

# Dialogue agents

Christopher Potts

CS 244U: Natural language understanding  
May 21



## Plan and goals

- ① Language in context, language as action
- ② The Switchboard Dialog-Act Corpus: modeling interactional linguistic patterns
- ③ The Cards Corpus: task-oriented dialogue
- ④ PLOW: dialogue agents for the Web browser
- ⑤ A framework for increasingly sophisticated dialogue agents:
  - a. Markov Decision Processes (MDP)
  - b. Partially Observable MDP (POMDP)
  - c. Decentralized POMDP (Dec-POMDP)
  - d. Approximations of b and c

## Plan and goals

- ① Language in context, language as action
- ② The Switchboard Dialog-Act Corpus: modeling interactional linguistic patterns
- ③ The Cards Corpus: task-oriented dialogue
- ④ PLOW: dialogue agents for the Web browser
- ⑤ A framework for increasingly sophisticated dialogue agents:
  - a. Markov Decision Processes (MDP)
  - b. Partially Observable MDP (POMDP)
  - c. Decentralized POMDP (Dec-POMDP)
  - d. Approximations of **b** and **c**

Deeply influenced by **Adam Vogel**, who successfully defended his Stanford PhD thesis this morning!



# Language is action

## Winograd (1986:170):

“all language use can be thought of as a way of activating procedures within the hearer. We can think of an utterance as a program – one that indirectly causes a set of operations to be carried out within the hearer’s cognitive system.”

## Eliza (Weizenbaum 1966)

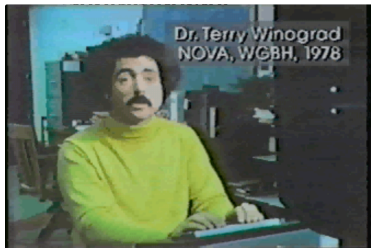
- 1 User: You are like my father in some ways.
- 2 ELIZA: What resemblance do you see
- 3 User: You are not very aggressive but I think you don't want me to notice that.
- 4 ELIZA: What makes you think I am not aggressive?
- 5 User: You don't argue with me.
- 6 ELIZA: Why do you think I don't argue with you?
- 7 User: You are afraid of me.
- 8 ELIZA: Does it please you to believe I am afraid of you?

Rewrite user's reply by (i) swapping 1st and 3rd person, (ii) interpolating stock phrases, and (iii) using scores to rank possible transformations.

## SHRDLU (Winograd 1972)

Full-grounded system that parses the user's input, maps it to logical form, interprets that logical form in its world, tries to take appropriate action, and generates responses.

<http://hci.stanford.edu/winograd/shrdlu/>



*One project did succeed. Terry Winograd's program SHRDLU could use English intelligently, but there was a catch: the only subject you could discuss was a micro-world of simulated blocks.*

# SHRDLU (Winograd 1972)

- 1 Person: WHAT DOES THE BOX CONTAIN?
- 2 Computer: THE BLUE PYRAMID AND THE BLUE BLOCK.
- 3 Person: WHAT IS THE PYRAMID SUPPORTED BY?
- 4 Computer: THE BOX.
- 5 Person: HOW MANY BLOCKS ARE NOT IN THE BOX?
- 6 Computer: FOUR OF THEM.
- 7 Person: IS AT LEAST ONE OF THEM NARROWER THAN THE ONE WHICH I TOLD YOU TO PICK UP?
- 8 Computer: YES, THE RED CUBE.
- 9 Person: IS IT SUPPORTED?
- 10 Computer: YES, BY THE TABLE.

# HAL

- In the 1967 Stanley Kubrick movie *2001: A Space Odyssey*, the spaceship's computer HAL can
  - display graphics;
  - play chess; and
  - conduct natural, open-domain conversations with humans.
- How well did the filmmakers do at predicting what computers would be capable in 2001?

