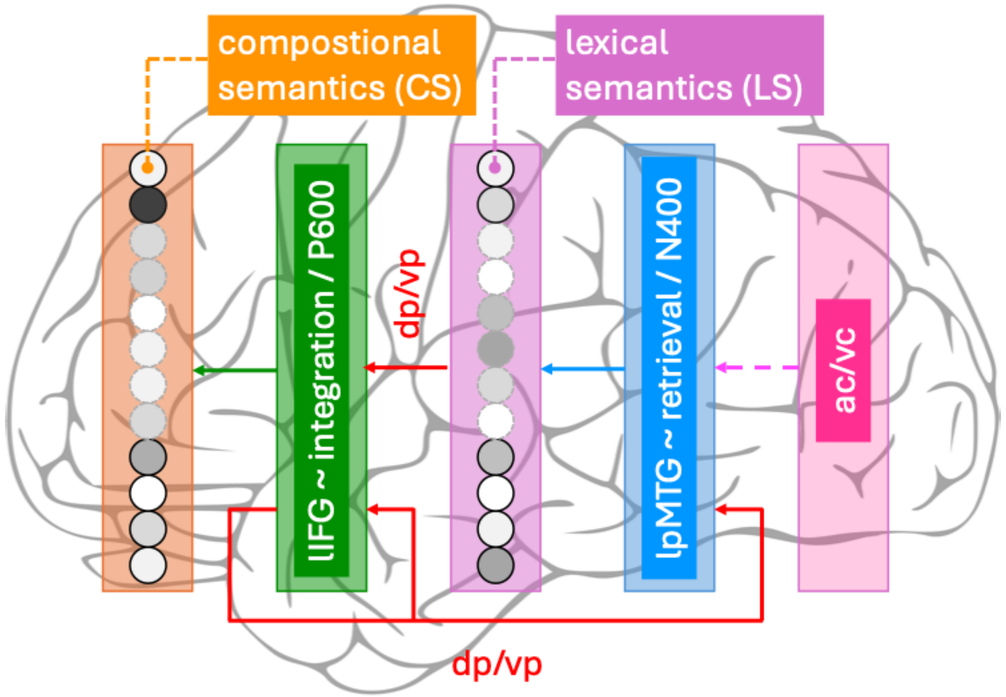

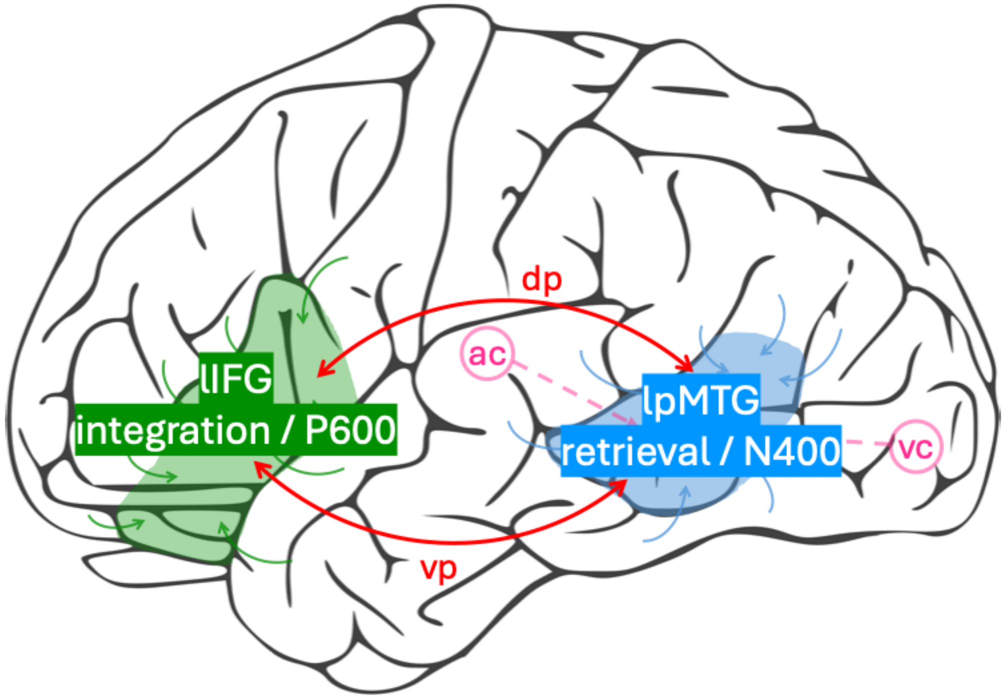

- Model trained to map sequences of words via LS to representation of CS
- LS derived from COALS (Rohde et al., 2009) / CS representations from DFS
- Aligns with functional-anatomical mapping of Retrieval-Integration theory

 Brouwer et al. (2017, *Cogn. Sci.*; 2021, *Frontiers Psych.)*

Incremental meaning construction in RI model

enter(mike,restaurant) ∧ order(mike,salad)

