Computing Dynamic Meanings: Building Integrated Competence-Performance Theories for Semantics

Days 4-5: Mechanistic processing models for formal semantics (DRT + ACT-R + Bayes)

Jakub Dotlačil & Adrian Brasoveanu

ESSLLI 2018, August 9-10 2018

Goals for this final part of the course

- introduce mechanistic processing models for formal semantics that integrate
  - dynamic semantics, specifically, Discourse Representation Theory (DRT, Kamp 1981; Kamp and Reyle 1993)
  - the ACT-R cognitive architecture
- show how to embed these mechanistic processing models into Bayesian models and fit them to experimental data
  - focus on the fan experiment in Anderson (1974); see also Anderson and Reder (1999)
  - briefly discuss experiments investigating the interaction of prononominal and presuppositional cataphora and conjunctions vs. conditionals reported in Brasoveanu and Dotlačil (2015)

An ACT-R based left-corner parser for DRT

Main idea:

- give a mechanistic account of the simultaneous syntactic and semantic parsing process
- using an independently motivated, general cognitive architecture (e.g., Lewis and Vasishth 2005 for syntax)

Once DRT+ACT-R mechanistic processing models available:

- fit them to experimental data by embedding them in Bayesian models
- estimate parameter values and quantify our uncertainty about them
- do quantitative model comparison for qualitative / symbolic theories of semantic phenomena

Why ACT-R? [quick recap]

- widely used; well developed goal structuring, memory structures and access, interaction with peripherals (vision and motor modules) etc.
- modular architecture: incorporate linguistics while largely maintaining previous body of cognitive psychology research
- hybrid architecture: symbolic (discrete / qualitative) and subsymbolic (continuous / quantitative) components
- symbolic structures to describe human behavior (Marr's algorithmic level)
- well suited to encode linguistic theories
- subsymbolic components to model performance, can be fit to experimental data (e.g., via embedding in Bayesian models), mainly latencies and accuracy, but also neurological data (see Anderson 2007 a.o.)

Why DRT as our semantic framework?

Meaning representations, i.e., Discourse Representation Structures (DRSs):

- are well understood mathematically and widely used in formal semantics
- can simultaneously function as:
  - meaning representations (logical forms)
  - their content (model structures)

    . . . at least for atomic DRSs (Kamp and Reyle, 1993, pp. 96-97)
- support rich theories of intra- and cross-sentential anaphora and presupposition resolution

DRSs can be thought of as mental models (Johnson-Laird, 1983, 2004; Johnson-Laird et al., 1989), with advantages: (i) richer array of representations and operations; (ii) comprehensive mathematical theory of their structure and interpretation (dynamic logic / dynamic semantics); (iii) used by linguists for a wide variety of semantic phenomena.

DRT: an example

Kamp and Reyle (1993)

(1) A delegate arrived. She registered.

- the indefinite *a delegate* in first sentence introduces a discourse referent (dref) x
- ... that can be picked up by the pronoun *she* in the second sentence
- resulting interpretation: there is an entity x that's a delegate, that arrived and registered
- cross-sentential pronoun binding of this sort: not immediately available in classical static semantics
- even more problematic: cross-clausal 'donkey' anaphora:
  If a farmer owns a donkey, she feeds it.
  Every farmer who owns a donkey feeds it.

Interpretation in DRT

Kamp (1981), Kamp and Reyle (1993)

Broadly:

1. build a syntactic structure for the current sentence
2. apply construction and interpretation rules (roughly) top-down to the syntactic structure to create a discourse representation structure (DRS)
3. interpret the DRS
4. move to next sentence, go to step 1

DRT: an example (ctd.)

Kamp and Reyle (1993)

(2) A delegate arrived. She registered.

DRT: an example (ctd.)

Kamp and Reyle (1993)

(2) A delegate arrived. She registered.

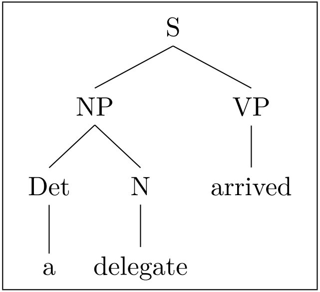

DRT: an example (ctd.)

Kamp and Reyle (1993)

(2) A delegate arrived. She registered.

