Explicit world-knowledge and distributional semantic representations

ESSLLI 2017 Day 4: representational conflicts in models of the lexicon

Asad Sayeed and Alessandra Zarcone

University of Gothenburg, Saarland University

We got the impression that some of the audience does not have a psycholinguistic background. . .

. . . which means that we'll have a little overview of some relevant points

Psycholinguistics

- A catch-all term for a large number of sub-fields.
- Unifying theme:
  - There is something called language.
  - Human minds do this thing called language.
  - What do humans minds do in order to "do" language?
- Example areas: child language acquisition, second language acquisition, adult (native) sentence processing, adult (second-language) sentence processing . . .

What follows is a very non-technical "crash course" in what some experimental psycholinguists do.

Adult sentence processing

We're going to narrow in on adult processing (even though there's lots of issues for modeling lexical semantics in e.g. child language research).

Adult sentence processing research

- Focus on "fully-acquired" language, e.g., adult native speaker.
- Main research question: what happens when a language user interprets an utterance?

The central issue is thus: the "time course" of an utterance in the processing system.

Adult sentence processing

Time course: focus on behaviour over time.

Donald put the cake in the ocean.

For example, we want *somehow* to measure *something* that happens at "ocean". Roughly, three kinds of measurements:

- Introspective - conscious human report, judgements, linguistic responses etc.
- Reaction-based - observed behavioural changes, time delays, etc., in response to stimulus.
- Physiological - some measurable aspect of the body (usu. brain) that reflects some biological "effort".

Introspective measurement

Introspective techniques are simple. E.g.:

- Linguist making "traditional" grammaticality judgement.
- Ratings collected over groups of non-linguists.
- Fill-in-the-blank/cloze responses.

How this might work:

Rate from 1-7

Donald put the cake in the ocean.

- Pro: often very fast/cheap to collect, *often does* lead to good insights, scientifically strong results.
- Con: not "real-time" and prone to dangers of subjectivity.

Reaction-based techniques

Usually requires some active or passive measurement. E.g.:

- Self-paced reading: subjects press button to see next word, measure per-word reaction time.
- Eye-tracking: observe a subject's eyes moving across a sentence.
- Timed decision tasks.

How this might work:

Expect slower self-paced reading reaction time at "ocean"

Donald put the cake in the . . . ocean.

- Pro: remove some amount of subjectivity, often fine-grained (ms-scale) rich data.
- Con: somewhat expensive equipment, very indirect, very sensitive to experimental conditions.

Physiological techniques

Measurement of signals from the body at varying degrees of depth. E.g.:

- Event-related potentials (ERP) – use EEGs to roughly measure changes in location of electrical output.
- Functional MRI – use giant magnets to measure minute differences in water flow in the brain.
- Pupillometry – measure pupil dilation changes to test neural activations.

How this might work:

Expect N400 ERP at "ocean"

Donald put the cake in the . . . ocean.

- Pro: somewhat more direct view of *something* happening in the brain.
- Con: very expensive equipment and software, sometimes limited time or spatial resolution.

Things psycholinguists like

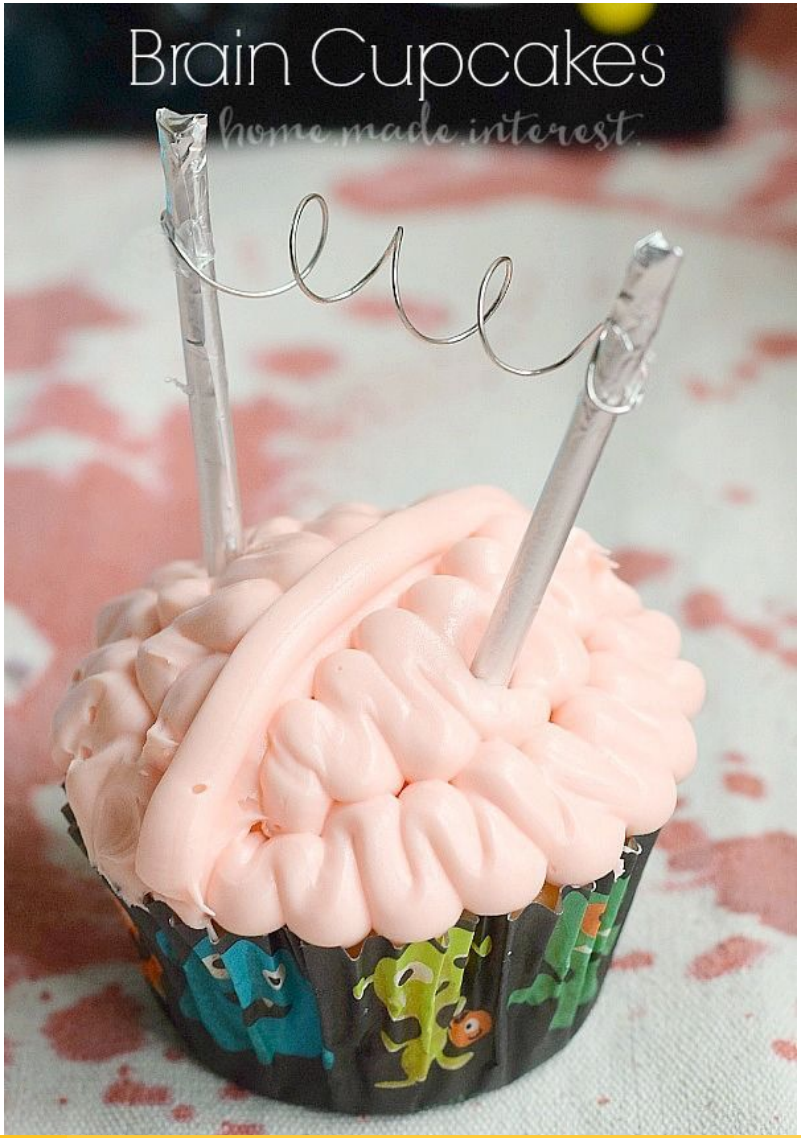

Why are we talking about this?

We aren't going to be able to talk about every technique.

- So that you can contextualize what we *are* going to talk about.
- Any model of the world-knowledge interactions with the lexicon ought to be observable in behaviour.
- We *do* find that many hypotheses can be validated* that way.

*For values of validation that include, "do you trust their way of calculating a p-value?"

Part 1: incongruities between distributional knowledge and world-knowledge in event structure representation

Well, apparent incongruities, anyway. . .

Let's look at the problem of logical metonymy.
(BEGIN PRETENDING TO BE ALESSANDRA)

Logical metonymy and covert events

Jack Kerouac began the book
around 1949 in New York

Logical metonymy and covert events

Jack Kerouac began the book around 1949 in New York
$\rightarrow$ writing

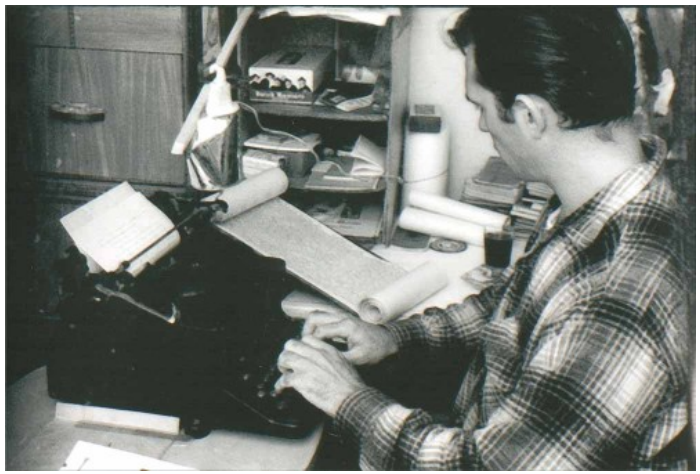

Logical metonymy and covert events

Jack Kerouac began the book around 1949 in New York
$\rightarrow$ writing

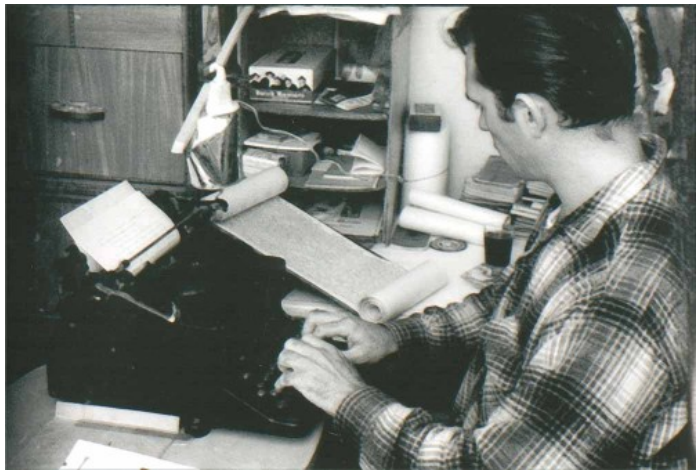

Logical Metonymies [Pustejovsky, 1995]

- involve *covert events* (*metonymy*: book $\rightarrow$ writing the book)

Logical metonymy and covert events

Jack Kerouac began the book around 1949 in New York
$\rightarrow$ writing

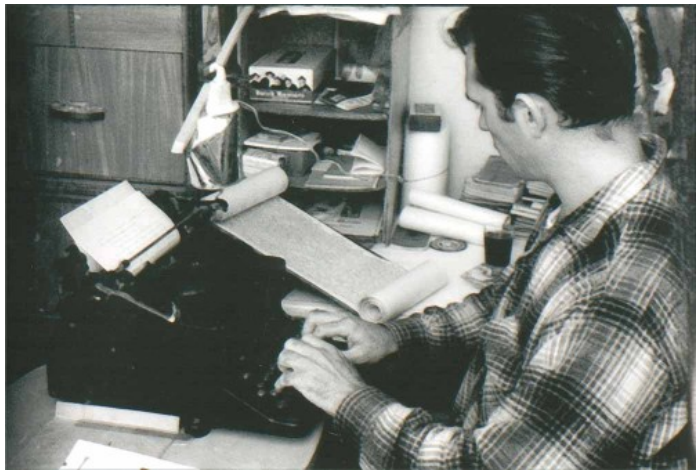

Logical Metonymies [Pustejovsky, 1995]

- involve *covert events* (*metonymy*: book $\rightarrow$ writing the book)
  - not realized on the surface, but understood
  - influence reading times
  - a challenge to compositionality

Logical metonymy and covert events

Jack Kerouac began the book around 1949 in New York
$\rightarrow$ writing

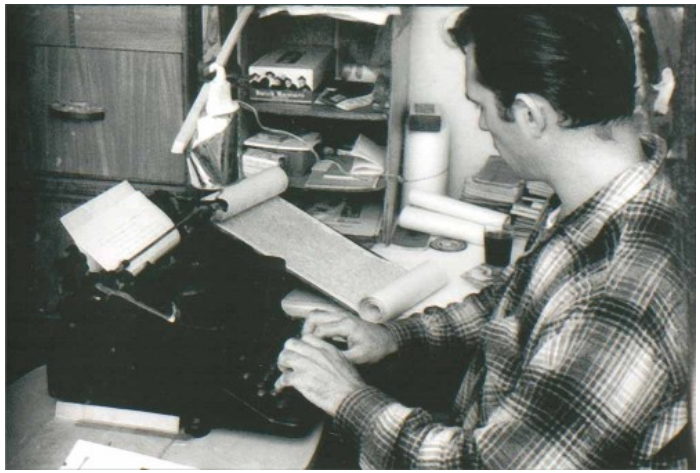

Logical Metonymies [Pustejovsky, 1995]

- involve *covert events* (*metonymy*: book $\rightarrow$ writing the book)
  - not realized on the surface, but understood
  - influence reading times
  - a challenge to compositionality
- **The Source Question**:
  What is the *source* of the covert event (lexicon, world knowledge)?

