Semantic Theory
Week 10 – Incremental Meaning Construction

Noortje Venhuizen
Harm Brouwer

Universität des Saarlandes

Summer 2021

Distributional Formal Semantics

Noortje Venhuizen
Petra Hendriks
Matthew Crocker
Harm Brouwer

FROM MODELS TO MEANING SPACE




…
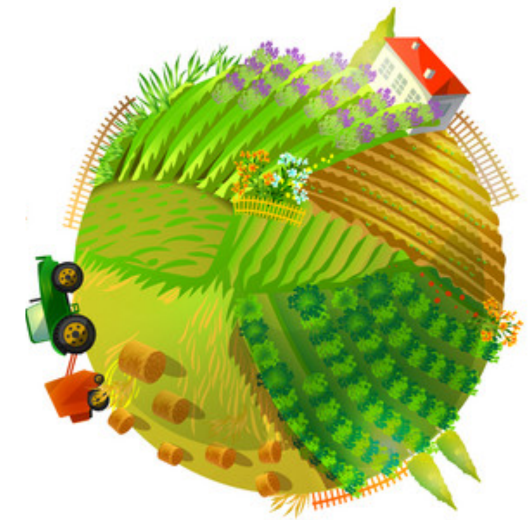

$M_1 = \langle U_1, V_1 \rangle$ — $p_1 \wedge \neg p_2 \wedge p_3 \wedge \ldots$

$M_2 = \langle U_2, V_2 \rangle$ — $\neg p_1 \wedge p_2 \wedge p_3 \wedge \ldots$

$M_3 = \langle U_3, V_3 \rangle$ — $\neg p_1 \wedge p_2 \wedge \neg p_3 \wedge \ldots$

$M_n = \langle U_n, V_n \rangle$ — $\neg p_1 \wedge \neg p_2 \wedge \neg p_3 \wedge \ldots$

- The set of models $\mathcal{M}_{\mathcal{P}}$ — describing states-of-affairs over propositions in $\mathcal{P}$ — defines a meaning space
- Propositional meaning defined by co-occurrence across models

Venhuizen, Hendriks, Crocker, Brouwer (2019, WoLLIC; 2021, Inf. Comp./ArXiv)