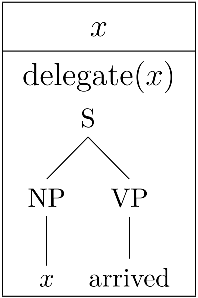

DRT: an example (ctd.)

Kamp and Reyle (1993)

(2) A delegate arrived. She registered.

x
delegate(x) arrive(x)

DRT: an example (ctd.)

Kamp and Reyle (1993)

(2) A delegate arrived. She registered.

$$\begin{array}{|c|}\hline x \\ \hline \text{delegate}(x) \\ \text{arrive}(x) \\ \hline \end{array} \quad ;$$

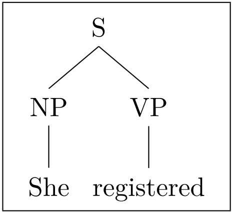

DRT: an example (ctd.)

Kamp and Reyle (1993)

(2) A delegate arrived. **She registered.**

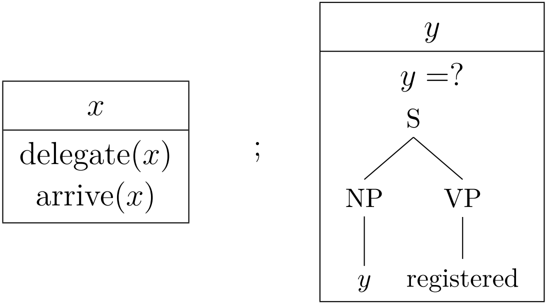

DRT: an example (ctd.)

Kamp and Reyle (1993)

(2) A delegate arrived. She registered.

$$\begin{array}{|c|}\hline x \\ \hline \text{delegate}(x) \\ \text{arrive}(x) \\ \hline \end{array} \quad ; \quad \begin{array}{|c|}\hline y \\ \hline y = x \\ \hline \end{array}$$

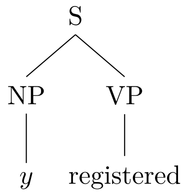

DRT: an example (ctd.)

Kamp and Reyle (1993)

(2) A delegate arrived. She registered.

$$\begin{array}{|c|}\hline x \\ \hline \text{delegate}(x) \\ \text{arrive}(x) \\ \hline \end{array} \quad ; \quad \begin{array}{|c|}\hline y \\ \hline y = x \\ \text{register}(y) \\ \hline \end{array}$$

Semantic interpretation and syntactic structure

DRS construction / semantic interpretation is defined in Kamp and Reyle (1993) to work off of / after a full syntactic structure is available.

- following the standard Montagovian / mathematical logic format

Need to modify it so that

- DRS construction
- syntactic structure building

proceed in parallel, incrementally.

DRT: synchronizing syntax and semantics

- A delegate arrived.

DRT: synchronizing syntax and semantics

- A delegate arrived.

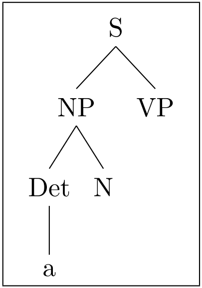

DRT: synchronizing syntax and semantics

- A delegate arrived.

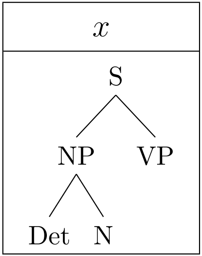

DRT: synchronizing syntax and semantics

- A **delegate** arrived.

DRT: synchronizing syntax and semantics

- A **delegate** arrived.

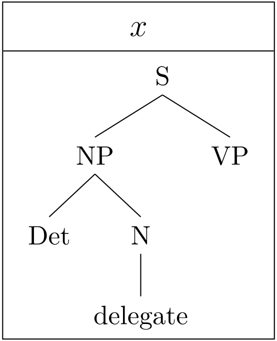

DRT: synchronizing syntax and semantics

- A **delegate** arrived.

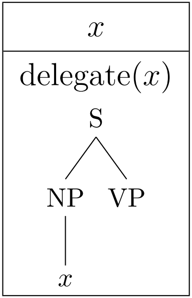

DRT: synchronizing syntax and semantics

- A delegate arrived.

DRT: synchronizing syntax and semantics

- A delegate arrived.