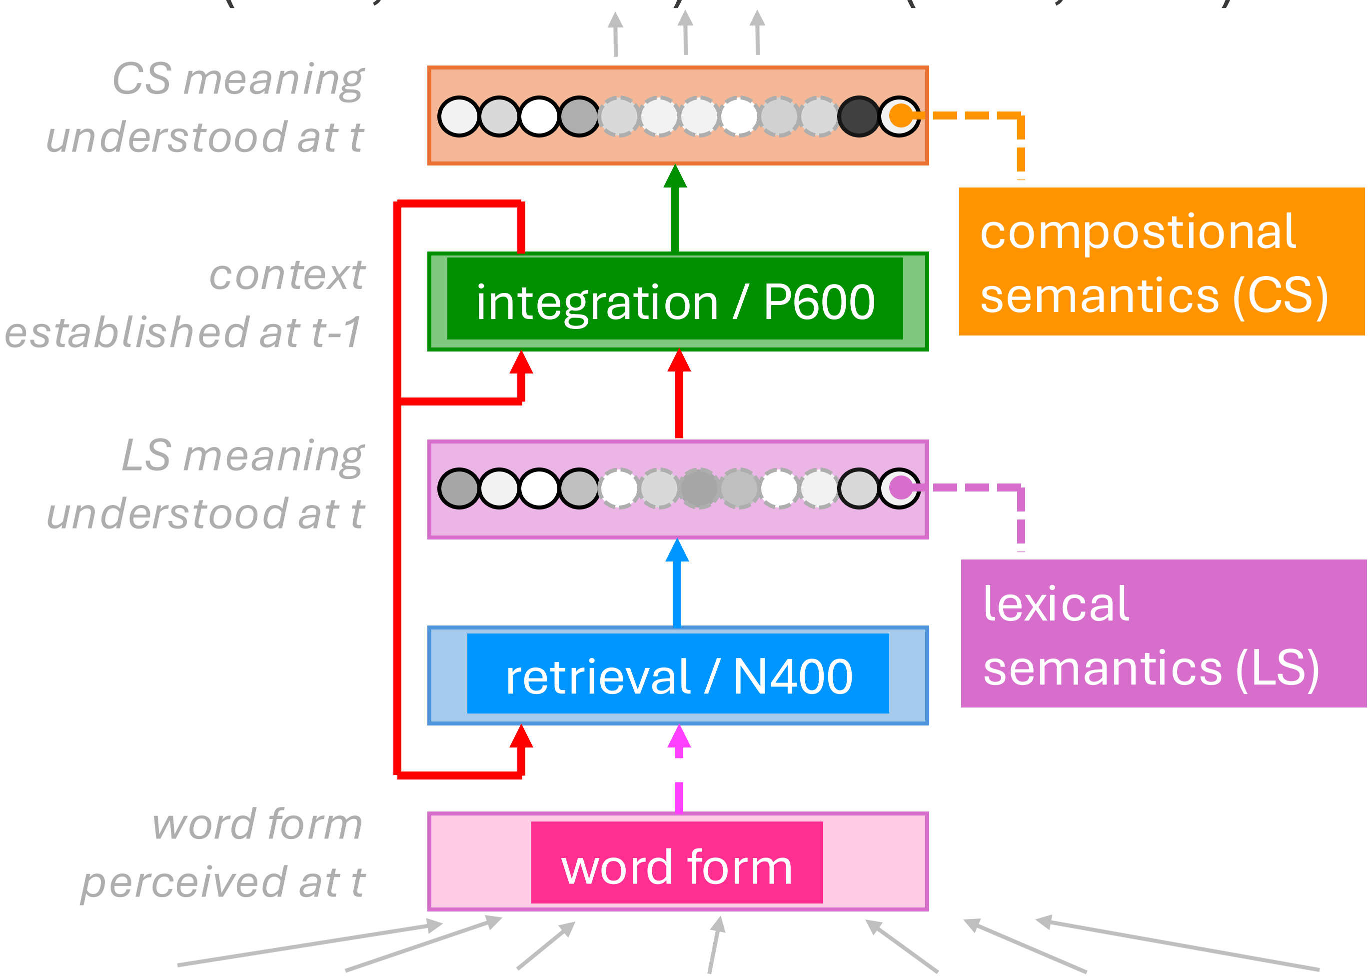

"Mike" "entered" "the" "restaurant" "and" "ordered" "salad"

Expectation-based processing based on:

- Sentence-semantics mappings in training
  → **linguistic experience**
- Co-occurrences in the meaning space
  → **world knowledge**

Meaning-space navigation

- 3-D representation of 150-D meaning space (using MDS)
- Propositions that co-occur frequently are positioned close in space

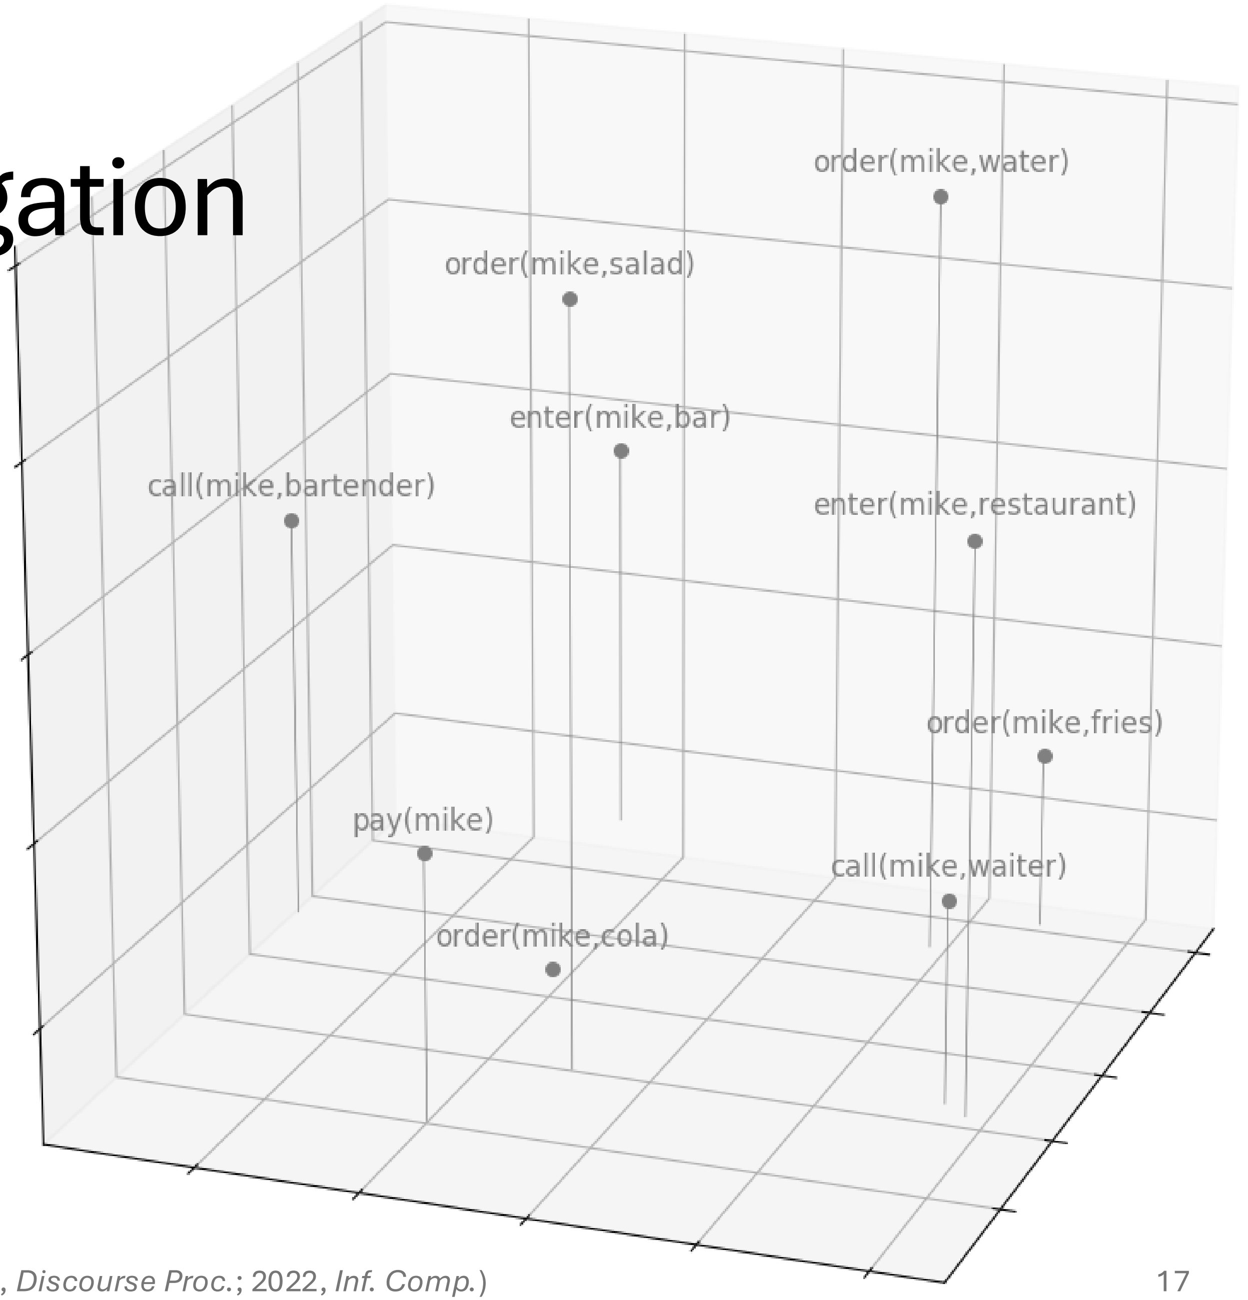

Venhuizen et al. (2019, *Discourse Proc.*; 2022, *Inf. Comp.*)