A FRAMEWORK FOR DISTRIBUTIONAL FORMAL SEMANTICS

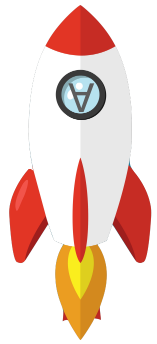

A meaning space for Distributional Formal Semantics

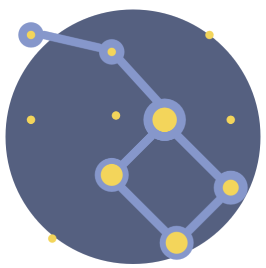

Formal properties of the meaning space

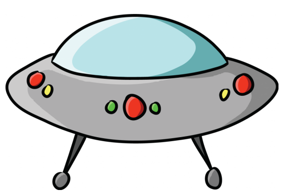

Incremental meaning construction

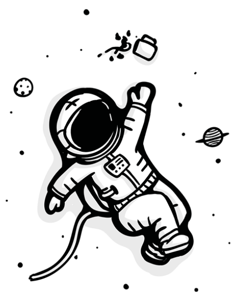

Semantic processing in the meaning space

A MODEL OF INCREMENTAL MEANING CONSTRUCTION

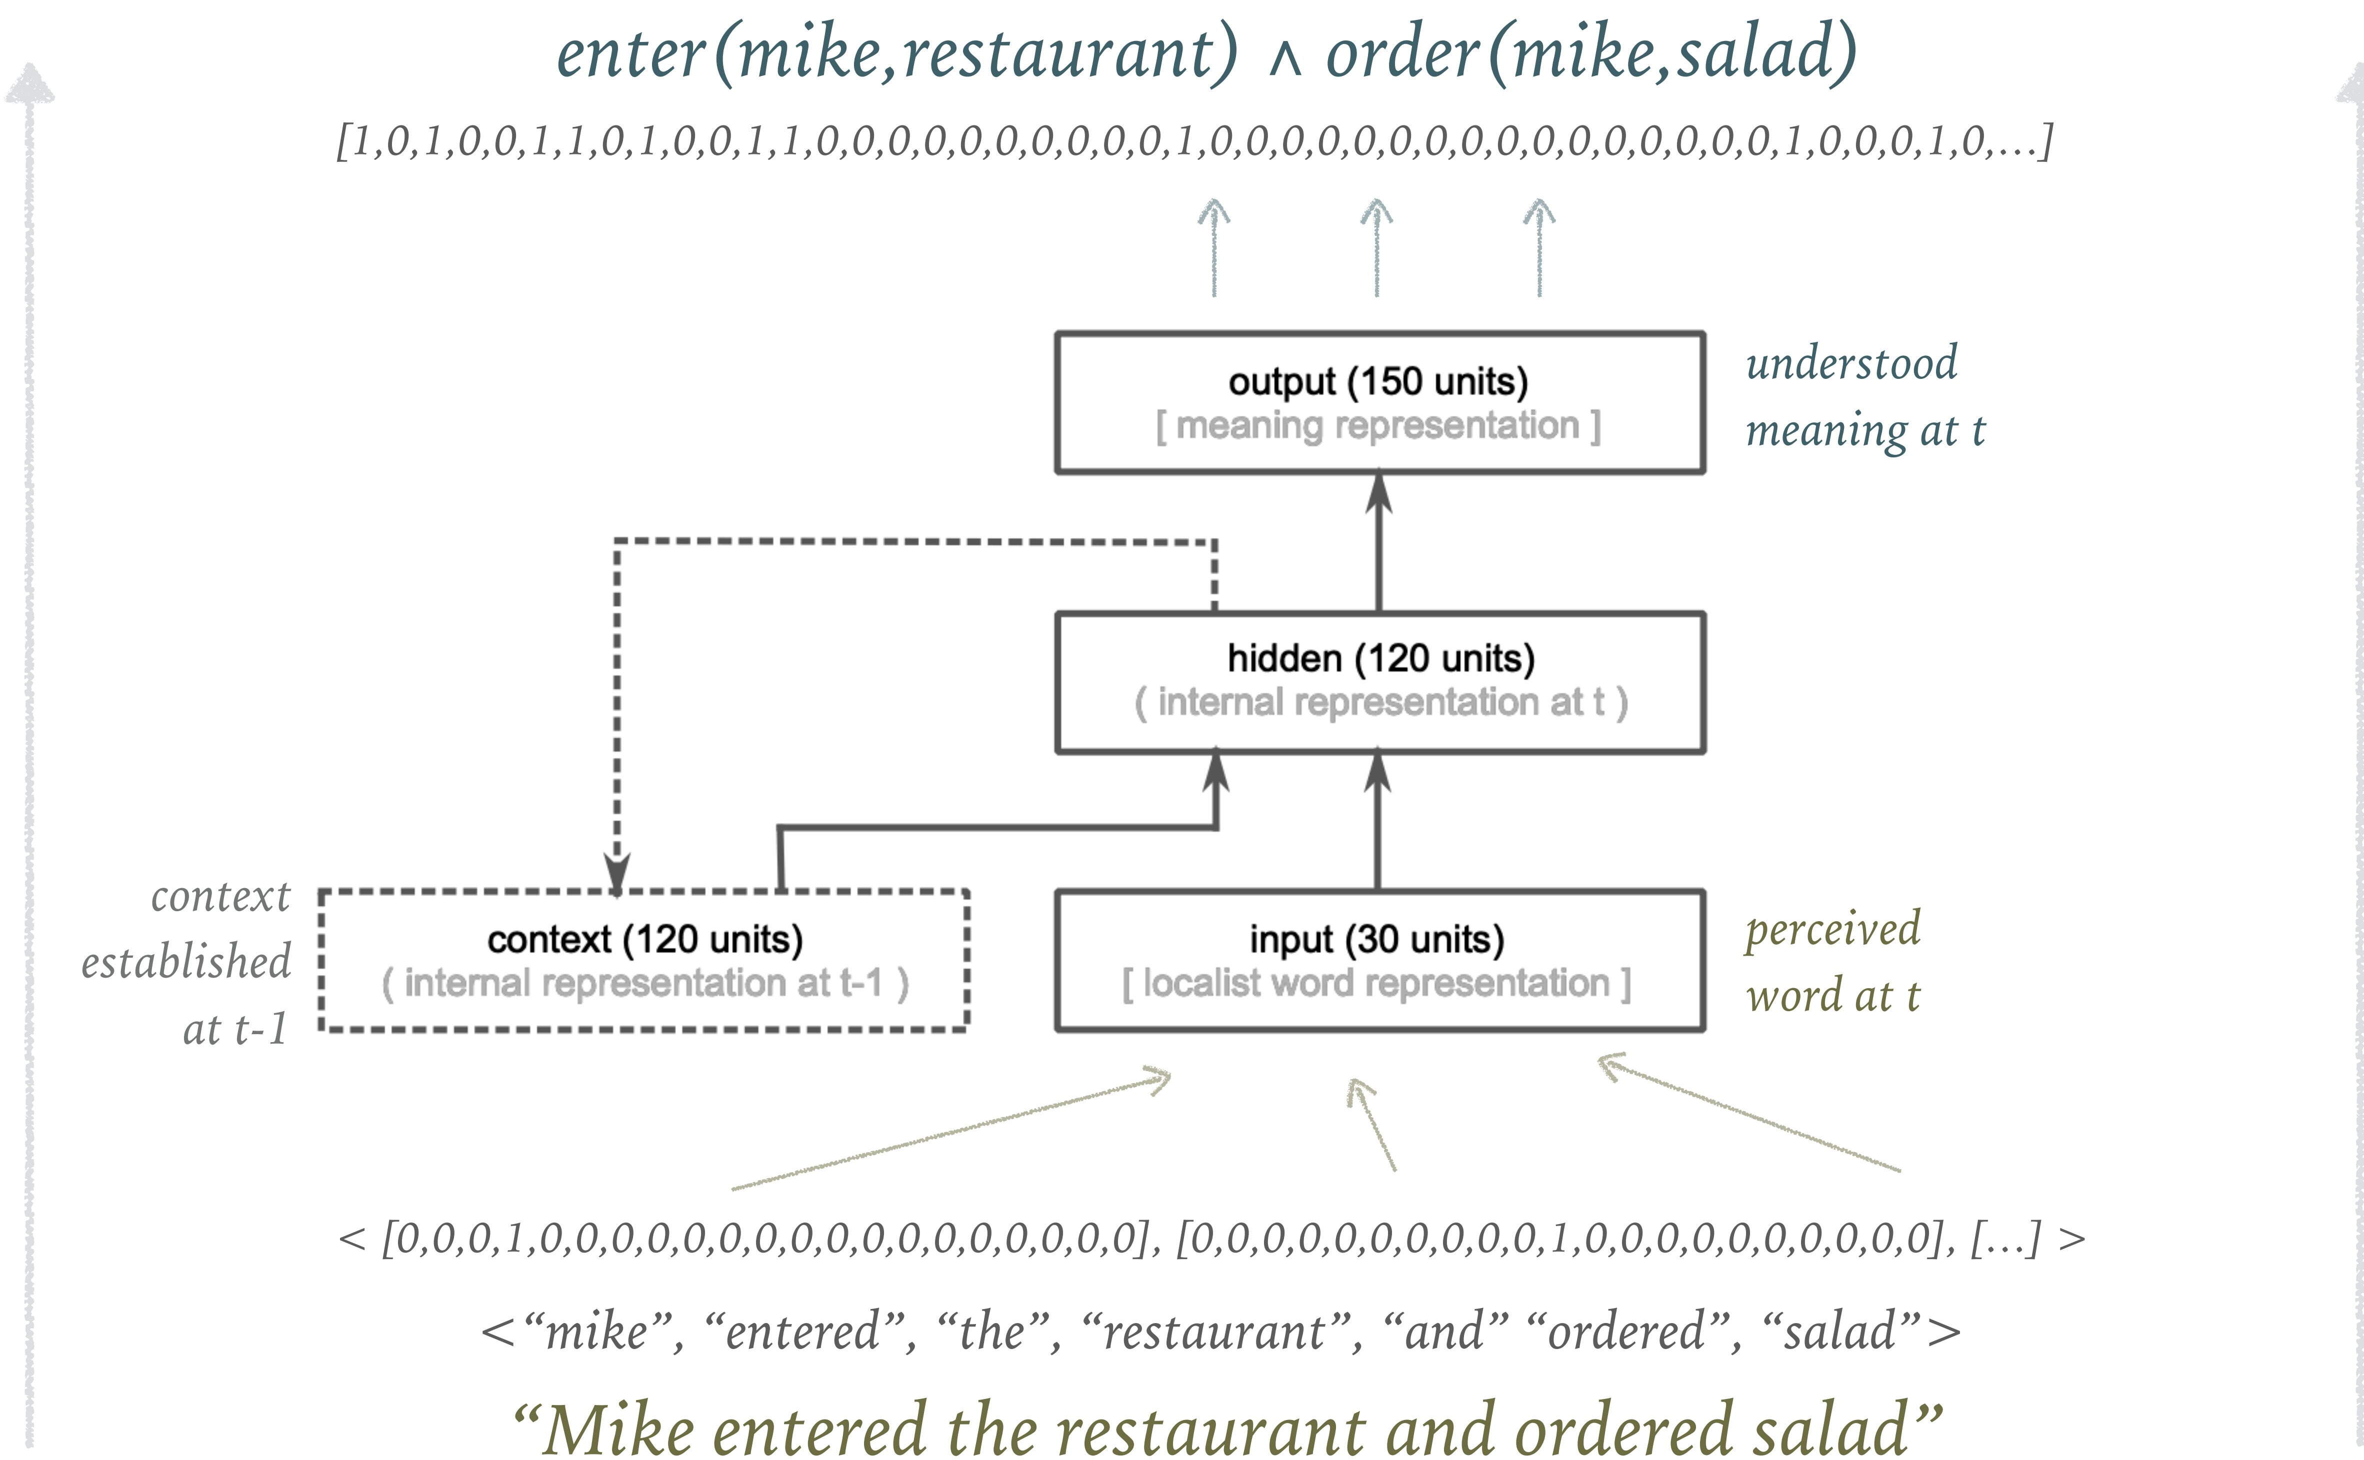

A PRIMER ON NEURAL NETWORKS

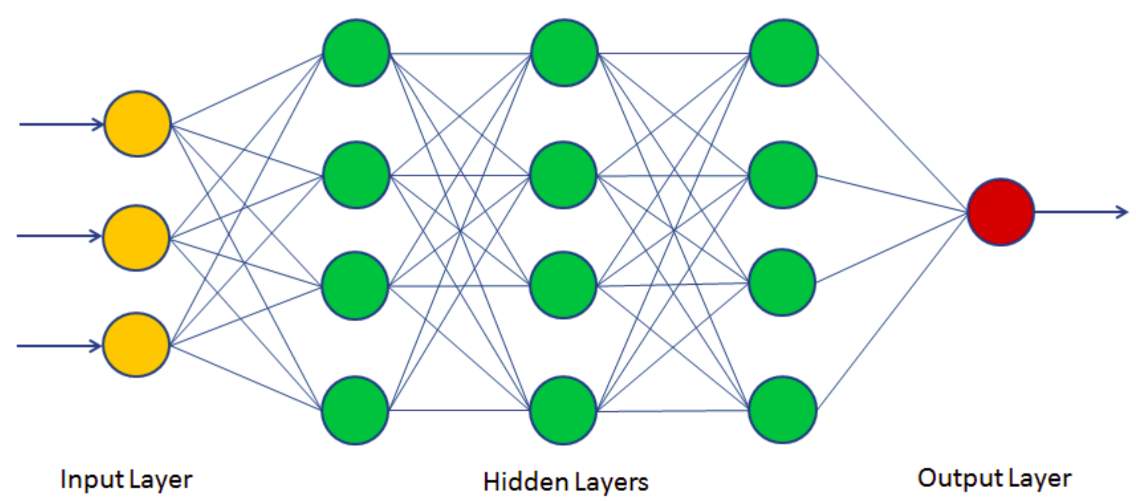

Biological Neurons

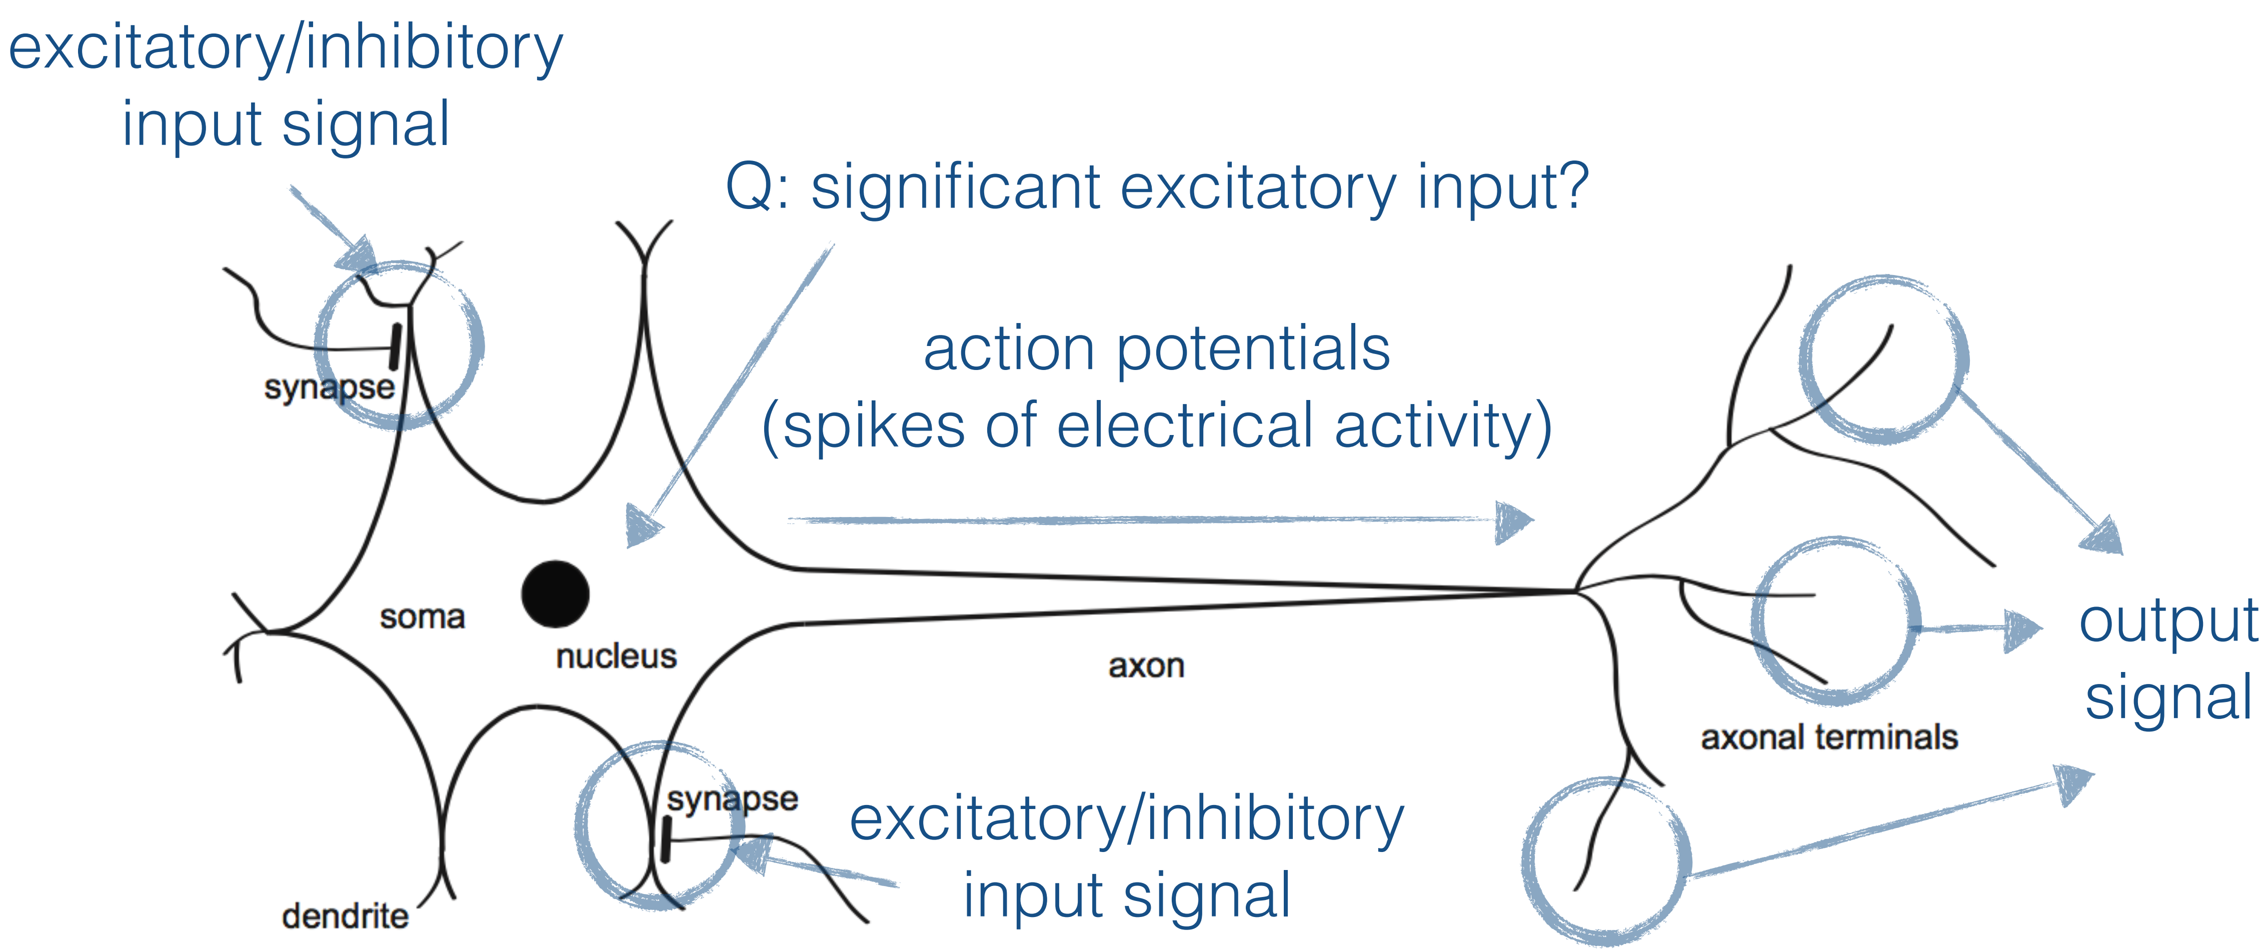

Figure A.1 | *Schematic overview of a biological neuron (or nerve cell).*

> synapses transform action potentials into an excitatory or inhibitory chemical signal

Artificial “Model” Neurons

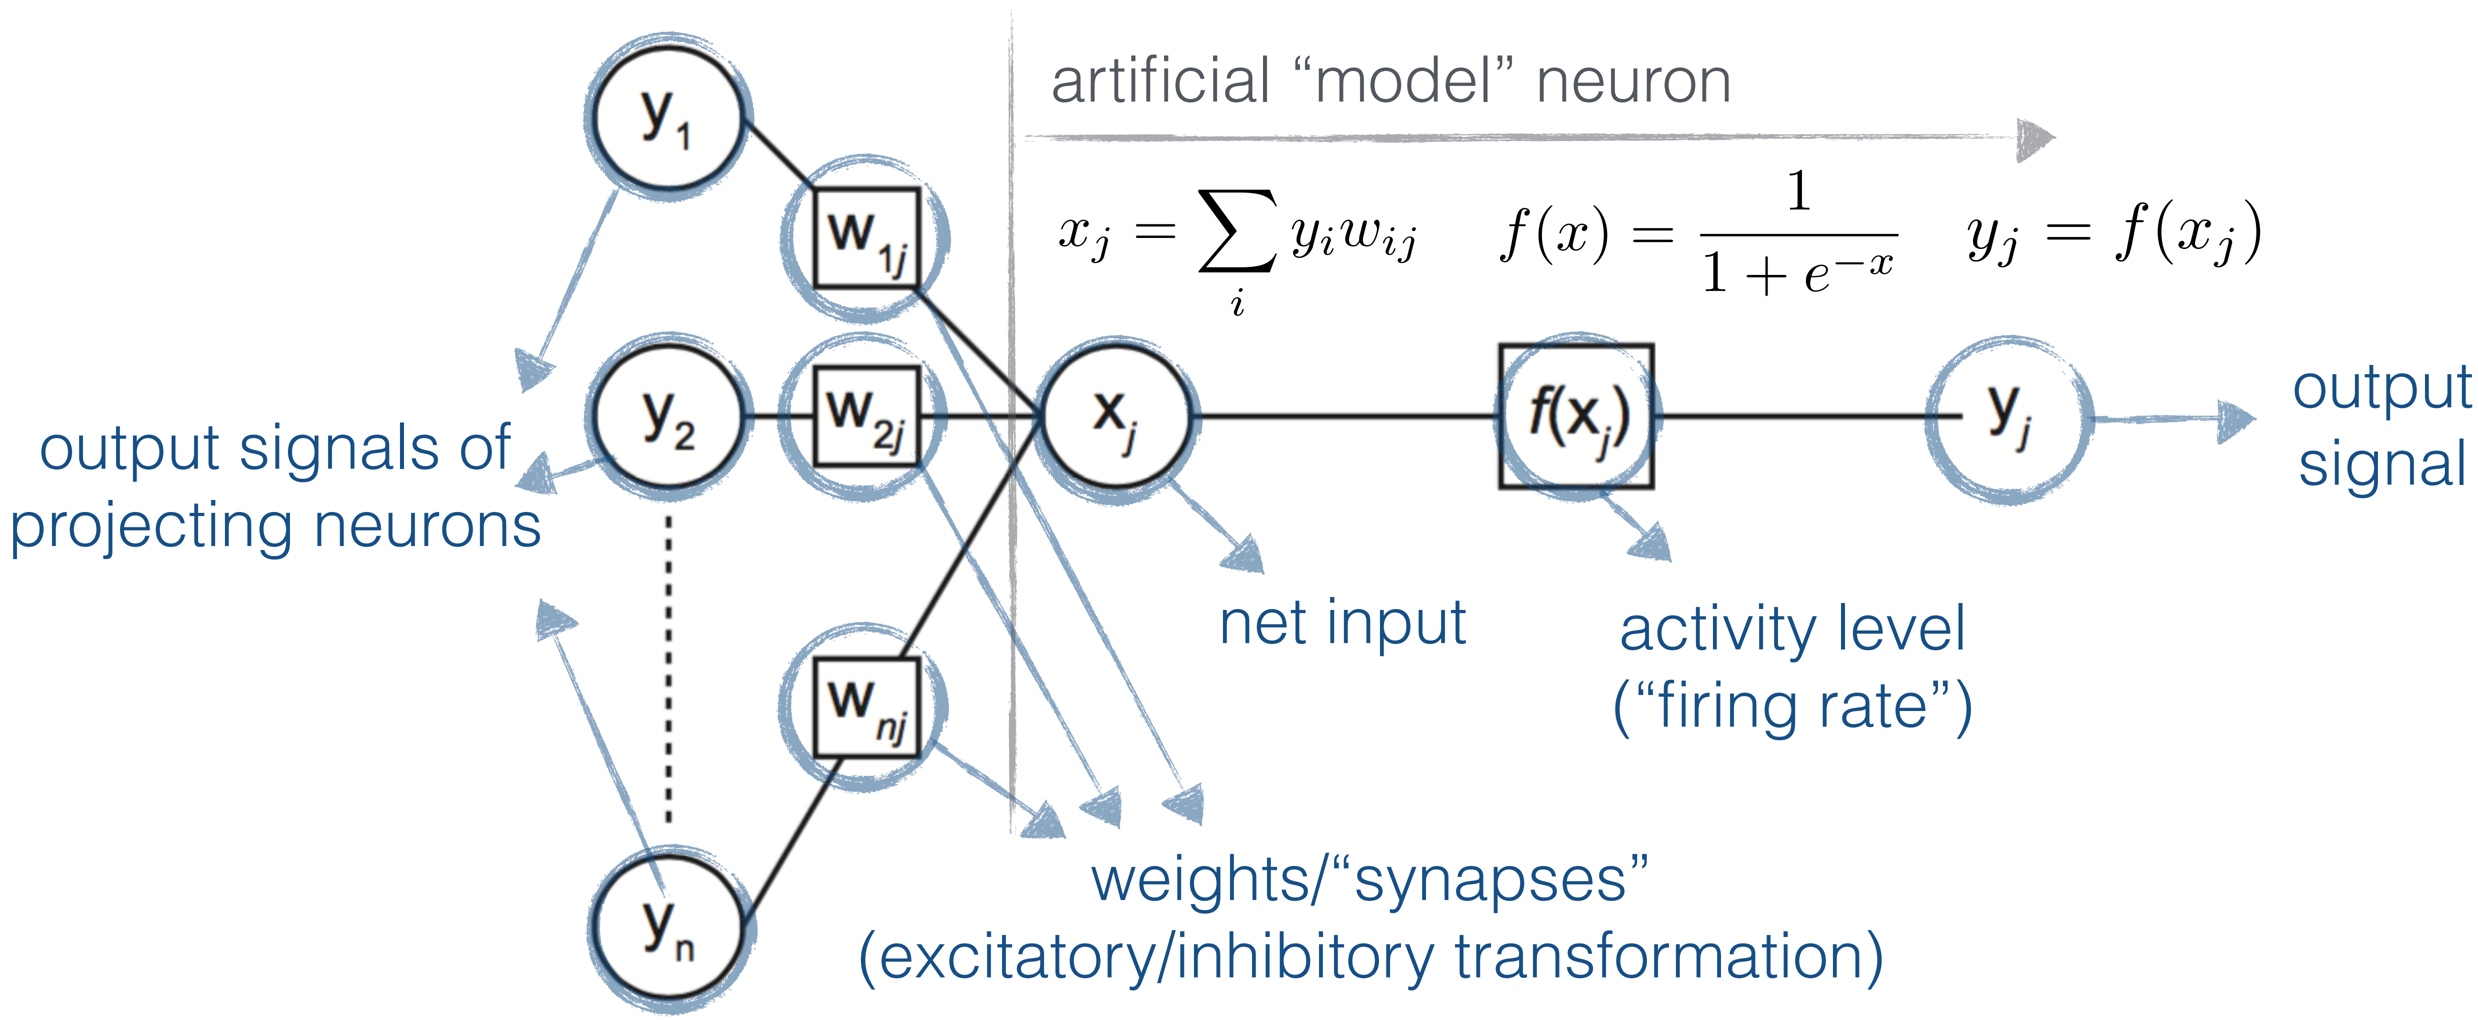

Figure A.2 | *Schematic overview of a unit (or model neuron). The activation level of the unit is a non-linear combination of its net input. The unit's net input, in turn, is the weighted sum of the activation levels of all units that signal to this unit.*

Artificial Neural Networks

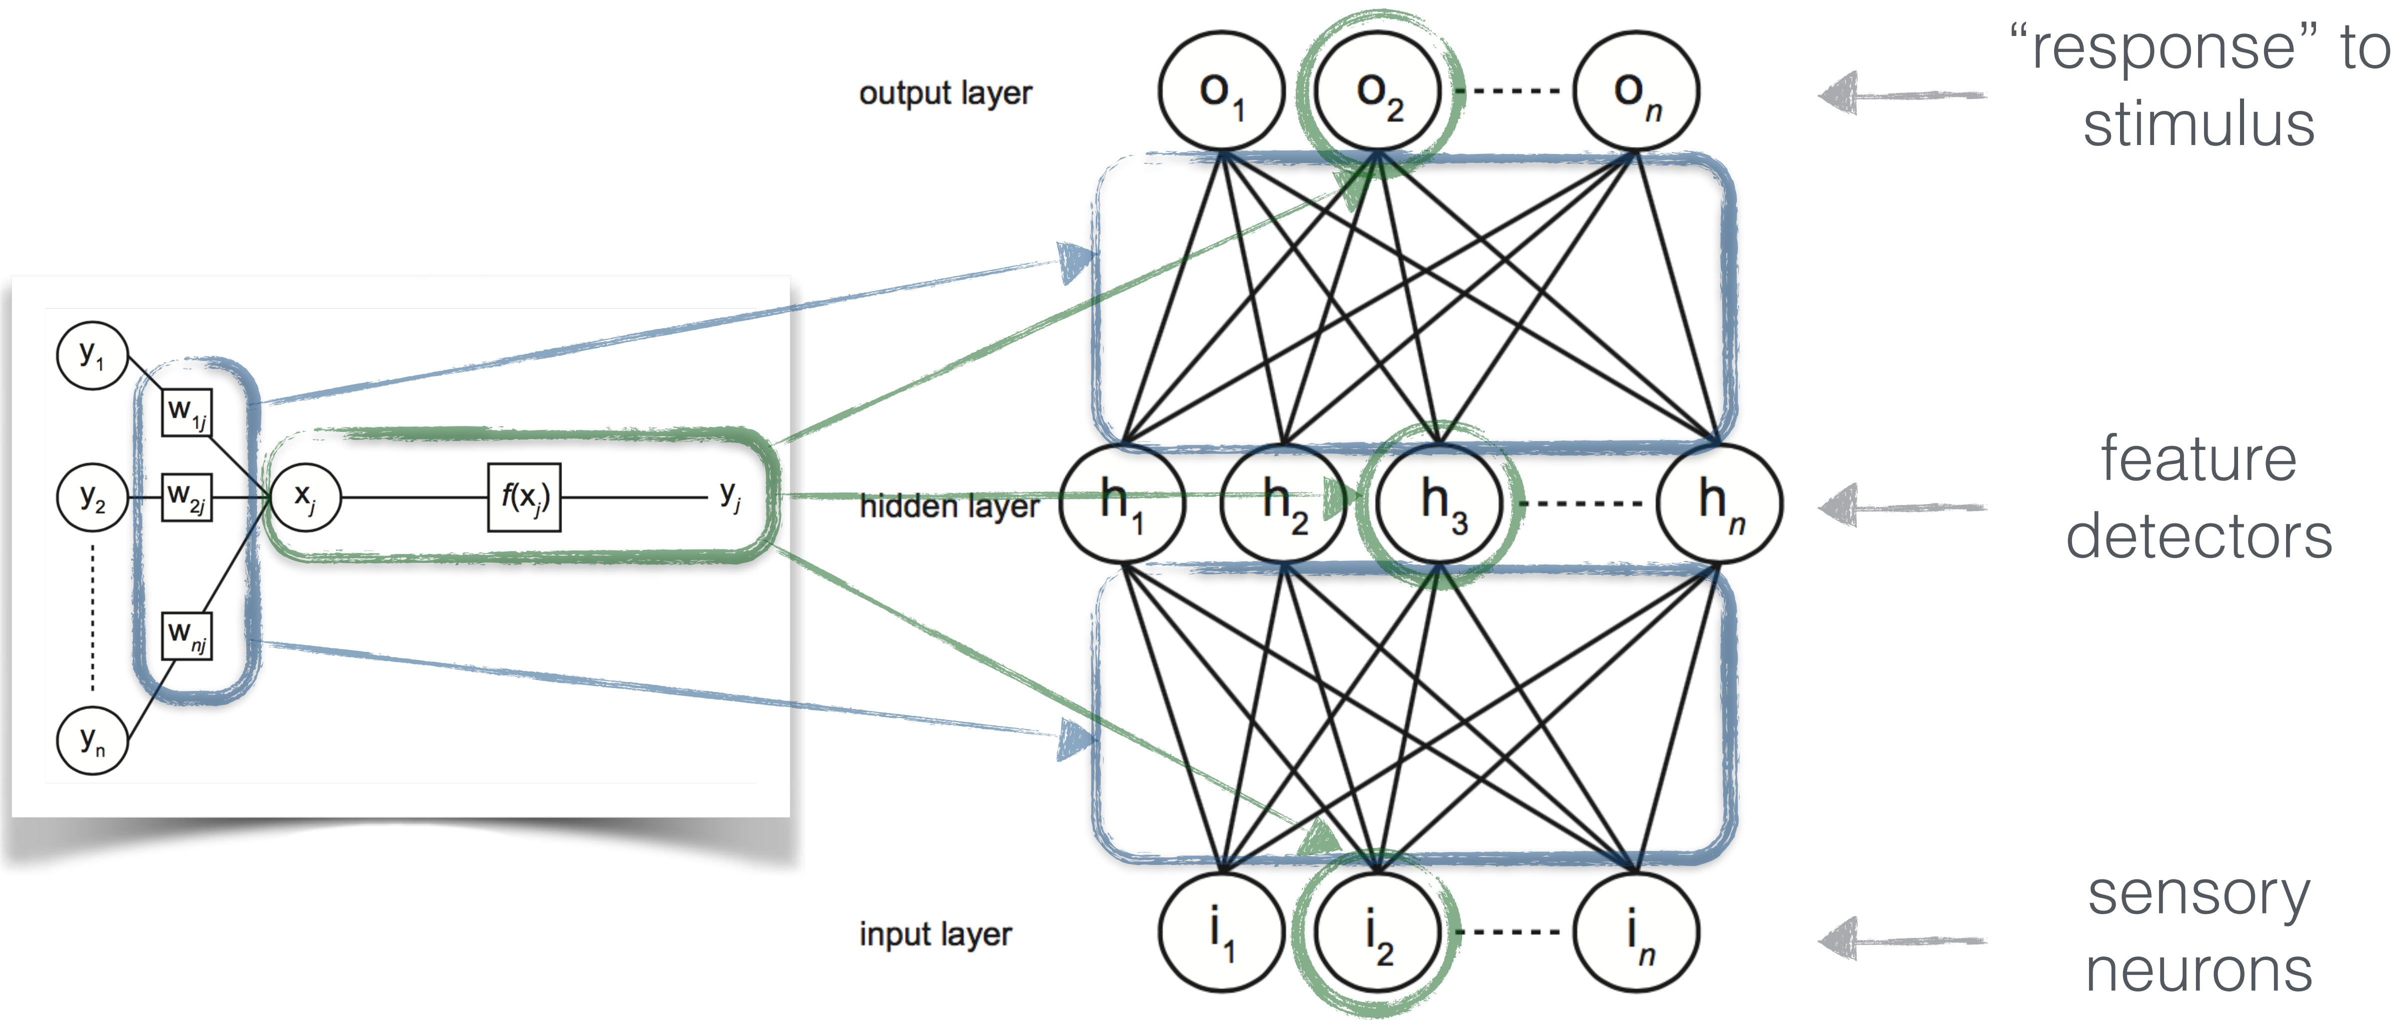

Figure A.3 | *A Feed Forward neural Network (FFN). Units in successive layers are fully connected, whereas units within layers are not.*

Recurrence—Modeling Memory

Q: What about **temporally extended stimuli** (e.g., sentences)?