Logical metonymy and covert events

Jack Kerouac began the book around 1949 in New York
$\rightarrow$ writing

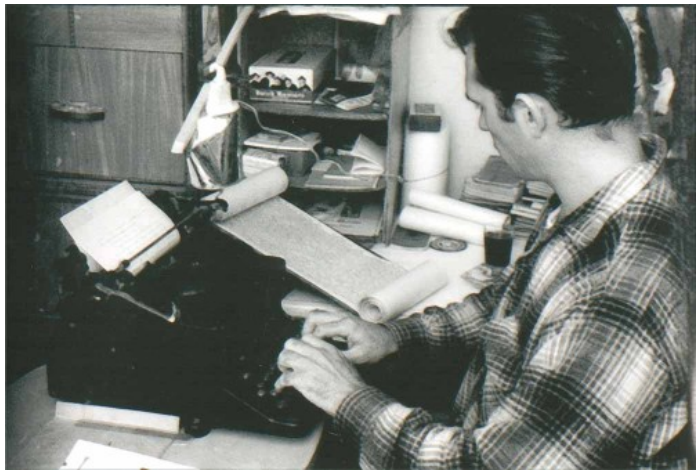

Logical Metonymies [Pustejovsky, 1995]

- EVent-selecting verb + ENtity-denoting object

Logical metonymy and covert events

Jack Kerouac began **the book**$_{EN}$ around 1949 in New York

$\rightarrow$ writing

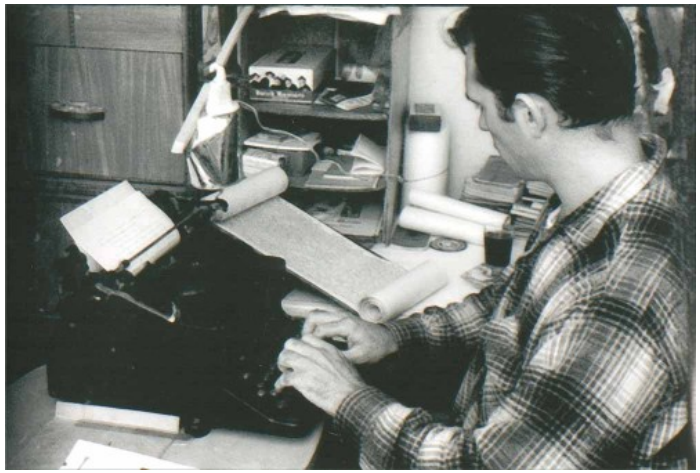

Logical Metonymies [Pustejovsky, 1995]

- EVent-selecting verb + ENtity-denoting object
  $\Leftrightarrow$ Jack Kerouac began **his journey**$_{EV}$ across America.

Logical metonymy and covert events

Jack Kerouac began **the book**$_{EN}$ around 1949 in New York
$\rightarrow$ writing

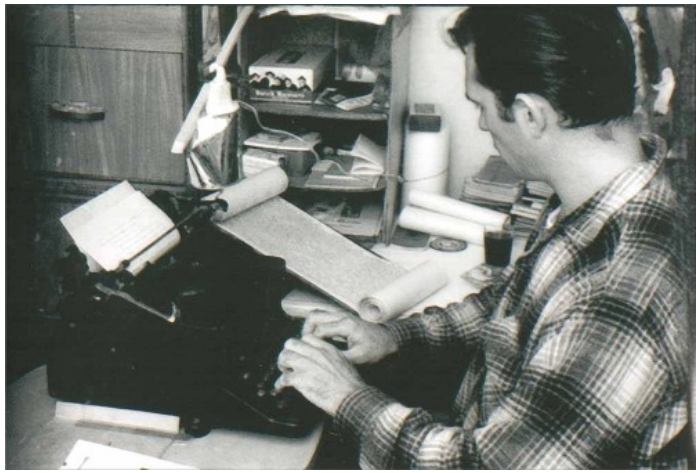

Logical Metonymies [Pustejovsky, 1995]

- EVent-selecting verb + ENtity-denoting object
  $\Leftrightarrow$ Jack Kerouac began **his journey**$_{EV}$ across America.
- **The Trigger Question**:
  What *triggers* the metonymy (and the covert event)?

Psycholinguistic motivation

wash car

wash hair

Psycholinguistic motivation

wash car

$\rightarrow$ *hose, sponge, outdoor*

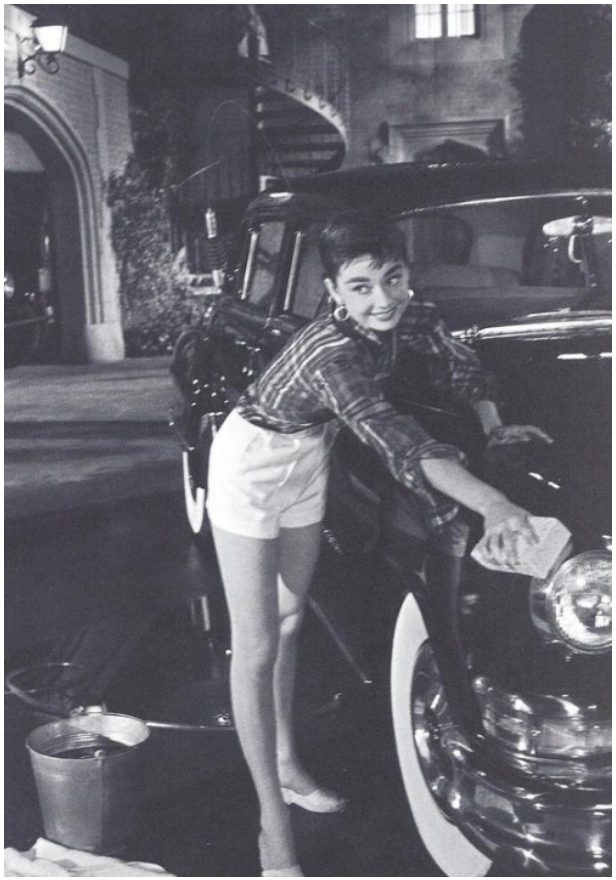

wash hair

$\rightarrow$ *shampoo, sink, bathroom*

Psycholinguistic motivation

Generalized event knowledge [McRae and Matsuki, 2009]:

Prototypical knowledge about events and their participants
(first and second-hand experience, available in our memory)

Psycholinguistic motivation

Generalized event knowledge [McRae and Matsuki, 2009]:

Prototypical knowledge about events and their participants
(first and second-hand experience, available in our memory)

- activated by words in isolation $\Rightarrow$ cue concepts from typical scenarios

$$\langle arrest \rangle \xrightarrow{agent} cop$$
$$\langle arrest \rangle \xrightarrow{patient} crook$$

Psycholinguistic motivation

Generalized event knowledge [McRae and Matsuki, 2009]:

Prototypical knowledge about events and their participants
(first and second-hand experience, available in our memory)

- activated by words in isolation $\Rightarrow$ cue concepts from typical scenarios

$$\langle arrest \rangle \xrightarrow{agent} cop$$
$$\langle arrest \rangle \xrightarrow{patient} crook$$

- words rapidly combine $\Rightarrow$ **expectations** about upcoming input
  [Bicknell et al., 2010; Matsuki et al., 2011]

Psycholinguistic motivation

Generalized event knowledge [McRae and Matsuki, 2009]:

Prototypical knowledge about events and their participants
(first and second-hand experience, available in our memory)

- activated by words in isolation $\Rightarrow$ cue concepts from typical scenarios

  $\langle arrest \rangle \xrightarrow{agent} cop$
  $\langle arrest \rangle \xrightarrow{patient} crook$

- words rapidly combine $\Rightarrow$ **expectations** about upcoming input
  [Bicknell et al., 2010; Matsuki et al., 2011]

  Donna used the hose to wash her filthy...

Psycholinguistic motivation

Generalized event knowledge [McRae and Matsuki, 2009]:

Prototypical knowledge about events and their participants
(first and second-hand experience, available in our memory)

- activated by words in isolation $\Rightarrow$ cue concepts from typical scenarios

$$\langle arrest \rangle \xrightarrow{agent} cop$$
$$\langle arrest \rangle \xrightarrow{patient} crook$$

- words rapidly combine $\Rightarrow$ **expectations** about upcoming input
  [Bicknell et al., 2010; Matsuki et al., 2011]

  Donna used the hose to wash her filthy... **car** / **hair**

Psycholinguistic motivation

Generalized event knowledge [McRae and Matsuki, 2009]:

Prototypical knowledge about events and their participants
(first and second-hand experience, available in our memory)

- activated by words in isolation ⇒ cue concepts from typical scenarios

  $\langle arrest \rangle \xrightarrow{agent} cop$
  $\langle arrest \rangle \xrightarrow{patient} crook$

- words rapidly combine ⇒ **expectations** about upcoming input
  [Bicknell et al., 2010; Matsuki et al., 2011]

  Donna used the hose to wash her filthy... **car** / **hair**

- Operationalize thematic role-based expectations
  ⇒ **thematic fit**: typicality of a filler for a given argument slot

The Words-as-Cues Hypothesis [Zarcone, 2014]

The **baker** finished the icing

The Words-as-Cues Hypothesis [Zarcone, 2014]

The **baker** finished the icing $\rightarrow$ **spreading**

The Words-as-Cues Hypothesis [Zarcone, 2014]

The **child** finished the icing → **eating**

The Words-as-Cues Hypothesis [Zarcone, 2014]

The Words-as-Cues Hypothesis:

The Words-as-Cues Hypothesis [Zarcone, 2014]

The Words-as-Cues Hypothesis:

- **The Source Question**: What is the *source* of the covert event?

