Situated Communication

Alex Lascarides
collaborations with Mattias Appelgren, Nicholas Asher and Julie Hunter

GDR LIFT, November 2021

Motivating example

I sent Charlotte to her room.

I nod towards a scratch on the wall

I was making dinner.

Features of the Scratches example

A nonlinguistic event appropriated as a discourse move

Nonlinguistic events must be ***conceptualized*** in a discourse-relevant way

- someone (or something) scratched the wall.
- Charlotte scratched the wall.

Different utterance ⇒ different conceptualisation

I moved the table into the living room this morning.
I nod towards scratch on the wall
I had to buy some new paint.

Features of the Scratches example

A nonlinguistic event appropriated as a discourse move

Linguistic units don't refer to or describe the scratching event.
The scratching event:

- is a part of the message
  - coherently related to linguistic moves
- affects how that message is **constructed**
- but is not produced as part of the discourse

Starting Point: Coherence-based Discourse Semantics

Some assumptions

- A discourse move contributes (an instance of) a proposition
- It must be semantically related to some part of the discourse context
  - Explanation, Elaboration, Narration, Contrast, Result, etc.
- Salience: only certain parts of that context are available.

Salience: The Right Frontier Constraint

Subordinating relations ('push down')

Develop a certain point; maintain its salience

- Explanation, Elaboration, Background...

Coordinating relations ('push to the right')

Push the discourse forward, shutting off accessibility of previous moves

- Continuation, Narration, Result...

The Right Frontier Constraint (RFC)

New moves must attach to **The Right Frontier**:

- most recent move;
  extended segment(s) it's part of;
  moves it's subordinate to.

Example

Max had a lovely evening.
He had a great meal.
He ate salmon.
He devoured lots of cheese.
He won a dancing competition.