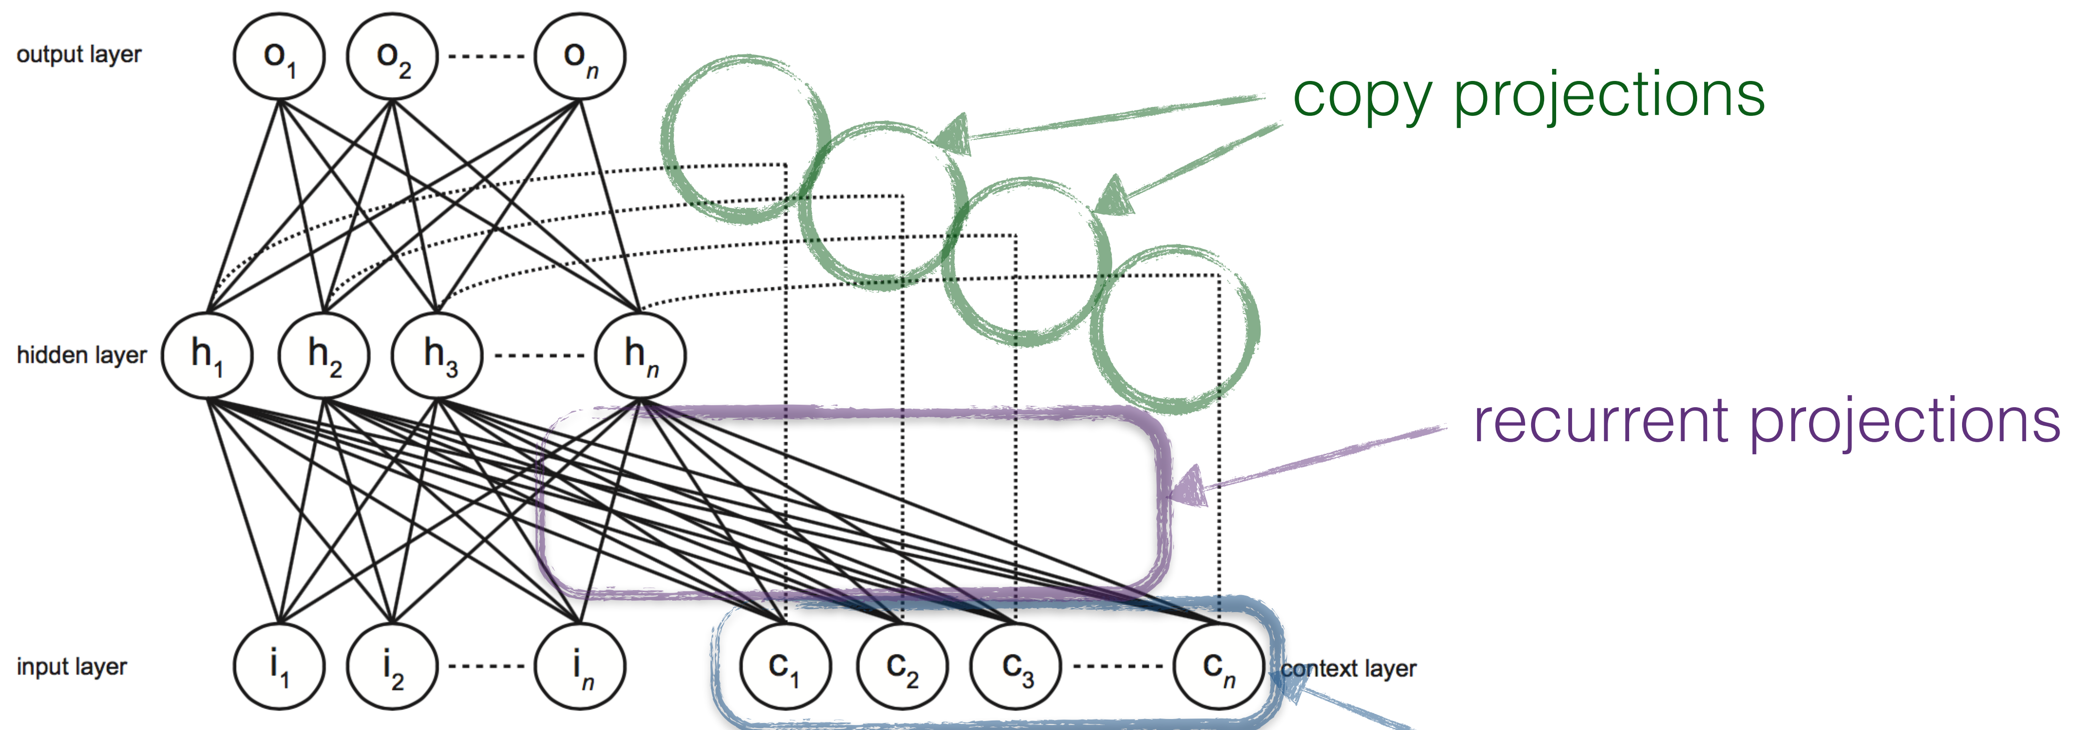

Figure A.5 | *A Simple Recurrent Network (SRN). Before being updated, the activation values of the units in the hidden layer are copied to their corresponding unit in the context layer (the fine dotted lines represent copy connections).*

> a **Simple Recurrent Network (SRN)** is a very powerful tool for cognitive modeling

Elman (1990), *Cogn. Sci.*

Learning in Neural Networks

> Neural Networks learn from experience (**training**)

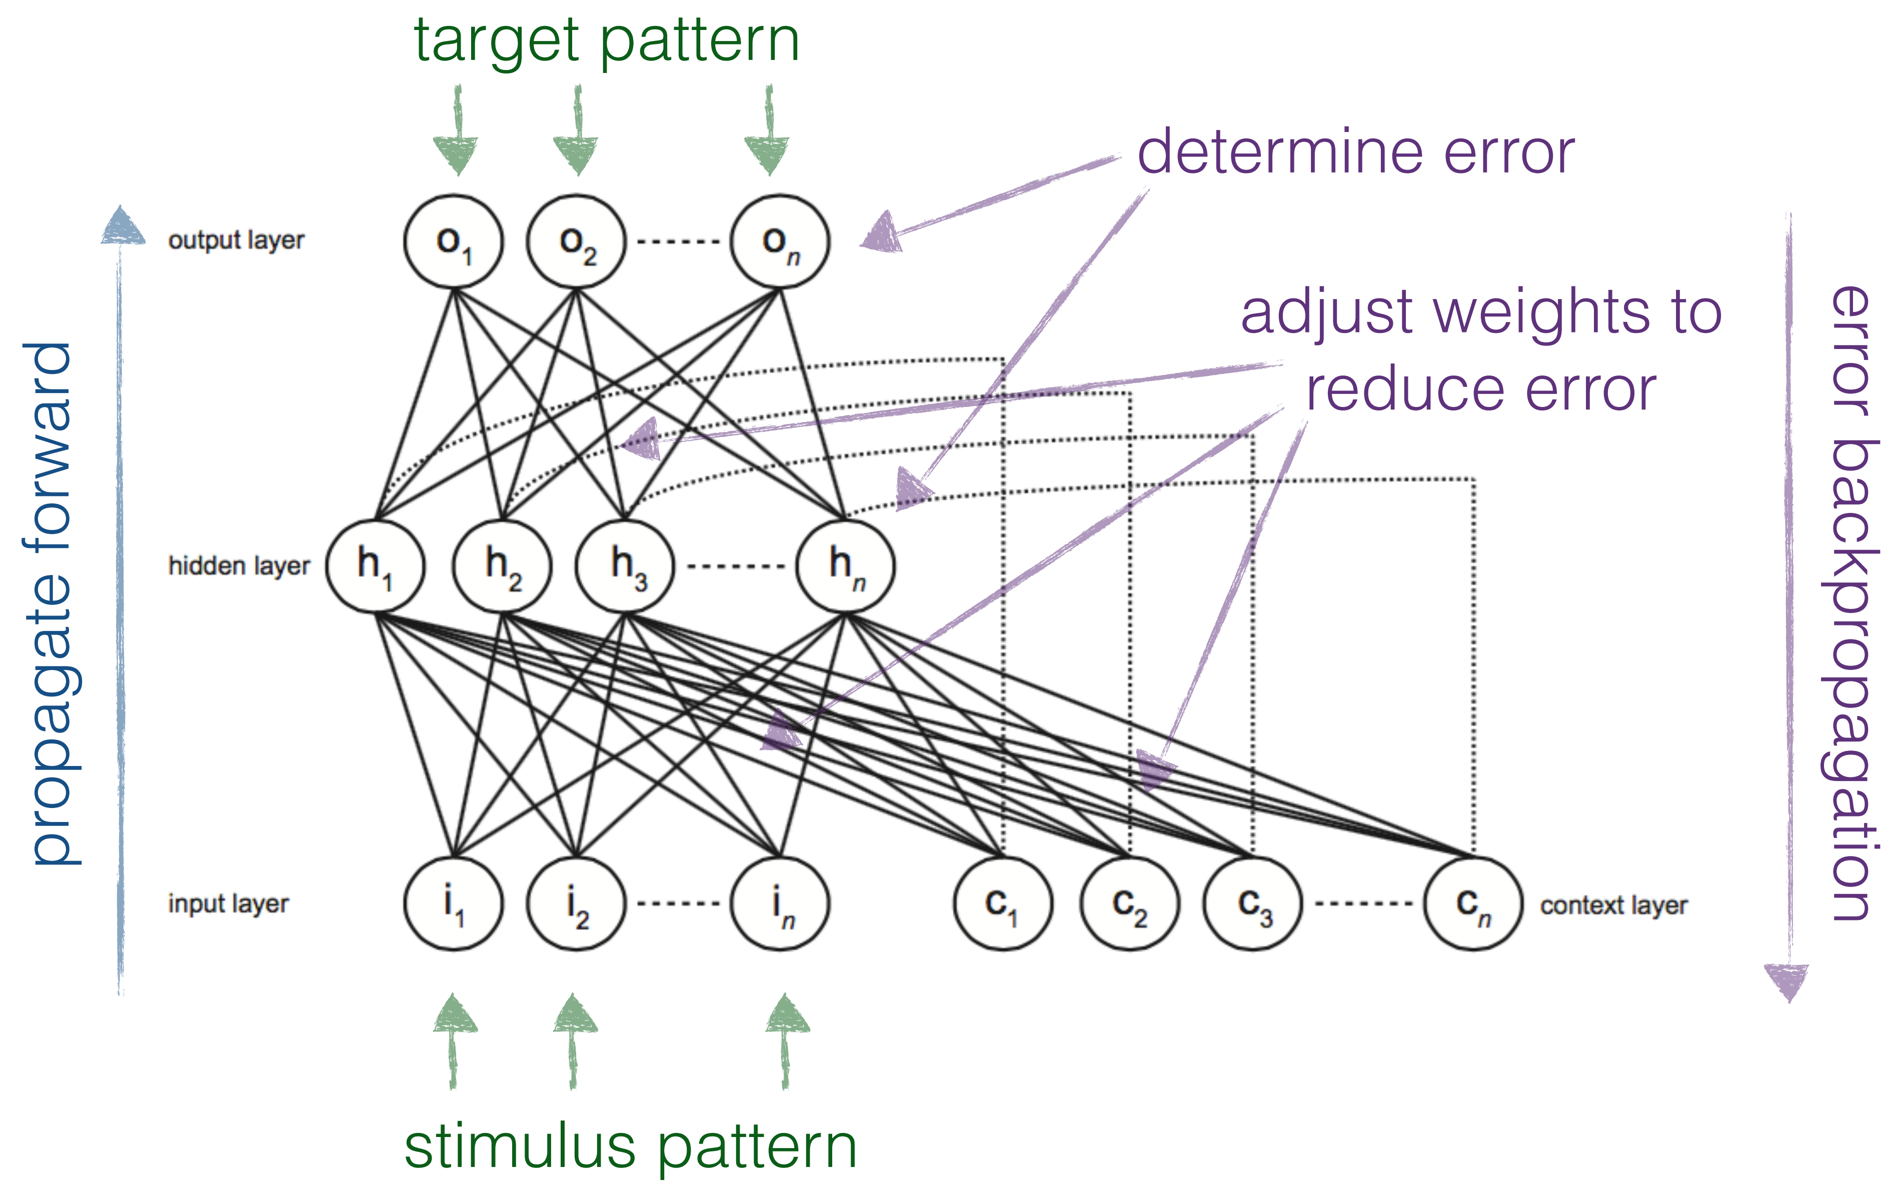

> challenge in neural network modeling is to **minimize error** for a set of stimuli

Rumelhart et al. (1986), *Nature*

INCREMENTAL MEANING CONSTRUCTION

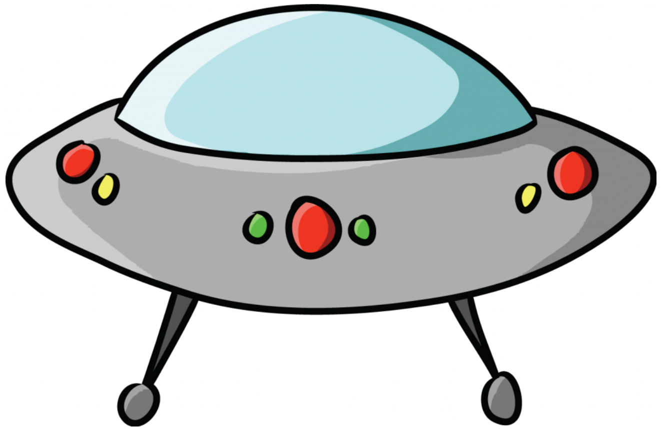

CONSTRUCTING THE MODEL: LANGUAGE

Utterance$_{[a/b]}$	**Semantics**$_a$	**Semantics**$_b$
P [entered/didn't enter] a L	$enter(P,L)$	$\neg(enter(P,L)) \wedge referent(P)$
P [entered/didn't enter] the L	$enter(P,L)$	$\neg(enter(P,L)) \wedge referent(P) \wedge referent(L)$
P [called/didn't call] the S	$call(P,S)$	$\neg(call(P,S)) \wedge referent(P) \wedge referent(S)$
P [ordered/didn't order] O	$order(P,O)$	$\neg(order(P,O)) \wedge referent(P)$
P [paid/didn't pay]	$pay(P)$	$\neg(pay(P)) \wedge referent(P)$
the S [arrived/didn't arrive]	$arrive(S)$	$\neg(arrive(S)) \wedge referent(S)$
the S [brought/didn't bring] O	$bring(S,O)$	$\neg(bring(S,O)) \wedge referent(S)$
[P/someone] entered the L he/she ordered O	$enter(P,L) \wedge order(P,O)$	$\exists x_{m/f}(enter(x,L) \wedge order(x,O))$
[P/someone] entered the L he/she called the S	$enter(P,L) \wedge call(P,S)$	$\exists x_{m/f}(enter(x,L) \wedge call(x,S))$
[P/someone] called the S he/she ordered O	$call(P,S) \wedge order(P,O)$	$\exists x_{m/f}(call(x,S) \wedge order(x,O))$
[P/someone] called the S he/she paid	$call(P,S) \wedge pay(P)$	$\exists x_{m/f}(call(x,S) \wedge pay(x))$
[P/someone] called the S he brought O	$call(P,S) \wedge bring(S,O)$	$\exists x(call(x,S) \wedge bring(S,O))$
[P/someone] called the S he arrived	$call(P,S) \wedge arrive(S)$	$\exists x(call(x,S) \wedge arrive(S))$