Meaning-space navigation

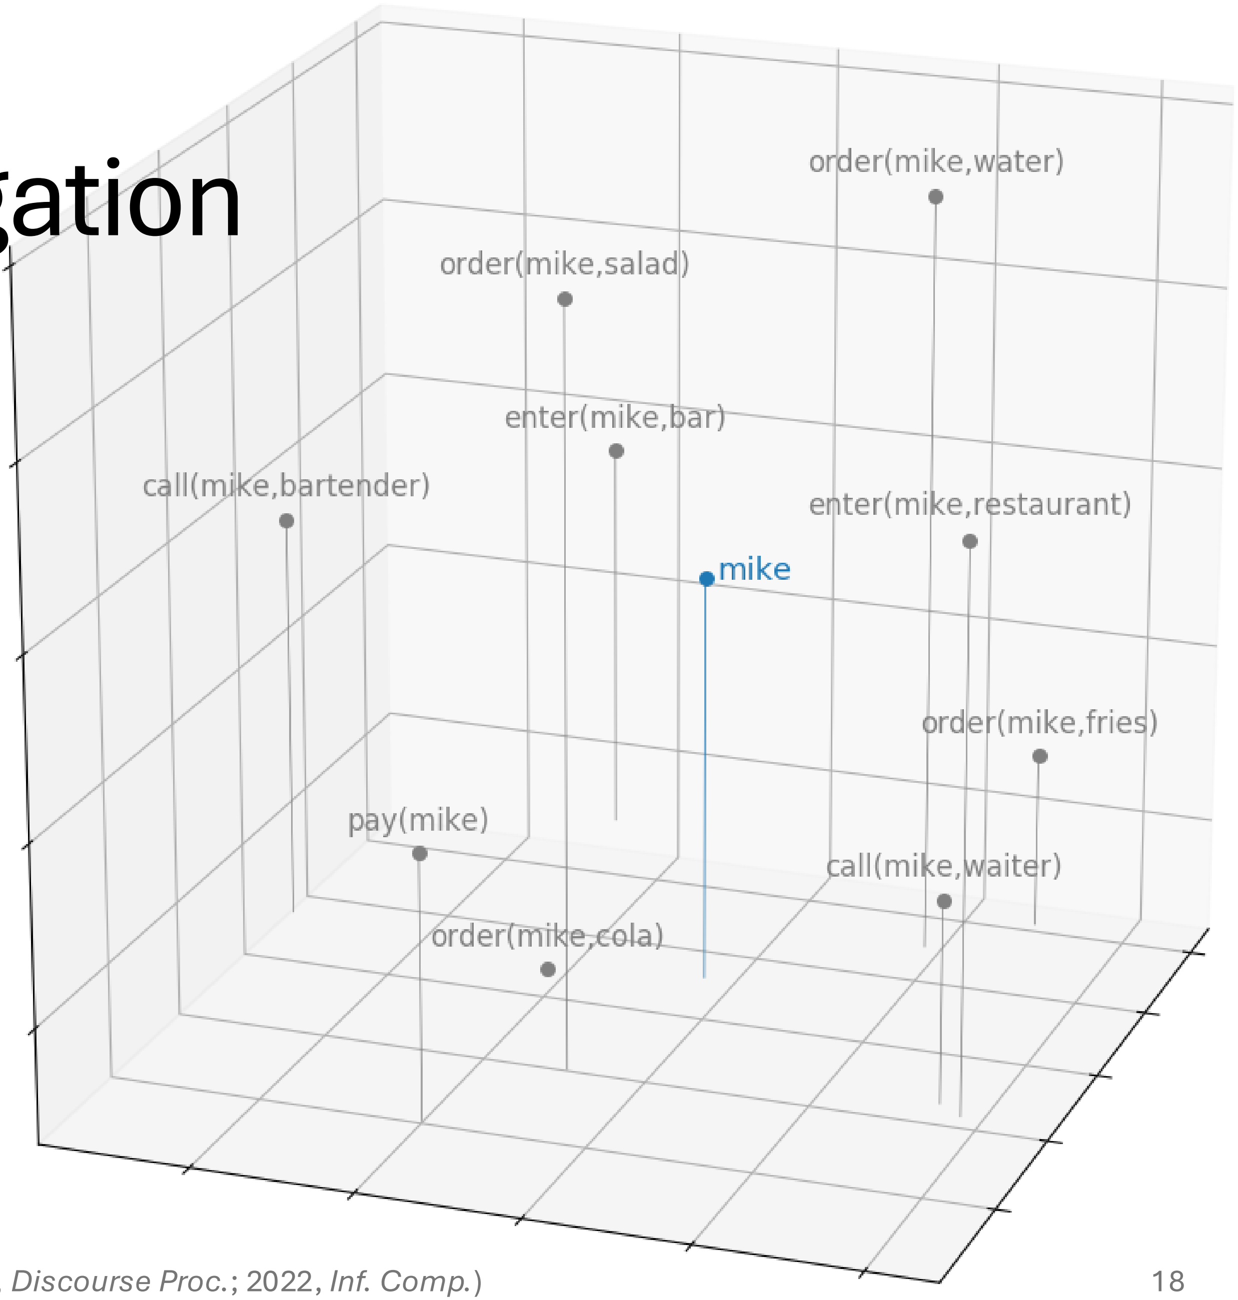

“Mike”

- Model-derived meaning abstracts over meanings of all propositions pertaining to *mike*

Venhuizen et al. (2019, *Discourse Proc.*; 2022, *Inf. Comp.*)

Meaning-space navigation

"Mike entered"

- Model navigates to a point that represents the contextualised meaning of "entered" given "mike"

