Coordinating on context and construal

Christopher Potts

Stanford Linguistics

Google, February 19, 2015

From sketch to detailed image (Levinson, 2000)

Figure 0.1
Rembrandt sketch

From sketch to detailed image (Levinson, 2000)

Figure 0.1
Rembrandt sketch

"So here is the miracle: from a merest, sketchiest squiggle of lines, you and I converge to find adumbration of a coherent scene."

From sketch to detailed image (Levinson, 2000)

"The problem of utterance interpretation is not dissimilar to this visual miracle. An utterance is not, as it were, a veridical model or "snapshot" of the scene it describes [. . .]. Rather, an utterance is just as sketchy as the Rembrandt drawing."

From sketch to detailed image (Levinson, 2000)

"The problem of utterance interpretation is not dissimilar to this visual miracle. An utterance is not, as it were, a veridical model or "snapshot" of the scene it describes [. . .]. Rather, an utterance is just as sketchy as the Rembrandt drawing."

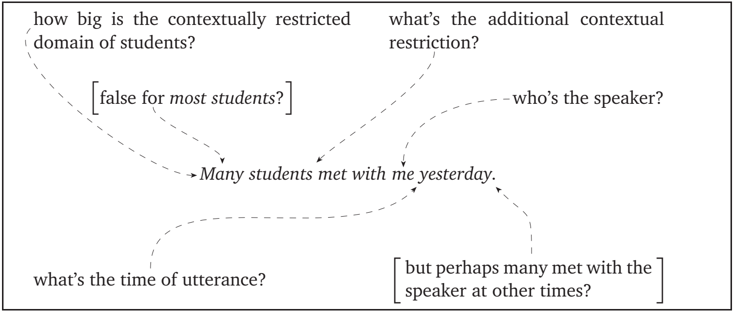

Gricean pragmatics

- **The cooperative principle (a super-maxim)**: Make your contribution as is required, when it is required, by the conversation in which you are engaged.
- **Quality**: Contribute only what you know to be true. Do not say false things. Do not say things for which you lack evidence.
- **Quantity**: Make your contribution as informative as is required. Do not say more than is required.
- **Relation (Relevance)**: Make your contribution relevant.
- **Manner**: (i) Avoid obscurity; (ii) avoid ambiguity; (iii) be brief; (iv) be orderly.
- **Politeness**: Be polite, so be tactful, respectful, generous, praising, modest, deferential, and sympathetic. (Leech)

Gricean pragmatics

- **The cooperative principle (a super-maxim)**: Make your contribution as is required, when it is required, by the conversation in which you are engaged.
- **Quality**: Contribute only what you know to be true. Do not say false things. Do not say things for which you lack evidence.
- **Quantity**: Make your contribution as informative as is required. Do not say more than is required.
- **Relation (Relevance)**: Make your contribution relevant.
- **Manner**: (i) Avoid obscurity; (ii) avoid ambiguity; (iii) be brief; (iv) be orderly.
- **Politeness**: Be polite, so be tactful, respectful, generous, praising, modest, deferential, and sympathetic. (Leech)

Conversational implicature

Definition

Speaker S saying U to listener L conversationally implicates q iff

1. S and L mutually, publicly presume that S is cooperative.
2. To maintain 1 given U, it must be supposed that S thinks q.
3. S thinks that both S and L mutually, publicly presume that L is willing and able to work out that 2 holds.

Conversational implicature

Definition

Speaker S saying U to listener L conversationally implicates q iff

1. S and L mutually, publicly presume that S is cooperative.
2. To maintain 1 given U, it must be supposed that S thinks q.
3. S thinks that both S and L mutually, publicly presume that L is willing and able to work out that 2 holds.

Example

Ann: What city does Paul live in?
Bob: Hmm . . . he lives in California.

Conversational implicature

Definition

Speaker S saying U to listener L conversationally implicates q iff

1. S and L mutually, publicly presume that S is cooperative.
2. To maintain 1 given U, it must be supposed that S thinks q.
3. S thinks that both S and L mutually, publicly presume that L is willing and able to work out that 2 holds.

Example

Ann: What city does Paul live in?
Bob: Hmm ... he lives in California.

(A) Assume Bob is cooperative.
(B) Bob supplied less information than was required, seemingly contradicting (A).
(C) Assume Bob does not know which city Paul lives in.
(D) Then Bob's answer is optimal given his evidence.

Conversational implicature

Definition

Speaker S saying U to listener L conversationally implicates q iff

1. S and L mutually, publicly presume that S is cooperative.
2. To maintain 1 given U, it must be supposed that S thinks q.
3. S thinks that both S and L mutually, publicly presume that L is willing and able to work out that 2 holds.

Example

Prosecutor:	Do you have a Swiss bank account?
Defendant:	My company had one years ago.

Conversational implicature

Definition

Speaker *S* saying *U* to listener *L* conversationally implicates *q* iff

1. *S* and *L* mutually, publicly presume that *S* is cooperative.
2. To maintain 1 given *U*, it must be supposed that *S* thinks *q*.
3. *S* thinks that both *S* and *L* mutually, publicly presume that *L* is willing and able to work out that 2 holds.

Example

Prosecutor: Do you have a Swiss bank account?
Defendant: My company had one years ago.

(A) Defendant is cooperative and an expert about his accounts.
(B) He failed to address the question, seemingly contradicting (A).
(C) The more relevant statement "I have a Swiss bank account" must be pragmatically inaccessible.
(D) By (A), falsity is the best explanation for its inaccessibility.

Conversational implicature

Definition

Speaker *S* saying *U* to listener *L* conversationally implicates *q* iff

1. *S* and *L* mutually, publicly presume that *S* is cooperative.
2. To maintain 1 given *U*, it must be supposed that *S* thinks *q*.
3. *S* thinks that both *S* and *L* mutually, publicly presume that *L* is willing and able to work out that 2 holds.

Example

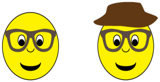

r_1	r_2	

	r_1	r_2
'glasses'	T	T
'hat'	F	T

Conversational implicature

Definition

Speaker S saying U to listener L conversationally implicates q iff

1. S and L mutually, publicly presume that S is cooperative.
2. To maintain 1 given U, it must be supposed that S thinks q.
3. S thinks that both S and L mutually, publicly presume that L is willing and able to work out that 2 holds.

Example

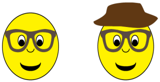

r_1 r_2

	r_1	r_2
'glasses'	T	T
'hat'	F	T

(A) Assume the speaker is cooperative.
(B) 'glasses' is less informative that 'hat'.
(C) To reconcile 'glasses' with (A), assume the speaker lacks evidence for 'hat'.
(D) By the nature of the game, the speaker lacks evidence for 'hat' iff 'hat' is false.

Conversational implicature

Definition

Speaker S saying U to listener L conversationally implicates q iff

1. S and L mutually, publicly presume that S is cooperative.
2. To maintain 1 given U, it must be supposed that S thinks q.
3. S thinks that both S and L mutually, publicly presume that L is willing and able to work out that 2 holds.

Implicature as social, interactional

Implicatures are inferences that listeners make to reconcile the speaker's linguistic behavior with the assumption that the speaker is cooperative.

Implicatures and cognitive complexity

The speaker must believe that the listener will infer that the speaker believes the implicature.

Language as a system of conventions

Convention (Lewis, 1969)

Regularity R in the behavior of members of population P is a convention iff

1. almost everyone prefers to conform to R on condition that almost everyone else does; and
2. almost everyone would just as happily defect to alternative regularity R' if everyone else did.

Language as a system of conventions

Convention (Lewis, 1969)

Regularity R in the behavior of members of population P is a convention iff

1. almost everyone prefers to conform to R on condition that almost everyone else does; and
2. almost everyone would just as happily defect to alternative regularity R' if everyone else did.

Smith et al. (2013)

As a convention-based communication agent, I assume

1. there is a single set of linguistic conventions $\mathcal{L}$
2. everyone knows $\mathcal{L}$
3. everyone else believes that I know $\mathcal{L}$
4. but (social anxiety!) I don't really know $\mathcal{L}$!