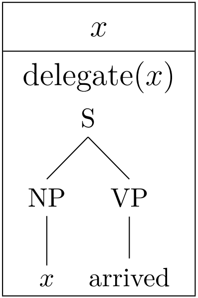

DRT: synchronizing syntax and semantics

- A delegate arrived.

$$\boxed{\begin{array}{c} x \\ \hline \text{delegate}(x) \\ \text{arrive}(x) \end{array}}$$

An ACT-R based left-corner parser for DRT

Distributing the parser over memory components:

1. declarative memory: lexical knowledge, knowledge of incrementally constructed syntactic structures and DRSs
2. procedural memory: knowledge of grammar (both syntax and semantics)
3. imaginal buffer: holds the currently constructed syntactic structure / parse state
4. goal buffer: holding stack of expected syntactic categories, argument information for semantics etc.
5. semantic buffer (`discourse_context`): holds the currently constructed DRS
6. a second semantic buffer will hold unresolved DRSs for pronominal / presuppositional anaphora / cataphora

An ACT-R based left-corner parser for DRT

Cognitive process of simultaneous syn / sem parsing per word:

1. scan word
2. lexical access: move lexical information about word into retrieval buffer
3. check goal buffer and (possibly) check imaginal buffer and `discourse_context` buffer for current syn / sem parse
4. based on buffer contents:
   - create syntactic structure and attach it to current parse state
   - create new semantic representation or update current one
   - store syntax and semantics in dec. mem. and/or
   - keep (parts of) current parse in buffer
   - update goals
5. go to step 1

Left-corner parser for DRT: an example

- A boy sleeps.

```
A --- ------.
```

Input

- Goals: S
- Found:
  a / Det

Output

- Goals: N, NP, VP
- Syntax:

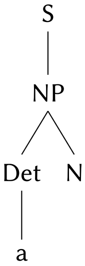

- Semantics:

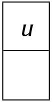

Left-corner parser for DRT: an example (ctd.)

- A boy sleeps.

`- boy ------.`

Input

- Goals: N NP S
- Found:
  boy / N

Output

- Goals: VP
- Syntax:

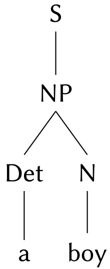

- Semantics:

u
$\mathrm{boy}(u)$

Left-corner parser for DRT: an example (ctd.)

- A boy sleeps.

```
- --- sleeps.
```

Input

- Goals: VP
- Found: *sleeps* / V

Output

- Goals: –
- Syntax:

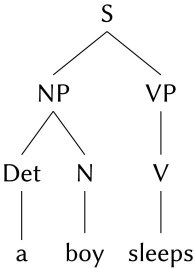

- Semantics:

u
$\text{boy}(u)$ $\text{sleep}(u)$

Fan experiment: updating & evaluating DRSs

Anderson (1974); Anderson and Reder (1999)

Model simultaneous syn/sem parsing in a classical exp. about:

- how basic propositional information of the kind encoded by atomic DRSs
- . . . is (stored and) retrieved from declarative memory

[not exactly how fan exp. originally conceptualized in '74/'99 b/c work at formal semantics/computational psycholinguistics interface more recent]

Involves essential components of real-time sem interpretation:

i. compose/integrate sem representations of new words with sem representation of previous discourse
ii. update discourse sem representation accordingly
iii. evaluate new sem representations relative to our mental model of the world
iv. integrate their content into our world knowledge database

Aside: atomic DRSs

Atomic DRSs:

- equivalent to
  - atomic first-order logic formulas
  - conjunctions thereof
  - atomic formulas or conjunctions thereof + a prefix of existential quantifiers
- multiple atomic DRSs can be merged into a single atomic DRS
- . . . with caveats for certain cases, usually requiring bound-variable renaming
- see Kamp (1981); Kamp and Reyle (1993); Groenendijk and Stokhof (1991); Muskens (1996) among many others

Sem frameworks and sem/psycholing interface

DRT – most obvious choice for psycholing models of sem:

- always had an explicit representational commitment
- motivated by goal of interfacing sem and cognitive science more closely
- the classic Kamp (1981) paper begins:

  "Two conceptions of meaning have dominated formal semantics of natural language. The first of these sees meaning principally as that which determines conditions of truth. [...] According to the second conception meaning is, first and foremost, that which a language user grasps when he understands the words he hears or reads. [...] these two conceptions [...] have remained largely separated for a considerable period of time. This separation has become an obstacle to the development of semantic theory [...] The theory presented here is an attempt to remove this obstacle. It combines a definition of truth with a systematic account of semantic representations."
  (Kamp 1981, p. 189)