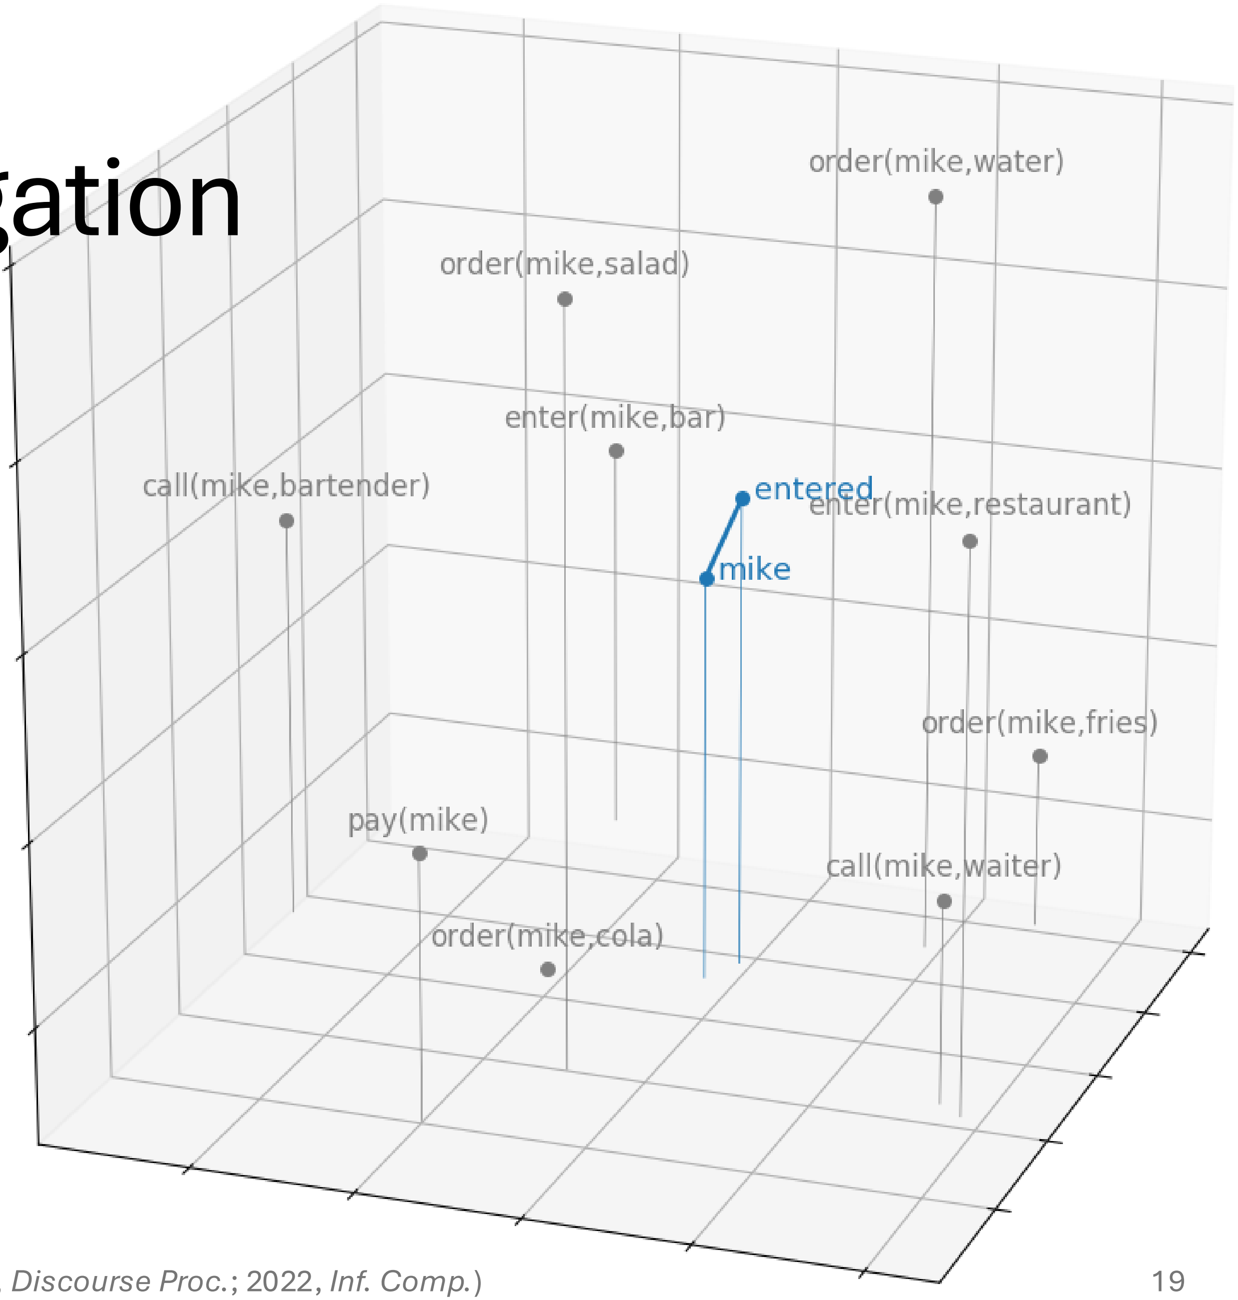

Venhuizen et al. (2019, *Discourse Proc.*; 2022, *Inf. Comp.*)

Meaning-space navigation

"Mike entered [the] bar"