Plan for today

1. The Rational Speech Acts (RSA) model
2. Training effective literal listeners
3. The joint inferences of deeply pragmatic listeners
4. The Cards task-oriented dialogue corpus
5. Language and action together

The Rational Speech Acts model

1. The Rational Speech Acts (RSA) model
2. Training effective literal listeners
3. The joint inferences of deeply pragmatic listeners
4. The Cards task-oriented dialogue corpus
5. Language and action together

Mike Frank

Noah Goodman

The Rational Speech Acts model

Definition (Literal listener)

$$L_0(world \mid msg, \mathcal{L}) \propto \frac{\mathbb{I}(world \in \mathcal{L}(msg))}{|\mathcal{L}(msg)|} P(world)$$

Definition (Pragmatic speaker)

$$S_1(msg \mid world, \mathcal{L}) \propto \exp \lambda \left(\log L_0(world \mid msg, \mathcal{L}) - C(msg)\right)$$

Definition (Pragmatic listener)

$$L_1(world \mid msg, \mathcal{L}) \propto S_1(msg \mid world, \mathcal{L}) P(world)$$

The origins of RSA

- Rosenberg & Cohen (1964): early Bayesian model of production and comprehension in reference games.
- Lewis (1969): signaling systems (H. Clark 1996)
- Rabin (1990): recursive strategic signaling
- Camerer et al. (2004): cognitive hierarchy models for games of conflict and coordination
- Franke (2008, 2009) and Jäger (2007, 2012): iterated best response
- Golland et al. (2010): $L_1(S_0)$ with semantic parsing
- Frank & Goodman (2012): $L_1(S_1(L_1(S_0)))$

An ad hoc conversational implicature

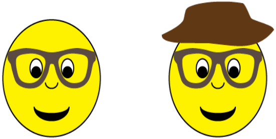

r_1 r_2

Referents

r_1	.5
r_2	.5

Prior

	r_1	r_2
'glasses'	T	T
'hat'	F	T

Messages

'glasses'	0
'hat'	0

Costs

Figure: Scenario

An ad hoc conversational implicature

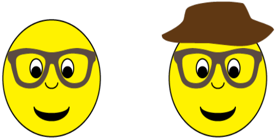

r_1 r_2

Referents

r_1	.5
r_2	.5

Prior

	r_1	r_2
'glasses'	T	T
'hat'	F	T

Messages

'glasses'	0
'hat'	0

Costs

Figure: Scenario

	r_1	r_2
'glasses'	.75	.25
'hat'	0	1

L_1

	'glasses'	'hat'
r_1	1	0
r_2	.33	.67

S_1

	r_1	r_2
'glasses'	.5	.5
'hat'	0	1

L_0

Figure: Reasoning

Experimental results

- Implicatures encourage mutual exclusivity, a.k.a., the pidgeon-hole principle (E. Clark 1987; Frank et al. 2009). This reasoning is pervasive in communication.
- Implicature reasoning in simple reference games is extremely well-supported (Vogel et al., 2014; Degen & Franke, 2012).
- Eye-tracking studies have illuminated the time-course of implicature reasoning during sentence processing (Grodner & Sedivy, 2008; Huang & Snedeker, 2009; Grodner et al., 2010).
- For first-language acquisition, simple reference games separate linguistic abilities from pragmatic abilities — and kids turn out to be pretty good at pragmatics (Stiller et al., 2011).

The role of context

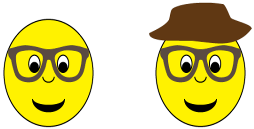

r_1 r_2

Referents

r_1	.5
r_2	.5

Prior

	r_1	r_2
'glasses'	T	T
'hat'	F	T

Messages

'glasses'	0
'hat'	0

Costs

Figure: Scenario

	r_1	r_2
'glasses'	.75	.25
'hat'	0	1

L_1

The role of context

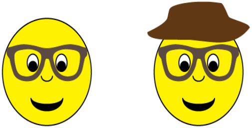

r_1	r_2

Referents

r_1	.5
r_2	.5

Prior

	r_1	r_2
'glasses'	.5	.5
'hat'	0	1

L_1

	r_1	r_2
'glasses'	T	T
'hat'	F	T

Messages

'glasses'	0
'hat'	0

Costs

Figure: Scenario

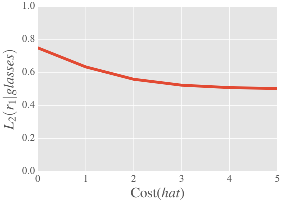

The role of context

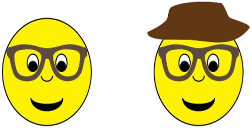

r_1	r_2

Referents

r_1	.5
r_2	.5

Prior

	r_1	r_2
'glasses'	.99	.01
'hat'	0	1

L_1

	r_1	r_2
'glasses'	T	T
'hat'	F	T

Messages

'glasses'	0
'hat'	0

Costs

Figure: Scenario

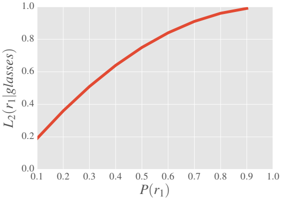

The role of context

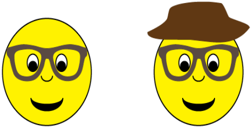

r_1 r_2

Referents

r_1	.5
r_2	.5

Prior

	r_1	r_2
‘glasses’	T	T
‘hat’	F	T

Messages

‘glasses’	0
‘hat’	0

Costs

Figure: Scenario

	r_1	r_2
‘glasses’	.95	.05
‘hat’	0	1

L_{10}

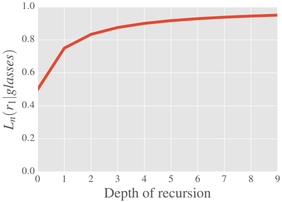

Bounded rationality

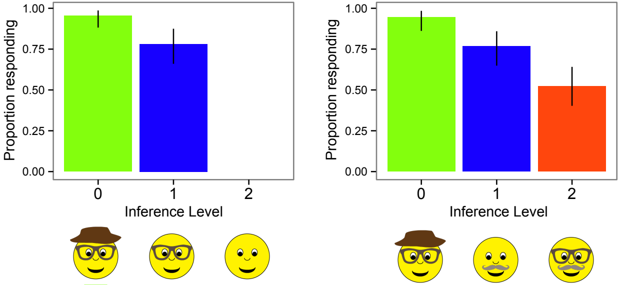

“hat”	1	0	0
“glasses”	1	1	0

“hat”	1	0	0
“glasses”	1	0	1
“mustache”	0	1	1

(Vogel et al., 2014)

Bounded rationality

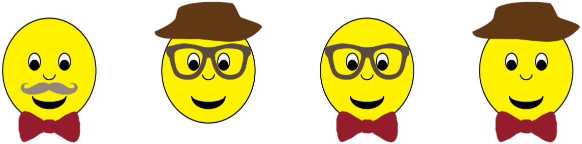

Self-trained discriminative RSA

	r_1	r_2
'glasses'	T	T
'hat'	F	T

Weights: $(1, 0)$

(Vogel et al., 2014)

Self-trained discriminative RSA

	r_1	r_2
'glasses'	T	T
'hat'	F	T

Weights: $(1, 0)$

	r_1	r_2	r_3
'glasses'	T	F	F
'hat'	T	F	T
'mustache'	F	T	T

Weights: ...