Example

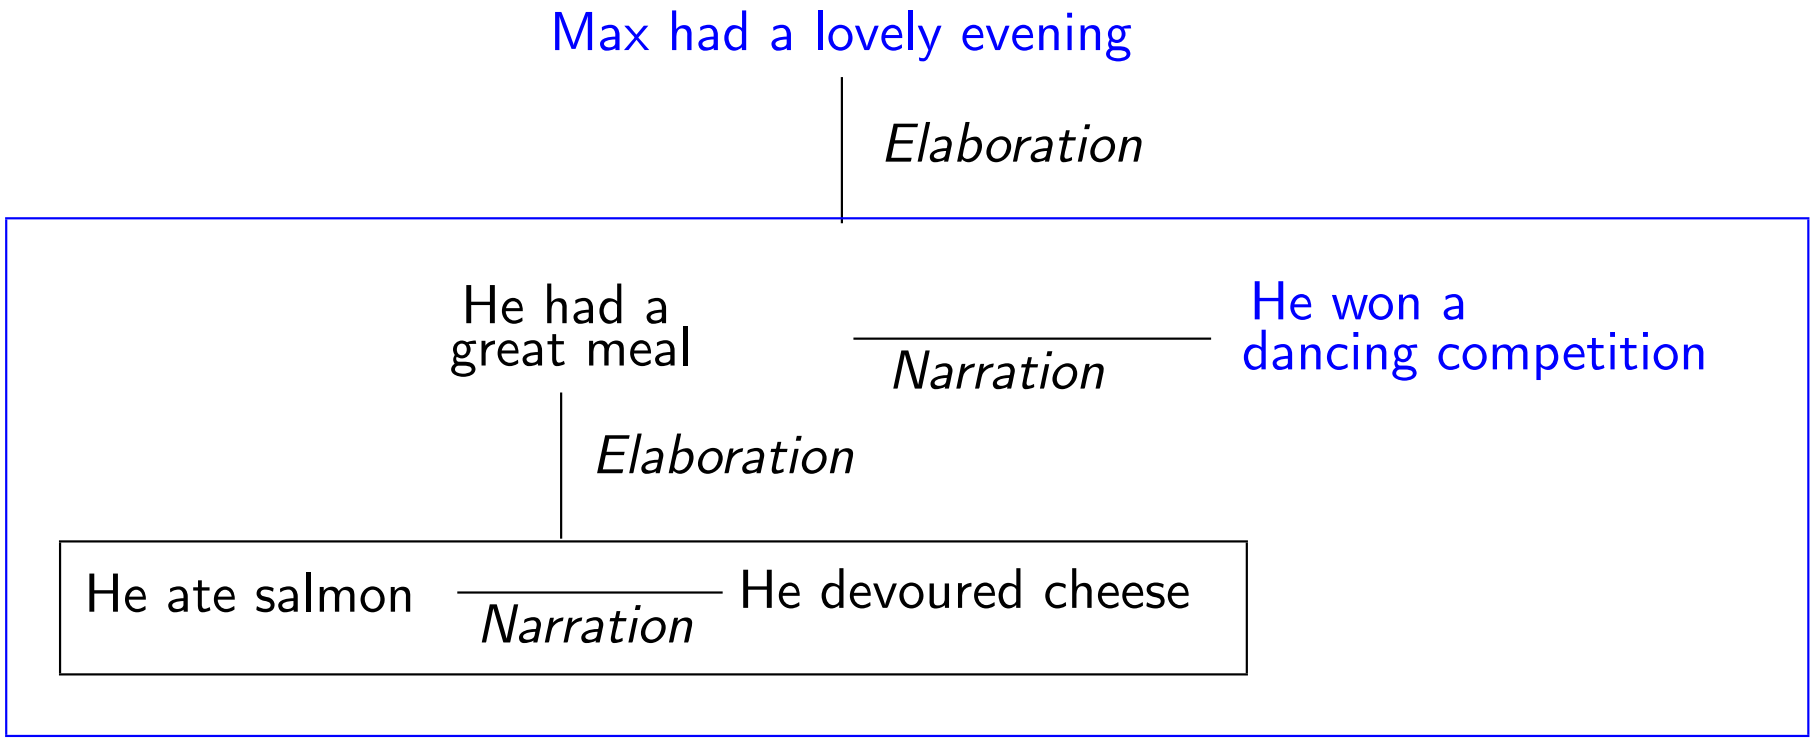

It was a lovely pink.

Interaction between coherence and conceptualization

(1) I sent Charlotte to her room. (2') She scratched the wall.

$\Rightarrow$ Explanation(1,2')
(2') $<_t$ (1)

Similarly for (1)+ non-linguistic:

$\Rightarrow$ Explanation(1,e)
e: *Charlotte created* $[\textit{the scratch on the wall}]_g$

Non-linguistic events disambiguate linguistic moves

Red blocks should be on blue blocks!

$\forall x \forall y(\mathit{red}(x) \wedge \mathit{on}(x,y) \rightarrow \mathit{blue}(y))$ $\forall x \forall y(\mathit{blue}(y) \wedge \mathit{on}(x,y) \rightarrow \mathit{red}(x))$

Non-linguistic events disambiguate linguistic moves

Red blocks should be on blue blocks!

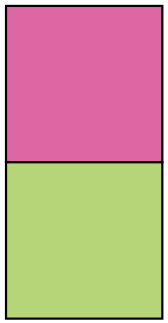

$\forall x \forall y (red(x) \wedge on(x,y) \rightarrow blue(y))$

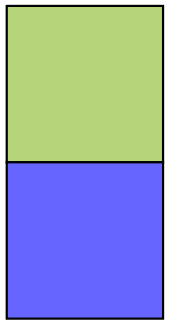

$\forall x \forall y (blue(y) \wedge on(x,y) \rightarrow red(x))$

Semantics of
$Correction(a,u)$
resolves the linguistic ambiguity

Two claims for *Situated Communication*

1. Nonlinguistic events affect **discourse structure, its evolution and its interpretation** in nontrivial ways.
   - Initial analysis via the STAC corpus
     joint work with Julie Hunter and Nicholas Asher.
2. **The semantics of coherence relations influences conceptualisation** of nonlinguistic events.
   - Experiments in Interactive Task Learning (ITL) involving corrective feedback.
     joint work with Mattias Appelgren.