Sem frameworks and sem/psycholing interface

In addition:

- DRT is a dynamic sem framework ⇒ notion of DRS merge: merge two representations into a larger representation of the same type
- this makes construction, maintenance and incremental update of sem representations more straightforward
- . . . and similar to construction, maintenance and incremental update of syntactic representations

  merging DRSs (when possible): consequence of facts about dynamic conjunction and update semantics associated with variable assignments and atomic lexical relations (Groenendijk and Stokhof 1991; Muskens 1996; Brasoveanu 2007 a.o.)
- DRT not the only possible choice: less "representational" systems, both dynamic and static, also possible – but less straightforwardly so, at least at a first glance

Fan experiment (Anderson 1974; Anderson and Reder 1999)

Fan effect:

- “the phenomenon that, as participants study more facts about a particular concept, their time to retrieve a particular fact about that concept increases.”
- “Fan effects have been found in the retrieval real-world knowledge […], face recognition […] [etc.]”
- “The fan effect is generally conceived of as having strong implications for how retrieval processes interact with memory representations.”
- “It has been used to study the representation of semantic information […] and of prior knowledge […].”

  (Anderson and Reder 1999, p. 186)

Fan experiment: training phase

Original fan experiment in Anderson (1974): fan effect in recognition memory.

Training phase: memorize 26 facts about people in locations

(3)
a. A lawyer is in a cave.
b. A debutante is in a bank.
c. A doctor is in a bank.
d. A doctor is in a shop.
e. A captain is in a church.
f. A captain is in a park.
g. A fireman is in a park.
h. A hippie is in a park.
i. A hippie is in a church.
j. A hippie is in a town.
k. ...

Fan experiment: test phase

Test phase: a series of target and foil sentences

- target sentences: studied in training; basically, true
  - A doctor is in a shop.
- foil sentences: not studied in training; basically, false
  foils: novel combinations of the same people and locations
  - A hippie is in a cave.

Participants had to:

- recognize target / true sentences
- reject foil / false sentences

Fan exp. & semantics/psycholinguistics interface

Fan experiment: clear example of work at interface between:

- formal semantics
- computational psycholinguistics

(although not conceptualized and modeled as such before)

Training phase: model building
Test phase: incremental syn/sem parsing + sem evaluation

So:

Let's build an end-to-end processing model of all these 3 stages and fit it to fan exp. data.

Fan exp.: model building & spreading activation

– the 10 items above form a minimal network of facts for the 9 conditions in the fan exp.
– conditions: how many studied facts are connected to each type of person and location
– to see this, represent the 10 facts as a network in which each fact is connected by an edge to the type of person and location it is about
– person types / common nouns listed on the left, location types / common nouns listed on the right

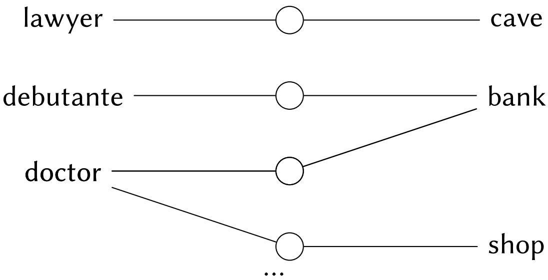

Fan exp.: model building & spreading activation

Network representation shows how:

- person and location nouns/concepts fan into 1, 2 or 3 sentences/facts

The 9 conditions in the fan exp.:

			location fan		
			1	**2**	**3**
person	fan	**1**	lawyer-cave	debutante-bank	fireman-park
		2	doctor-shop	captain-church	captain-park
		3	hippie-town	hippie-church	hippie-park

- fan: the number of facts/sentences associated with each concept/common noun.

Fan exp.: model building & spreading activation

Fan reflects spreading activation:

- in addition to base activation, a chunk's activation depends on the specific attributes and values it contains
- a chunk i for a fact is connected to/associated with the person/location concepts it contains as values
- $\ldots$ in the same way nodes are connected in a neural network
- activation spreads from values to the containing chunk
- our fact chunk i is associatively activated by the activation of the concepts it contains (person and location)

Spreading activation: influence of cognitive context on dec. mem.