(Vogel et al., 2014)

Self-trained discriminative RSA

SELFTRAIN(Games G)

1 Initialize $S = S_0$
2 Repeat:
3 $\quad L = \text{TRAINLISTENER}(G, S)$ `# Train on S's production prefs.`
4 $\quad S = \text{TRAINSPEAKER}(G, L)$ `# Train on L's construal prefs.`
5 Return (S, L)

(Vogel et al., 2014)

Self-trained discriminative RSA

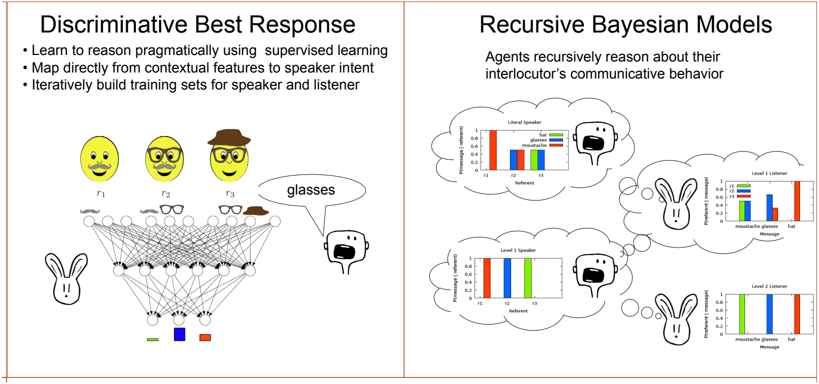

(Vogel et al., 2014)

Self-trained discriminative RSA

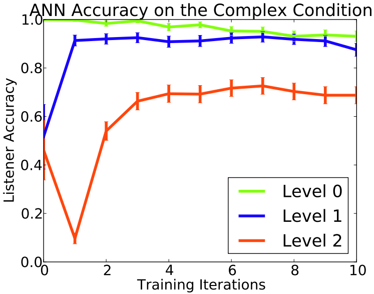

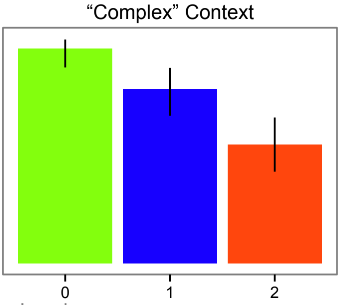

(Vogel et al., 2014)


Angel Chang


Will Monroe


Sam Bowman


Chris Manning

Scene generation

Show me an original 3d scene of a home office . . .

Scene generation

Show me an original 3d scene of a home office . . .

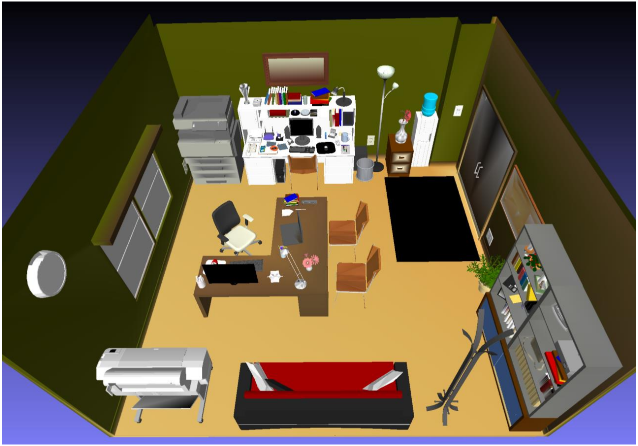

Scene as denotations

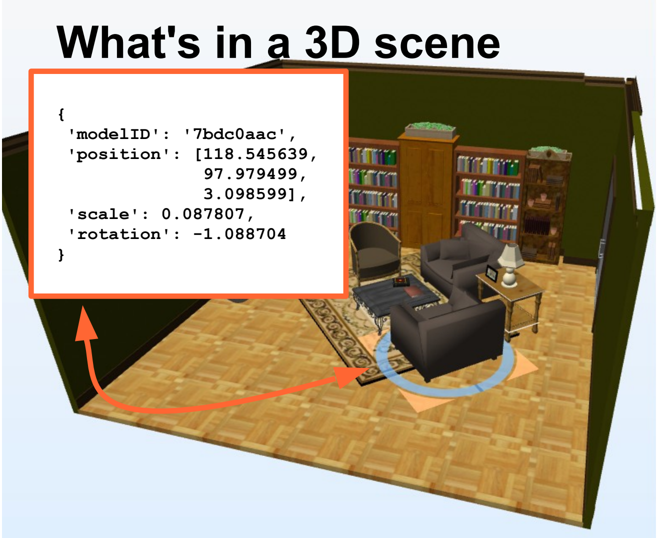

Scene as denotations

```
{
 'modelID': '7bdc0aac',
 'position': [118.545639,
              97.
              3.0
 'scale': 0.08780
 'rotation': -1.0
}
```

Field	Value
name	ellington armchair
id	7bdc0aac
tags	armchair, chair, ellington, haughton, sam, seating, woodmark
category	Chair
wnlemmas	armchair
unit	0.028974
up	[0, 0, 1]
front	[0, -1, 0]

Scene generation corpus

There is a bed and there is a chair next to the bed.

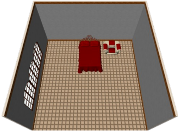

The room has three windows on one wall. There is a red bed in the back of the room. Along side the bed is a side chair that is red and white.

This room has a bed with red bedding against the wall. Next to the bed is a chair.

there is a antique looking bed with red covers and pillows in a room. next to it is a recliner chair with red padding. also there are windows.

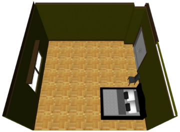

there is a bed with five pillows on it, and next to it is a chair

There is a bed in the room with two pillows and a small chair near to the right side of it.

There is a large grey bed in the bottom right corner of the room. Above the bed is a small black chair.

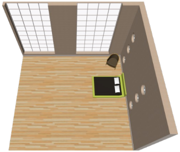

Floor to ceiling windows on back wall. Green bed with two pillows and black blanket. Lights recessed into right side wall. Light wood flooring. A chair is in the upper right hand corner

There is a bed on the side of the room. There is a chair in the corner, next to the windows.

I see a bed and a chair.

Scene generation as semantic interpretation

There is a 3 person couch and table in the center of the room.

THE GOOD

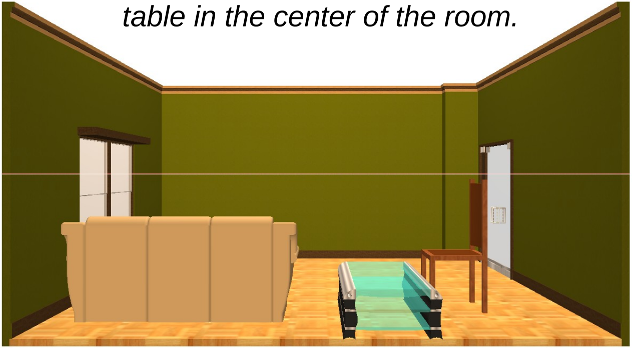

Scene generation as semantic interpretation

An L shaped couch with a vase on the corner.

THE BAD

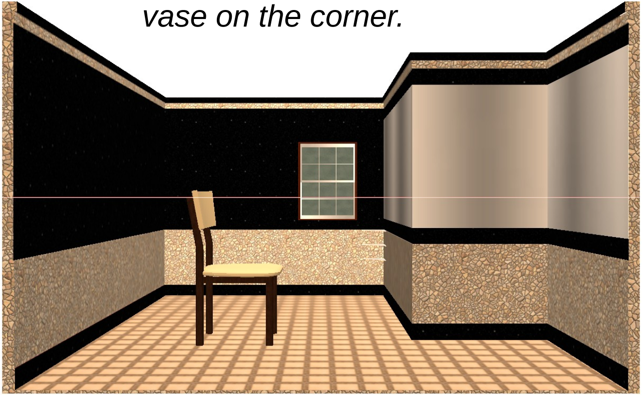

Scene generation as semantic interpretation

It is a square-shaped room with a wooden floor covered by a tan rug and an intricate wallpaper. There is a tall window in the corner with a small ceiling and desk-type object. In the middle of the room there is a gray-and-black carefully furnished bed with a simplistic gray cupboard and lamp on the opposite side of it in relation to the corner window.