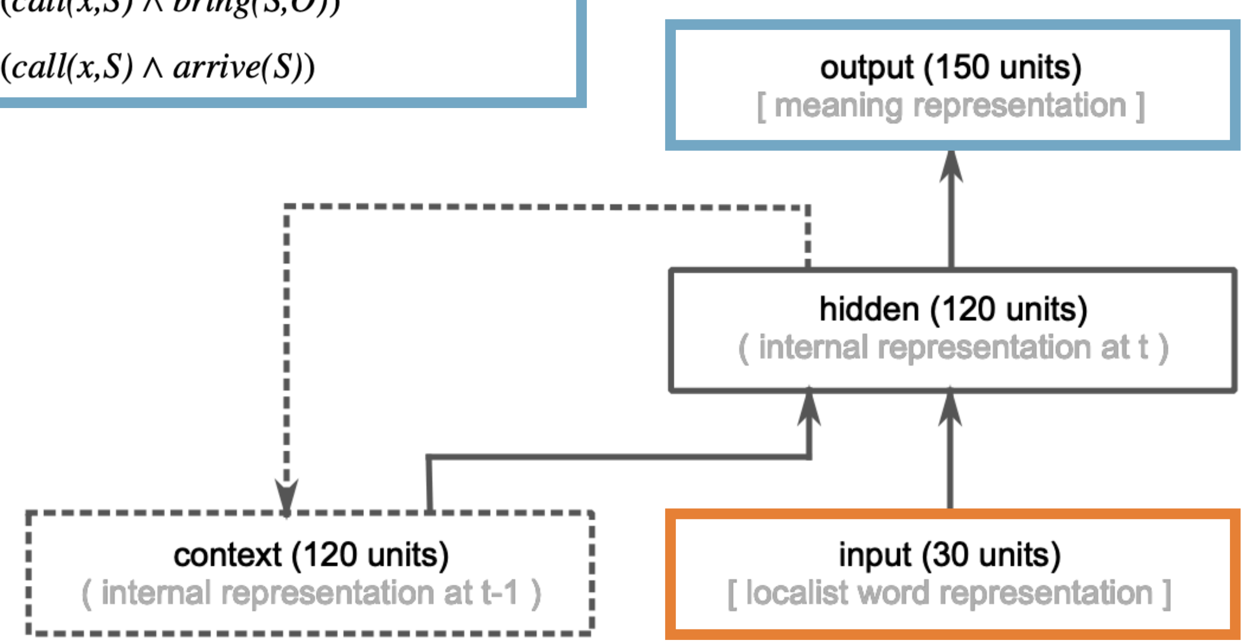

Venhuizen, Hendriks, Crocker, Brouwer (2021, Inf. Comp./ArXiv)

MEANING SPACE NAVIGATION

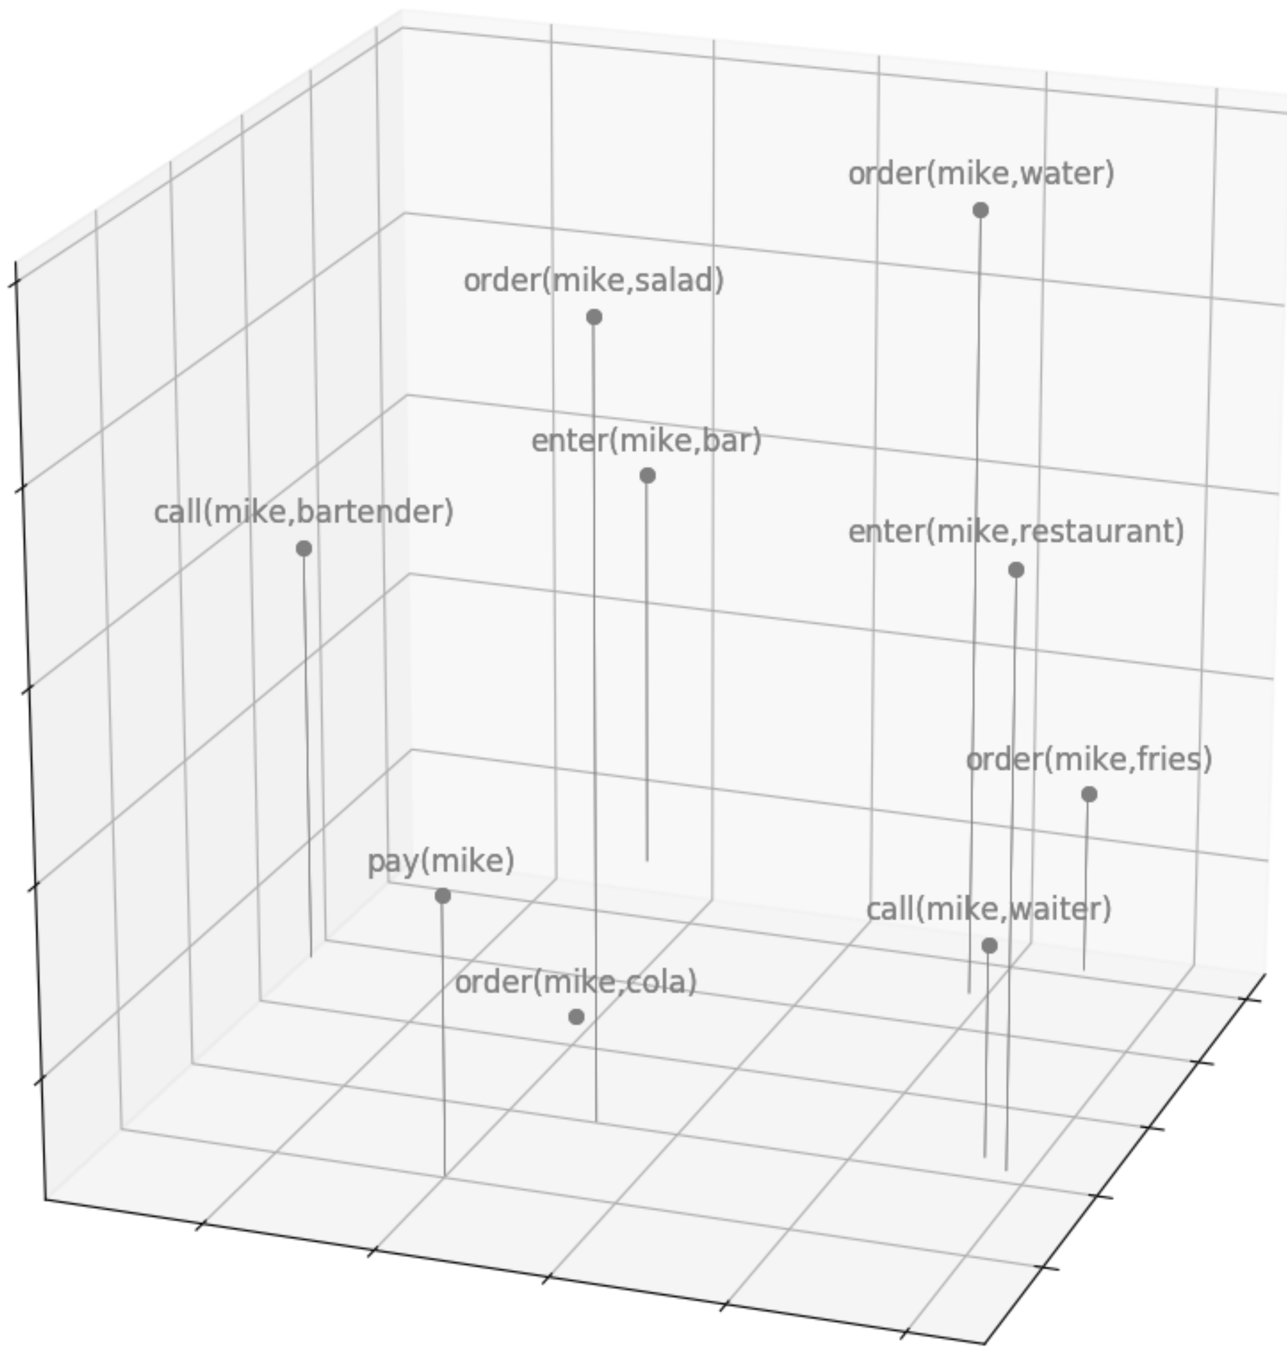

- Propositions that co-occur frequently in $\mathcal{M}$ are positioned close to each other in space

MEANING SPACE NAVIGATION

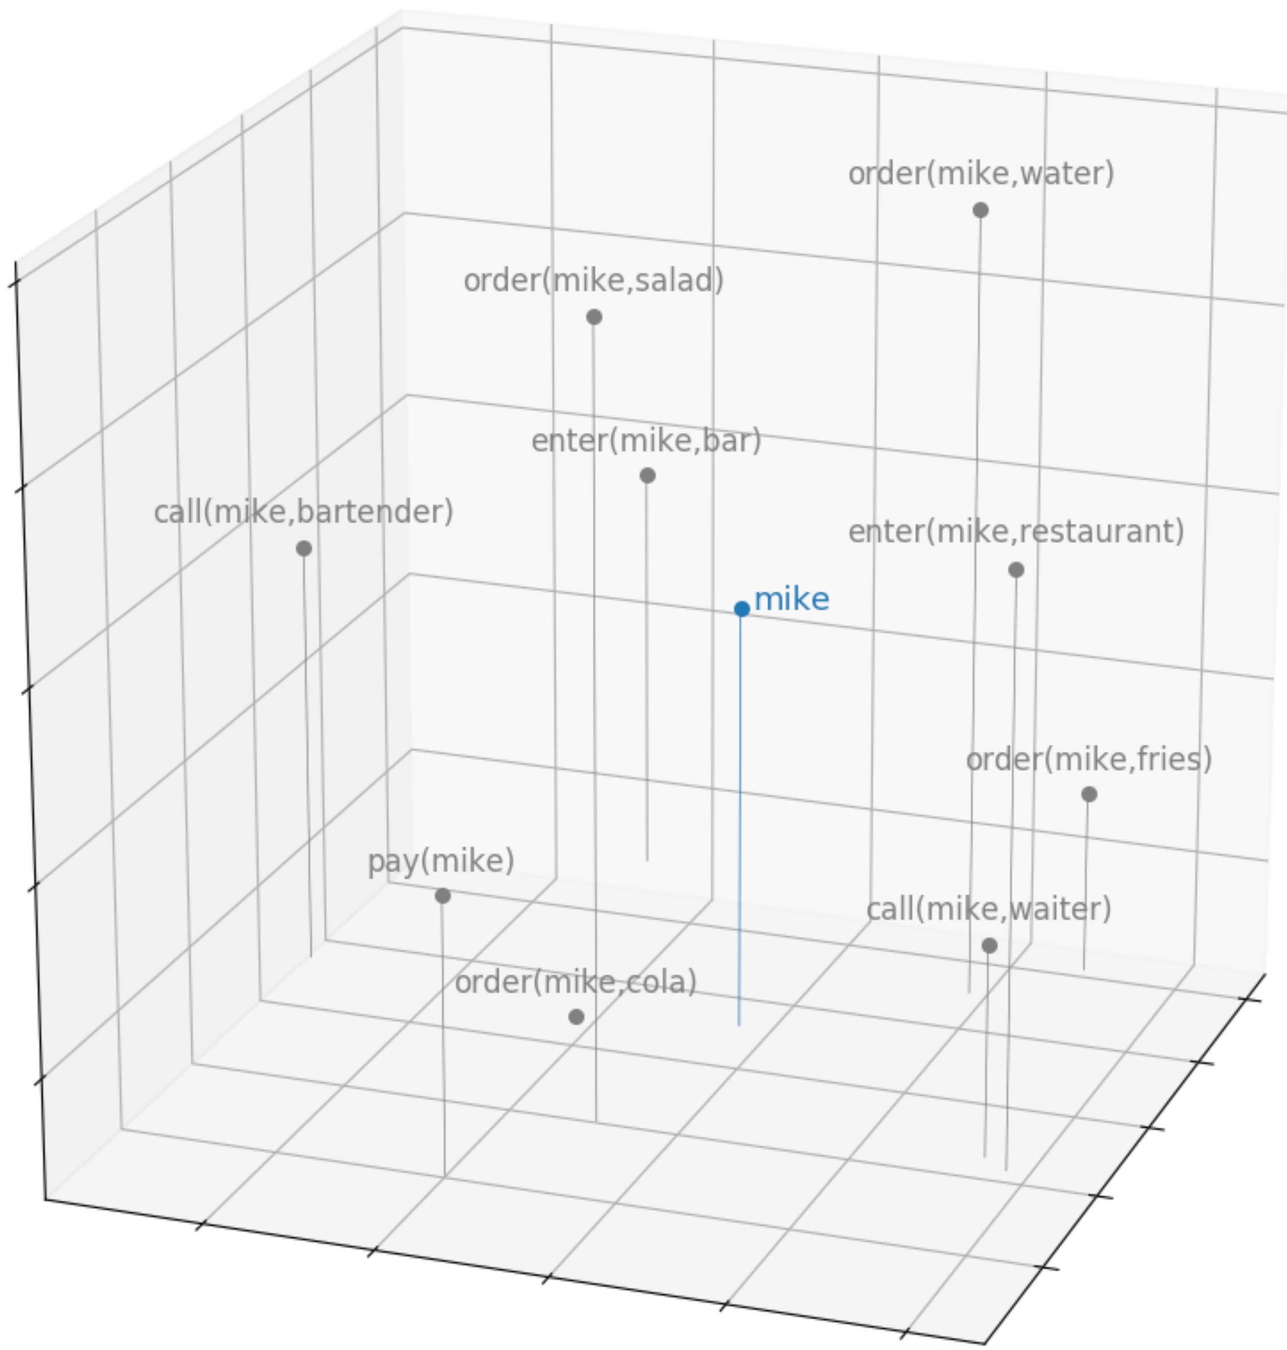

- Model-derived meaning of '*mike*' abstracts over the meanings of all propositions pertaining to *mike*

MEANING SPACE NAVIGATION

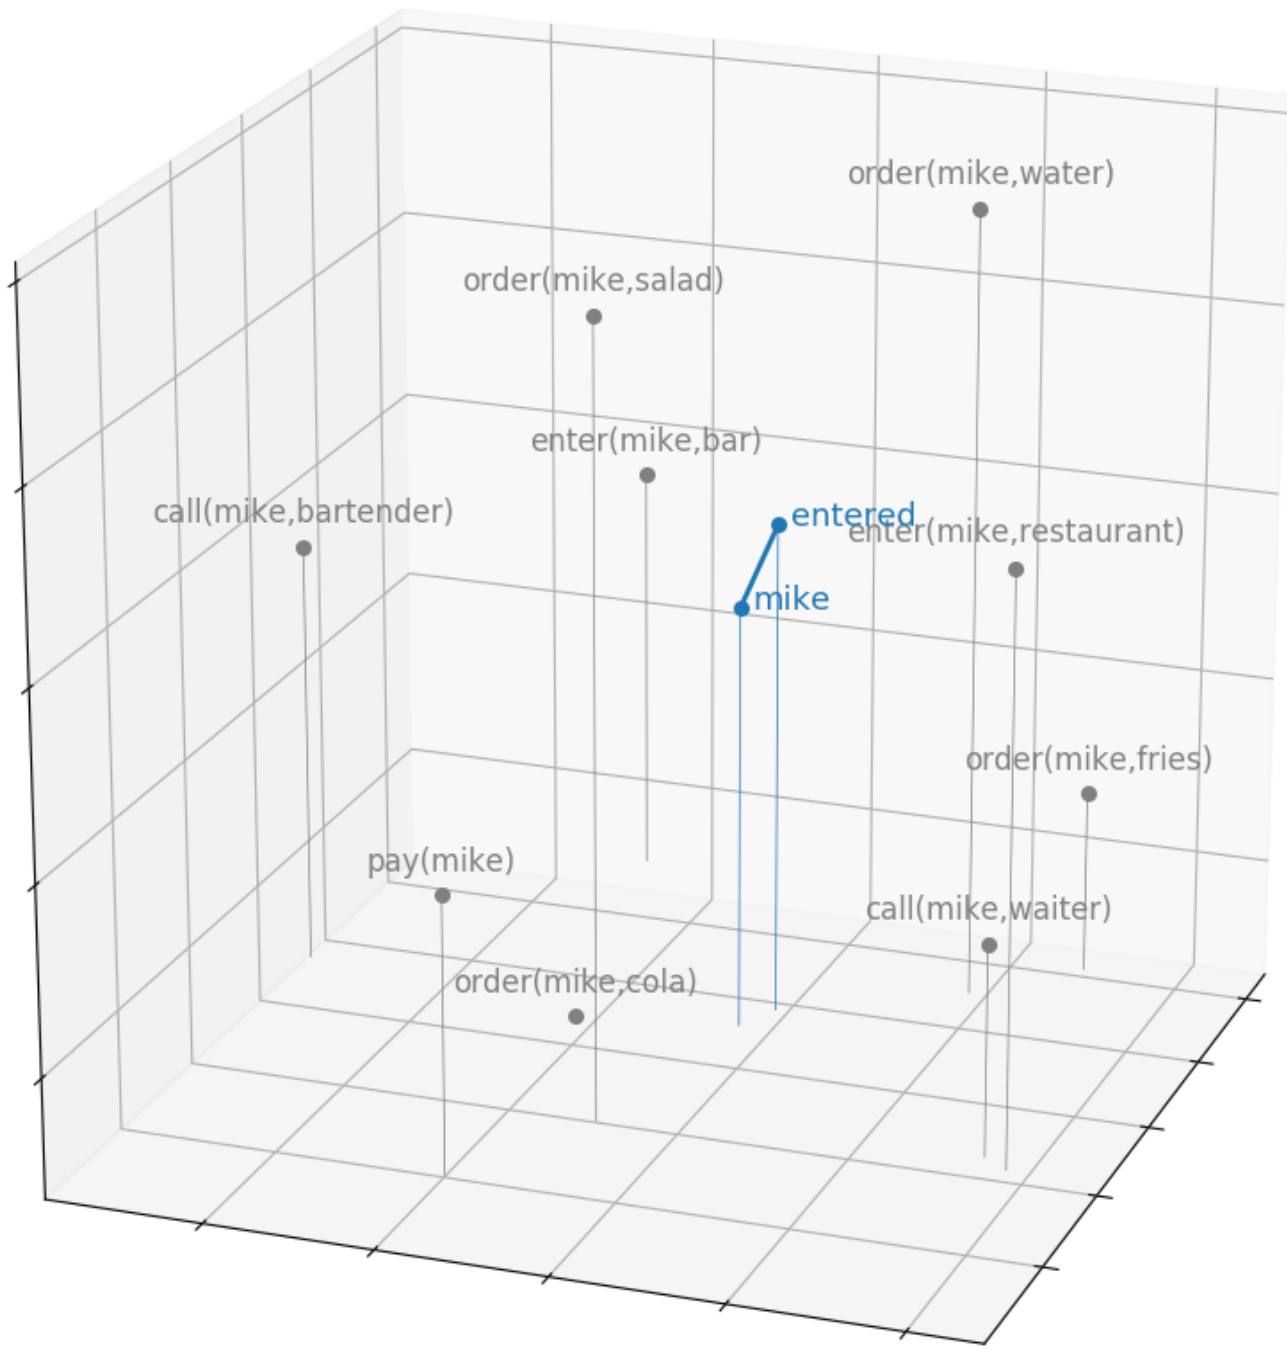

- At the word "entered", the model navigates to a point that represents the contextualised meaning "mike entered"