Theoretical Issues

Salience

Is the RFC still valid?

Semantics

Dynamic update with real world events

Use the The STAC Corpus as evidence.

(https://www.irit.fr/STAC/corpus.html)

A corpus

Settlers of Catan

- multi-party, win-lose game
- players use resources (wood, clay, ...) to build roads and settlements
- board: multiple regions, each assigned a resource and number (2 - 12)
- players get resources by rolling dice, trading, or stealing
- robber: roll of a 7; discard, steal, move

The game board

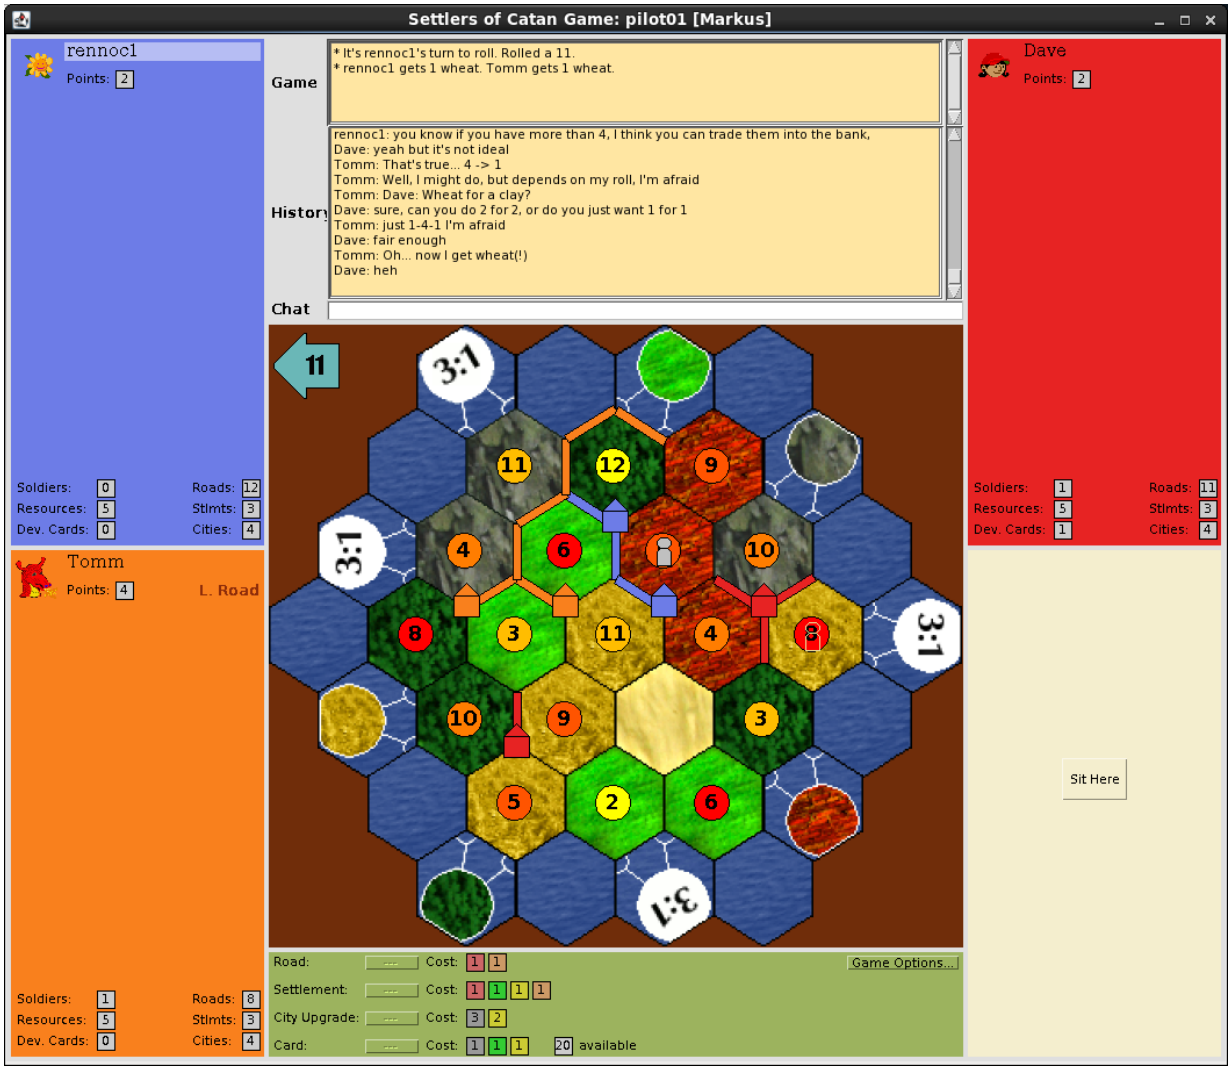

Annotation

- 59 games, each with dozens of dialogues with 1-30+ turns
- Annotation in the style of SDRT (Asher & Lascarides 2003)
  - Also who offers what to whom (defined in a FS)
- Annotation was tackled in two phases:
  chat-only vs. chat + game
  20% of chat-only annotations were wrong!!

`https://www.irit.fr/STAC/corpus.html`.

Is the RFC still valid?

- What are the constraints on how a speaker can exploit the world around her to accomplish her discourse purposes?
- And what happens when the world changes while she's talking? Not a static set of indices (à la Kaplan)

$\Rightarrow$ Start by looking at what kind of structures are allowed.