THE UGLY

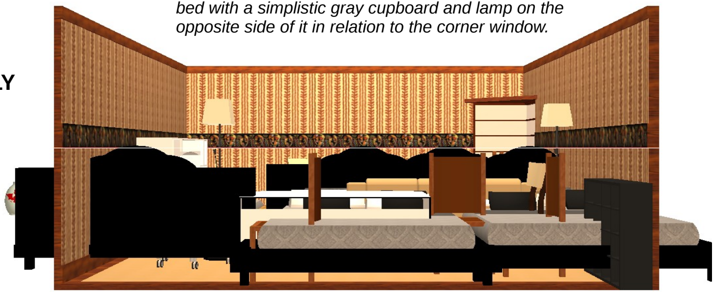

Recursive neural networks for natural logic

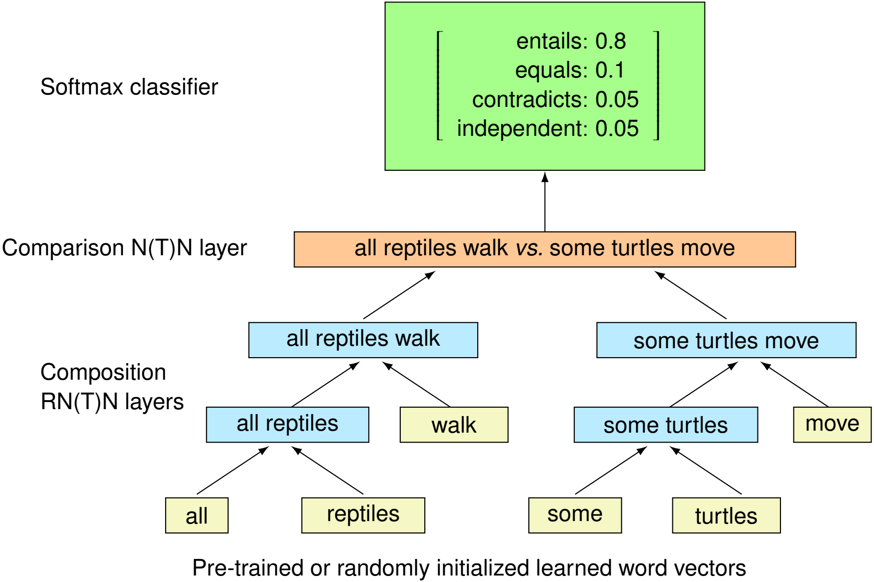

Experiments

Simulated data

- Learning the natural logic relational algebra ✓
- Learning propositional logic theorem provers ✓
- Learning to reason with quantifiers and negation ✓

(Bowman et al., 2014a,b)

Experiments

Simulated data

- Learning the natural logic relational algebra ✓
- Learning propositional logic theorem provers ✓
- Learning to reason with quantifiers and negation ✓

Naturalistic data

- WordNet relations 95% test training on 33% of the data
- The SICK textual entailment challenge 76.9% test

(Bowman et al., 2014a,b)

A new natural language inference corpus

To date: entailment, contradiction, and independence sentences for 15.5k ImageFlickr pictures/captions.

Image caption	Entailment	Contradiction	Independent
Three people with political signs.	People have signs displaying political themes.	Three people have signs promoting their football team.	Men and women are holding up political placards at a rally.
A person working for the city begins cutting down a tree.	A city employee is working outdoors.	The town sheriff is sitting on a tree swing.	A woman who works for the city is using a chainsaw.

1. The Rational Speech Acts (RSA) model
2. Training effective literal listeners
3. The joint inferences of deeply pragmatic listeners
4. The Cards task-oriented dialogue corpus
5. Language and action together

Lexical uncertainty

(Bergen et al., 2012, 2014; Potts et al., 2015)

Lexical uncertainty

1. It's a sofa, not a couch.

(Bergen et al., 2012, 2014; Potts et al., 2015)

Lexical uncertainty

1. It's a sofa, not a couch.
2. synagogues and other churches

(Bergen et al., 2012, 2014; Potts et al., 2015)

Lexical uncertainty

1. It's a sofa, not a couch.
2. synagogues and other churches
3. superb but not outstanding

(Bergen et al., 2012, 2014; Potts et al., 2015)

Lexical uncertainty

1. It's a sofa, not a couch.
2. synagogues and other churches
3. superb but not outstanding
4. $\mathbf{L}(\textit{world}, \mathcal{L} \mid \textit{msg}) \propto P(\textit{world})P(\mathcal{L})S_1(\textit{msg} \mid \textit{world}, \mathcal{L})$

(Bergen et al., 2012, 2014; Potts et al., 2015)

Lexical uncertainty

1. It's a sofa, not a couch.
2. synagogues and other churches
3. superb but not outstanding
4. $\mathbf{L}(\mathit{world}, \mathcal{L} \mid \mathit{msg}) \propto P(\mathit{world})P(\mathcal{L})S_1(\mathit{msg} \mid \mathit{world}, \mathcal{L})$
5. $\mathbf{L}(\mathit{world} \mid \mathit{msg}) \propto P(\mathit{world}) \sum_{\mathcal{L} \in \mathbb{L}} P(\mathcal{L})S_1(\mathit{msg} \mid \mathit{world}, \mathcal{L})$

(Bergen et al., 2012, 2014; Potts et al., 2015)

Anxious experts

(Levy & Potts, 2015)

Anxious experts

1. *oenophile* means *wine lover*

(Levy & Potts, 2015)

Anxious experts

1. *oenophile* means *wine lover*
2. the bow lute, such as the Bambara ndang, (Hearst, 1992)

(Levy & Potts, 2015)

Anxious experts

1. *oenophile* means *wine lover*
2. the bow lute, such as the Bambara ndang, (Hearst, 1992)
3. wine lover or oenophile

(Levy & Potts, 2015)

Anxious experts

1. *oenophile* means *wine lover*
2. the bow lute, such as the Bambara ndang, (Hearst, 1992)
3. wine lover or oenophile
4. synagogues and other churches

(Levy & Potts, 2015)

Anxious experts

1. *oenophile* means *wine lover*
2. the bow lute, such as the Bambara ndang, (Hearst, 1992)
3. wine lover or oenophile
4. synagogues and other churches
5. synagogues or churches

(Levy & Potts, 2015)

Anxious experts

1. *oenophile* means *wine lover*
2. the bow lute, such as the Bambara ndang, (Hearst, 1992)
3. wine lover or oenophile
4. synagogues and other churches
5. synagogues or churches
6. $S_2(msg \mid world, \mathcal{L}) \propto$
   $\exp\left(\alpha \log\left(L_1(world \mid msg, \mathcal{L})\right) - \beta \log\left(L_1(\mathcal{L} \mid msg)\right) - C(msg)\right)$

(Levy & Potts, 2015)

Contextual uncertainty

Contextual uncertainty

1. Chris has to miss class today.

Contextual uncertainty

1. Chris has to miss class today.
2. A friend tweeting about bread-baking and soccer:
   "Who could have predicted that?!"

Contextual uncertainty

1. Chris has to miss class today.
2. A friend tweeting about bread-baking and soccer:
   “Who could have predicted that?!”
3. Hand me the fork.

Contextual uncertainty

1. Chris has to miss class today.
2. A friend tweeting about bread-baking and soccer: “Who could have predicted that?!”
3. Hand me the fork.

4. $\mathbf{L}(\mathit{world}, \mathit{context} \mid \mathit{msg}, \mathcal{L}) \propto$
$P(\mathit{context})\, S_1(\mathit{msg}, \mid \mathit{world}, \mathit{context}, \mathcal{L})\, P_{\mathit{context}}(\mathit{world})$

Joint emotional and informational goals

1. Hyperbole
   a. I told you a thousand times already.
   b. It took a million years to get the waiter to our table.
   c. The watch cost $5000.
2. Sarcasm
   a. Oh, that's wonderful! (it's terrible)
   b. Yeah, delicious. (disgusting)
   c. Sounds great. (sounds terrible)
3. Metaphor
   a. Juliet is the sun.
   b. I feel sick as a dog.
   c. Our new boss is a shark.

(Kao et al., 2014a,b)

The Cards corpus

The Cards world

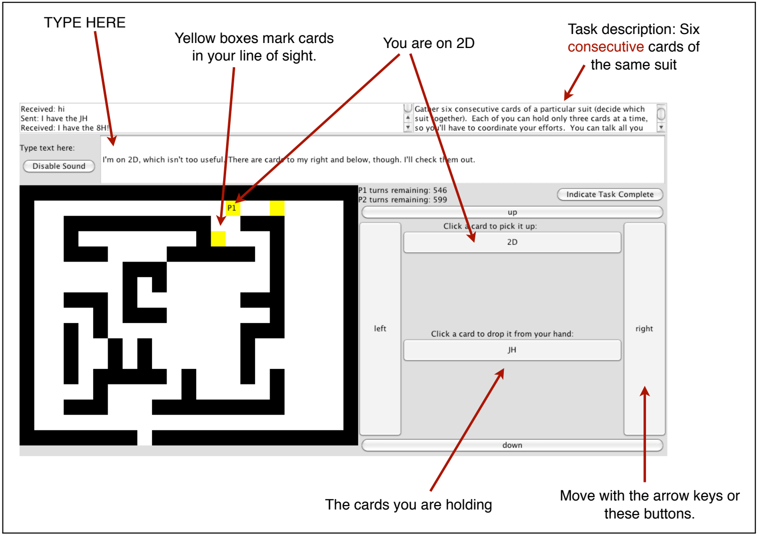

The Cards world

Gather six consecutive cards of a particular suit (decide which suit together), or determine that this is impossible. Each of you can hold only three cards at a time, so you'll have to coordinate your efforts. You can talk all you want, but you can make only a limited number of moves.