MEANING SPACE NAVIGATION

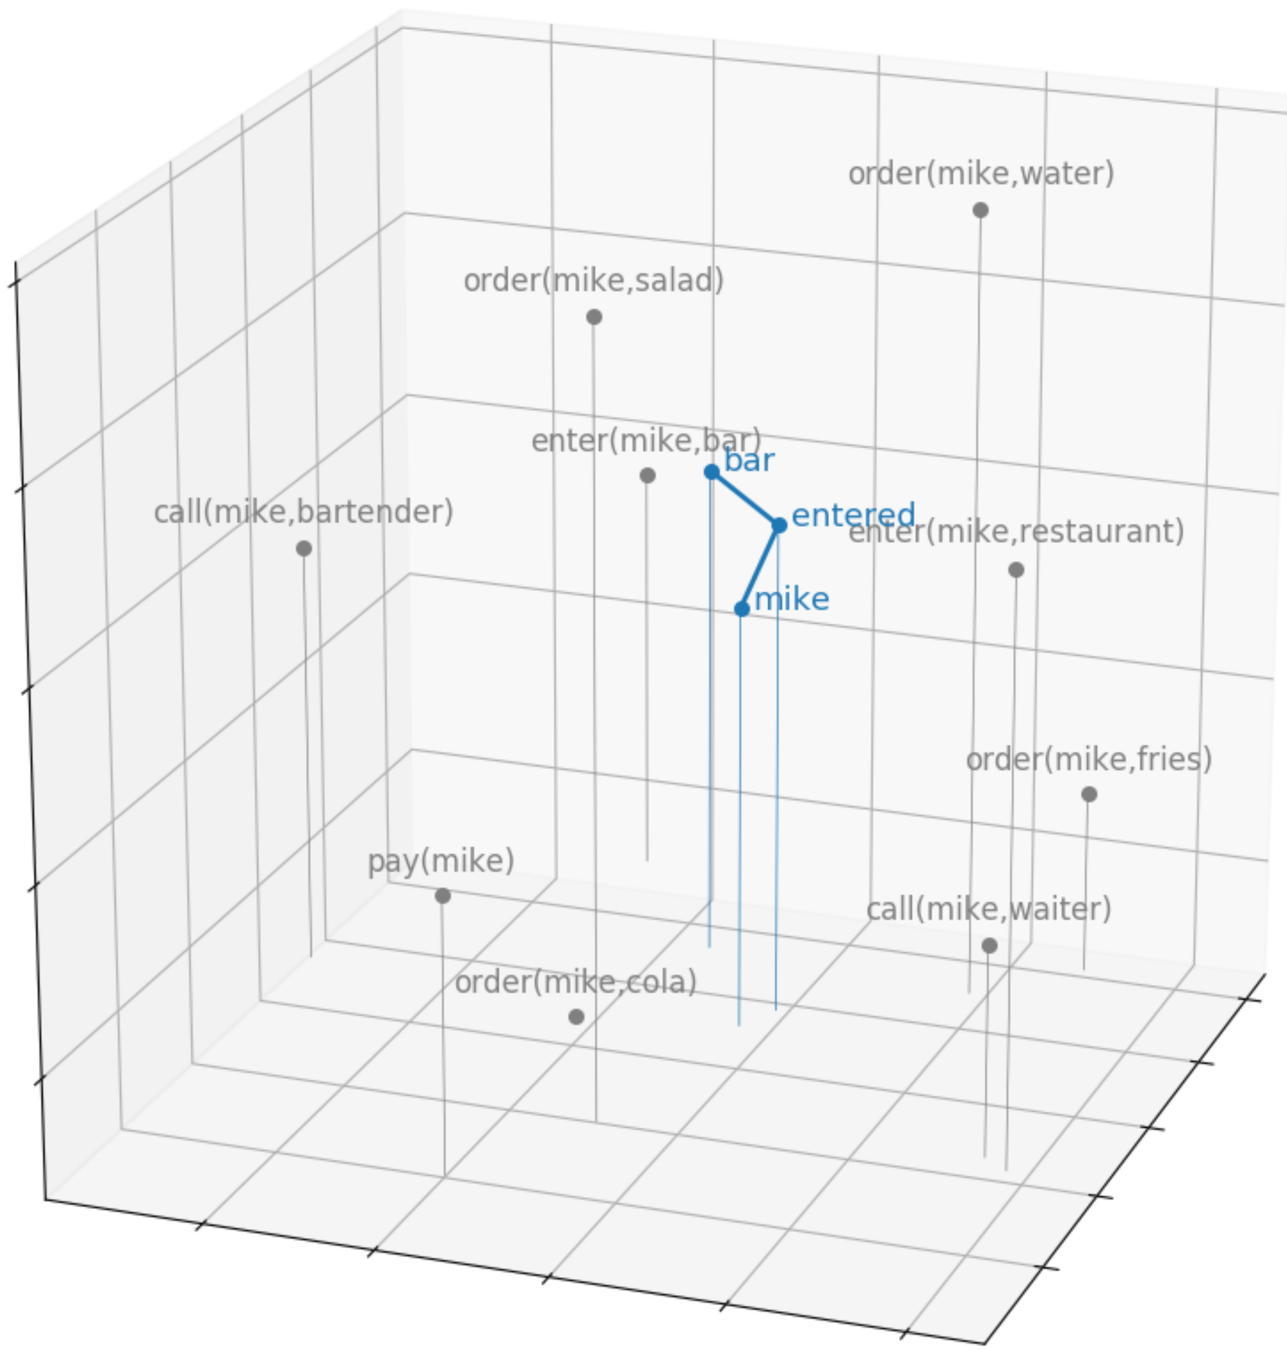

- The utterance "*mike entered the bar*" approximates the propositional meaning vector for *enter(mike,bar)*

MEANING SPACE NAVIGATION

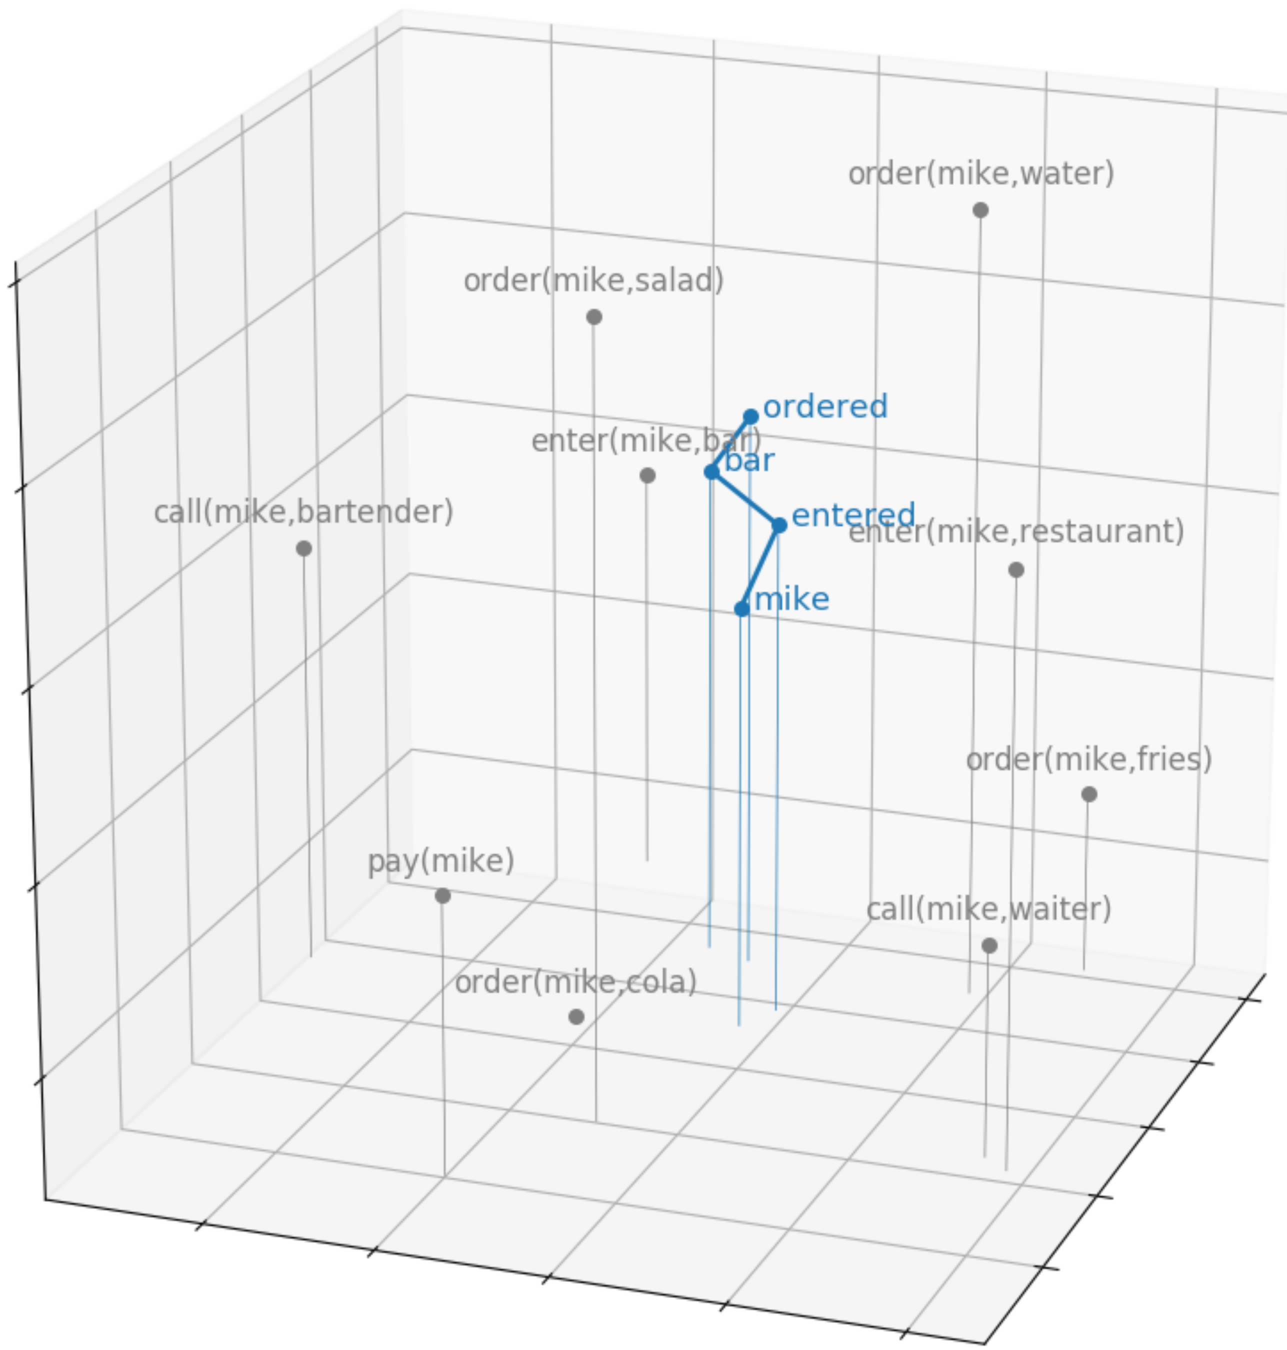

➤ The meaning vector after processing *"mike entered the bar [.] he ordered"* is close to *order* propositions that are typical given *enter(mike,bar)*

MEANING SPACE NAVIGATION

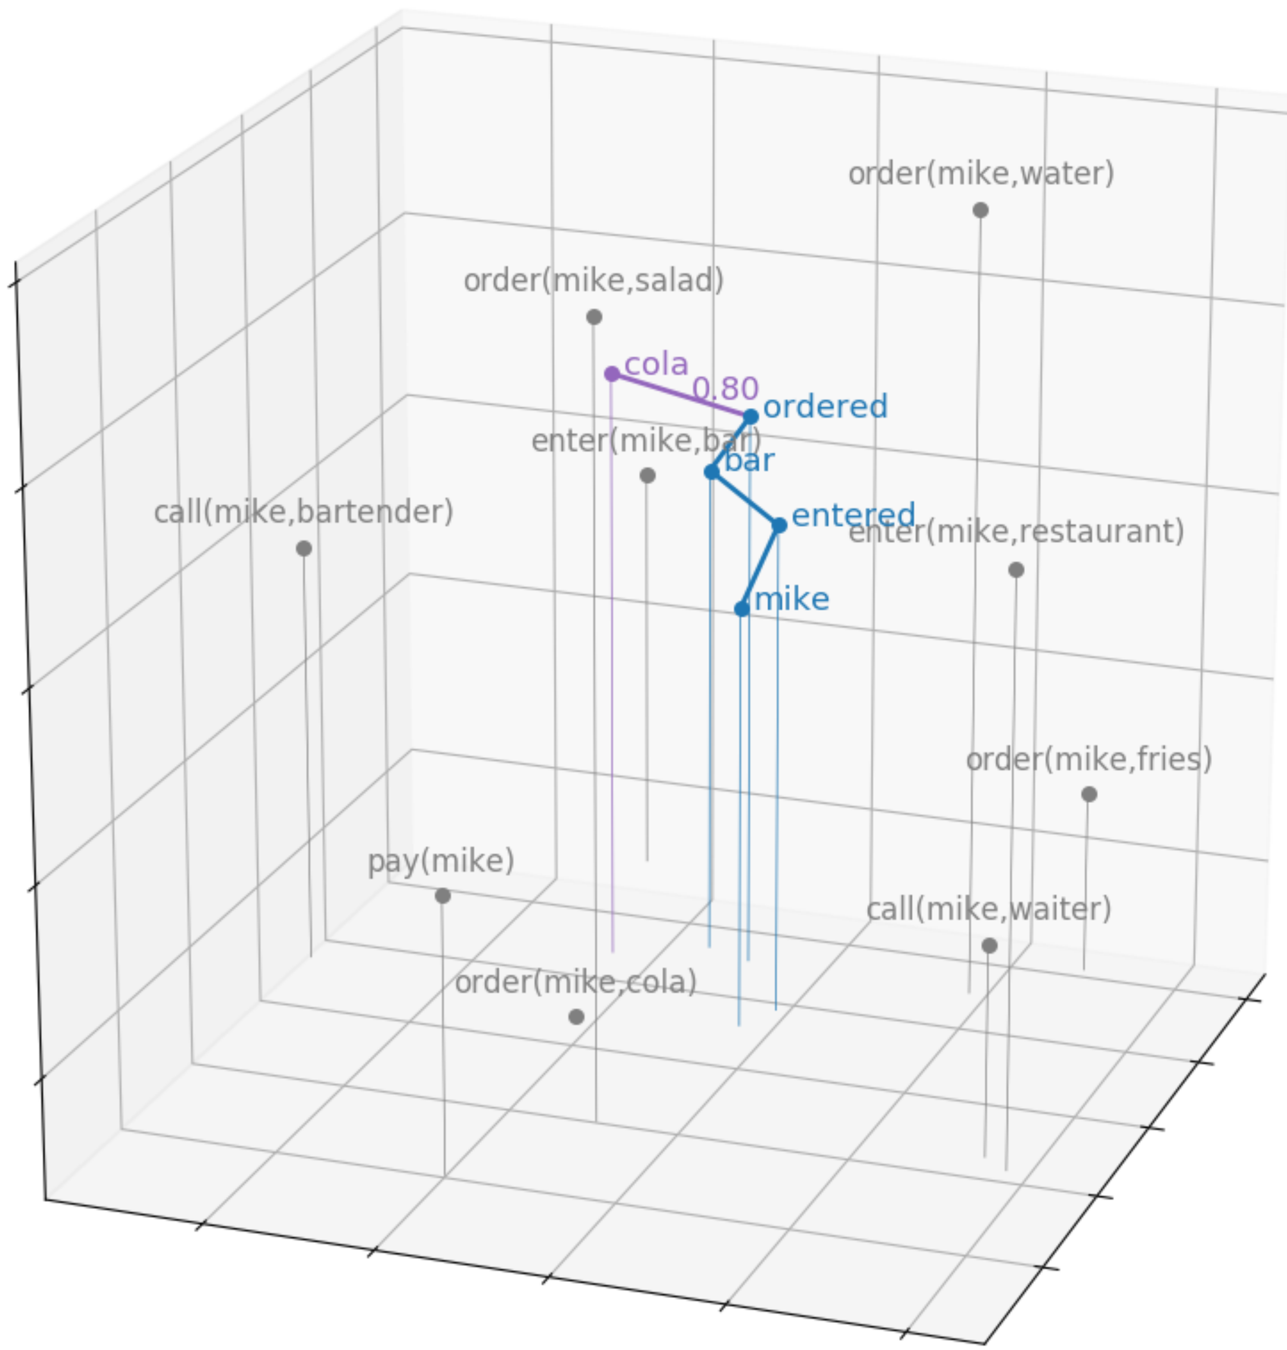

- When the utterance is continued with *"cola"*, the model approximates the conjunctive meaning vector *enter(mike,bar)* ∧ *order(mike,cola)*