- Approximates the propositional meaning vector for *enter(mike,bar)*

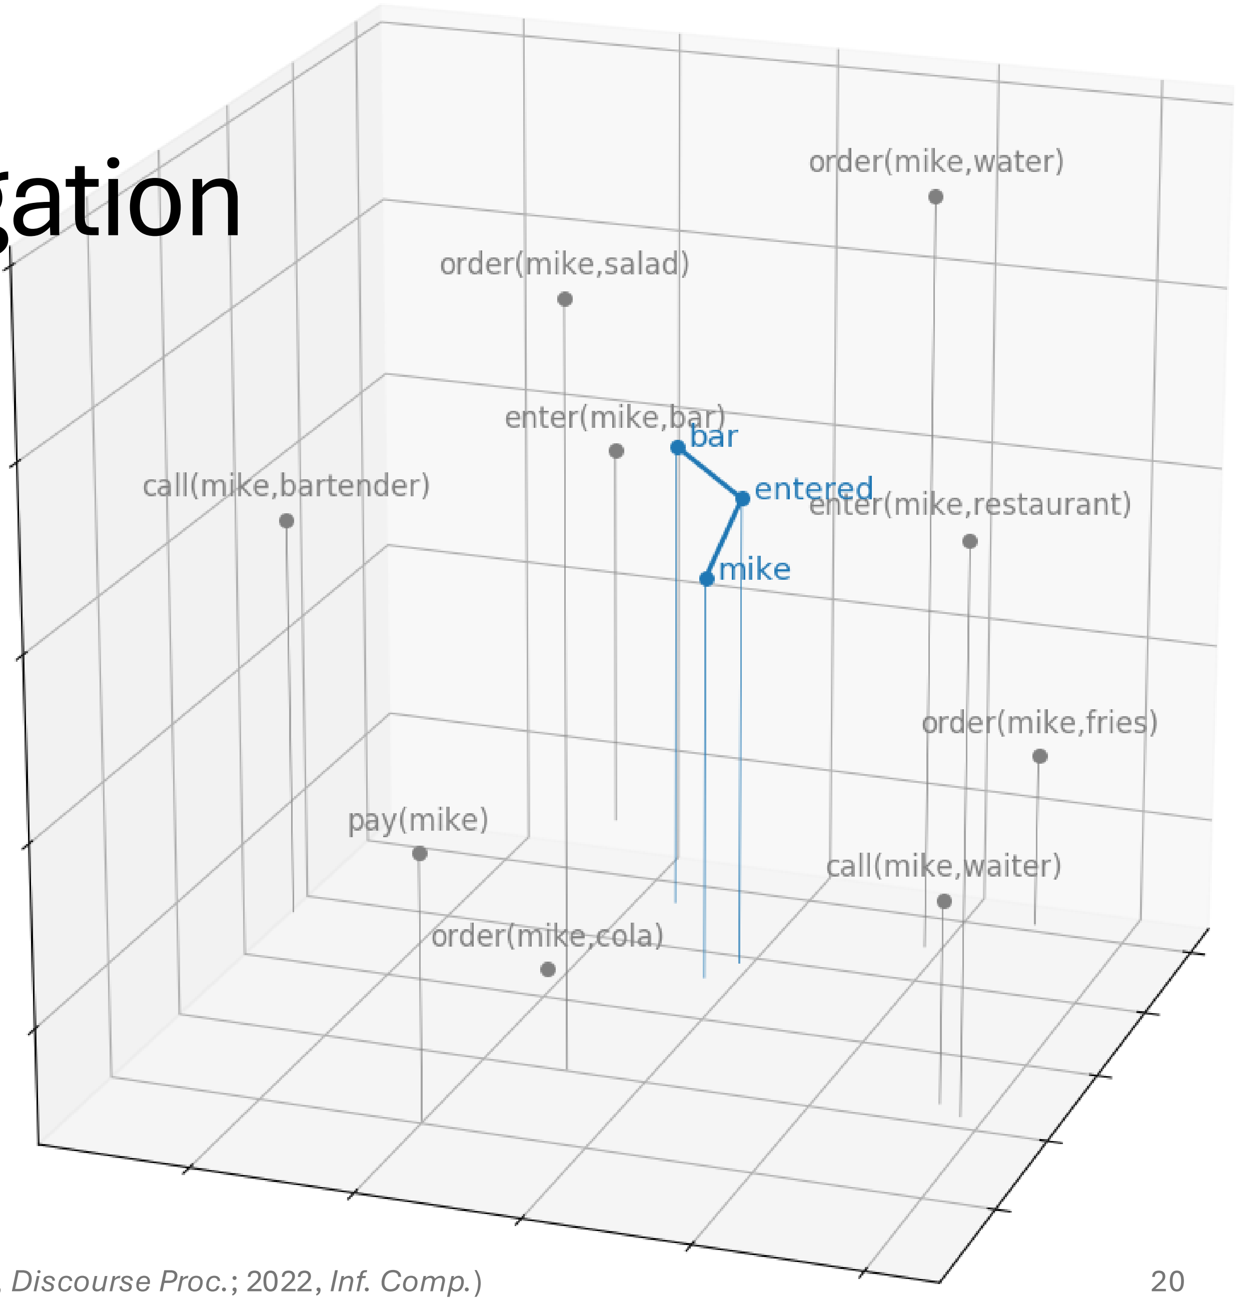

Venhuizen et al. (2019, *Discourse Proc.*; 2022, *Inf. Comp.*)

Meaning-space navigation

"Mike entered [the] bar [he] ordered"

- Close to *order* propositions that are typical given *enter(mike,bar)*

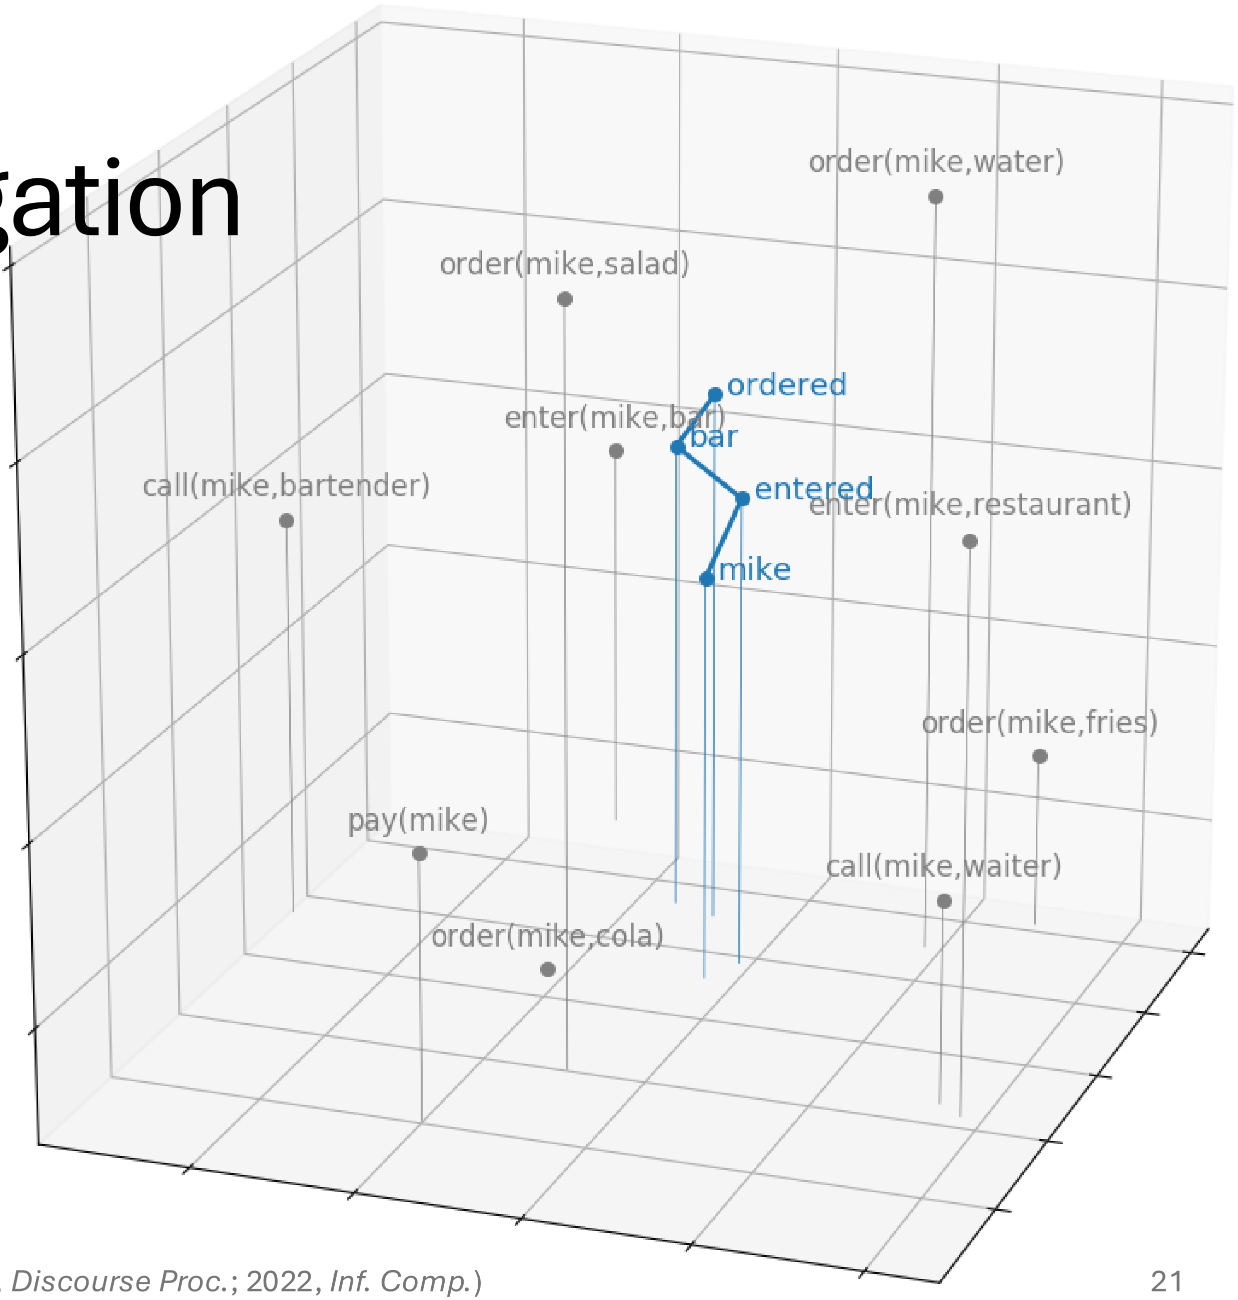

Venhuizen et al. (2019, *Discourse Proc.*; 2022, *Inf. Comp.*)