- if we are currently attending to a specific concept/value
- we are associatively activating chunks containing that value

Spreading activation

- fact/chunk i does not simply additively inherit all the activations $W_1, W_2, \ldots$ of the concepts/values it contains
- $W_1, W_2, \ldots$ are scaled by the associative strengths between chunk i and the values it contains

Example:

- e.g., suppose the value/concept *doctor* has weight W_1
- we don't simply add that to base activation B_i of our chunk/fact *A doctor is in a shop*
- we scale it by the associative strength S_{1i}: this is associated with the attribute/slot PERSON of chunk i

$$\left|\begin{array}{llll} \textsc{person} & (S_{1i}): & \textit{doctor} & (W_1) \\ \textsc{location} & (S_{2i}): & \textit{shop} & (W_2) \end{array}\right|$$

Spreading activation

Total activation A_i for chunk i:

$$\boxed{A_i = B_i + (W_1 S_{1i} + W_2 S_{2i} + \dots)}$$

- B_i: base activation
- $W_1 S_{1i} + W_2 S_{2i} + \dots$: spreading activation from all values in chunk i
  - $W_1 S_{1i}$: activation spreading from value 1
  - $W_2 S_{2i}$: activation spreading from value 2
  - $\dots$: activation spreading from any other values chunk i might contain

Spreading activation

- associative strengths $S_{1i}, S_{2i}, \ldots$: like connection strengths in a neural network
- the higher the strength of the connection, the more one point in the network (a value/concept) will influence another (a chunk/fact)
- values/concepts with higher connections strengths are more prominent/important for that chunk,
- ... so the more activation the chunk overall will inherit from the activation of these values/concepts

Main intuition about activation

To understand the additive relation between base activation and spreading activation, imagine declarative memory as sea of darkness, with small rafts, i.e., chunks, floating everywhere on it.

- each raft has a small light, and the brightness of that light indicates its total activation
- the brighter that light is, the easier the raft is to find and grab – that is, we can retrieve it more accurately and more quickly
- the light on each raft is powered by two power sources

Main intuition about activation (ctd.)

i. base activation: a rechargeable battery stored on the raft itself
   - this keeps track of the history of previous usages of a chunk
   - every time we use/retrieve a raft/chunk, we plug its ‘local battery’ in for a quick charge
   - immediately after that, the battery will have more power, so the light will be brighter

Main intuition about activation (ctd.)

ii. spreading activation: current cognitive context can increase the brightness of the light on a raft
   - sources: the values held in the goal and/or imaginal buffer
   - if these values are also stored on some of the rafts in declarative memory (they are values in some slots), they act as wires delivering extra power to the lights on the rafts
   - each value in the goal/imaginal buffer has a set amount of battery power: these are the W_j values
   - that power gets distributed to all the rafts in declarative memory that also store that value

Prediction: the more rafts a value is connected with, i.e., the higher the fan of a value, the less power will be transmitted to each individual raft.

$\Rightarrow$ This is the fan effect.

Fan experiment: results

Target RTs			location fan 1	2	3
person	fan	**1**	1.11	1.17	1.15
		2	1.17	1.20	1.23
		3	1.22	1.22	1.36

Foil RTs			location fan 1	2	3
person	fan	**1**	1.20	1.25	1.26
		2	1.22	1.36	1.47
		3	1.26	1.29	1.47

Fan experiment: results (ctd.)

- averaging over targets and foils, the effect of 1-fan (both person and location) ≈ 1.2 s,
  increased by about 50 ms for each additional fan (≈ 1.25 s for 2-fan, ≈ 1.3 s for 3-fan)
- the min effect: latency is a function of minimum fan: participants respond more slowly to the 2-2 fan items than to the 1-3 or 3-1 items (repeatedly replicated)
  $\Rightarrow$ evidence for parallel access to memory from the two cues / concepts, with search being more determined by the lower fan concept
- approx. equal fan effects for targets and foils, only slightly larger effect for foils
  $\Rightarrow$ foil rejection not done by serial (exhaustive) search of the facts in memory

ACT-R account of the fan effect

- crucially relies on spreading activation, i.e., second term in the sum below (Anderson and Reder 1999)
- $A_i = B_i + \sum_j W_j S_{ji}$
- specifically, fan exp. conditions manipulate strengths of association S_{ji} between concepts j and facts i
- let's zoom in and see how these strengths are usually formalized

(4) $$S_{ji} = S + \log(P(i|j))$$

- S: constant (baseline strength), free param.
- $P(i|j)$: probability of i when j is present; estimate of how predictive concept j is of fact i

ACT-R account of the fan effect (ctd.)