The Cards world

Gather six consecutive cards of a particular suit (decide which suit together), or determine that this is impossible. Each of you can hold only three cards at a time, so you'll have to coordinate your efforts. You can talk all you want, but you can make only a limited number of moves.

What's going on?

$\Downarrow$

Which suit should we pursue?

$\Downarrow$

Which sequence should we pursue?

$\Downarrow$

Where is card X?

By the numbers

- 1,266 transcripts
- Game length mean: 373.21 actions (median 305, sd 215.20)
- Actions:
  - Card pickup: 19,157
  - Card drop: 12,325
  - Move: 371,811
  - Utterance: 45,805
    - Utt. length mean: 5.69 words (median 5, sd 4.74)
    - Total word count: 260,788
    - Total vocabulary: ≈4,000

Task-oriented dialogue corpora

Corpus	Task type	Domain	Task-orient.	Docs.	Format
Switchboard	discussion	open	very loose	2,400	aud/txt
SCARE	search	3d world	tight	15	aud/vid/txt
TRAINS	routes	map	tight	120	aud/txt
Map Task	routes	map	tight	128	aud/vid/txt
Columbia Games	games	maps	tight	12	aud/txt
Cards	search	2d grid	tight	1,266	txt **in context**

Chief selling points for Cards:

- Pretty large.
- Controlled enough that similar things happen often.
- Very highly structured — the only corpus whose release version allows the user to replay all games with perfect fidelity.

Grounded semantics (literal listeners)

"in the bottom you see the opening on the bottom row" ⇓ BOARD(entrance & bottom); H: 5.48

"in the top right of the middle part of the board" ⇓ middle(top & right); H: 5.27

"i'm in the center" ⇓ BOARD(middle); H: 7.37

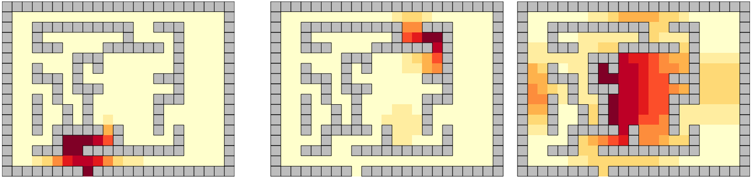

Utterances as bags of words. No preprocessing for spelling correction, lemmatization, etc. Assign semantic tags using log-linear classifiers trained on the corpus data.

Language and action, language as action


Adam Vogel

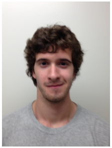
Max Bodoia


Dan Jurafsky

Simplified Cards scenario

Both players must find the ace of spades.

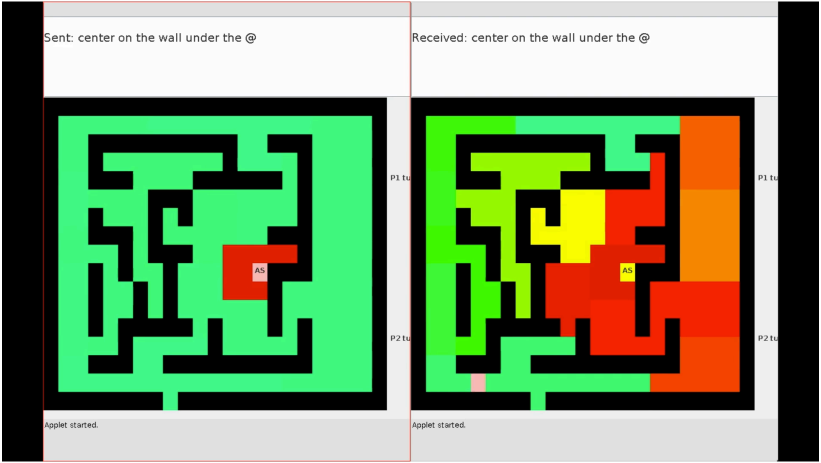

[DialogBot home movie]

Agent framework

We want our agent to:

- Make moves that are likely to lead it to the card.
- Change its behavior based on observations it receives.
- Respond to advice from the other player.
- Give advice to the other player.

Modeling the problem as a POMDP allows us to train agents that have these properties.

(Vogel et al., 2013a,b)

Approximate solutions take us only part of the way

- Even approximate solutions tractable only for $<$ 10K states.

Approximate solutions take us only part of the way

- Even approximate solutions tractable only for $<$ 10K states.
Card loc.		Agent loc.		Partner loc.		Partner's card beliefs
231	$\times$	231	$\times$	231	$\times$	231
		$\approx$ 50K		$\approx$12M		$\approx$3B

Approximate solutions take us only part of the way

- Even approximate solutions tractable only for $<$ 10K states.

- Card loc. Agent loc. Partner loc. Partner's card beliefs

  $231 \times 231 \times 231 \times 231$

  $\approx 50K$ $\approx 12M$ $\approx 3B$

- Language as a representation for planning:

231 (columns) × 130 (rows)

	(0,0)	(0,1)	...	(9,6)	(9,7)	(9,8)	(9,9)	...	(17,23)
BOARD(bottom;entrance):	0	0	.01	.03	.04	.05	.04	.12	0
BOARD(middle;top):	0	0	.02	.02	.03	.02	.02	.07	0
middle(middle):	0	0	.02	.07	.08	.09	.10	.02	0
BOARD(bottom;corner;right):	0	0	.03	.06	.05	.06	.07	.03	0
...	0	0	.11	.02	.01	.01	.01	.06	0
	0	0	.03	.13	.13	.14	.12	.01	0
	0	0	.07	.04	.05	.06	.07	.01	0
	0	0	.01	.01	.02	.01	.01	.04	0
	...	...	...	...	...	...	...	...	...

Approximate solutions take us only part of the way

- Even approximate solutions tractable only for $<$ 10K states.
Card loc.		Agent loc.		Partner loc.		Partner's card beliefs
231	$\times$	231	$\times$	231	$\times$	231
		$\approx$ 50K		$\approx$12M		$\approx$3B
- Language as a representation for planning:

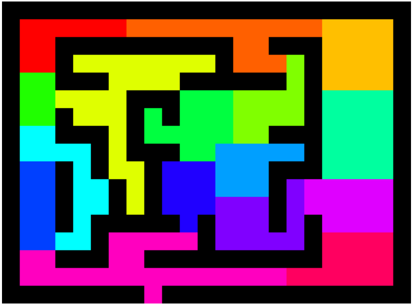

Belief-state approximation

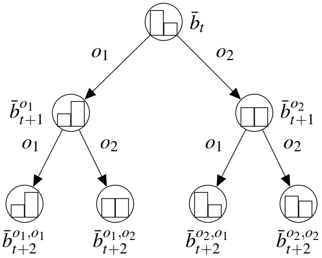

(a) Exact multi-agent belief tracking

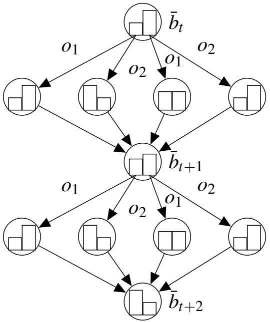

(b) Approximate multi-agent belief tracking

ListenerBot

- A POMDP agent that learns to navigate its world and interpret language.
- Driven by its small negative reward for not having the card and its large positive reward for finding it.
- No sensitivity to the other player.
- Literal listeners: each message *msg* denotes $P(\mathit{world} \mid \mathit{msg})$
- Bayes rule to incorporate these as observations.