Meaning-space navigation

"Mike entered [the] bar [he] ordered cola"

- Transition in meaning space quantifies expectancy of continuation in terms of Surprisal (Shannon, 1948)

$$S(m_{ab}) = -\log P(b|a)$$

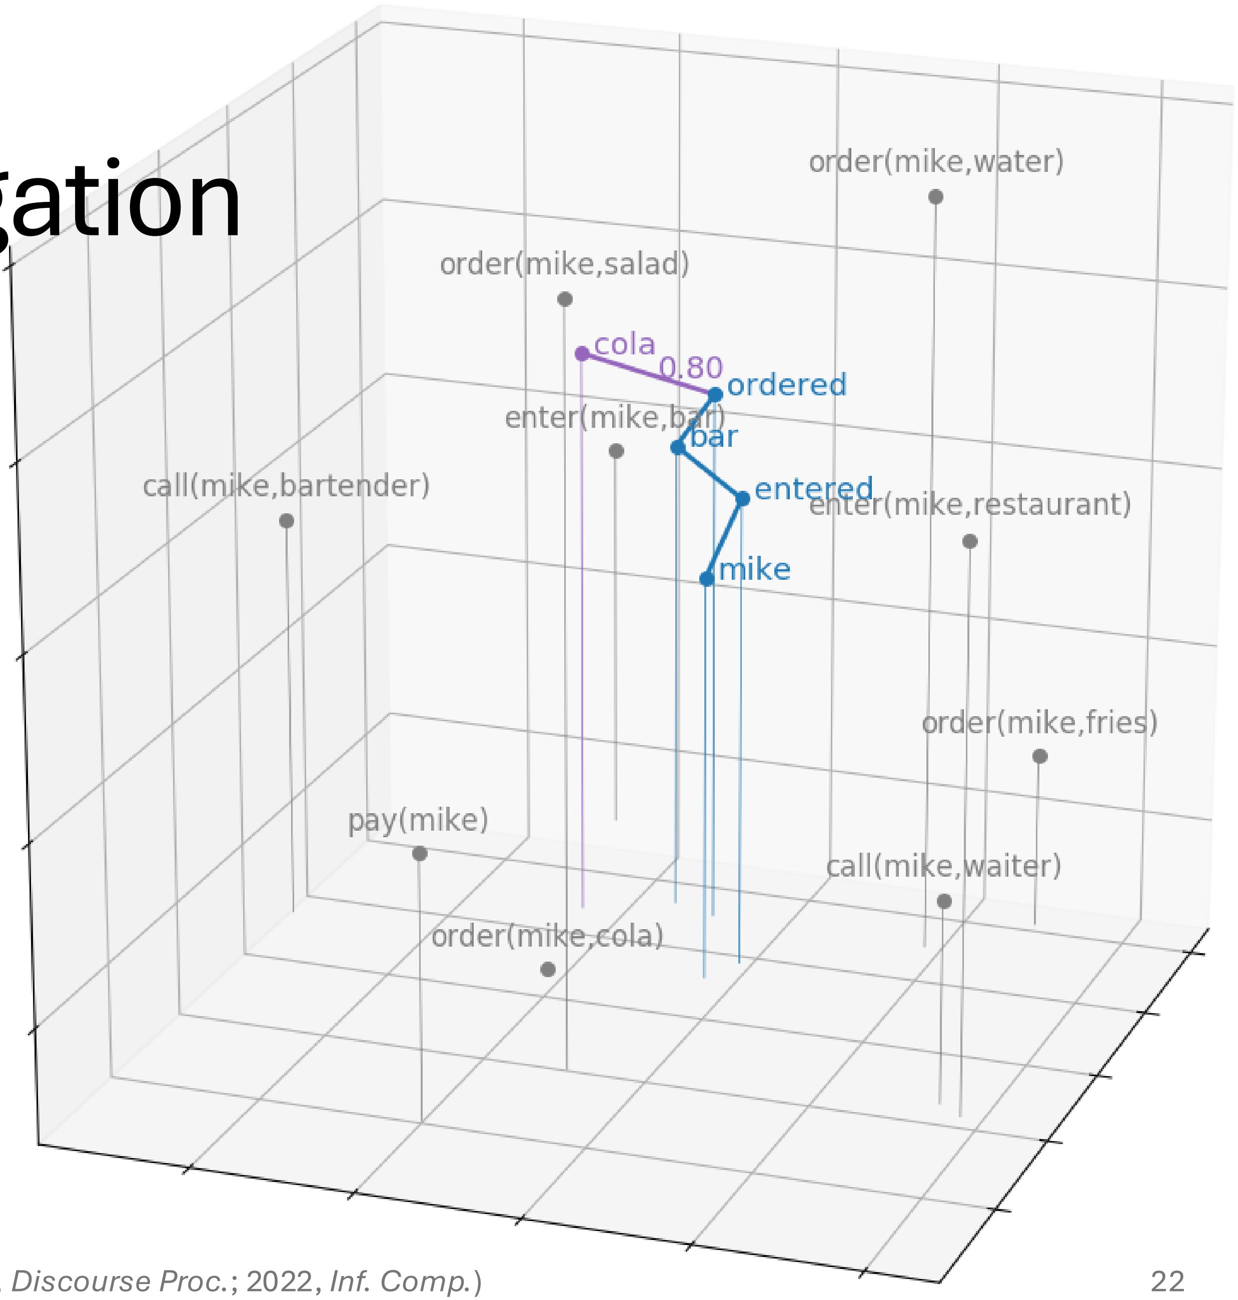

Venhuizen et al. (2019, *Discourse Proc.*; 2022, *Inf. Comp.*)

Meaning-space navigation

“Mike entered [the] bar [he] ordered fries”

- Larger (more surprising) transition reflects less expected continuation
- Expectancy derives from **linguistic experience** and/or **world knowledge**