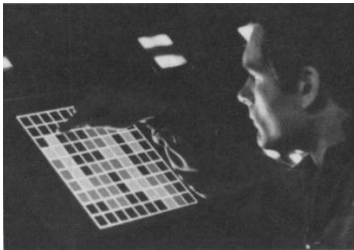
(Slide idea from Andrew McCallum)



## HAL

## Graphics

## HAL



## Jurassic Park (1993)



(Slide idea from Andrew McCallum)

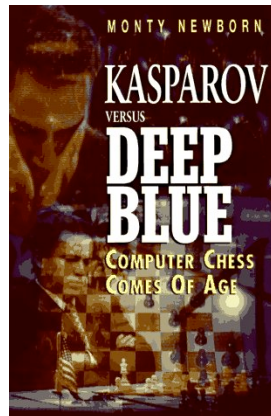
# HAL

## Chess

HAL



Deep Blue (1997)



(Slide idea from Andrew McCallum)

# HAL

## Dialogue

HAL

2014

David Bowman: Open the pod bay doors, HAL.

HAL: I'm sorry, Dave, I'm afraid I can't do that.

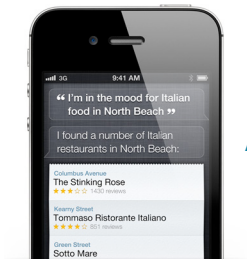
David: What are you talking about, HAL?

HAL: I know that you and Frank were planning to disconnect me, and I'm afraid that's something I cannot allow to happen.



(Slide idea from Andrew McCallum)

# Siri



**You:** Any good burger joints around here?

**Siri:** I found a number of burger restaurants near you.

**You:** Hmm. How about tacos?

**Apple:** [Siri remembers that you asked about restaurants. so it will look for Mexican restaurants in the neighborhood. And Siri is proactive, so it will question you until it finds what you're looking for.]

(Slide from Marie de Marneffe)

# Siri

**Colbert:** For the love of God, the cameras are on, give me something?

**Siri:** What kind of place are you looking for?  
Camera stores or churches?

[...]

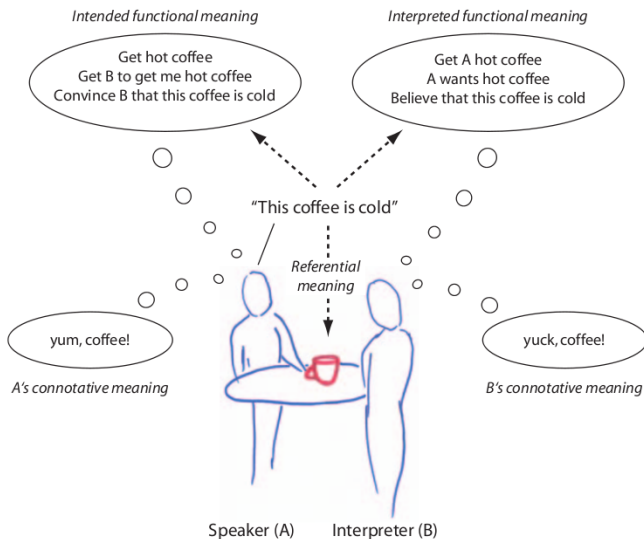
**Colbert:** I don't want to search for anything! I want to write the show!

**Siri:** Searching the Web for "search for anything. I want to write the shuffle."



(Slide from Marie de Marneffe)

# Language in context



## Discourse models

Discourse is structured by an evolving set of abstract, implicit, issues about which the participants have only partial knowledge:

### Questions under Discussion (QUDs)

McCarthy 1980; Groenendijk and Stokhof 1984; Rooth 1985; Lewis 1988; Ginzburg 1996a; Roberts 1996; Büring 1999

### Decision problems (games)

Lewis 1969; Clark 1996; Merin 1997; Blutner 1998; Parikh 2001; Beaver 2002; van Rooy 2003; Benz et al. 2005; Franke 2009

### Goal-orientation

Perrault and Allen 1980; Allen 1991; Hobbs et al. 1993; Graff 2000; Allen et al. 2007; Stone et al. 2007

For much more: <http://www.ling.ohio-state.edu/~croberts/QUDbib/>

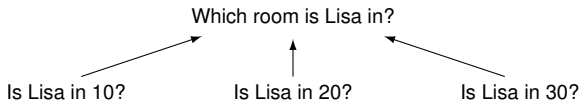
## Pragmatically required over-answering

**Context:** Homer calls a hotel.

Homer: Is Lisa Simpson in Room 10?