ListenerBot example

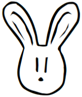

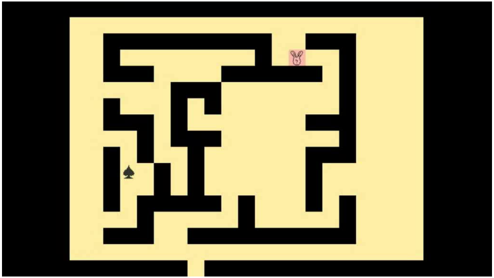

[ListenerBot home movie]

ListenerBot example

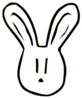

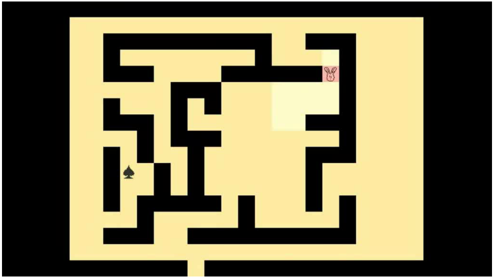

[ListenerBot home movie]

ListenerBot example

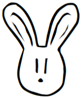

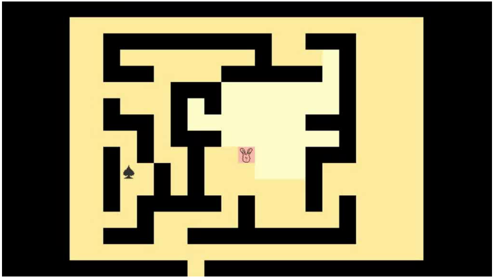

[ListenerBot home movie]

ListenerBot example

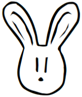

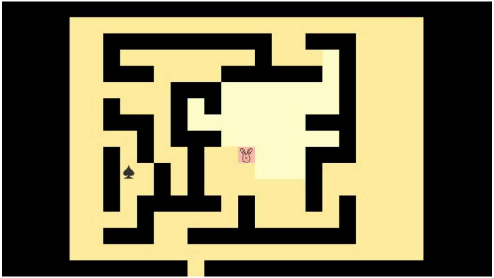

"it's on the left side"

$\Downarrow$

$\textsc{board}(\mathbf{left})$

$\Downarrow$

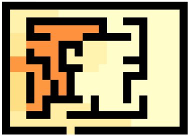

[ListenerBot home movie]

ListenerBot example

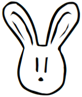

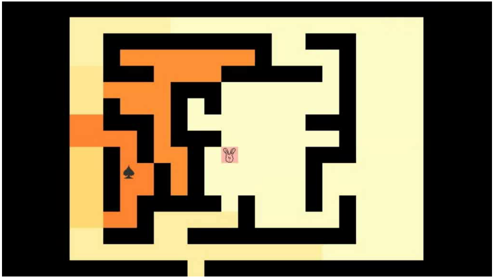

"it's on the left side"

$\Downarrow$

BOARD(left)

$\Downarrow$

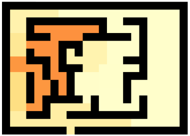

[ListenerBot home movie]

DialogBot

DialogBot is a strict extension of Listener Bot:

- The set of states is now all combinations of
  - both players' positions
  - the card's region
  - the region the other player believes the card to be in
- The set of actions now includes dialogue actions.
- (The player assumes that) a dialogue action *U* alters the other player's beliefs in the same way that *U* would impact his own.
- Same basic reward structure as for Listenerbot, except now also sensitive to whether the other player has found the card.
- Approximate RSA is a special case.

How the agents relate to each other

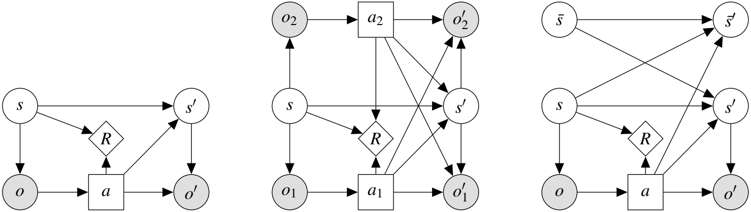

(a) ListenerBot POMDP (b) Full Dec-POMDP (c) DialogBot POMDP

DialogBot and ListenerBot play together

DialogBot beliefs

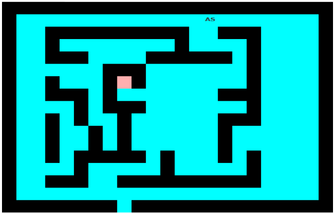

ListenerBot beliefs

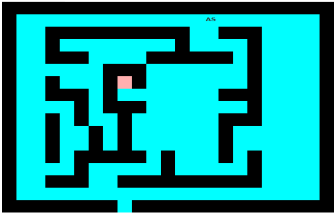

DialogBot beliefs: ListenerBot's position

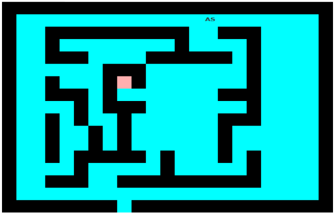

DialogBot beliefs: ListenerBot's beliefs

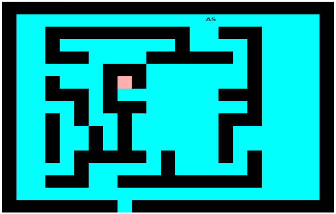

DialogBot and ListenerBot play together

DialogBot beliefs

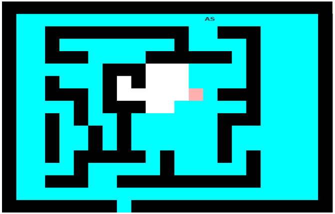

ListenerBot beliefs

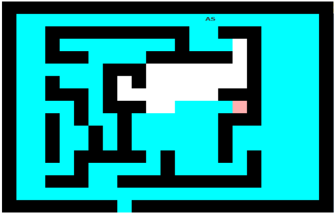

DialogBot beliefs: ListenerBot's position

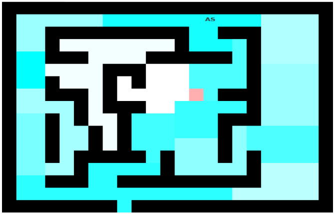

DialogBot beliefs: ListenerBot's beliefs

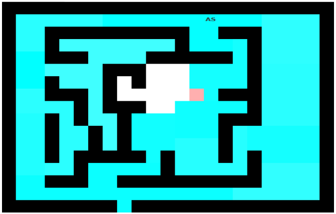

DialogBot and ListenerBot play together

Dialogbot: “Top”

DialogBot beliefs

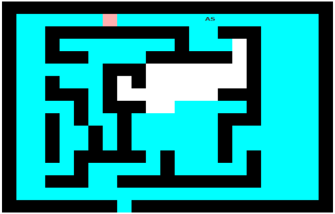

ListenerBot beliefs

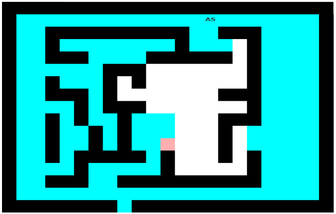

DialogBot beliefs: ListenerBot’s position

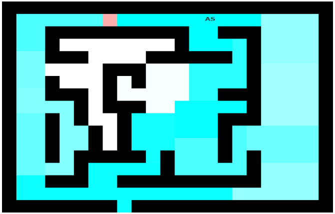

DialogBot beliefs: ListenerBot’s beliefs

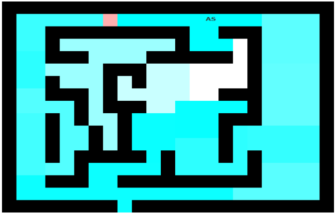

DialogBot and ListenerBot play together

Dialogbot: "Top"