Is the RFC still valid?

Game events moving forward do not shut off salience of previous events (in contrast to linguistic moves)
unless we comment on them.

Example

159.1	Server	ljay played a Soldier card.
159.4	Server	ljay stole a resource from gwfs
160	Server	ljay rolled a 4 and a 4.
161	Server	gwfs gets 2 wheat.
163	**gwfs**	**touché**

Is the RFC still valid?

Game events moving forward do not shut off salience of previous events (in contrast to linguistic moves)
unless we comment on them.

Example

159.1	Server	ljay played a Soldier card.
159.4	Server	ljay stole a resource from gwfs
160	Server	ljay rolled a 4 and a 4.
161	Server	gwfs gets 2 wheat.
162	gwfs	the wheat's growing again!
163	**gwfs**	**# touché**

Asymmetric structures

Similar to multiparty threads but one thread depends on moves in another for its interpretation (asymmetry)

An Example

341	Server	gwfs rolled a 6 and a 3.
342	Server	inca gets 2 wheat. dmm gets 1 wheat.
344	gwfs	**9 nooo!**
344.0.1	UI	gwfs ended their turn.
344.0.2	Server	It's inca's turn to roll the dice.
345	Server	inca rolled a 1 and a 3.
346	Server	CheshireCatGrin gets 1 ore, 1 wood. gwfs gets 2 wood.
347	gwfs	**4 better :)**