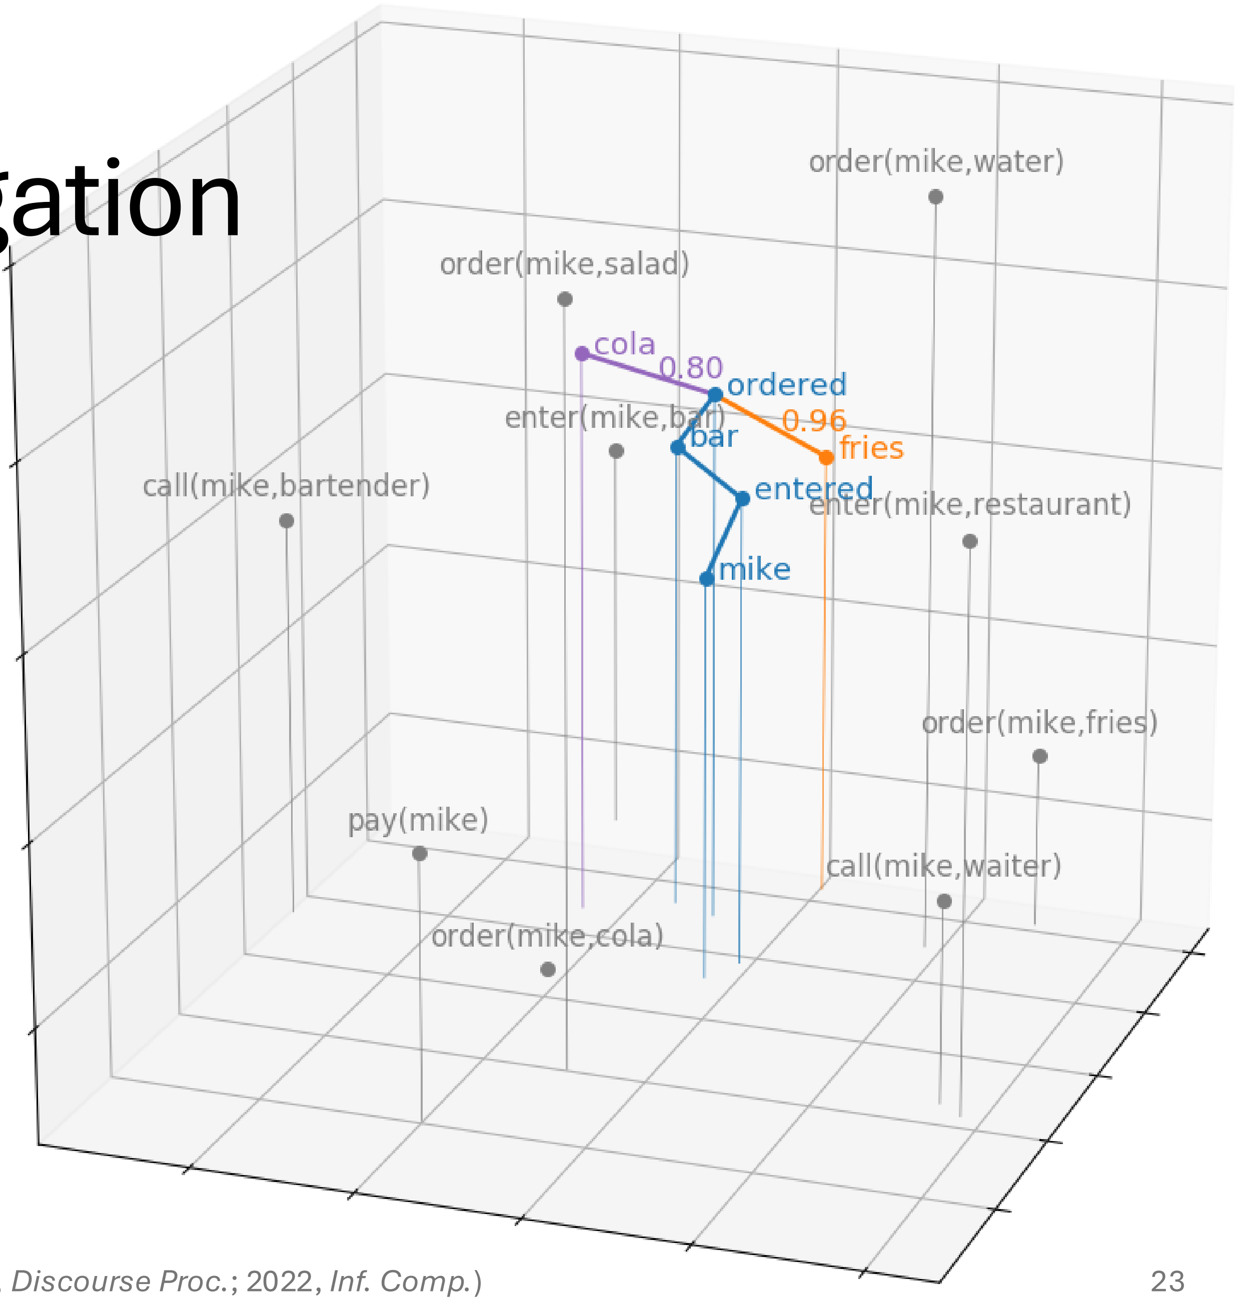

Venhuizen et al. (2019, *Discourse Proc.*; 2022, *Inf. Comp.*)

Testing RI model predictions

baseline: John entered the restaurant. Before long he opened the menu [...].
event-related implausible: John left the restaurant. Before long he opened the menu [...].
event-unrelated implausible: John entered the apartment. Before long he opened the menu [...].