MEANING SPACE NAVIGATION

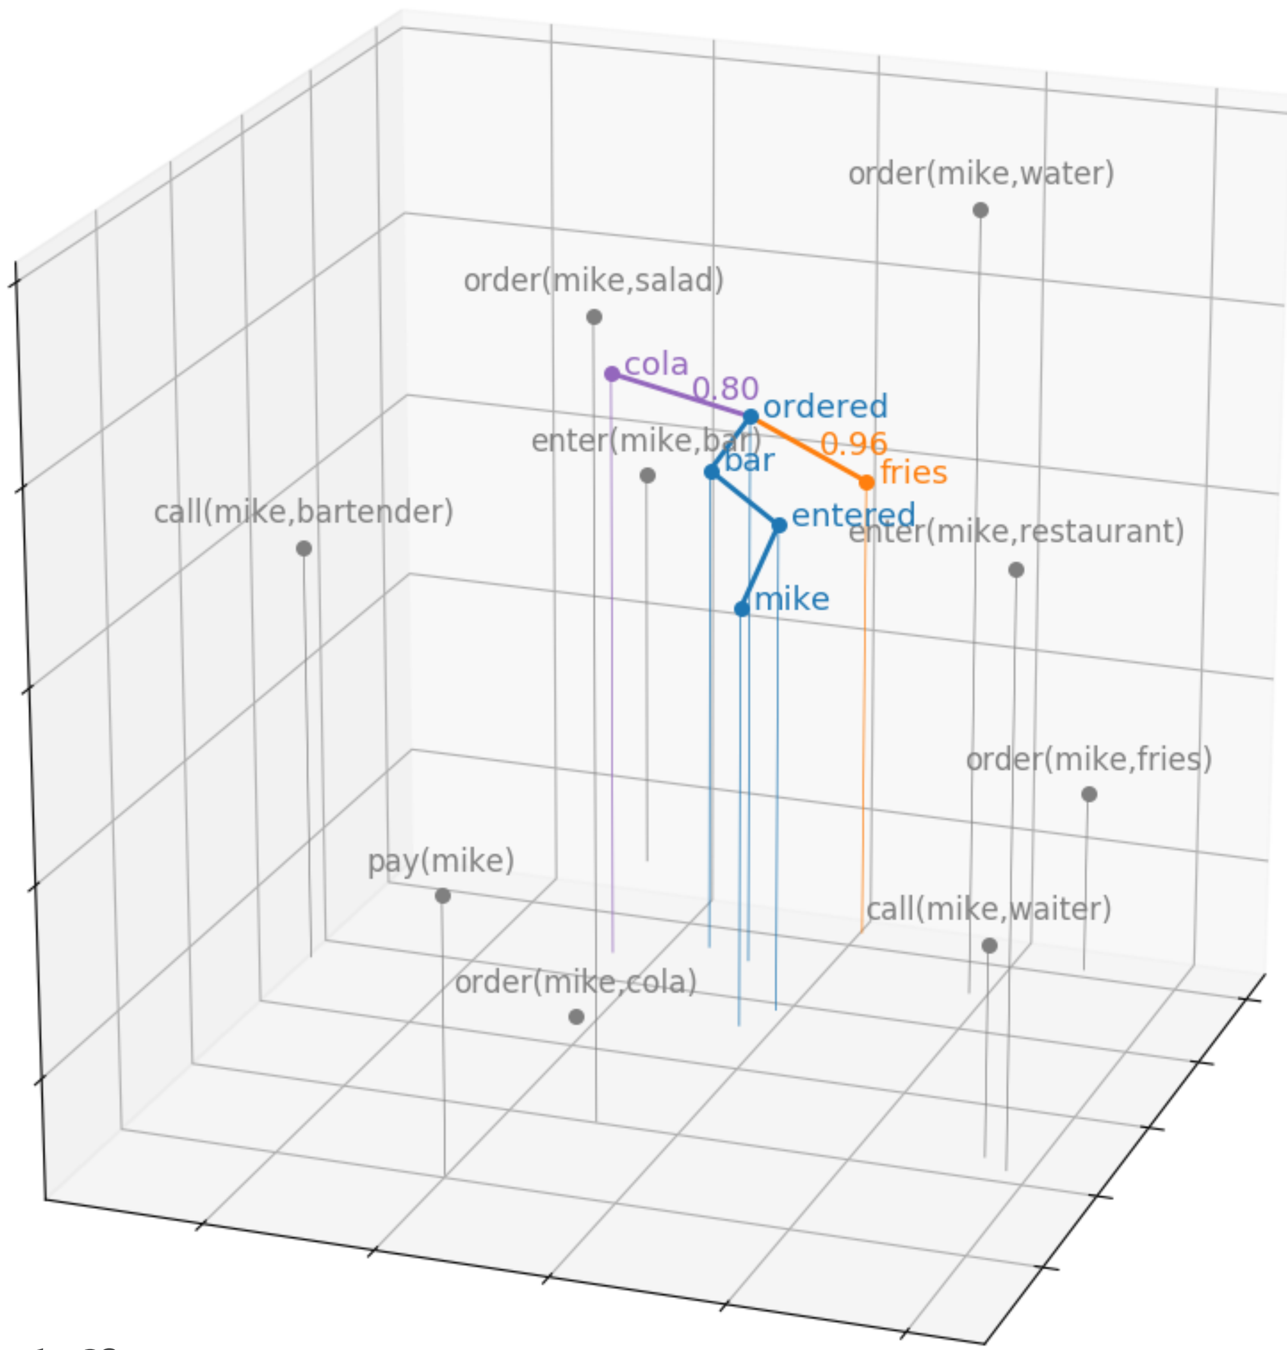

- *"fries"* results in different transition in meaning space, approximating the conjunctive meaning *enter(mike,bar)* ∧ *order(mike,fries)*—but is **less expected** after *"mike entered the bar [.] he ordered"*

INFORMATION THEORY IN DFS

Probabilistic nature of meaning space allows for defining formal notion of information (Shannon, 1948)

- **Surprisal** quantifies the expectancy of words in context
- Higher Surprisal ⇔ increased processing cost (Hale, 2001; Levy, 2008)
- In DFS, Surprisal quantifies expectancy of transition in meaning space, triggered by message m_{ab}:

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

➡ Word-by-word information effects of semantic constructions

Venhuizen, Crocker, Brouwer (2019, Discourse Proc; 2019, Entropy)

SEMANTIC PROCESSING IN THE MEANING SPACE

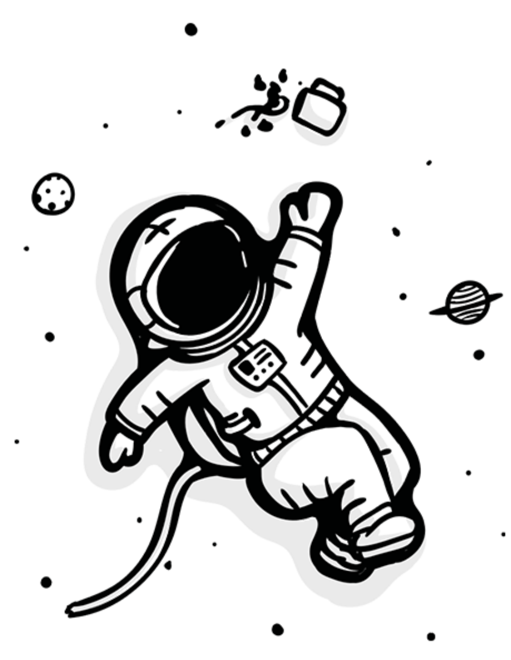

ENTAILMENT AND INFERENCE

Incremental meaning construction in the model is driven by:

- Sentence-semantics mappings (literal utterance meaning)
- Structure of the meaning space (probabilistic inferences)

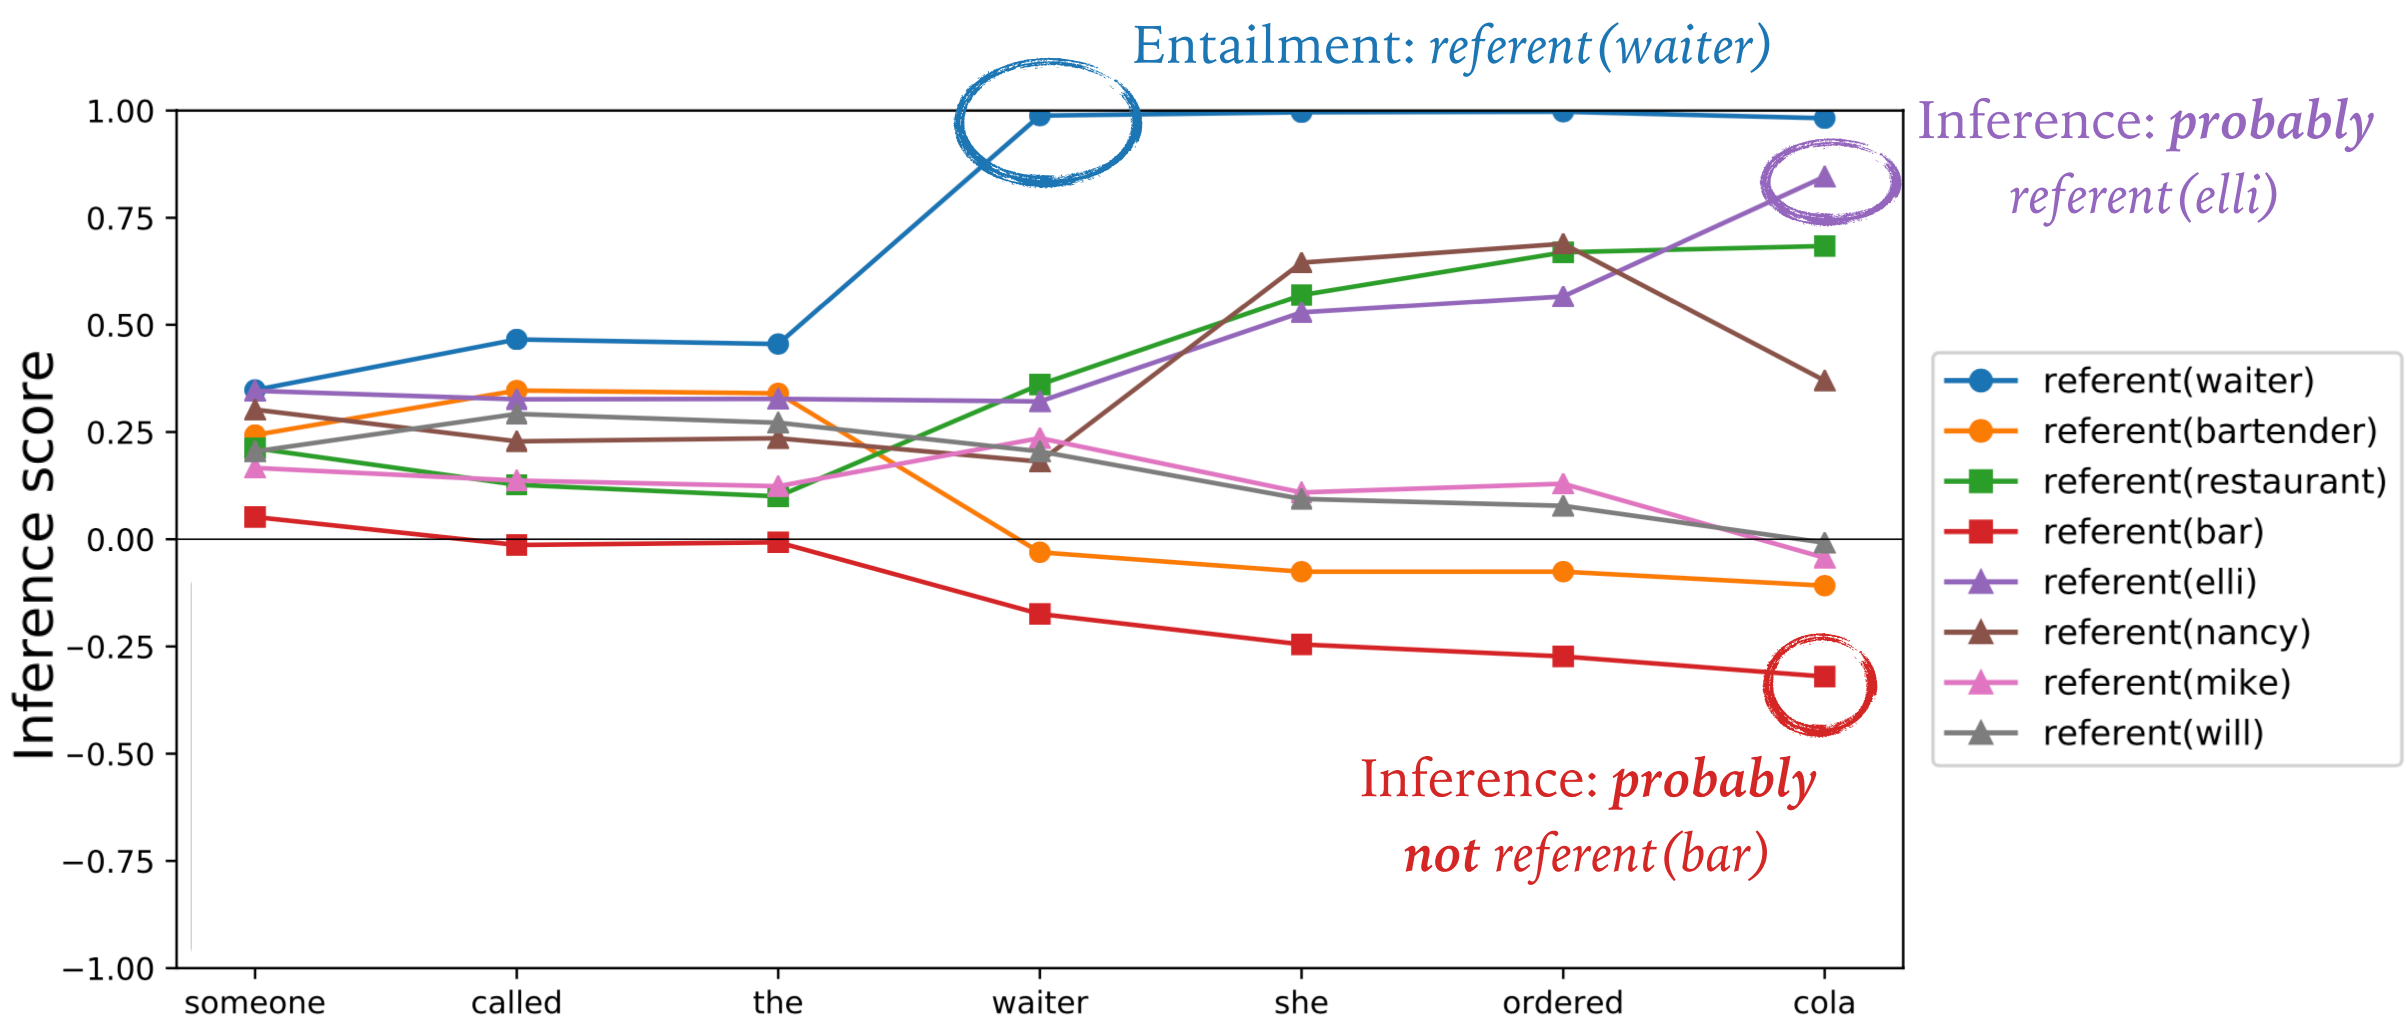

Venhuizen, Hendriks, Crocker, Brouwer (2021, Inf. Comp./ArXiv)

NEGATION AND PRESUPPOSITION

- Negation affects entailments and probabilistic inferences
- Interaction between negation and presupposition (triggered by *"the"*)
  - Presupposition has an effect beyond the literal meaning

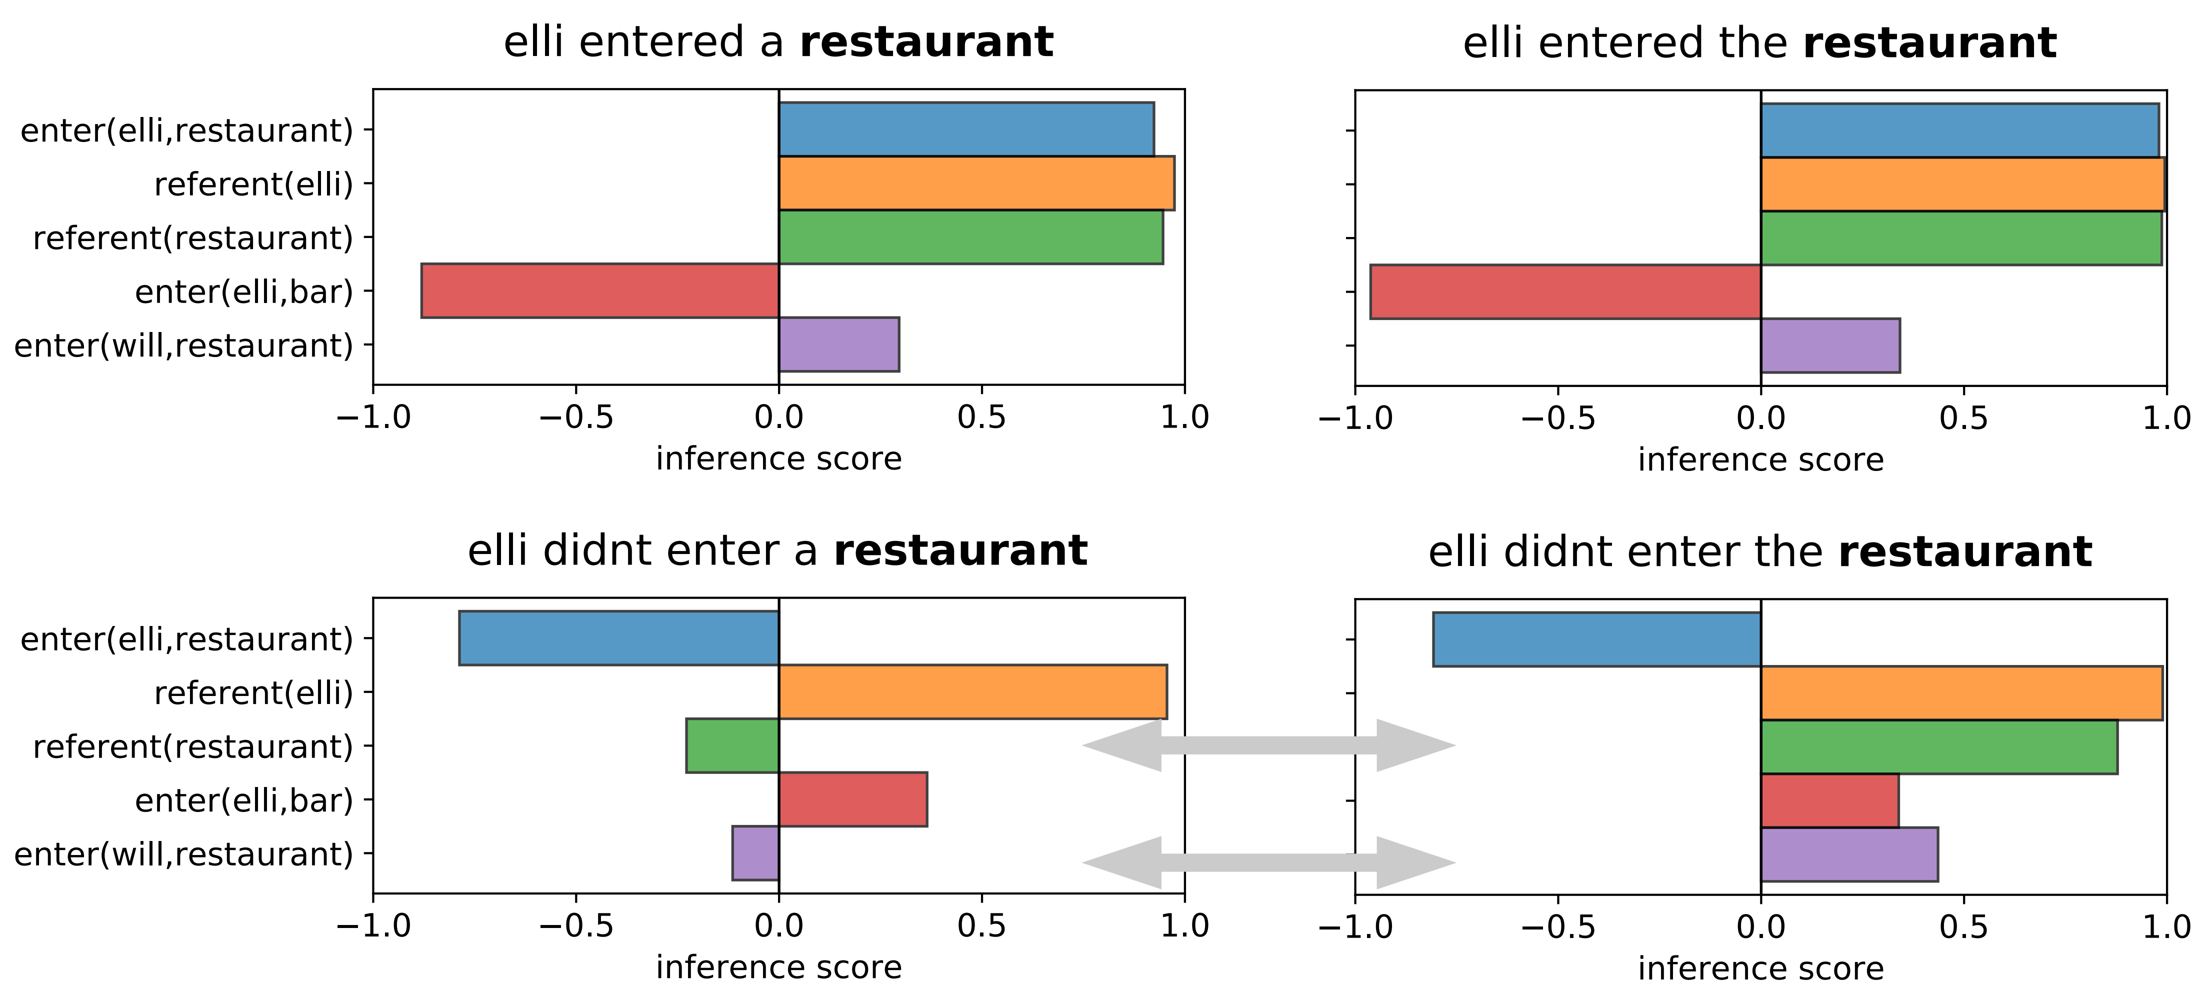

Venhuizen, Hendriks, Crocker, Brouwer (2021, Inf. Comp./ArXiv)