  $\Rightarrow$ generalized knowledge of events and their participants:
  covert events relevant to typical event scenarios are retrieved

The Words-as-Cues Hypothesis [Zarcone, 2014]

The Words-as-Cues Hypothesis:

- **The Source Question**: What is the *source* of the covert event?

  ⇒ generalized knowledge of events and their participants:
  covert events relevant to typical event scenarios are retrieved

  The baker finished the icing (**spreading** vs. **eating**)

The Words-as-Cues Hypothesis [Zarcone, 2014]

The Words-as-Cues Hypothesis:

- **The Source Question**: What is the *source* of the covert event?

  ⇒ generalized knowledge of events and their participants:
  covert events relevant to typical event scenarios are retrieved

  The baker finished the icing (**spreading** vs. **eating**)

- **The Trigger Question**: What *triggers* the metonymy?

  ⇒ low thematic fit between the verb and the object
  event-denoting nouns are better fillers for metonymic verbs

The Words-as-Cues Hypothesis [Zarcone, 2014]

The Words-as-Cues Hypothesis:

- **The Source Question**: What is the *source* of the covert event?

  ⇒ generalized knowledge of events and their participants:
  covert events relevant to typical event scenarios are retrieved

  The baker finished the icing (**spreading** vs. **eating**)

- **The Trigger Question**: What *triggers* the metonymy?

  ⇒ low thematic fit between the verb and the object
  event-denoting nouns are better fillers for metonymic verbs

  ✓ begin **the journey** ⇔ × begin **the book**

The Words-as-Cues Hypothesis [Zarcone, 2014]

The Words-as-Cues Hypothesis:

- **The Source Question**: What is the *source* of the covert event?

  ⇒ generalized knowledge of events and their participants:
  covert events relevant to typical event scenarios are retrieved

  The baker finished the icing (**spreading** vs. **eating**)

- **The Trigger Question**: What *triggers* the metonymy?

  ⇒ low thematic fit between the verb and the object
  event-denoting nouns are better fillers for metonymic verbs

  ✓ begin **the journey** ⇔ × begin **the book**

✓ ranked (testable) set of interpretations, determined by context

The source question

What is the source of the covert event?

The source question

What is the source of the covert event?

$\Downarrow$

generalized event knowledge: high thematic fit covert events, relevant to typical scenarios

The Source Question: Psycholinguistic evidence

Der Konditor	/	**das Kind**	hörte auf,	die	Glasur	aufzutragen	und	fing mit..
The baker	/	**the child**	finished	the	icing	to spread	and	started with...

The Source Question: Psycholinguistic evidence

Der Konditor	/	**das Kind**	hörte auf,	die Glasur	aufzutragen	und fing mit..
The baker	/	**the child**	finished	the icing	to spread	and started with...

$$\mathbf{RT}_{high} < \mathbf{RT}_{low}$$

The Source Question: Psycholinguistic evidence

Der Konditor / **das Kind** hörte auf, die Glasur aufzutragen und fing mit..
The baker / **the child** finished the icing to spread and started with...

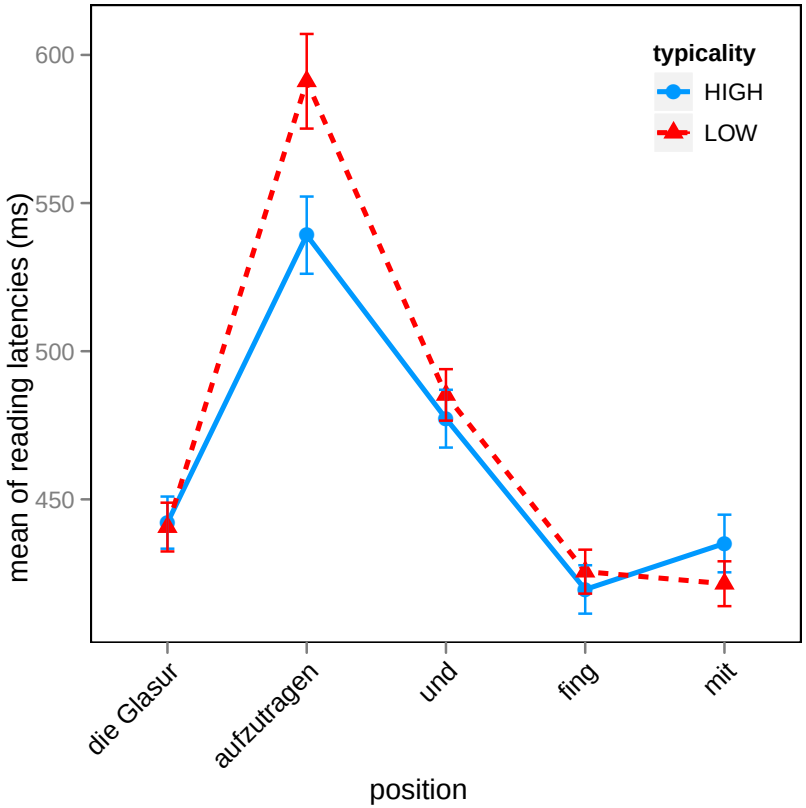

$\mathbf{RT}_{high} < \mathbf{RT}_{low}$

$\Downarrow$

facilitation effect for the **high typicality** condition

[Zarcone et al., 2014]

The Source Question: Computational modeling

A computational model of covert event interpretation for the Words-as-Cues Hypothesis:

The Source Question: Computational modeling

A computational model of covert event interpretation for the Words-as-Cues Hypothesis:

- **similarity-based**: ranked set of interpretations
  $\Rightarrow$ similar verbs, similar expectations

The Source Question: Computational modeling

A computational model of covert event interpretation for the Words-as-Cues Hypothesis:

- **similarity-based**: ranked set of interpretations
  $\Rightarrow$ similar verbs, similar expectations
- **compositional**: typical arguments $\rightarrow$ expectations for covert events
  $\Rightarrow$ integration of contextual cues

The Source Question: Computational modeling

A computational model of covert event interpretation for the Words-as-Cues Hypothesis:

- **similarity-based**: ranked set of interpretations
  $\Rightarrow$ similar verbs, similar expectations
- **compositional**: typical arguments $\rightarrow$ expectations for covert events
  $\Rightarrow$ integration of contextual cues
- **thematic-fit based** model
  $\Rightarrow$ the event with the best thematic fit is chosen

The Source Question: Computational modeling

A computational model of covert event interpretation for the Words-as-Cues Hypothesis:

- **similarity-based**: ranked set of interpretations
  $\Rightarrow$ similar verbs, similar expectations
- **compositional**: typical arguments $\rightarrow$ expectations for covert events
  $\Rightarrow$ integration of contextual cues
- **thematic-fit based** model
  $\Rightarrow$ the event with the best thematic fit is chosen

Konditor	aufhören	Glasur	auftragen	essen
Kind	aufhören	Glasur	essen	auftragen

Task: choose the high-typicality event over the low-typicality event (dataset from the psycholinguistic experiments)

The Source Question: a similarity-based model

Distributional Memory (DM) [Baroni and Lenci, 2010]

The Source Question: a similarity-based model

Distributional Memory (DM) [Baroni and Lenci, 2010]

- corpus-extracted weighted *word-link-word* tuples

	$\langle verb, bomb \rangle$	$\langle subj, kill \rangle$	$\langle verb, gun \rangle$	$\langle subj, shoot \rangle$	$\langle verb, book \rangle$	$\langle subj, read \rangle$
marine	40.0	82.1	85.3	44.8	3.2	3.3
teacher	5.2	7.0	9.3	4.7	48.4	53.6

The Source Question: a similarity-based model

Distributional Memory (DM) [Baroni and Lenci, 2010]

- corpus-extracted weighted *word-link-word* tuples

	$\langle verb, bomb \rangle$	$\langle subj, kill \rangle$	$\langle verb, gun \rangle$	$\langle subj, shoot \rangle$	$\langle verb, book \rangle$	$\langle subj, read \rangle$
marine	40.0	82.1	85.3	44.8	3.2	3.3
teacher	5.2	7.0	9.3	4.7	48.4	53.6

- Weighted expectations (**thematic fit**):
  $marine \xrightarrow{object} gun, bomb, \ldots$

The Source Question: a similarity-based model

Distributional Memory (DM) [Baroni and Lenci, 2010]

- corpus-extracted weighted *word-link-word* tuples

	⟨*verb,bomb*⟩	⟨*subj,kill*⟩	⟨*verb,gun*⟩	⟨*subj,shoot*⟩	⟨*verb,book*⟩	⟨*subj,read*⟩
marine	40.0	82.1	85.3	44.8	3.2	3.3
teacher	5.2	7.0	9.3	4.7	48.4	53.6

- Weighted expectations (**thematic fit**):
  $marine \xrightarrow{object} gun, bomb, \ldots$
- English DM [Baroni and Lenci, 2010] and German DM [Padoó and Utt, 2012]

The Source Question: a similarity-based model

A **compositional** model inspired by the ECU model [Lenci, 2011]:

The Source Question: a similarity-based model

A **compositional** model inspired by the ECU model [Lenci, 2011]:

- subject's expectations: *brewer* $\xrightarrow{subj}$ *event*

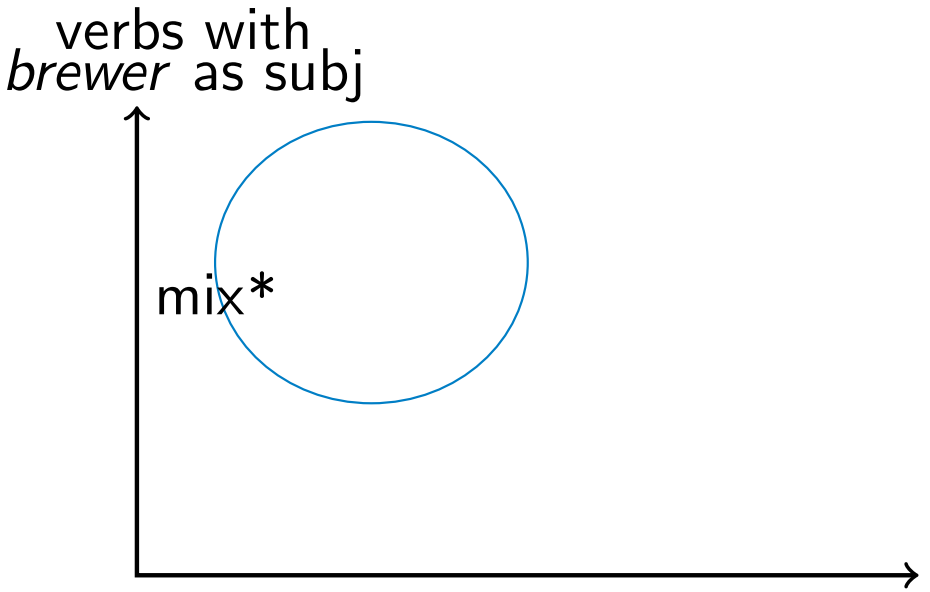

The Source Question: a similarity-based model

A **compositional** model inspired by the ECU model [Lenci, 2011]:

- subject's expectations: *brewer* $\xrightarrow{subj}$ *event*
- object's expectations: *beer* $\xrightarrow{obj}$ *event*

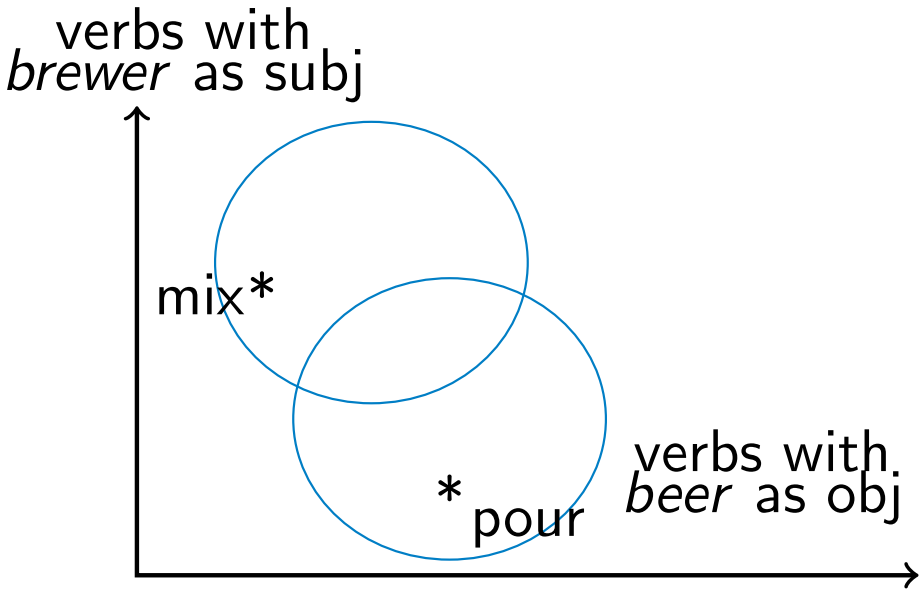

The Source Question: a similarity-based model

A **compositional** model inspired by the ECU model [Lenci, 2011]:

- subject's expectations: *brewer* $\xrightarrow{subj}$ *event*
- object's expectations: *beer* $\xrightarrow{obj}$ *event*
- metonymic verb's expectations: *finish* $\xrightarrow{comp^{-1}}$ *event*

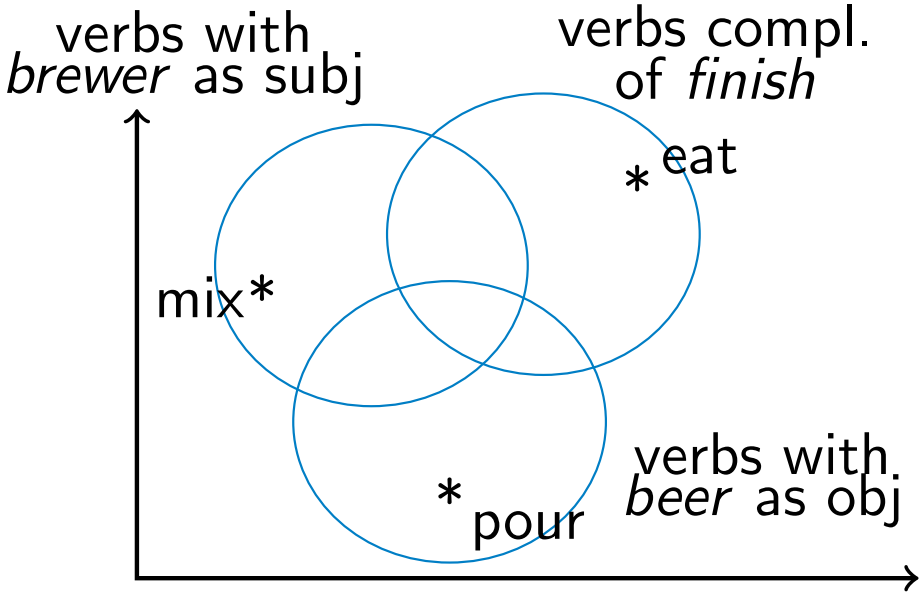

The Source Question: a similarity-based model

A **compositional** model inspired by the ECU model [Lenci, 2011]:

- subject's expectations: *brewer* $\xrightarrow{subj}$ *event*
- object's expectations: *beer* $\xrightarrow{obj}$ *event*
- metonymic verb's expectations: *finish* $\xrightarrow{comp^{-1}}$ *event*
  $\Rightarrow$ composed expectations: $\langle brewer, finish, beer \rangle \xrightarrow{covert\ \ event}$ *prototype*

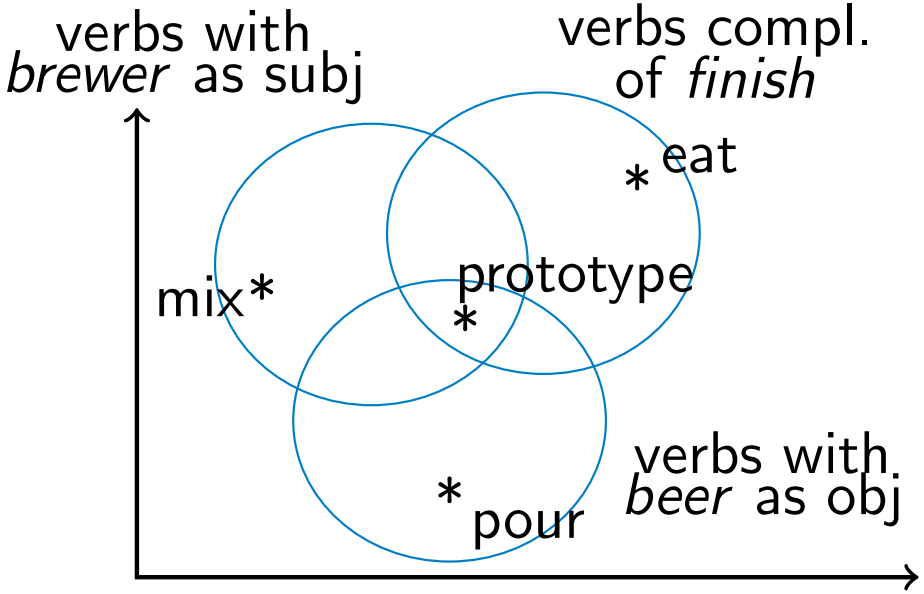

The Source Question: a similarity-based model

A **compositional** model inspired by the ECU model [Lenci, 2011]:

- subject's expectations: *brewer* $\xrightarrow{subj}$ *event*
- object's expectations: *beer* $\xrightarrow{obj}$ *event*
- metonymic verb's expectations: *finish* $\xrightarrow{comp^{-1}}$ *event*
  $\Rightarrow$ composed expectations: $\langle brewer, finish, beer \rangle \xrightarrow{covert\ \ event}$ *prototype*

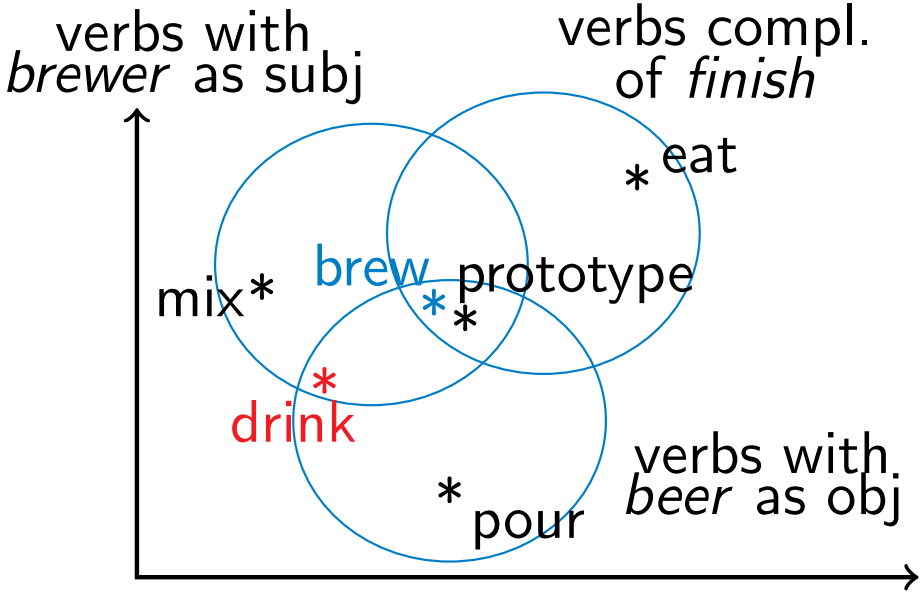

Thematic fit of an event: **similarity** to the prototype

$\langle brewer, finish, beer \rangle \rightarrow$ brew, drink

The Source Question: a similarity-based model

A **compositional** model inspired by the ECU model [Lenci, 2011]:

- subject's expectations: *brewer* $\xrightarrow{subj}$ *event*
- object's expectations: *beer* $\xrightarrow{obj}$ *event*
- metonymic verb's expectations: *finish* $\xrightarrow{comp^{-1}}$ *event*
  $\Rightarrow$ composed expectations: $\langle brewer, finish, beer \rangle \xrightarrow{covert\ \ event}$ *prototype*