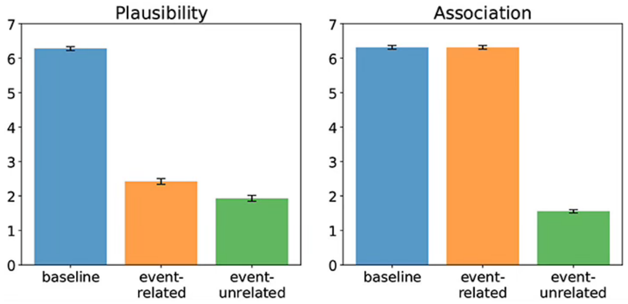

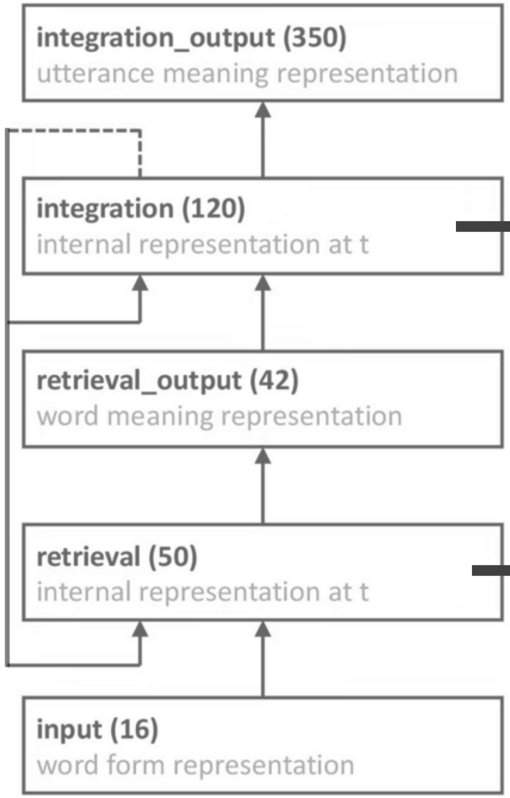

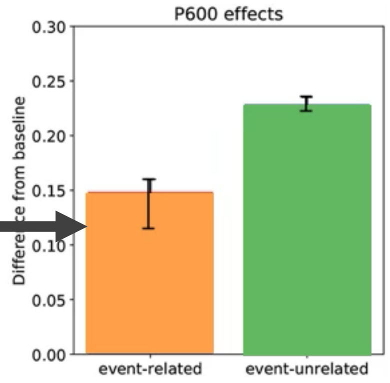

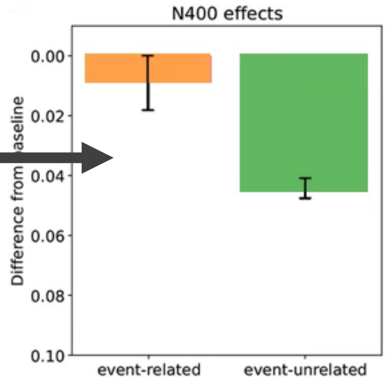

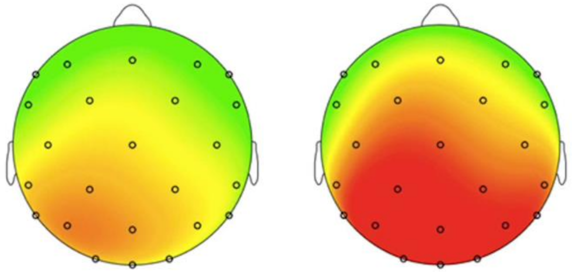

event-related *minus* baseline

event-unrelated *minus* baseline

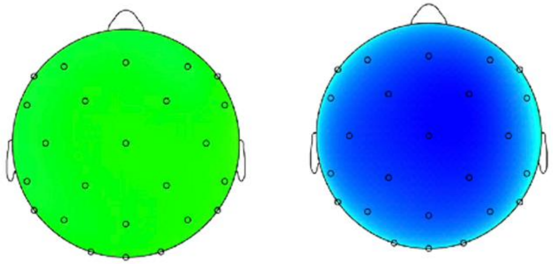

✅ **P600**

✅ **N400**

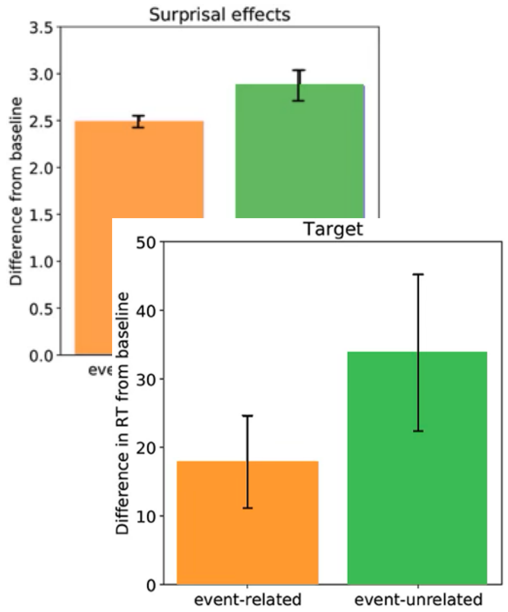

✅ **Reading times**

Referential Retrieval-Integration theory

- **Referential ambiguity effect: Nref**

  Sustained negativity observed in response to referential processing difficulty

- **Minimal RI extension**

  Retrieved **word meaning** = conceptual knowledge + referential knowledge

- **Referential knowledge**

  Binding vs. Accommodation

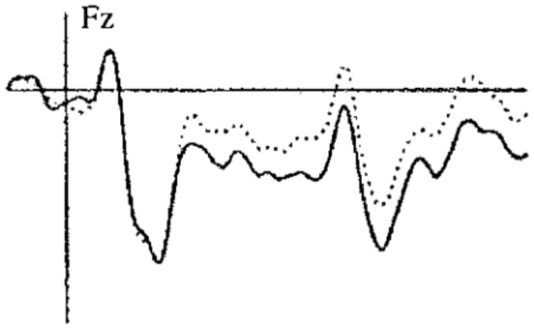

a. David told the two girls to clean up their room. … David told the girl …

x Y
x = David′
girls(Y)
tell(x, Y,“clean room”)

\+

v
x = David′
tell(x, v, …)
girl(v)
v=?

b. David told the boy and the girl to clean up their room. … David told the girl …

x y z
x = David′
boy(y)
girl(z)
tell(x, y,“clean room”)
tell(x, z,“clean room”)

\+

v
x = David′
tell(x, v, …)
girl(v)
v = z

Conclusion: From words to utterances

How does the meaning of complex utterances derive from the meaning of the individual words that constitute these utterances?

- **Linguistic perspective:** Lexical and compositional semantics offer complementary models of meaning
- **Neural perspective:** incremental language comprehension as Retrieval-Integration cycles

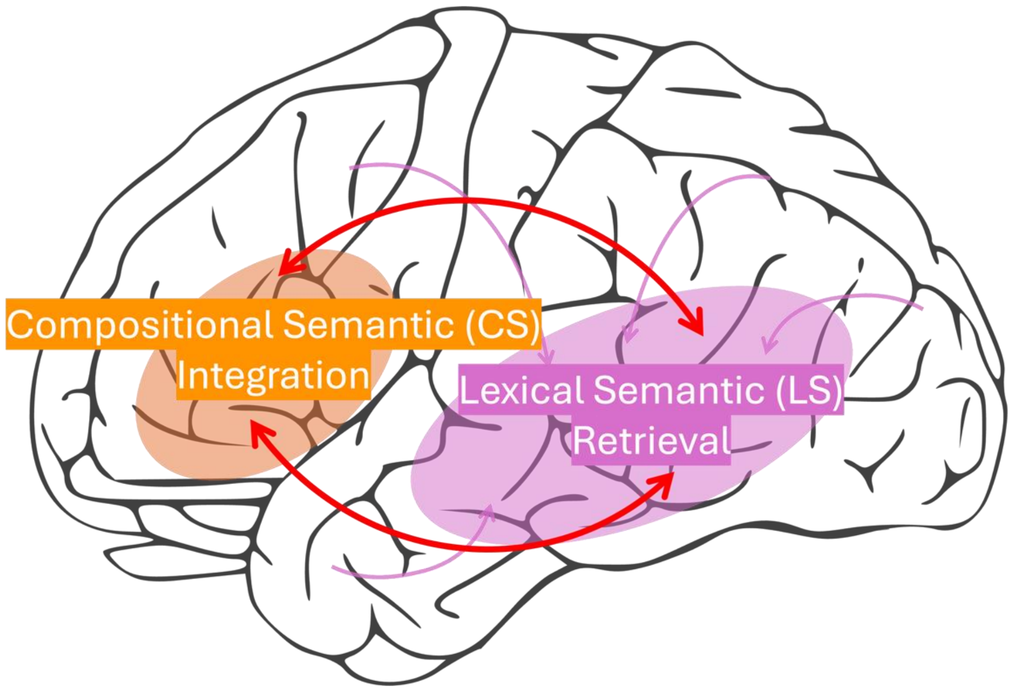

> Two models of meaning

Revisiting the principle of compositionality

Compositional integration is …

- **a non-linear mapping**
  - from LS meaning space into distinct CS space
- **continuous**
  - graded inferences in LS and CS spaces
- **expectation-based**
  - world knowledge and linguistic experience
- **spatiotemporally extended**
  - retrieval and integration are spatially extended as well as temporally overlapping