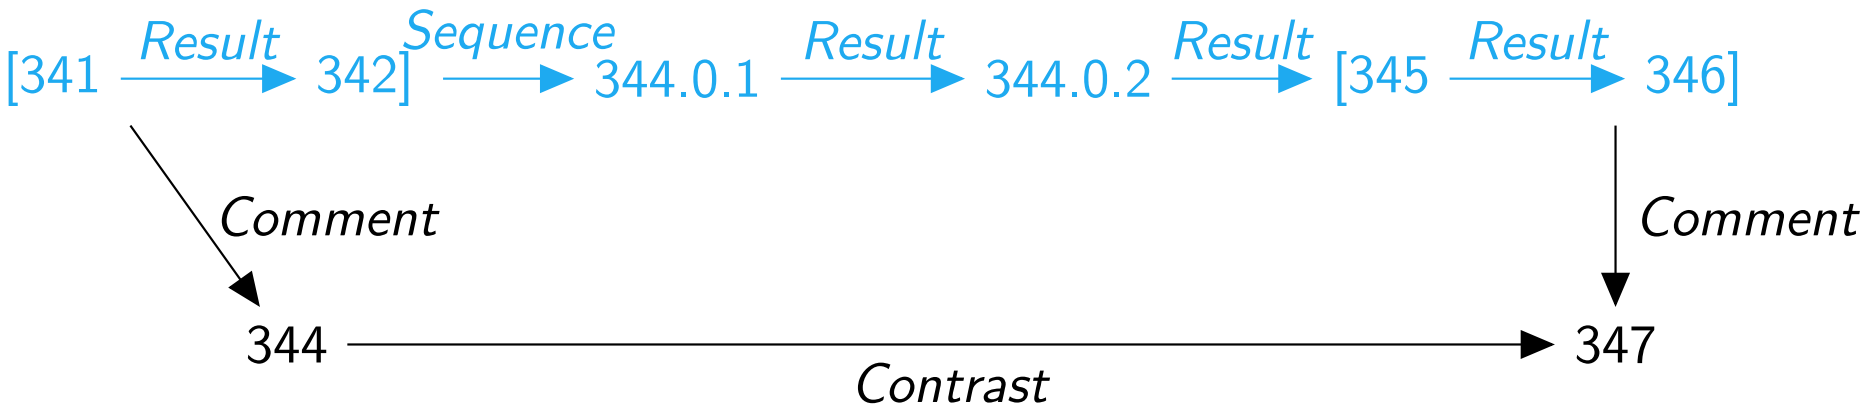

Contrast(344,347) violates RFC.

An example with multiple dependencies

534	gwfs	**anyone want to trade their ore for my wood?**
535	ljay	**nope**
538	gwfs	**it may prove a prudent trade, lj. . .**
539	ljay	**nope**
539.1	Server	gwfs played a Soldier card.
539.4	Server	gwfs stole a resource from ljay
540	gwfs	**apologies. . .**