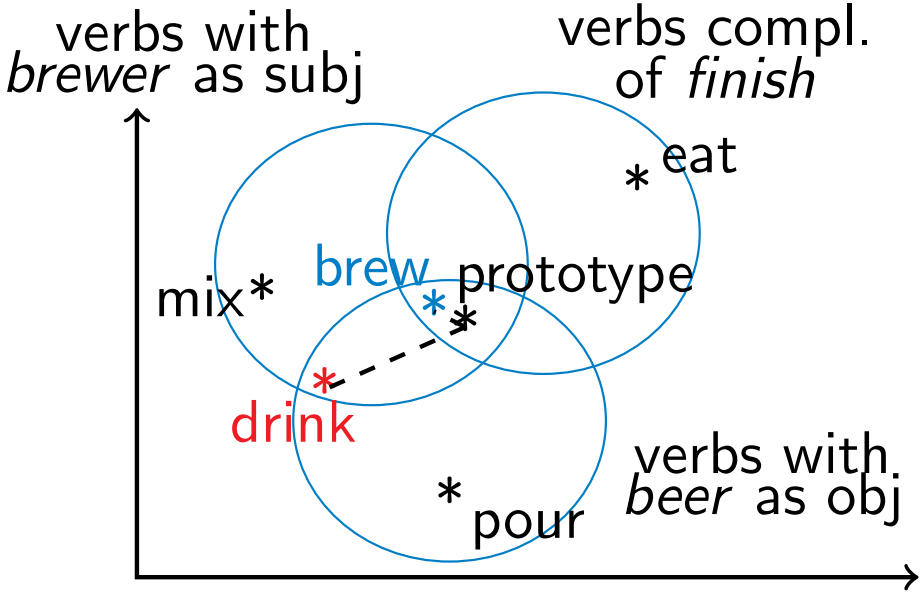

Thematic fit of an event: **similarity** to the prototype

$\langle brewer, finish, beer \rangle \rightarrow$ brew, drink

Sim(pr,brew) $>$ **Sim(pr,drink)**

The Source Question: Computational modeling

- Comparison with **Probabilistic Models** [Lapata et al., 2003]: covert event in a given context maximizes $P(s, v, o, e)$

The Source Question: Computational modeling

- Comparison with **Probabilistic Models** [Lapata et al., 2003]: covert event in a given context maximizes $P(s, v, o, e)$

		Probabilistic Models		Similarity-based Models	
	BL	**SOV**	**SO**	**SOV**	**SO**
Accuracy	50%	62%	75%	53%	68%
Coverage	100%	44%	75%	94%	98%
Backoff Acc.	50%	55%	**69%**	53%	**68%**

The Source Question: Computational modeling

- Comparison with **Probabilistic Models** [Lapata et al., 2003]: covert event in a given context maximizes $P(s, v, o, e)$

		Probabilistic Models		Similarity-based Models	
	BL	**SOV**	**SO**	**SOV**	**SO**
Accuracy	50%	62%	75%	53%	68%
Coverage	100%	44%	75%	94%	98%
Backoff Acc.	50%	55%	**69%**	53%	**68%**

- Both classes outperform the baselines (BL)

The Source Question: Computational modeling

- Comparison with **Probabilistic Models** [Lapata et al., 2003]: covert event in a given context maximizes $P(s, v, o, e)$

		Probabilistic Models		Similarity-based Models	
	BL	SOV	SO	SOV	SO
Accuracy	50%	62%	75%	53%	68%
Coverage	100%	44%	75%	94%	98%
Backoff Acc.	50%	55%	**69%**	53%	**68%**

- Both classes outperform the baselines (BL)
- **Similarity-based Models**: comparable accuracy to **Probabilistic Models** while guaranteeing higher coverage

The Source Question: Computational modeling

- Comparison with **Probabilistic Models** [Lapata et al., 2003]: covert event in a given context maximizes $P(s, v, o, e)$

		Probabilistic Models		Similarity-based Models	
	BL	SOV	SO	SOV	SO
Accuracy	50%	62%	75%	53%	68%
Coverage	100%	44%	75%	94%	98%
Backoff Acc.	50%	55%	**69%**	53%	**68%**

- Both classes outperform the baselines (BL)
- **Similarity-based Models**: comparable accuracy to **Probabilistic Models** while guaranteeing higher coverage
- SO models perform better than SOV models: the metonymic verb not very informative

The Trigger Question

What is the source of the covert event?

The Trigger Question

What is the source of the covert event?

$\Downarrow$

✓ generalized event knowledge: high thematic fit covert events, relevant to typical scenarios

The Trigger Question

What is the source of the covert event?

✓ generalized event knowledge: high thematic fit covert events, relevant to typical scenarios

What triggers the metonymy?

The Trigger Question

What is the source of the covert event?

⇓

✓ generalized event knowledge: high thematic fit covert events, relevant to typical scenarios

What triggers the metonymy?

⇓

low thematic fit between the verb and the object: event-denoting nouns are better fillers for metonymic verbs

The Trigger Question: Computational Modeling

	EV	EN
meton. v.	✓ The boy **started** the **fight**	× The boy **started** the **puzzle**
non-meton. v.	✓ The boy **saw** the **fight**	✓ The boy **saw** the **puzzle**

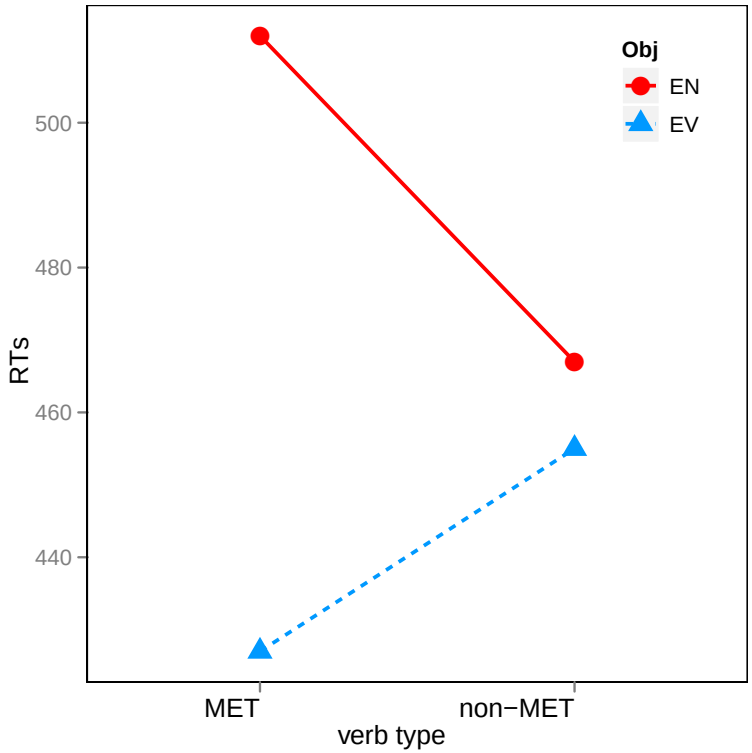

[Traxler et al., 2002]

The Trigger Question: Computational Modeling

	EV	EN
meton. v.	✓ The boy **started** the **fight**	× The boy **started** the **puzzle**
non-meton. v.	✓ The boy **saw** the **fight**	✓ The boy **saw** the **puzzle**

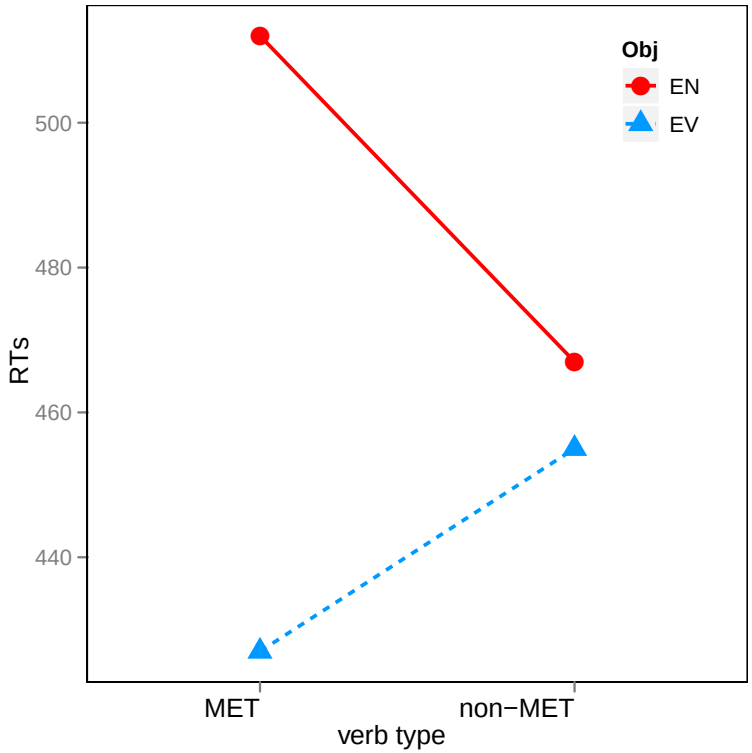

- longest reading times for metonymic verb + entity-denoting object

[Traxler et al., 2002]

The Trigger Question: Computational Modeling

	EV	EN
meton. v.	✓ The boy **started** the **fight**	× The boy **started** the **puzzle**
non-meton. v.	✓ The boy **saw** the **fight**	✓ The boy **saw** the **puzzle**

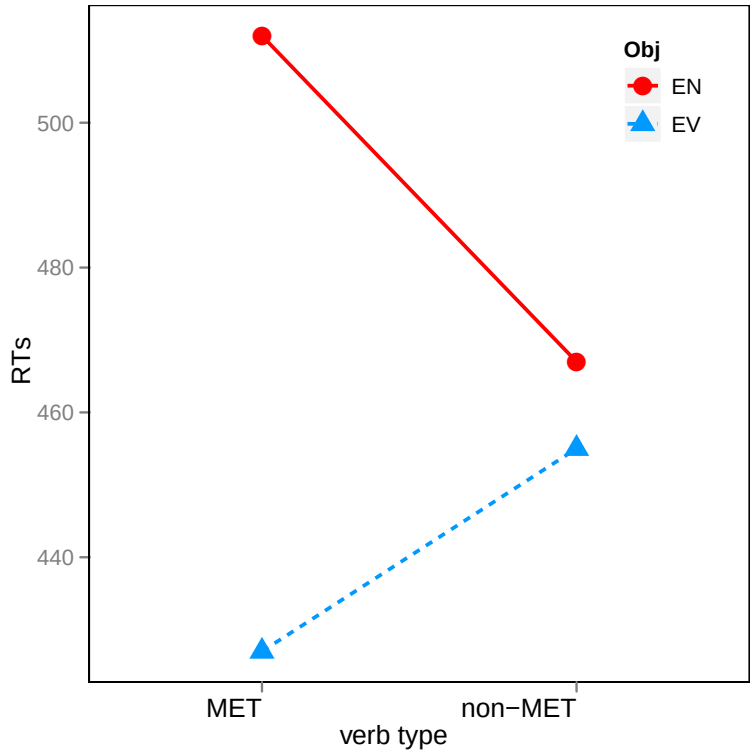

- longest reading times for metonymic verb + entity-denoting object
  - type-clash or lower thematic fit?