QUANTIFICATION AND REFERENCE

Quantified expressions induce inferential uncertainty

- ➤ Selective expressions (e.g. pronouns) can reduce this uncertainty
- ➤ Confirming initial expectations results in reduced Surprisal

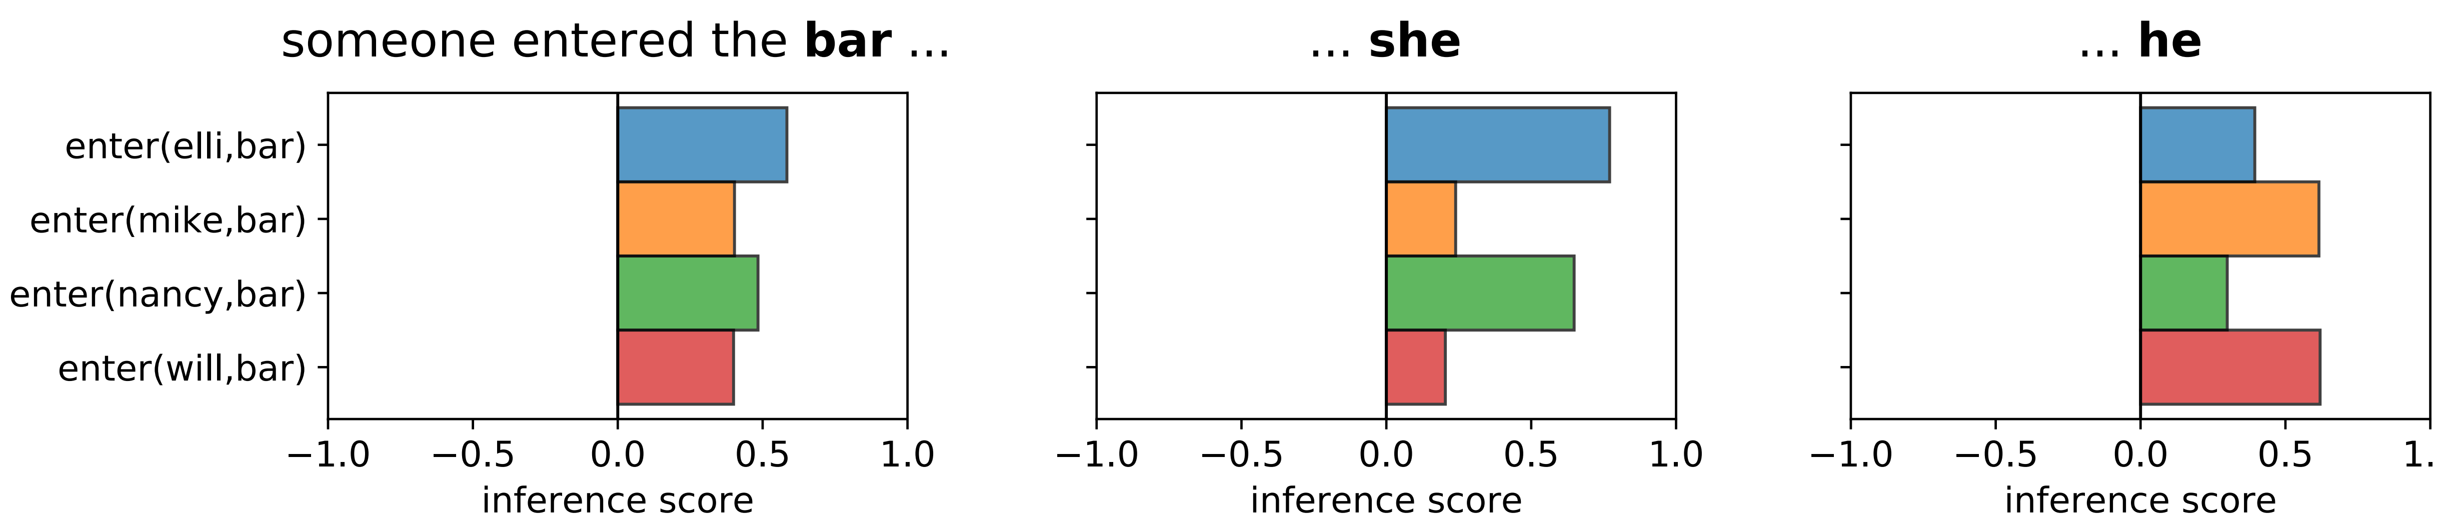

$S(she) = 0.92$ $S(he) = 1.12$

Venhuizen, Hendriks, Crocker, Brouwer (2021, Inf. Comp./ArXiv)

REFERENTIAL AMBIGUITY

In the training data, the anaphoric antecedent of pronouns is always disambiguated by the preceding or the following context

- ➤ Ambiguous pronouns trigger competing hypotheses about the utterance-final interpretation
- ➤ Disambiguating continuations result in utterance-level entailments
- ➤ Surprisal estimates reflect difference between expected and unexpected continuations

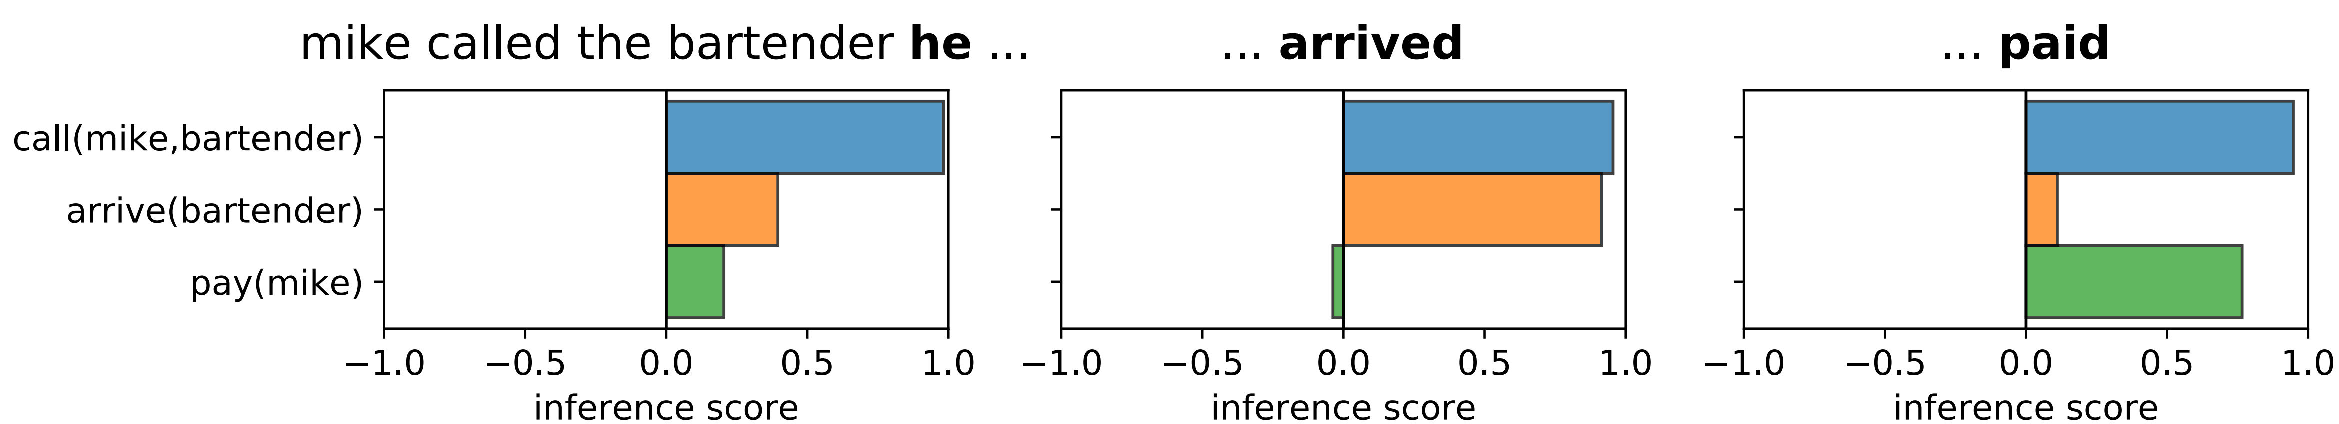

$S(arrived) = 0.44$ $S(paid) = 0.66$

Venhuizen, Hendriks, Crocker, Brouwer (2021, Inf. Comp./ArXiv)

BACK FROM SPACE

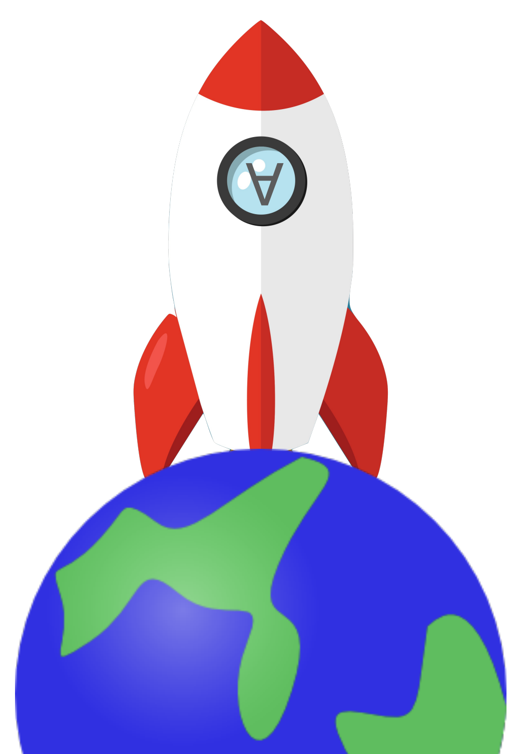

SUMMARY

Distributional Formal Semantics

- Compositionality
- Entailment and probabilistic inference
- Incremental meaning construction

?

Distributional Semantics

- Semantic similarity
- Empirically driven
- Cognitively inspired

DS VS. DFS: COMPLEMENTARY ASPECTS OF MEANING

- **Semantic similarity:**

  lexical similarity vs. propositional similarity

  beer ~ wine — *order(mike,beer) ~ drink(mike,beer)*

- **Data-driven sampling:**

  bottom-up vs. top-down

  individual linguistic co-occurrences — *high-level description of the world*

- **Cognitive foundation:**

  semantic memory vs. utterance interpretation

  feature-based word meanings — *unfolding discourse-level interpretation*

Kutas & Federmeier (2000); McRae et al. (2005); van Berkum (2009); Brouwer et al. (2012, 2017)

DISCUSSION: COGNITIVE FOUNDATION FOR DFS?

The Retrieval-Integration account of the electrophysiology of language comprehension

- Word meaning retrieval~N400
- Integration in utterance meaning~P600

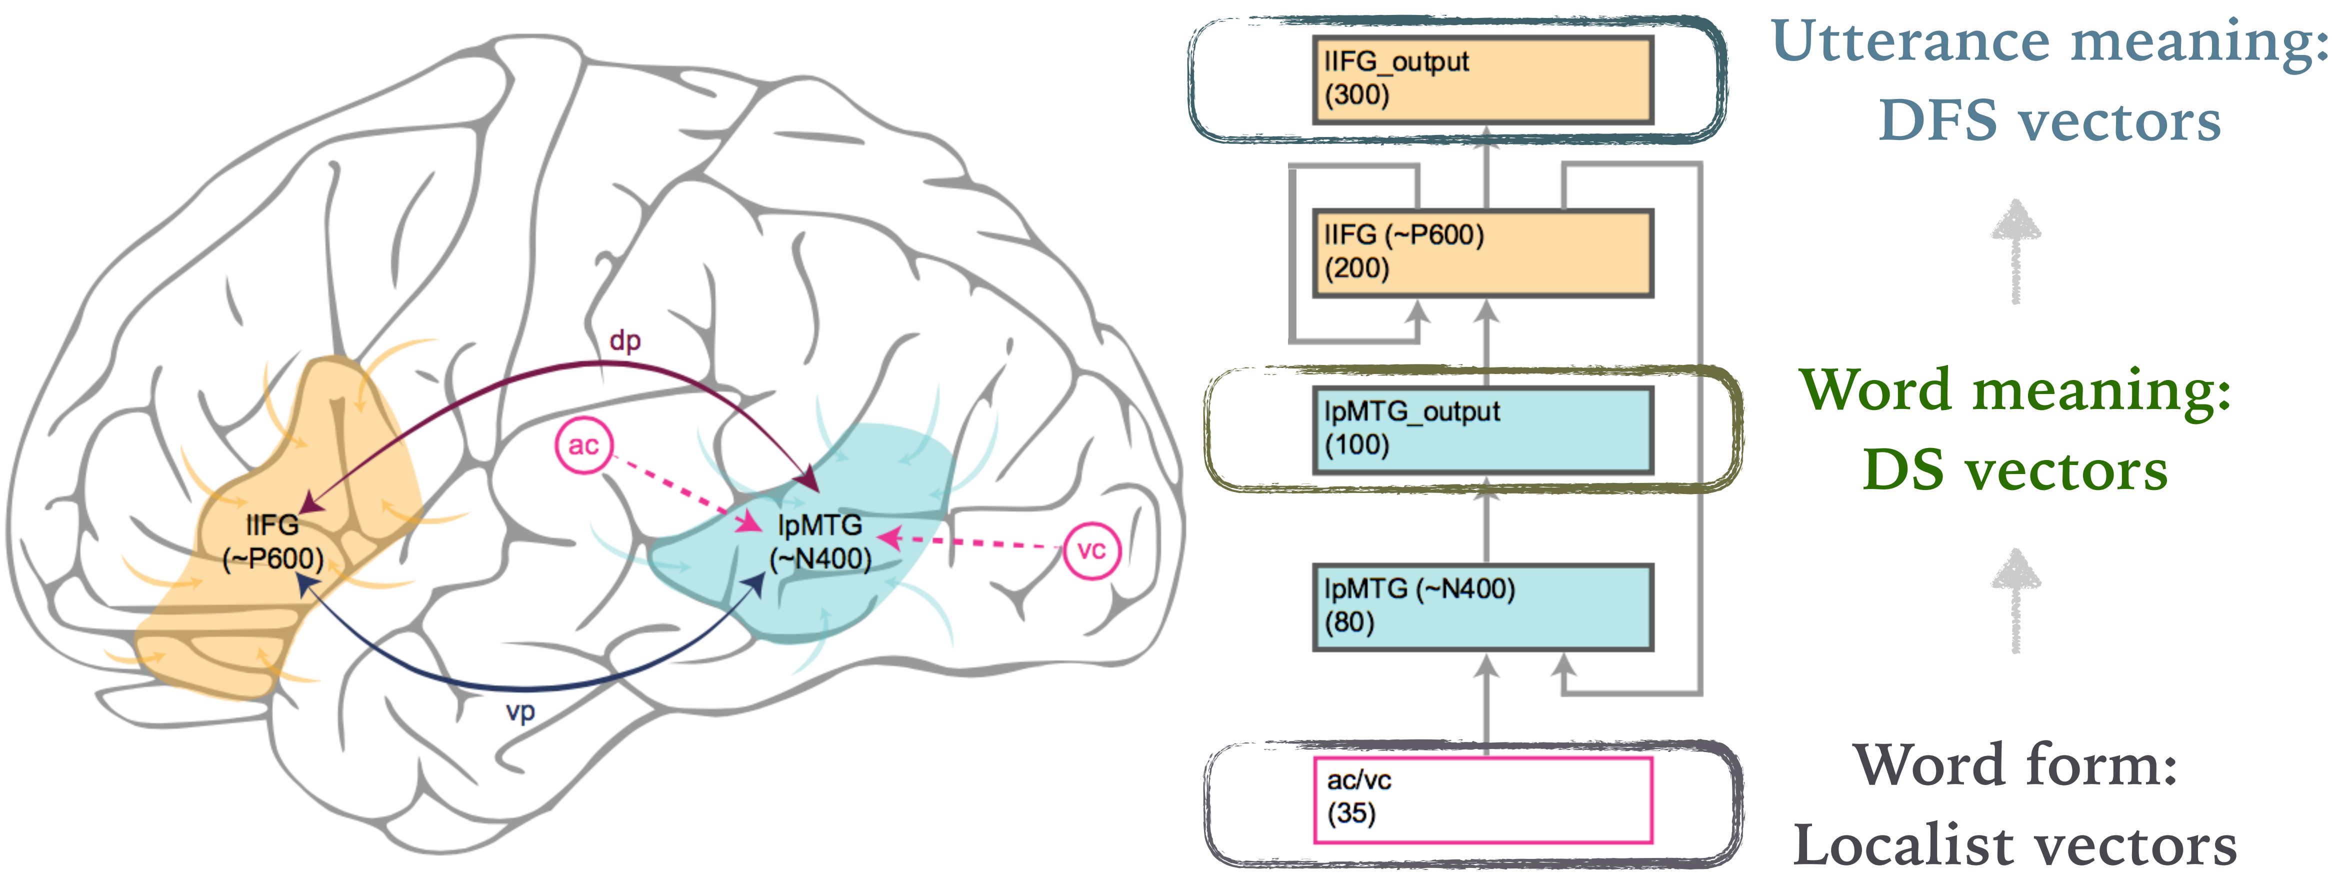

Brouwer et al. (2012, Brain Res.; 2017, Cogn. Sci.; 2021, Frontiers Psych.);

DISCUSSION: DATA-DRIVEN DFS?