DialogBot beliefs

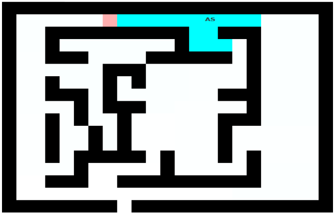

ListenerBot beliefs

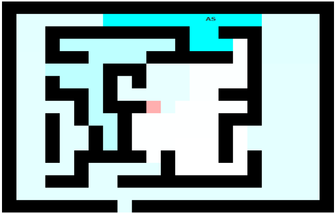

DialogBot beliefs: ListenerBot's position

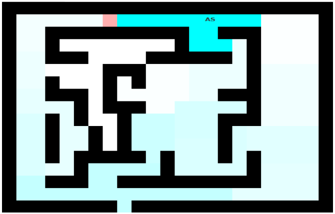

DialogBot beliefs: ListenerBot's beliefs

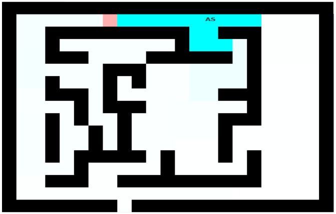

DialogBot and ListenerBot play together

DialogBot beliefs

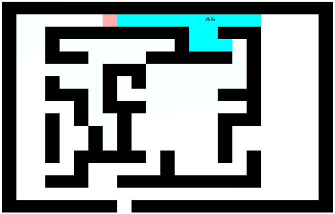

ListenerBot beliefs

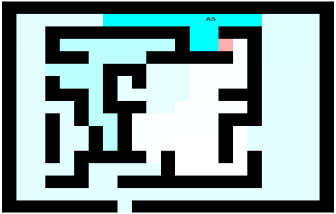

DialogBot beliefs: ListenerBot's position

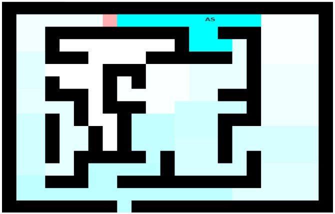

DialogBot beliefs: ListenerBot's beliefs

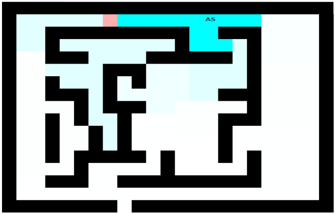

DialogBot and ListenerBot play together

DialogBot beliefs

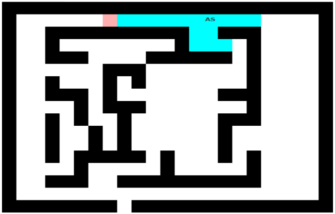

ListenerBot beliefs

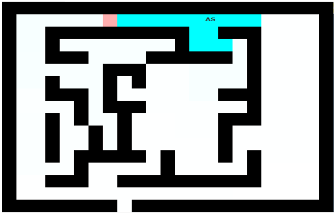

DialogBot beliefs: ListenerBot’s position

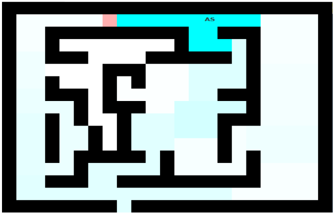

DialogBot beliefs: ListenerBot’s beliefs

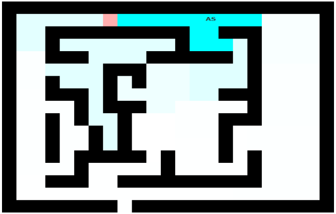

DialogBot and ListenerBot play together

DialogBot beliefs

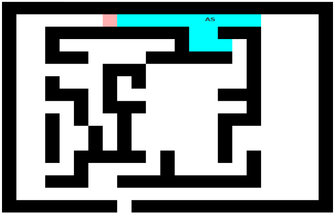

ListenerBot beliefs

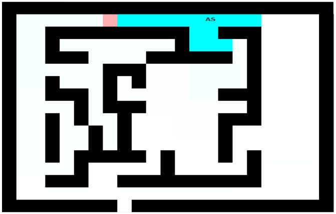

DialogBot beliefs: ListenerBot's position

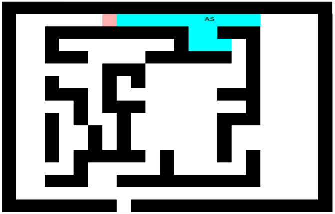

DialogBot beliefs: ListenerBot's beliefs

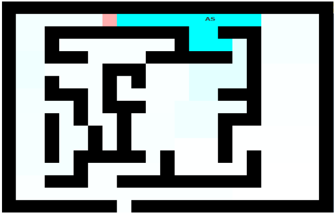

Emergent pragmatics

Quality

- The Gricean maxim of quality says roughly “Be truthful”.
- For DialogBot, this emerges from the decision problem: false information is (typically) more costly.
- DialogBot would lie if he thought it would move them toward the objective.

Quantity and Relevance

- The Gricean maxims of quantity and relevance for informative, timely contributions.
- When DialogBot finds the card, he communicates the information, not because he is hard-coded to do so, but rather because it will help the other player find it.

Grown-up DialogBots (a week of policy exploration)

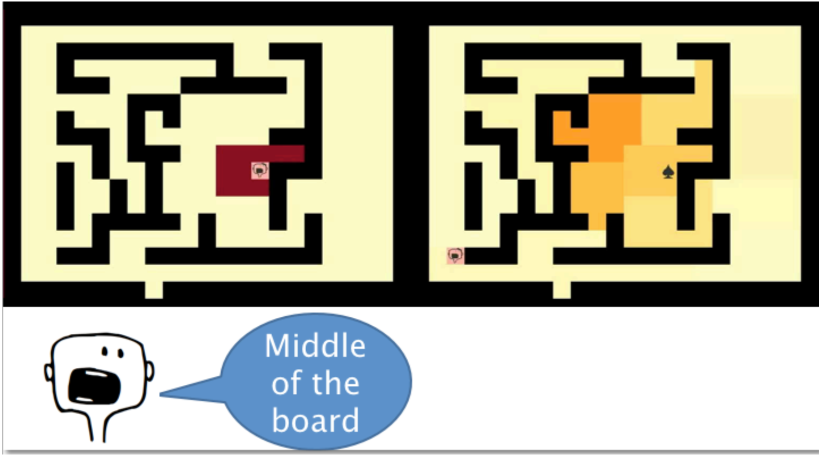

Baby DialogBots (a few hours of policy exploration)

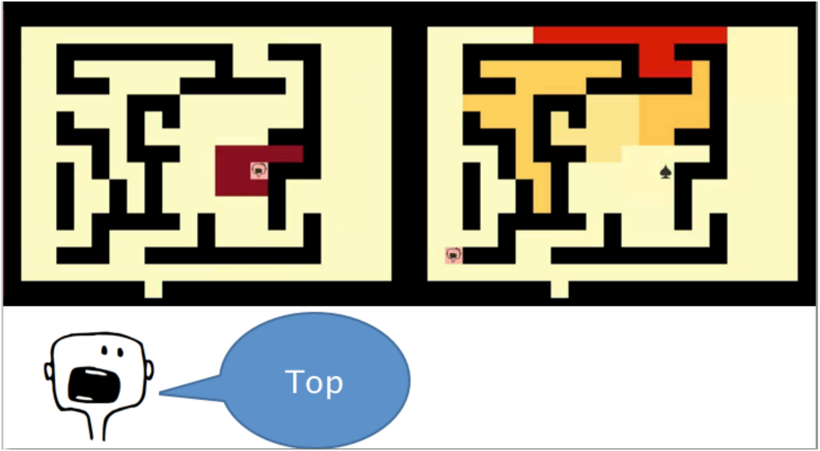

Experimental results

Agents	% Success	Average Moves
ListenerBot & ListenerBot	84.4%	19.8
ListenerBot & DialogBot	87.2%	17.5
DialogBot & DialogBot	90.6%	16.6

Table: 500 random initial states per agent combination.

Literal interpretations in the Cards world

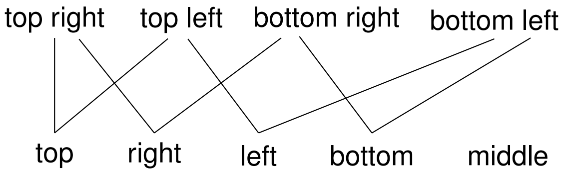

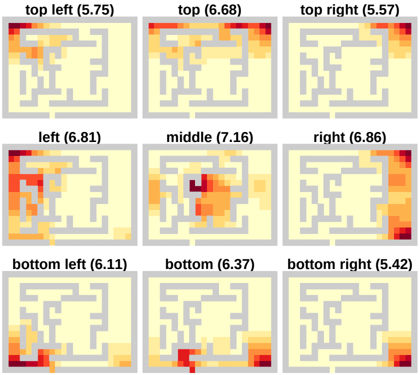

Implicature in the Cards world

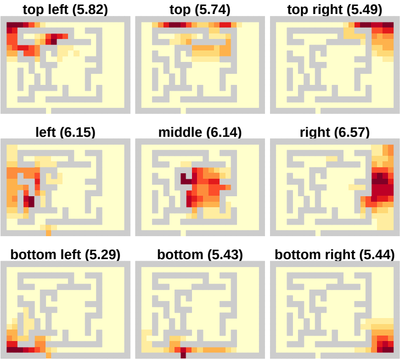

Figure: Human

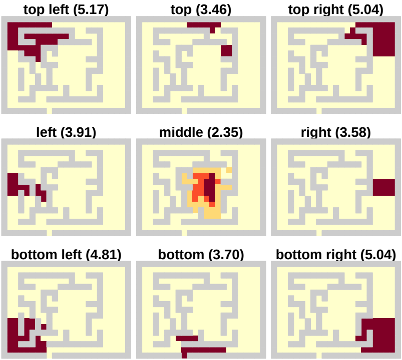

Figure: DialogBot

1. Literal speaker S: finds the cards and utters the message with the highest literal probability for his location.
2. Level-one listener $L(S)$: interprets each message as the set of beliefs that S must have to produce it.