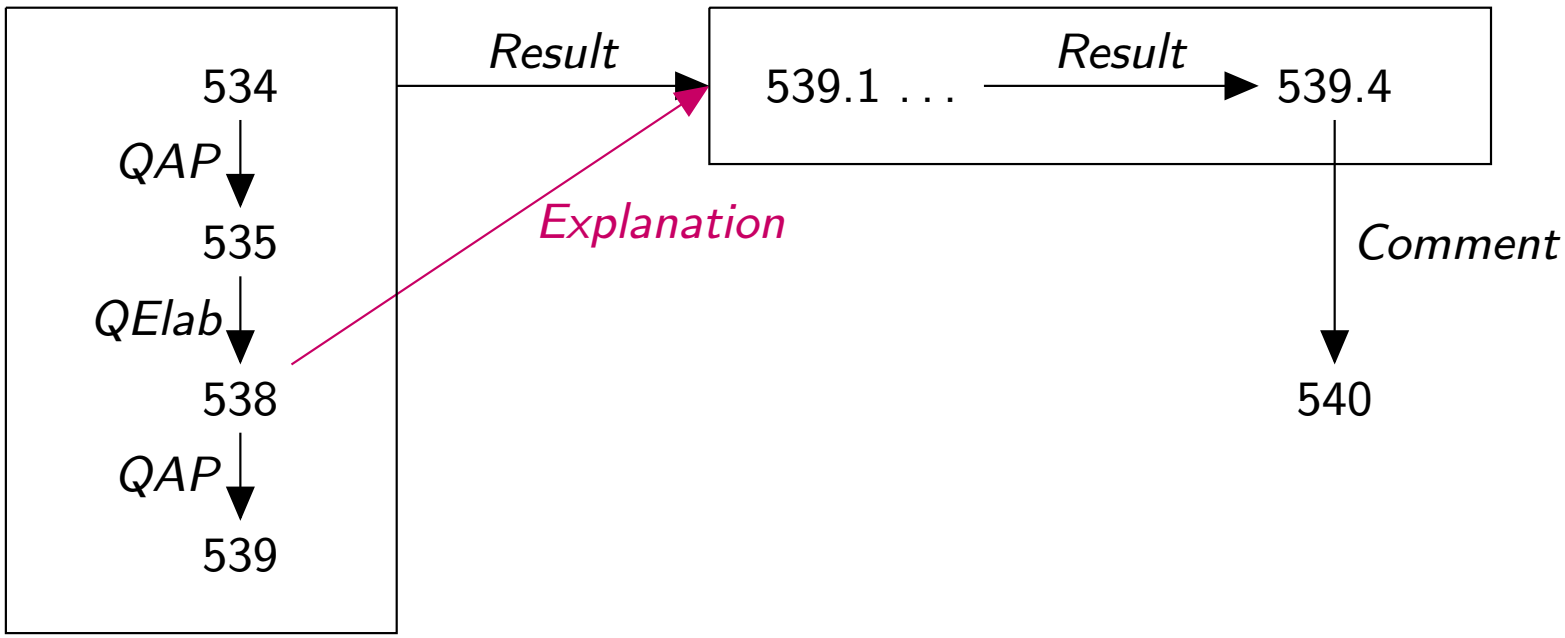

Semantics

Build on existing models of dynamic semantics

- represent dialogue with commitment slates for each speaker, etc.
- add a representation of the sequence of events in the actual world through the course of the conversation
- speakers can take on commitments to actual events, and this limits discourse continuations

Coherence and Conceptualisation: Experiments in ITL

Blocks world: learning agent can put one block on another, and must build a single tower from blocks on the table.

- Ignorant of goal constraints:
  e.g., *each red block must be on a blue block*
- Unaware of the domain-level concepts that define the goal:
  Can observe RGB values, but doesn't know the partition.
- Unaware of the colour terms: red, blue etc are neologisms.

The Interactive Learning Task

Learner's Task

- Learn to solve the planning problem by learning:
  - the goal description
  - how to conceptualise the domain
    (i.e., how to partition the RGB spectrum into colour categories)

using own experience and expert's corrective feedback as evidence.

Evidence via Corrections

Correction(a, u)

a: agent puts green block on blue block

u: No! Red blocks should be on blue blocks.

Challenge: Message is ambiguous and learner may not know the blocks are green and blue...

Formalising the set up

Signal	(Hidden) Message	Rule name
Red blocks should be on blue blocks	$\forall x \forall y (red(x) \wedge on(x,y) \rightarrow blue(y))$	$r_1^{r,b}$
	$\forall x \forall y (blue(y) \wedge on(x,y) \rightarrow red(x))$	$r_2^{r,b}$

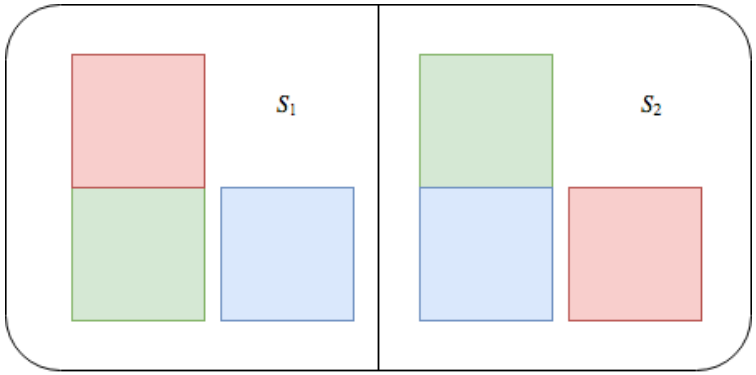

- Can't tell which message is intended in S_1 and S_2, even if you know the blocks' colours!
- So signal includes pointing to the tower or the block on table to disambiguate the message (if you know the blocks' colours).