[Traxler et al., 2002]

The Trigger Question: Computational Modeling

	EV	EN
meton. v.	✓ The boy **started** the **fight**	× The boy **started** the **puzzle**
non-meton. v.	✓ The boy **saw** the **fight**	✓ The boy **saw** the **puzzle**

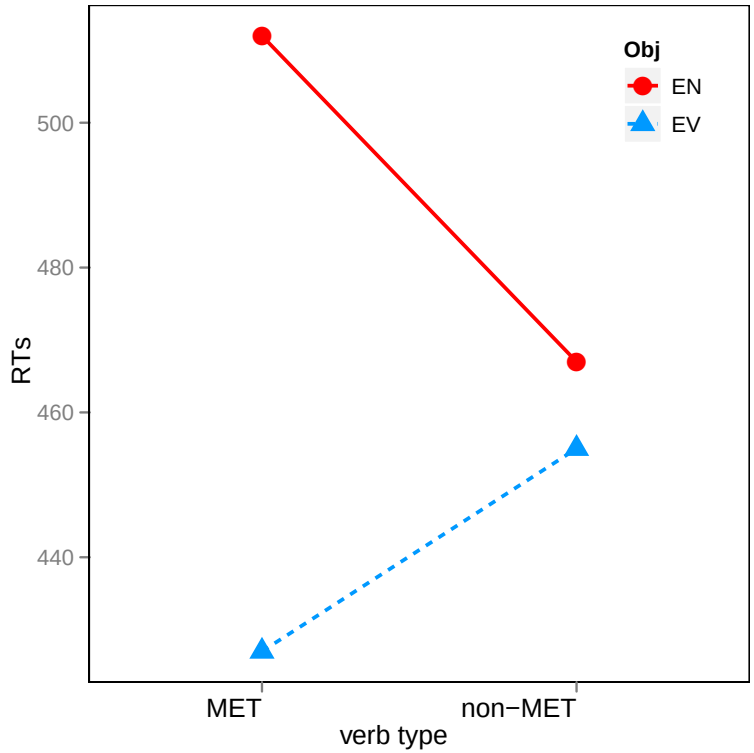

- longest reading times for metonymic verb + entity-denoting object
  - type-clash or lower thematic fit?
  - computational model of thematic fit (no explicit type information)

[Traxler et al., 2002]

The Trigger Question: Computational Modeling

The boy [started / saw] **the fight**$_{EV}$ / **the puzzle**$_{EN}$

RT Study

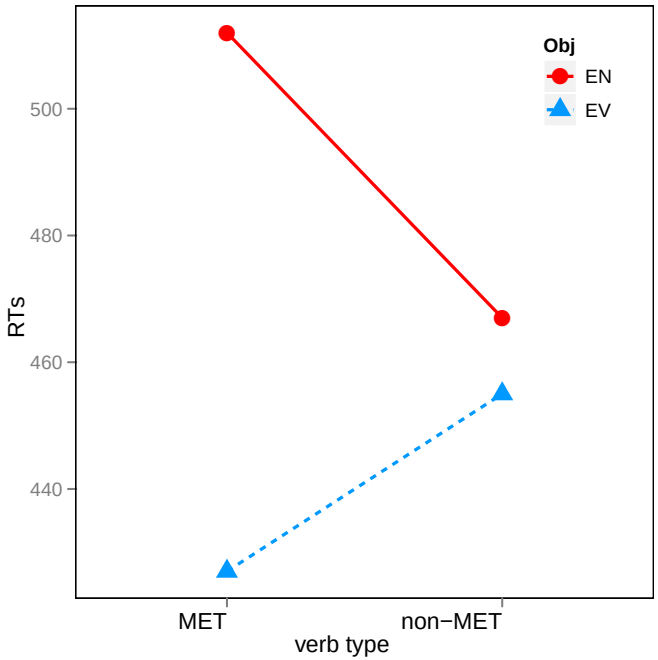

Thematic-fit Model

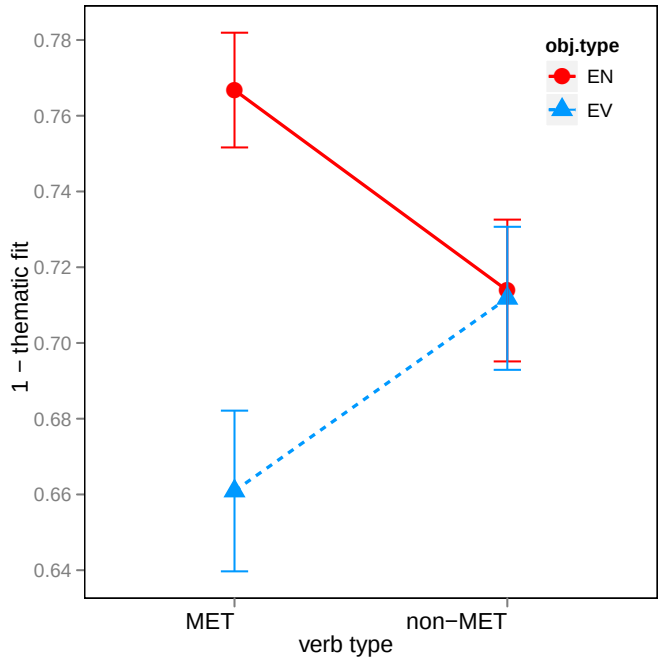

The Trigger Question: Computational Modeling

The boy [started / saw] **the fight**$_{EV}$ / **the puzzle**$_{EN}$

RT Study

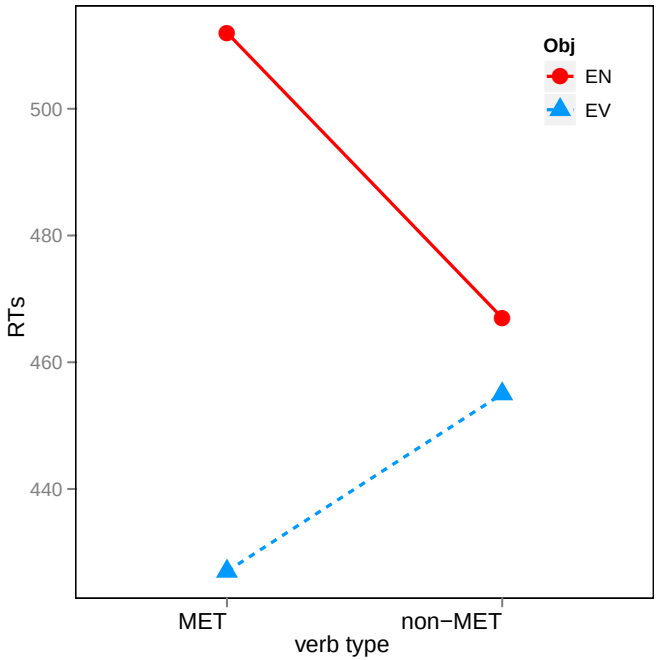

Thematic-fit Model

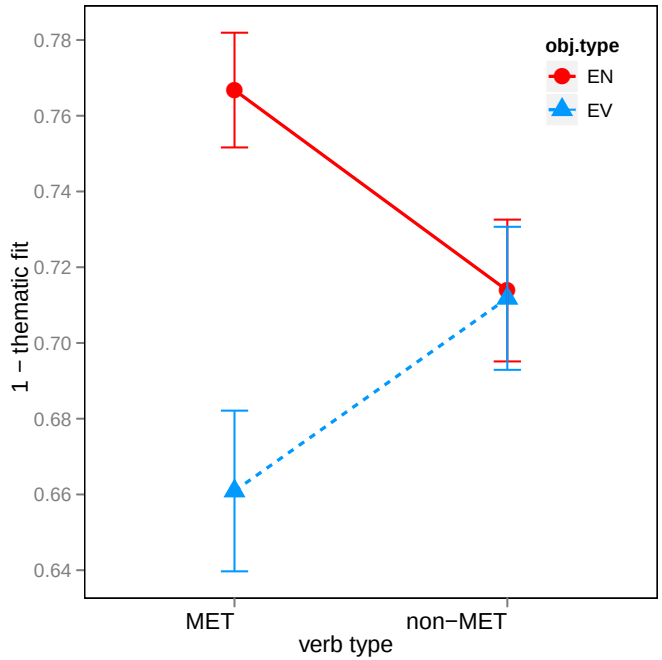

highest processing costs

[Traxler et al., 2002]

The Trigger Question: Computational Modeling

The boy [started / saw] **the fight**$_{EV}$ / **the puzzle**$_{EN}$

RT Study

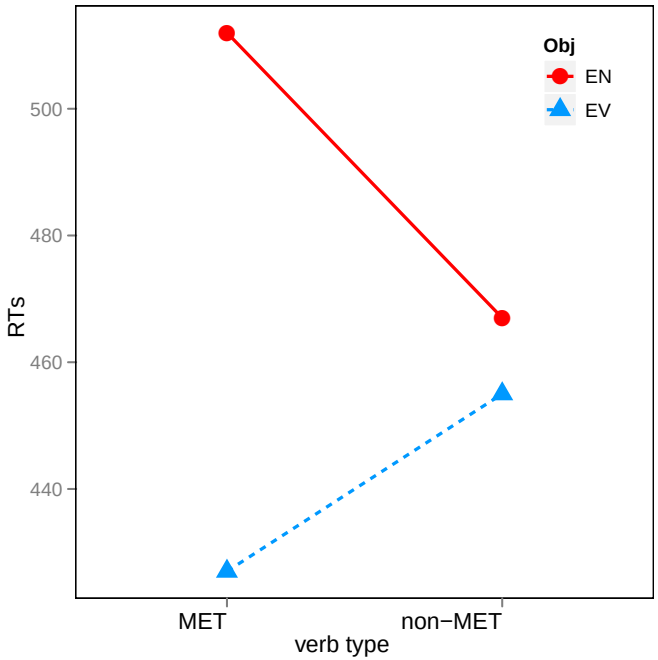

highest processing costs

[Traxler et al., 2002]

Thematic-fit Model

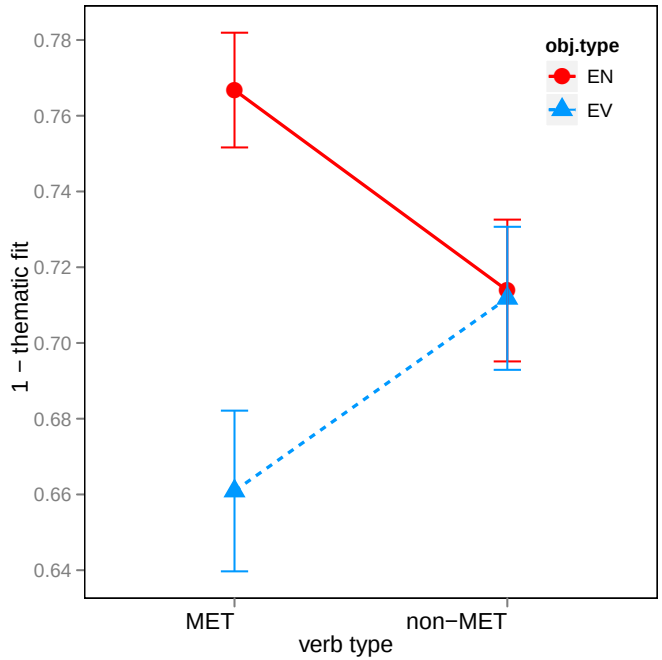

highest $(1 - th.fit)$ scores

[Zarcone et al., 2013]

The Trigger Question: Computational Modeling

The Trigger Question: Computational Modeling

What triggers the metonymy?