Looking ahead

Literal listeners

- Continued data collection of grounded interpretations.
- More compositionality, richer semantic representations.

Deep understanding of short-form communications

- Crowdsourced annotation for literal listeners.
- Joint inferences about context and construal.

CRF agents bringing language and action together

- A model-free approach to developing communication agents.
- So far: competitive listener bots in a fraction of the time.

Looking ahead

Literal listeners

- Continued data collection of grounded interpretations.
- More compositionality, richer semantic representations.

Deep understanding of short-form communications

- Crowdsourced annotation for literal listeners.
- Joint inferences about context and construal.

CRF agents bringing language and action together

- A model-free approach to developing communication agents.
- So far: competitive listener bots in a fraction of the time.

Thanks!

References I

Bergen, Leon, Noah D. Goodman & Roger Levy. 2012. That's what she (could have) said: How alternative utterances affect language use. In Naomi Miyake, David Peebles & Richard P. Cooper (eds.), *Proceedings of the thirty-fourth annual conference of the Cognitive Science Society*, 120–125. Austin, TX: Cognitive Science Society.

Bergen, Leon, Roger Levy & Noah D. Goodman. 2014. Pragmatic reasoning through semantic inference. Ms., MIT, UCSD, and Stanford.

Bowman, Samuel R., Christopher Potts & Christopher D. Manning. 2014a. Learning distributed word representations for natural logic reasoning. arXiv manuscript 1410.4176.

Bowman, Samuel R., Christopher Potts & Christopher D. Manning. 2014b. Recursive neural networks for learning logical semantics. arXiv manuscript 1406.1827.

Camerer, Colin F., Teck-Hua Ho & Juin-Kuan Chong. 2004. A cognitive hierarchy model of games. *The Quarterly Journal of Economics* 119(3). 861–898.

Clark, Eve V. 1987. The principle of contrast: A constraint on language acquisition. In Brian MacWhinney (ed.), *Mechanisms of language acquisition*, 1–33. Hillsdale, NJ: Erlbaum.

Clark, Herbert H. 1996. *Using language*. Cambridge: Cambridge University Press.

Degen, Judith & Michael Franke. 2012. Optimal reasoning about referential expressions. In *Proceedings of SemDIAL 2012*, Paris.

Frank, Michael C. & Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998.

Frank, Michael C., Noah D. Goodman & Joshua B. Tenenbaum. 2009. Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science* 20(5). 578–585.

Franke, Michael. 2008. Interpretation of optimal signals. In Krzysztof R. Apt & Robert van Rooij (eds.), *New perspectives on games and interaction*, vol. 4 Texts in Logics and Games, 297–310. Amsterdam University Press.

Franke, Michael. 2009. *Signal to act: Game theory in pragmatics* ILLC Dissertation Series. Institute for Logic, Language and Computation, University of Amsterdam.

Golland, Dave, Percy Liang & Dan Klein. 2010. A game-theoretic approach to generating spatial descriptions. In *Proceedings of the 2010 conference on empirical methods in natural language processing*, 410–419. Stroudsburg, PA: ACL. http://www.aclweb.org/anthology/D10-1040.

Grice, H. Paul. 1975. Logic and conversation. In Peter Cole & Jerry Morgan (eds.), *Syntax and semantics*, vol. 3: Speech Acts, 43–58. New York: Academic Press.

References II

Grodner, Daniel J., Natalie M. Klein, Kathleen M. Carbary & Michael K. Tanenhaus. 2010. "Some," and possibly all, scalar inferences are not delayed: Evidence for immediate pragmatic enrichment. *Cognition* 116(1). 42–55.

Grodner, Daniel J. & Julie Sedivy. 2008. The effects of speaker-specific information on pragmatic inferences. In Edward A. Gibson & Neal J. Pearlmutter (eds.), *The processing and acquisition of reference*, 239–272. Cambridge, MA: MIT Press.

Hearst, Marti A. 1992. Automatic acquisition of hyponyms from large text corpora. In *Proceedings of COLING 1992*, 539–545. Nantes: Association for Computational Linguistics.

Huang, Ti Ting & Jesse Snedeker. 2009. Online interpretation of scalar quantifiers: Insight into the semantics–pragmatics interface. *Cognitive Psychology* 58(3). 376–415.

Jäger, Gerhard. 2007. Game dynamics connects semantics and pragmatics. In Ahti-Veikko Pietarinen (ed.), *Game theory and linguistic meaning*, 89–102. Amsterdam: Elsevier.

Jäger, Gerhard. 2012. Game theory in semantics and pragmatics. In Claudia Maienborn, Klaus von Heusinger & Paul Portner (eds.), *Semantics: An international handbook of natural language meaning*, vol. 3, 2487–2425. Berlin: Mouton de Gruyter.

Kao, Justine T., Leon Bergen & Noah D. Goodman. 2014a. Formalizing the pragmatics of metaphor understanding. In *Proceedings of the 36th annual meeting of the cognitive science society*, 719–724. Wheat Ridge, CO: Cognitive Science Society.

Kao, Justine T., Jean Y. Wu, Leon Bergen & Noah D. Goodman. 2014b. Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences* 111(33). 12002–12007. doi:10.1073/pnas.1407479111.

Levinson, Stephen C. 2000. *Presumptive meanings: The theory of generalized conversational implicature*. Cambridge, MA: MIT Press.

Levy, Roger & Christopher Potts. 2015. Negotiating lexical uncertainty and expertise with disjunction. Poster presented at the 89th Annual Meeting of the Linguistic Society of America.

Lewis, David. 1969. *Convention*. Cambridge, MA: Harvard University Press. Reprinted 2002 by Blackwell.

Potts, Christopher, Daniel Lassiter, Roger Levy & Michael C. Frank. 2015. Embedded implicatures as pragmatic inferences under compositional lexical uncertainty. Ms., Stanford and UCSD.

Rabin, Matthew. 1990. Communication between rational agents. *Journal of Economic Theory* 51(1). 144–170. doi:10.1016/0022-0531(90)90055-O.

Rosenberg, Seymour & Bertram D. Cohen. 1964. Speakers' and listeners' processes in a word communication task. *Science* 145. 1201–1203.

References III

Smith, Nathaniel J., Noah D. Goodman & Michael C. Frank. 2013. Learning and using language via recursive pragmatic reasoning about other agents. In *Advances in neural information processing systems 26*, 3039–3047.

Stiller, Alex, Noah D. Goodman & Michael C. Frank. 2011. Ad-hoc scalar implicature in adults and children. In Laura Carlson, Christoph Hoelscher & Thomas F. Shipley (eds.), *Proceedings of the 33rd annual meeting of the Cognitive Science Society*, 2134–2139. Austin, TX: Cognitive Science Society.

Vogel, Adam, Max Bodoia, Christopher Potts & Dan Jurafsky. 2013a. Emergence of Gricean maxims from multi-agent decision theory. In *Human language technologies: The 2013 annual conference of the North American chapter of the Association for Computational Linguistics*, 1072–1081. Stroudsburg, PA: Association for Computational Linguistics.

Vogel, Adam, Andrés Gómez Emilsson, Michael C. Frank, Dan Jurafsky & Christopher Potts. 2014. Learning to reason pragmatically with cognitive limitations. In *Proceedings of the 36th annual meeting of the Cognitive Science Society*, 3055–3060. Wheat Ridge, CO: Cognitive Science Society.

Vogel, Adam, Christopher Potts & Dan Jurafsky. 2013b. Implicatures and nested beliefs in approximate Decentralized-POMDPs. In *Proceedings of the 2013 annual conference of the Association for Computational Linguistics*, 74–80. Stroudsburg, PA: Association for Computational Linguistics.