Graphical Model for "No! Red blocks should be on blue blocks" *points to tower*

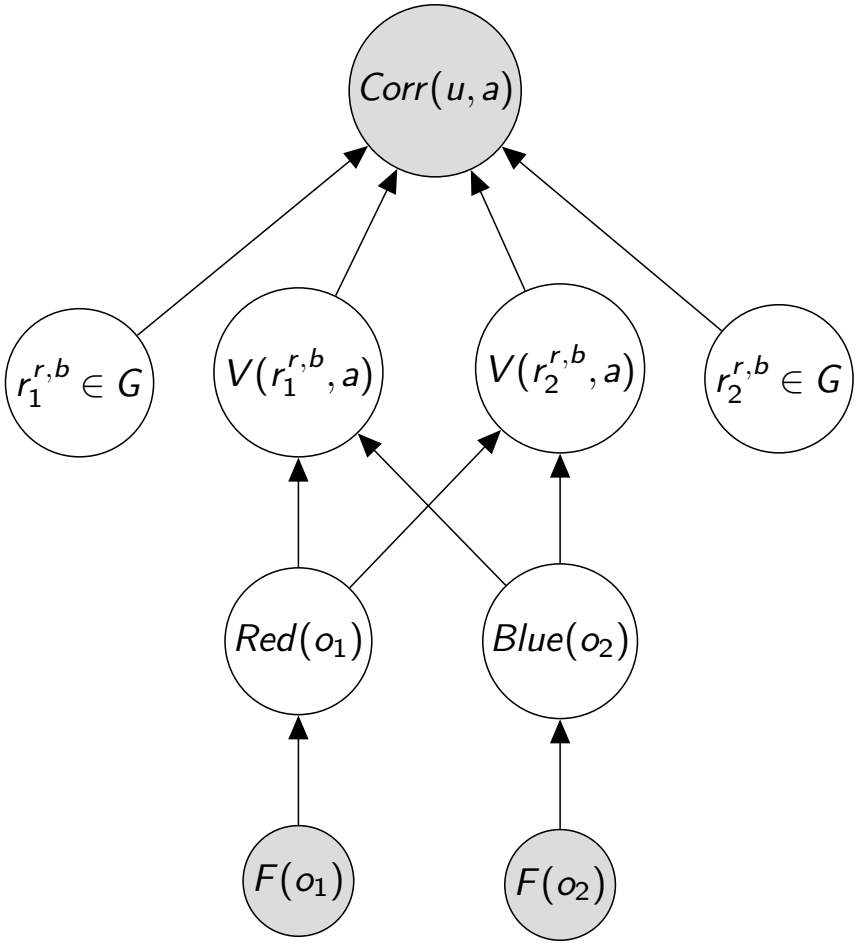

Observed: grey
Latent: white
$V(r, a)$: Boolean;
r violated by a
$r \in G$: Boolean
r is a goal constraint
$Red(o)$: Boolean; o is red
$F(o)$: o's RGB values
Posteriors for
$r \in G$, $P(Red(o)|F(o))$
become priors for next move

Semantics of correction imposes constraints on combination of values of the random variables

Experiments

(Hidden) Goals

Two rules: $r_1^{red,blue}$, $r_2^{green,maroon}$

Three rules: $r_1^{red,blue}$, $r_2^{red,blue}$, $r_1^{purple,orange}$

Evidence

anaphor: that's not red either, that's wrong for the same reason or not

full: Learn from actions that *aren't* corrected

Experimental run

- 50 trials per G (two rules, or three rules)
- 10 blocks on table (chosen so that G constraints are relevant)
- Learner is **full** or **simple**
  Expert is **anaphor** or **no-anaphor**

Results against Baselines: Cumulative Regret

Simple+no-anaphor: our simplest model.

Naive: Doesn't learn between trials;
simply avoids repeating corrected action.

No Language: Uses only "no": blocks with similar RGB values to o_1 can't be put on blocks with similar RGB values to o_2.

Results comparing models: Cumulative Regret

Two Rules:
$r_1^{red,blue}$, $r_2^{green,maroon}$

Three Rules:
$r_1^{red,blue}$, $r_2^{red,blue}$, $r_1^{purple,orange}$

Conclusion

- A formal semantics of embodied discourse calls for a radical overhaul of the types of structures and semantics we countenance.
- We need empirical embodied linguistic data to guide those revisions.
- The STAC corpus provides a basis for testing hypotheses about salience and model-theoretic semantics of embodied conversation.
- But the corpus doesn't address how language influences conceptualisation of the domain.
- Experiments in ITL demonstrate that reasoning about discourse coherence (semantics, saliance and anaphora) speeds up learning of: domain conceptualisation, symbol grounding, the planning problem