The Trigger Question: Computational Modeling

What triggers the metonymy?

low thematic fit between the verb and the object:
event-denoting nouns are
better fillers for metonymic verbs

The Trigger Question: Computational Modeling

What triggers the metonymy?

$\Downarrow$

low thematic fit between the verb and the object:
event-denoting nouns are
better fillers for metonymic verbs

- metonymic combinations distinguished in terms of **thematic fit**
- distributional characterization of metonymic verbs in terms of their **selectional behavior**

[Zarcone et al.2013; Utt et al., 2013]

The Trigger Question: Computational Modeling

What triggers the metonymy?

$\Downarrow$

low thematic fit between the verb and the object:
event-denoting nouns are
better fillers for metonymic verbs

- metonymic combinations distinguished in terms of **thematic fit**
- distributional characterization of metonymic verbs in terms of their **selectional behavior**

[Zarcone et al.2013; Utt et al., 2013]

Is thematic fit or type responsible for triggering the logical metonymy?

The Trigger Question: Psycholinguistic Evidence

Das Geburtstagskind hat mit	/	/	/	**angefangen**.
The birthday boy has with	/	/	/	**begun**.

[Zarcone and Padó, 2013]

The Trigger Question: Psycholinguistic Evidence

Das Geburtstagskind hat mit	**den Geschenken**	/	**der Suppe**	/	/	**angefangen**.
The birthday boy has with	**the presents**	/	**the soup**	/	/	**begun**.

2 EN-denoting objects,

[Zarcone and Padó, 2013]

The Trigger Question: Psycholinguistic Evidence

Das Geburtstagskind hat mit	/	/	**der Feier**	/ **der Schicht**	**angefangen**.
The birthday boy has with	/	/	**the party**	/ **the shift**	**begun**.

2 EN-denoting objects,
2 EV-denoting objects

[Zarcone and Padó, 2013]

The Trigger Question: Psycholinguistic Evidence

Das Geburtstagskind hat mit **den Geschenken** / / **der Feier** / **angefangen**.

The birthday boy has with **the presents** / / **the party** / **begun**.

2 high thematic fit objects (1 EN, 1 EV),

[Zarcone and Padó, 2013]

The Trigger Question: Psycholinguistic Evidence

Das Geburtstagskind hat mit	/	**der Suppe**	/	**der Schicht** **angefangen**.
The birthday boy has with	/	**the soup**	/	**the shift** **begun**.

2 high thematic fit objects
(1 EN, 1 EV),
2 low thematic fit objects
(1 EN, 1 EV)

[Zarcone and Padó, 2013]

The Trigger Question: Psycholinguistic Evidence

Das Geburtstagskind hat mit **den Geschenken** / **der Suppe** / **der Feier** / **der Schicht** **angefangen**.

The birthday boy has with **the presents** / **the soup** / **the party** / **the shift** **begun**.

Prediction from a type account:

$$\mathbf{RT}_{EV,HIGH} < \mathbf{RT}_{EN,HIGH}$$
$$\mathbf{RT}_{EV,LOW} < \mathbf{RT}_{EN,LOW}$$

[Zarcone and Padó, 2013]

The Trigger Question: Psycholinguistic Evidence

Das Geburtstagskind hat mit **den Geschenken** / **der Suppe** / **der Feier** / **der Schicht** **angefangen**.

The birthday boy has with **the presents** / **the soup** / **the party** / **the shift** **begun**.

Prediction from a thematic fit account:

$\mathbf{RT}_{EV,HIGH} < \mathbf{RT}_{EV,LOW}$

$\mathbf{RT}_{EN,HIGH} < \mathbf{RT}_{EN,LOW}$

[Zarcone and Padó, 2013]

The Trigger Question: Psycholinguistic Evidence

Das Geburtstagskind hat mit **den Geschenken** / **der Suppe** / **der Feier** / **der Schicht** **angefangen**.

The birthday boy has with **the presents** / **the soup** / **the party** / **the shift** **begun**.

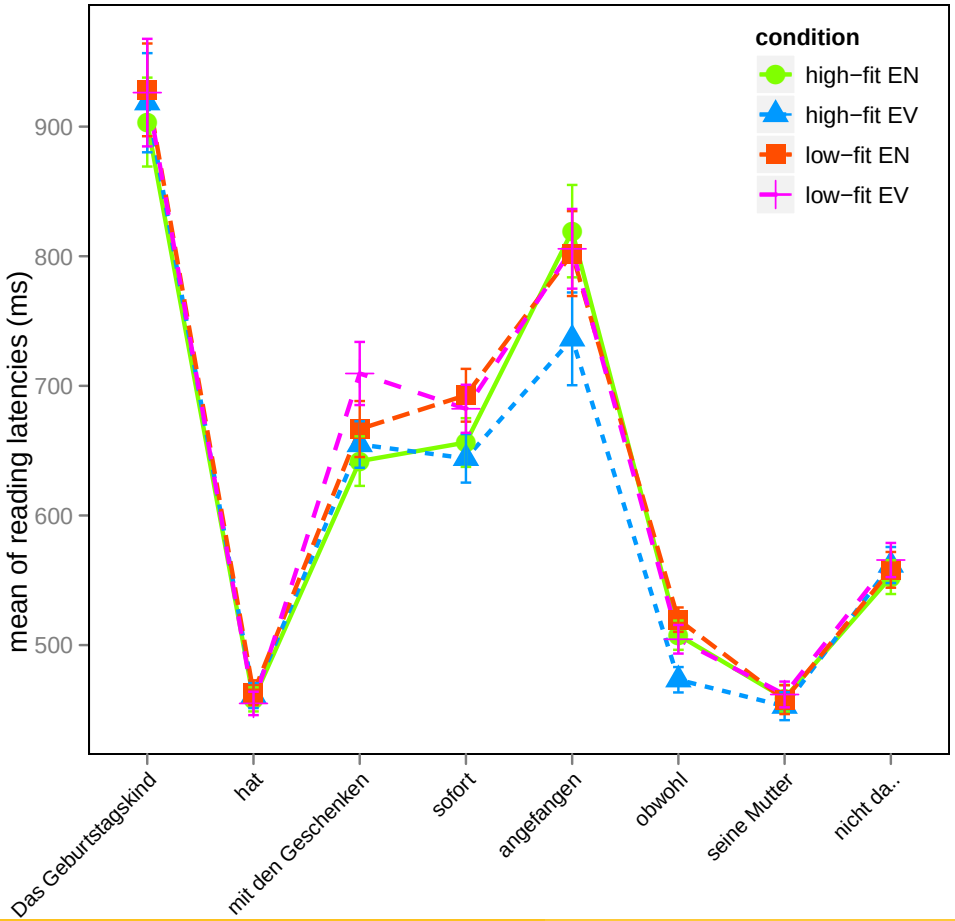

Results:

$\mathbf{RT}_{EV,HIGH} < \mathbf{RT}_{EV,LOW}$,
$\mathbf{RT}_{EN,HIGH}$, $\mathbf{RT}_{EN,LOW}$

[Zarcone and Padó, 2013]

The Trigger Question: Psycholinguistic Evidence

Das Geburtstagskind hat mit **den Geschenken** / **der Suppe** / **der Feier** / **der Schicht** **angefangen**.

The birthday boy has with **the presents** / **the soup** / **the party** / **the shift** **begun**.

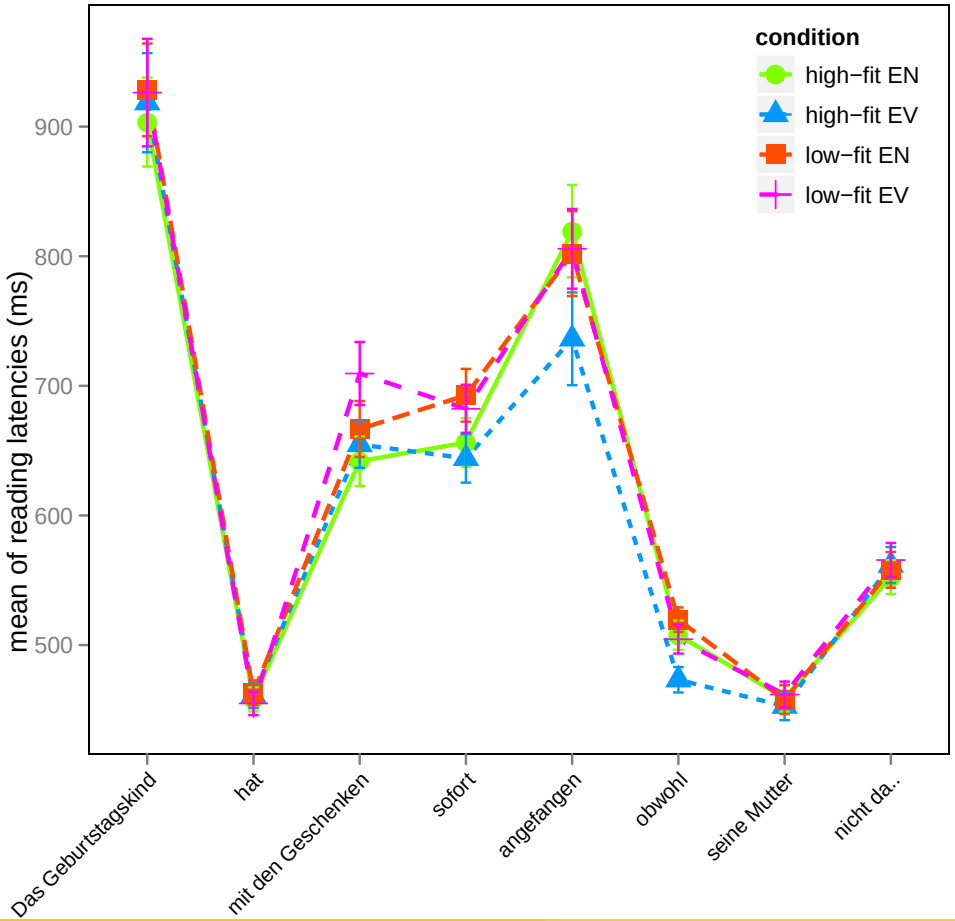

Results:

$\mathbf{RT}_{EV,HIGH} < \mathbf{RT}_{EV,LOW}$,

$\mathbf{RT}_{EN,HIGH}$, $\mathbf{RT}_{EN,LOW}$

$\Downarrow$

Both

type and **thematic fit**

are necessary

[Zarcone and Padó, 2013]

The Words-as-Cues Hypothesis revisited

What is the source of the covert event?