- when all facts studied / tested with equal frequency: $P(i|j) = \frac{1}{f_j}$, where f_j: fan of concept j
- if a concept has a fan of 1, e.g., *lawyer*, then probability of associated fact is 1
- if concept has a fan of 3, e.g., *hippie*, all 3 facts associated with the concept are equiprobable with prob. $1/3$

(5) $$S_{ji} = S + \log\left(\frac{1}{f_j}\right) = S - \log(f_j)$$ [final form]

Basic account of the fan effect:

- S_{ji}, hence activation, decreases as a log function of fan
- latency of fact retrieval inversely related to activation ⇒ increases as concept fan increases
- see Anderson and Reder (1999) (also Brasoveanu and Dotlačil in prep.) for account of other generalizations

Our proposal & model: main idea

The main idea behind our account & model of the fan effect:

> fan effect reflects the way DRSs are organized in declarative memory

Reformulate notion of fan and network of facts and concepts as relation between:

- main DRS contributed by a sentence
- the sub-DRSs contributed by its three parts
  - the person indefinite
  - the location indefinite
  - the relational predicate *in*

DRT + ACT-R

A lawyer is in a cave.

- DRSs (meaning representations) of the three major components *a lawyer*, *a cave*, binary predicate *in*: composed/combined together to form the DRS/meaning representation for full sentence
- exact nature of the three meaning components and the composition method vary from semantic framework to semantic framework
- do not need to fully specify a semantic framework to reformulate the fan experiment in formal semantics/DRT terms

DRT + ACT-R (ctd.)

A lawyer is in a cave.

- sufficient: main DRS contributed by the sentence formed out of three sub-DRSs contributed by the three sentential components
- this partitioning into 3 sub-DRSs matches:
  - the rough compositional skeleton generally assumed in the formal semantics literature
  - the real-time incremental comprehension process the ACT-R architecture imposes on us

DRT + ACT-R (ctd.)

A lawyer is in a cave.

$$
\text{MAIN-DRS}\;\left|\;
\begin{array}{ll}
\text{SUB-DRS}_1 : & \begin{array}{|c|}\hline x \\ \hline \text{lawyer}(x) \\ \hline\end{array} \\[2em]
\text{SUB-DRS}_2 : & \begin{array}{|c|}\hline y \\ \hline \text{cave}(y) \\ \hline\end{array} \\[2em]
\text{SUB-DRS}_3 : & \begin{array}{|c|}\hline \\ \hline \text{in}(x,\ y) \\ \hline\end{array}
\end{array}\;\right|
$$

DRT + ACT-R (ctd.)

A lawyer is in a cave.

- note: because of the seriality imposed in ACT-R by
  - one production firing at a time
  - buffers being able to hold only one chunk at a time
- ... we never have a full view of:
  - the syntactic tree representation
  - the DRS semantic representation
- representation assembled one sub-DRS at a time
- main DRS only implicitly available in declarative memory
- ... just like the full syntactic tree of the sentence is only implicitly available in declarative memory

DRT + ACT-R (ctd.)

A lawyer is in a cave.

- were we to merge/dynamically conjoin the 3 sub-DRSs into one DRS
- would obtain the DRS below – precisely the semantic representation assigned in DRT

x, y
lawyer(x) cave(y) in$(x,\ y)$

DRT + ACT-R (ctd.)

A lawyer is in a cave.

Final DRT + ACT-R representation:

$$
\text{MAIN-DRS}\left|\begin{array}{l}
\text{SUB-DRS}_1 : \left|\begin{array}{lll} \text{DREF} : & 1 \\ \text{PRED} & \textit{lawyer} \\ \text{ARG1} : & 1 \end{array}\right| \\[2ex]
\text{SUB-DRS}_2 : \left|\begin{array}{lll} \text{DREF} : & 2 \\ \text{PRED} & \textit{cave} \\ \text{ARG1} : & 2 \end{array}\right| \\[2ex]
\text{SUB-DRS}_3 : \left|\begin{array}{lll} \text{PRED} & \textit{in} \\ \text{ARG1} : & 1 \\ \text{ARG2} : & 2 \end{array}\right|
\end{array}\right|
$$

DRT + ACT-R: incremental interpretation

- A lawyer is in a cave.