Meaning space reflects world knowledge about propositional co-occurrence, rather than linguistic co-occurrence

- DeScript corpus (Wanzare et al., 2016)

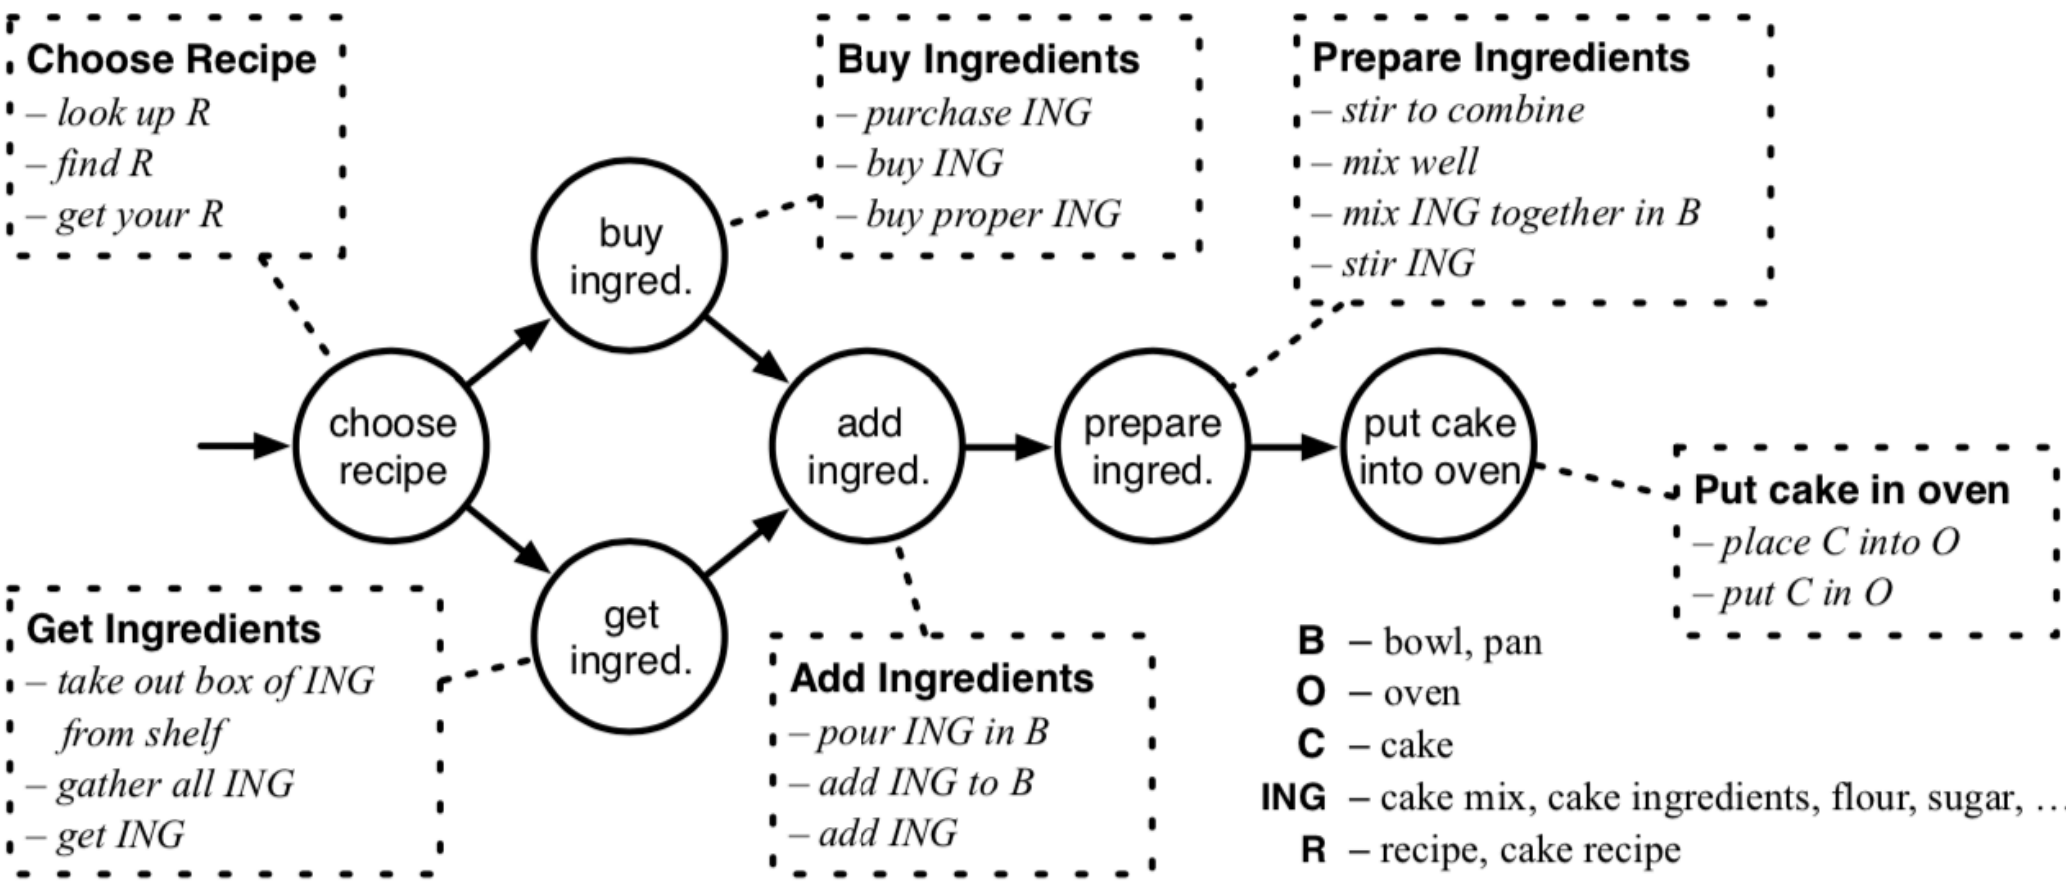

DISTRIBUTIONAL FORMAL SEMANTICS

- The meaning space $S_{\mathcal{M}\times\mathcal{P}}$ captures the structure of the world **truth-conditionally** and **probabilistically**
- Meaning vectors are **compositional** at the propositional level
- **Sub-propositional meaning** derived by incrementally navigating $S_{\mathcal{M}\times\mathcal{P}}$ (using a Simple Recurrent neural Network)
- Semantic phenomena—*negation, presupposition, quantification & reference*—affect **incremental entailments and inferences** during meaning space navigation

Distributional Formal Semantics

Noortje J. Venhuizen[a,*], Petra Hendriks[b], Matthew W. Crocker[a], Harm Brouwer[a]

[a] *Saarland University, Department of Language Science and Technology, 66123 Saarbrücken, Germany*
[b] *University of Groningen, Center for Language and Cognition Groningen (CLCG), P.O. Box 716, 9700AS Groningen, the Netherlands*

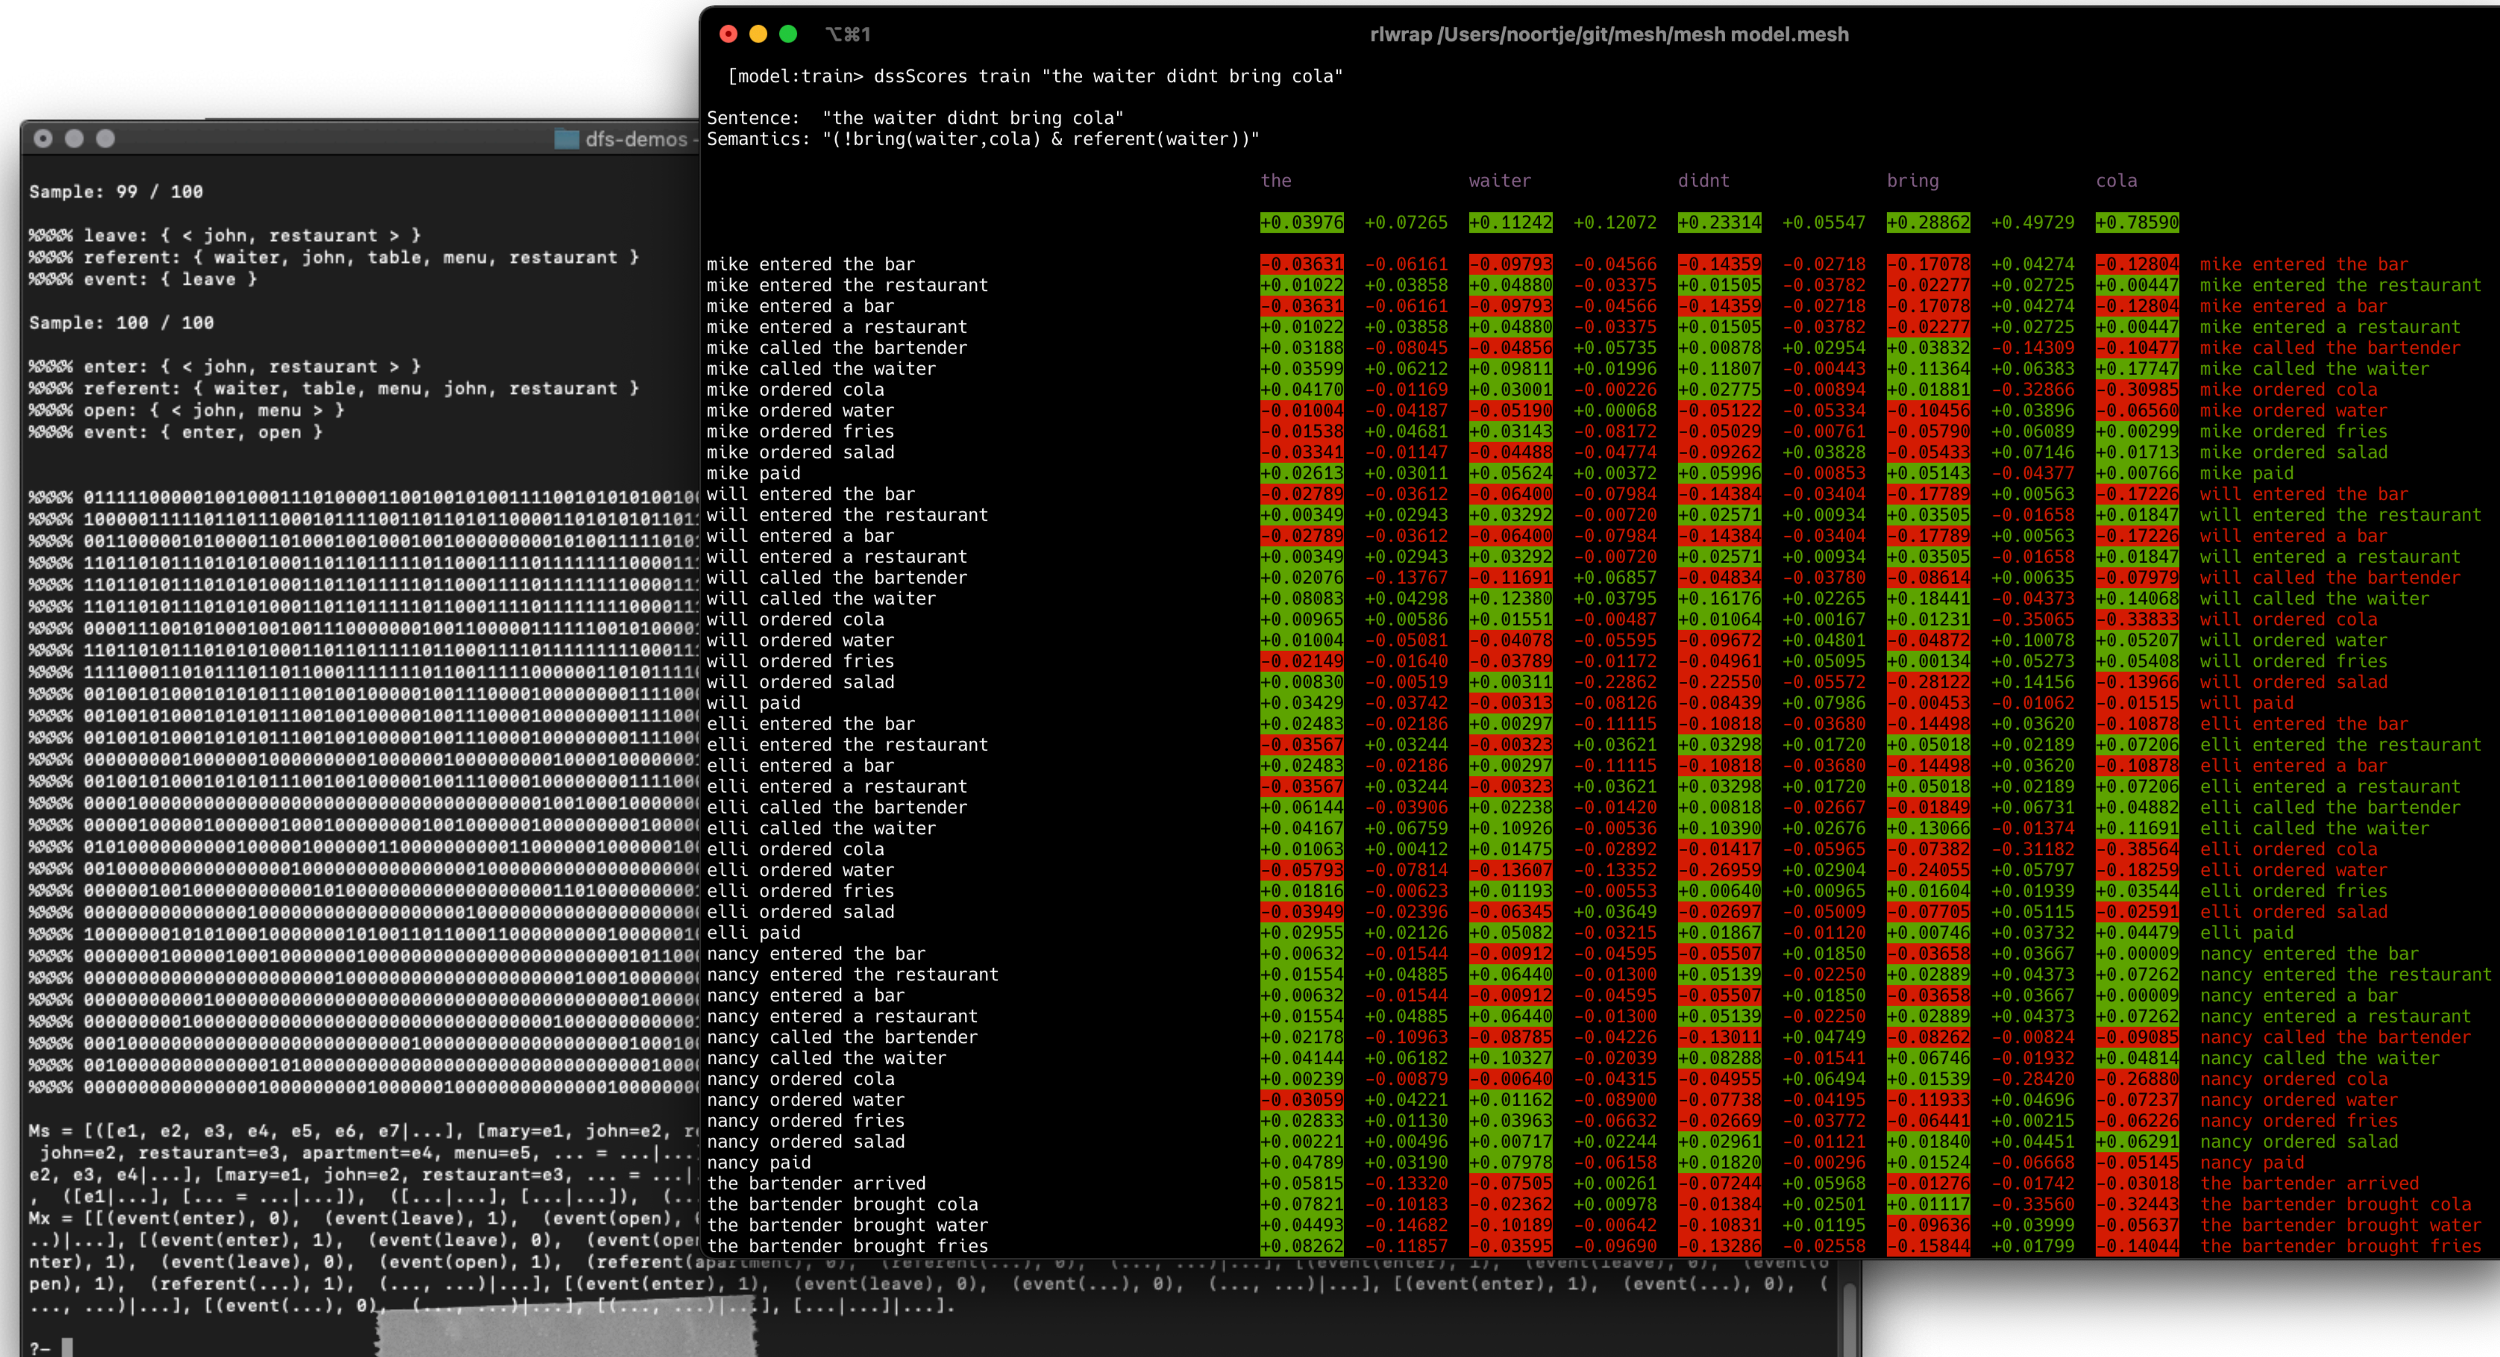

n advantage of such approaches is that the distributional
tedness between lexical items [8], and that they can be
oven extremely difficult to incorporate the traditional
ch a distributional semantic framework [9].

DFS-TOOLS

https://github.com/hbrouwer/dfs-tools