The Words-as-Cues Hypothesis revisited

What is the source of the covert event?

$\Downarrow$

✓ generalized event knowledge: high thematic fit covert events, relevant to typical scenarios

The Words-as-Cues Hypothesis revisited

What is the source of the covert event?

$\Downarrow$

✓ generalized event knowledge: high thematic fit covert events, relevant to typical scenarios

What triggers the metonymy?

The Words-as-Cues Hypothesis revisited

What is the source of the covert event?

$\Downarrow$

✓ generalized event knowledge: high thematic fit covert events, relevant to typical scenarios

What triggers the metonymy?

$\Downarrow$

low thematic fit between the verb and the object:

The Words-as-Cues Hypothesis revisited

What is the source of the covert event?

$\Downarrow$

✓ generalized event knowledge: high thematic fit covert events, relevant to typical scenarios

What triggers the metonymy?

$\Downarrow$

× low thematic fit between the verb and the object:

The Words-as-Cues Hypothesis revisited

What is the source of the covert event?

⇓

✓ generalized event knowledge: high thematic fit covert events, relevant to typical scenarios

What triggers the metonymy?

⇓

× low thematic fit between the verb and the object:

✓ type + thematic fit: expectations for high-typicality event-denoting objects

The Words-as-Cues Hypothesis revisited

Incremental, context-driven and expectation-driven specification process:

The Words-as-Cues Hypothesis revisited

Incremental, context-driven and expectation-driven specification process:

- metonymic verbs' selectional properties
  $\rightarrow$ expectations for high-typicality event-denoting objects

The Words-as-Cues Hypothesis revisited

Incremental, context-driven and expectation-driven specification process:

- metonymic verbs' selectional properties
  $\rightarrow$ expectations for high-typicality event-denoting objects
- contextual cues
  $\rightarrow$ update expectations for high-typicality covert events

The Words-as-Cues Hypothesis revisited

Incremental, context-driven and expectation-driven specification process:

- metonymic verbs' selectional properties
  $\rightarrow$ expectations for high-typicality event-denoting objects
- contextual cues
  $\rightarrow$ update expectations for high-typicality covert events
- semantic type as yet another constraint contributing to the expectation building process

The Words-as-Cues Hypothesis revisited

Incremental, context-driven and expectation-driven specification process:

- metonymic verbs' selectional properties
  $\rightarrow$ expectations for high-typicality event-denoting objects
- contextual cues
  $\rightarrow$ update expectations for high-typicality covert events
- semantic type as yet another constraint contributing to the expectation building process

A hybrid model [Resnik, 1996; Schulte im Walde, 2006]:

The Words-as-Cues Hypothesis revisited

Incremental, context-driven and expectation-driven specification process:

- metonymic verbs' selectional properties
  $\rightarrow$ expectations for high-typicality event-denoting objects
- contextual cues
  $\rightarrow$ update expectations for high-typicality covert events
- semantic type as yet another constraint contributing to the expectation building process

A hybrid model [Resnik, 1996; Schulte im Walde, 2006]:

- semantic type emerging from observed distributional behavior

The Words-as-Cues Hypothesis revisited

Incremental, context-driven and expectation-driven specification process:

- metonymic verbs' selectional properties
  $\rightarrow$ expectations for high-typicality event-denoting objects
- contextual cues
  $\rightarrow$ update expectations for high-typicality covert events
- semantic type as yet another constraint contributing to the expectation building process

A hybrid model [Resnik, 1996; Schulte im Walde, 2006]:

- semantic type emerging from observed distributional behavior
- selectional preferences as distributions over classes of fillers

The Words-as-Cues Hypothesis revisited

Incremental, context-driven and expectation-driven specification process:

- metonymic verbs' selectional properties
  $\rightarrow$ expectations for high-typicality event-denoting objects
- contextual cues
  $\rightarrow$ update expectations for high-typicality covert events
- semantic type as yet another constraint contributing to the expectation building process

A hybrid model [Resnik, 1996; Schulte im Walde, 2006]:

- semantic type emerging from observed distributional behavior
- selectional preferences as distributions over classes of fillers
- encoding both thematic fit and type

Lexicon and world knowledge

Work on generalized event knowledge:

Lexicon and world knowledge

Work on generalized event knowledge:

- it is possible to make predictions and verify hypotheses regarding the role of world knowledge in linguistic processing

Lexicon and world knowledge

Work on generalized event knowledge:

- it is possible to make predictions and verify hypotheses regarding the role of world knowledge in linguistic processing
- evidence for early use of rich knowledge about typical events and their participants

Lexicon and world knowledge

Work on generalized event knowledge:

- it is possible to make predictions and verify hypotheses regarding the role of world knowledge in linguistic processing
- evidence for early use of rich knowledge about typical events and their participants
  - $\Rightarrow$ during processing of explicit input
  - $\Rightarrow$ in covert event interpretation

(END PRETENDING TO BE ALESSANDRA)

So we trace some aspects of covert event reconstruction to semantic type. . .

. . . but what happens for events that violate our expectations?

Part 2: psycholinguistic evidence for distinct processing modalities

World knowledge and affordances

I'm specifically referring to the idea of affordances:

- From Gibson (1977): action possibilities latent in an environment.
  - E.g., we know that a knife can be used to cut or that a cake can be eaten.
- The problem of affordance knowledge in language:
  - What does it mean to know that some referents have particular potentials?
  - What knowledge is required to put that knowledge into use over time?

So meta

There's knowledge, and then there's knowledge.

People update representations over the course of a sentence.

- Sato et al. (2013): exploit Japanese word order to test whether object representations change over time.
- Paradigm: self-paced reading with picture selection task.
  - Sentences with canonical vs. non-canonical object shape expectation.
  - Final verb forces change from canonical to non-canonical expectation and vice versa.
  - Image verification – canonical or non-canonical object displayed, "yes" answer expected (even if false) either before (exp 1) or after verb (exp 2)
  - Fillers have possible "no" answers.

There's knowledge, and then there's knowledge.

What does this look like?

(1) a. Nana-ga reezooko-nonakani tamago-o subayaku otoshita
Nana-NOM refrigerator-LOC instantly dropped
'Nana dropped an egg in the refrigerator instantly.'
(canonical → non-canonical)

b. Nana-ga huraipan-nouede tamago-o subayaku korogashita
Nana-NOM pan-LOC egg-ACC instantly rolled
'Nana rolled an egg in the pan instantly.'
(non-canonical → canonical)

Then picture of broken or unbroken egg. (sometimes before verb)

There's knowledge, and then there's knowledge.

Results:

- Picture before verb: subjects respond "yes" faster when image matches canonicity.
- Picture after verb: subjects respond "yes" faster when image matches verb-induced state change.
- **No effect of original canonicity.**
- Reading time: faster in critical region for canonical fragments.

⇒ **mental representation updated dynamically, despite effort of representing non-canonicity**

There's knowledge, and then there's knowledge.

This suggests two levels of "affordance knowledge":

First-order affordance knowledge

Direct knowledge of object potentials, object state canonicity, etc.
⇒ Roughly analogous to **probabilistic/distributional** knowledge of object state?

Higher-order affordance knowledge

Knowledge of relationships between object states and potentials, instrument use, how to reanalyze object state given cues, etc.
⇒ Roughly analogous to knowledge of **plausibility**?

How do we represent "affordance knowledge"?

For first-order, you already know the answer: distributional models are getting pretty good at this.

- Distributional Memory, neural networks, etc.
- Correlations with static human judgements of what objects are *for*, what you can do *to* them, etc.
- Potential bias towards highly "salient" event combinations (through feature selection, use of mutual info stats).

For higher-order, this is what we do currently:

- Logical inference, reasoning with knowledge-bases, etc.
- Ever-open problems: empirical grounding, handling unexpectedness, representing effort.

Distributional knowledge is not necessarily complete.

Truism: the most effective learning algorithms are statistical/distributional.

- Surprisal, entropy, mutual information, etc: seem to kill two birds with one stone. (See what I did there?)
  - Apparently easy to obtain from corpus statistics.
  - Surprisal, at least, is a promising proxy for processing effort.
- But is plausibility (aka higher-order affordance knowledge?) directly captured here?
  - That was a rhetorical question to which the answer is NO.
  - But it does appear to have a separate effect on the brain...

Is plausibility really a thing?

Amsel et al. (2015), perceptuomotor features vs. event-relatedness. EEG experiment with the following sort of stimuli:

(2) During the African safari, Javier though he spotted a giraffe off in the distance. He quickly reached for his pair of . . . in the back.

a. binoculars (expected)
b. slippers (perceptuomotor-related anomaly)
c. zebras (event-related anomaly)
d. bubbles (unrelated)

This example for action affordances – but also tested several other event relations (e.g. taste, sound) from previous norming study.

But maybe giraffes LIKE bubbles!

Is plausibility really a thing?

Result summary:

- N400 amplitudes: expected $<$ perceptuomotor/event-related anomaly $<$ unrelated
- Perceptuomotor N400 effect – stronger frontal; event-related N400 effect – stronger occipital.
- Is plausibility the main driver of the N400 effect?
  - Plausibility Likert study – only marginal difference between perceptuomotor and event-related anomaly.
  - Reanalyzed N400 data across closely plausibility matched event vs. perceptuomotor vs. unrelated anomaly
    - Unrelated significantly stronger effect than event or perceptuomotor – probably not plausibility.

So plausibility is not "real" – is it?

Now we edge into speculations we'll get to tomorrow. . .