ISA :	*drs*
DREF :	1
ARG1 :	1

1
STILL-UNSPECIFIED-PREDICATE(1)

DRT + ACT-R: incremental interpretation (ctd.)

- A **lawyer** is in a cave.

$$\left|\begin{array}{ll} \textsc{isa}: & \textit{drs} \\ \textsc{dref}: & 1 \\ \textsc{pred} & \textit{lawyer} \\ \textsc{arg1}: & 1 \end{array}\right|$$

1
$lawyer(1)$

DRT + ACT-R: incremental interpretation (ctd.)

- A lawyer is **in** a cave.

$$\left|\begin{array}{ll} \textsc{isa}: & \textit{drs} \\ \textsc{pred} & \textit{in} \\ \textsc{arg1}: & 1 \end{array}\right|$$

$$\begin{array}{|c|} \hline \\ \hline in(1,\ \_) \\ \hline \end{array}$$

DRT + ACT-R: incremental interpretation (ctd.)

- A lawyer is in **a** cave.

$$\left|\begin{array}{ll} \text{ISA}: & \textit{drs} \\ \text{DREF}: & 2 \\ \text{PRED} & \textit{in} \\ \text{ARG1}: & 1 \\ \text{ARG2}: & 2 \end{array}\right|$$

$$\begin{array}{|c|}\hline 2 \\ \hline in(1,\ 2) \\ \hline \end{array}$$

$$\left|\begin{array}{ll} \text{ISA}: & \textit{drs} \\ \text{ARG1}: & 2 \end{array}\right|$$

$$\begin{array}{|c|}\hline \\ \hline \text{UNSPEC-PRED}(2) \\ \hline \end{array}$$

[other update options available]

DRT + ACT-R: incremental interpretation (ctd.)

- A lawyer is in a **cave**.

$$\left|\begin{array}{ll} \textsc{isa}: & \textit{drs} \\ \textsc{pred} & \textit{cave} \\ \textsc{arg1}: & 2 \end{array}\right|$$

$$\begin{array}{|c|} \hline \\ \hline cave(2) \\ \hline \end{array}$$

DRT+ACT-R+Bayes: sem eval. as mem. retrieval

- once the test sentence is fully parsed, the 3 sub-DRSs are recalled and stored in goal buffer
- for target / true sentences: they spread activation to the fact in declarative memory that verifies the sentence
- for foil / false sentences: they spread less activation because no fact is a perfect match
- we embed the DRT + ACT-R model in a Bayesian model and fit to data

Fan model parameters

- estimate 4 subsymbolic parameters:
  - `"buffer_spreading_activation"` (`"bsa"` for short), which is the W parameter (assume all W_j are equal)
  - `"strength_of_association"` (`"soa"` for short), which is the S parameter
  - `"rule_firing"` (`"rf"` for short), which is by default set to 50 ms
  - `"latency_factor"` (`"lf"` for short)

Fan model estimates

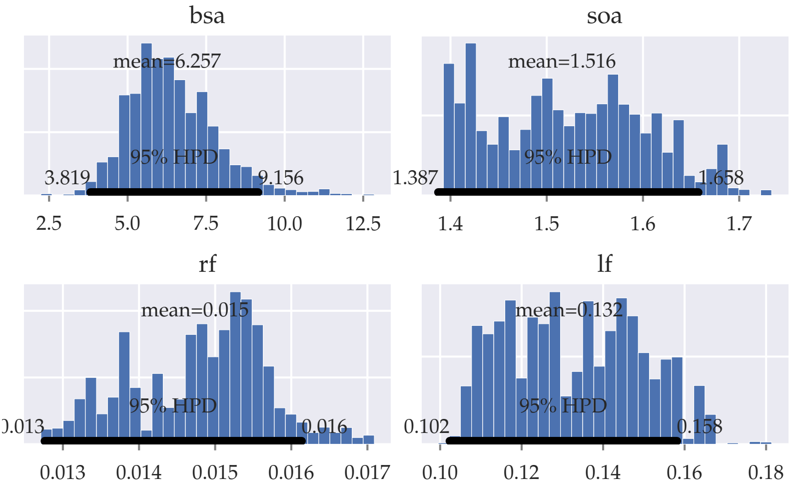

Fan model: observed vs. predicted RTs

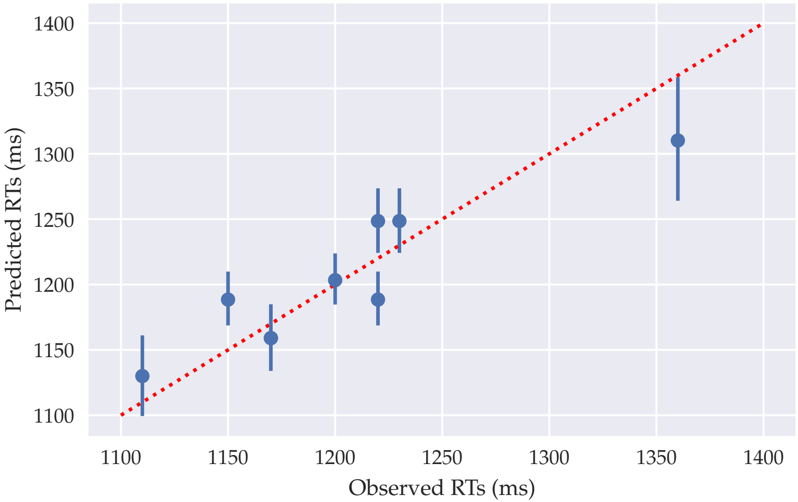

Conclusion

- model fits data fairly well
- better fit if we run model for longer than 5000 iterations, and with a burn-in larger than 500

Main messages

- we can bring together ACT-R and formal semantics theories in formally and computationally explicit way
- the resulting incremental interpreters can be fit to data $\Rightarrow$ different sem. and/or processing theories can be quantitatively compared
- fan effect: important insights into memory structures and cognitive processes that underlie incremental semantic processing and evaluation

Next: pronoun / presupposition resolution, which involves similar memory structures and retrieval processes.

Anderson, John R. 2007. How can the human mind occur in the physical universe?. Oxford University Press.

Anderson, John R, and Lynne M Reder. 1999. The Fan Effect: New Results and New Theories. Journal of Experimental Psychology: General 128:186–197.

Anderson, John Robert. 1974. Retrieval of propositional information from long-term memory. Cognitive Psychology 6:451 – 474.

Brasoveanu, Adrian. 2007. Structured nominal and modal reference. Doctoral Dissertation, Rutgers Univ.

Brasoveanu, Adrian, and Jakub Dotlačil. 2015. Incremental and predictive interpretation: Experimental evidence and possible accounts. In Proceedings of Semantics and Linguistic Theory (SALT) 25, 57–81.

Brasoveanu, Adrian, and Jakub Dotlačil. in prep. Formal Linguistics and Cognitive Architecture. Language, Cognition, and Mind (LCAM) Series. Dordrecht: Springer. The pyactr library (Python3 ACT-R) is available here: `https://github.com/jakdot/pyactr`.

Groenendijk, Jeroen, and Martin Stokhof. 1991. Dynamic predicate logic. Linguistics and Philosophy 14:39–100.

Johnson-Laird, Philip N. 1983. Mental models: Towards a cognitive science of language, inference, and consciousness. Harvard University Press.

Johnson-Laird, Philip N. 2004. The history of mental models. Psychology of reasoning: Theoretical and historical perspectives 179–212.

Johnson-Laird, Philip N., Ruth M. J. Byrne, and Patrizia Tabossi. 1989. Reasoning by model: The case of multiple quantification. Psychological Review 96:658–673.

Kamp, Hans. 1981. A theory of truth and semantic representation. In Formal methods in the study of language, ed. Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, 277–322. Amsterdam: Mathematical Centre Tracts.

Kamp, Hans, and Uwe Reyle. 1993. From discourse to logic. introduction to model theoretic semantics of natural language, formal logic and Discourse Representation Theory. Dordrecht: Kluwer.

Lewis, Richard, and Shravan Vasishth. 2005. An activation-based model of sentence processing as skilled memory retrieval. Cognitive Science 29:1–45.

Muskens, Reinhard A. 1996. Combining Montague Semantics and Discourse Representation. Linguistics and Philosophy 19:143–186.