Clerk A: She's in room 20.

Clerk B: #No.



(Roberts 1996; Ginzburg 1996a; Champollion 2008)



# Domain restriction

*I didn't see any.*

(Roberts 1996; Ginzburg 1996a; Malamud 2006)

## Domain restriction

- Are there typos in my slides?

*I didn't see any.*

(Roberts 1996; Ginzburg 1996a; Malamud 2006)

## Domain restriction

- Are there typos in my slides?
- Are there bookstores downtown?

*I didn't see any.*

(Roberts 1996; Ginzburg 1996a; Malamud 2006)

## Domain restriction

- Are there typos in my slides?
- Are there bookstores downtown?
- Are there cookies in the cupboard?

*I didn't see any.*

(Roberts 1996; Ginzburg 1996a; Malamud 2006)

## Domain restriction

- Are there typos in my slides?
- Are there bookstores downtown?
- Are there cookies in the cupboard?
- ...

*I didn't see any.*

(Roberts 1996; Ginzburg 1996a; Malamud 2006)

# Granularity

*Where are you from?*

(Groenendijk and Stokhof 1984; Ginzburg 1996b)

# Granularity

*Where are you from?*

- *Connecticut.*

(Issue: birthplaces)

(Groenendijk and Stokhof 1984; Ginzburg 1996b)

# Granularity

*Where are you from?*

- *Connecticut.* (Issue: birthplaces)
- *The U.S.* (Issue: nationalities)

(Groenendijk and Stokhof 1984; Ginzburg 1996b)



# Granularity

*Where are you from?*

- *Connecticut.* (Issue: birthplaces)
- *The U.S.* (Issue: nationalities)
- *Stanford.* (Issue: affiliations)

(Groenendijk and Stokhof 1984; Ginzburg 1996b)

# Granularity

*Where are you from?*

- *Connecticut.* (Issue: birthplaces)
- *The U.S.* (Issue: nationalities)
- *Stanford.* (Issue: affiliations)
- *Planet earth.* (Issue: intergalactic meetings)

(Groenendijk and Stokhof 1984; Ginzburg 1996b)

# Summary of corpus resources

- SwDA:  
<http://www.stanford.edu/~jurafsky/ws97/>
- SwDA with Treebank3 alignment:  
<http://comp Prag.christopherpotts.net/swda.html>
- Edinburgh Map Corpus:  
<http://groups.inf.ed.ac.uk/maptask/>
- TRIPS:  
<http://www.cs.rochester.edu/research/cisd/projects/trips/>
- TRAINS:  
<http://www.cs.rochester.edu/research/cisd/projects/trains/>
- Cards:  
<http://CardsCorpus.christopherpotts.net/>
- SCARE:  
<http://slate.cse.ohio-state.edu/quake-corpora/scare/>
- The Carnegie Mellon Communicator Corpus (human-computer):  
<http://www.speech.cs.cmu.edu/Communicator/>

# A decision-theoretic framework for dialogue agents

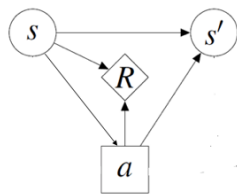


Figure: MDP

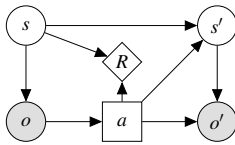


Figure: POMDP

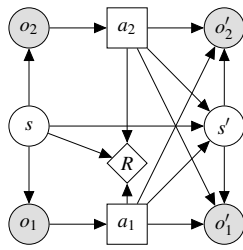


Figure: Dec-POMDP

## The Switchboard Dialog-Act Corpus

- The SwDA extends the Switchboard-1 Telephone Speech Corpus, Release 2, with turn/utterance-level dialog-act tags.
- The tags summarize syntactic, semantic, and pragmatic information about the associated turn.
- It is freely available:

<http://www.stanford.edu/~jurafsky/ws97/>

- The SwDA is not inherently linked to the Penn Treebank 3 parses of Switchboard, and it is far from straightforward to align the two resources (Calhoun et al. 2010).
- In addition, the SwDA is not distributed with the Switchboard's tables of metadata about the conversations and their participants.
- I created a CSV version of the corpus that pools all of this information to the best of my ability, thereby allowing study of the correlations among dialog tags, conversational metadata, and full syntactic structures:

<http://compprag.christopherpotts.net/swda.html>

## Example dialogue

```

^h   A.1 utt1:  {F Uh, } let's see. /
%    A.1 utt2:  How [ about, + {F uh, } let's see, about ] ten
qo   A.1 utt3:  {F uh, } what do you think was different ten ye
sv   B.2 utt1:  {D Well, } I would say as, far as social change
sv   B.2 utt2:  [ They, + they ] did more things together. /
b    @A.3 utt1: Uh-huh <>. /
sv   B.4 utt1:  {F Uh, } they ate dinner at the table together.
sv   B.4 utt2:  {F Uh, } the parents usually took out [ time, +
b    A.5 utt1:  Uh-huh. /
sv   B.6 utt1:  {F Uh, } although I'm not a mother, [ I, + I ]
qo   B.6 utt2:  {F Uh, } what # do you # --
%    A.7 utt1:  # We, # -/
+    B.8 utt1:  -- think about that? /
.
.
.

```

Table: FILENAME: 4360\_1599\_1589

# DAMSL tags

There are over 200 tags in the SwDA, most used only a few times. It is more common to work with a collapsed version involving just 44 tags.

	name	act tag	example	train count	full count
1	Statement-non-opinion	sd	Me, I'm in the legal department.	72824	75145
2	Acknowledge (Backchannel)	b	Uh-huh.	37096	38298
3	Statement-opinion	sv	I think it's great	25197	26428
4	Agree/Accept	aa	That's exactly it.	10820	11133
5	Abandoned or Turn-Exit	%	So, -	10569	15550
6	Appreciation	ba	I can imagine.	4633	4765
7	Yes-No-Question	qy	Do you have to have any special training?	4624	4727
8	Non-verbal	x	[Laughter], [Throat-clearing]	3548	3630
9	Yes answers	ny	Yes.	2934	3034
10	Conventional-closing	fc	Well, it's been nice talking to you.	2486	2582
11	Uninterpretable	%	But, uh, yeah	2158	15550
12	Wh-Question	qw	Well, how old are you?	1911	1979
13	No answers	nn	No.	1340	1377
14	Response Acknowledgement	bk	Oh, okay.	1277	1306
15	Hedge	h	I don't know if I'm making any sense or not.	1182	1226
16	Declarative Yes-No-Question	qy^d	So you can afford to get a house?	1174	1219
17	Other fo_o_fw_by_bc		Well give me a break, you know.	1074	883
18	Backchannel in question form	bh	Is that right?	1019	1053
19	Quotation	^q	You can't be pregnant and have cats	934	983
20	Summarize/reformulate	bf	Oh, you mean you switched schools for the kids.	919	952
21	Affirmative non-yes answers	na	It is.	836	847
22	Action-directive	ad	Why don't you go first	719	746

## DAMSL tags

There are over 200 tags in the SwDA, most used only a few times. It is more common to work with a collapsed version involving just 44 tags.

	name	act tag	example	train count	full count
23	Collaborative Completion	^2	Who aren't contributing.	699	723
24	Repeat-phrase	b^m	Oh, fajitas	660	688
25	Open-Question	qo	How about you?	632	656
26	Rhetorical-Questions	qh	Who would steal a newspaper?	557	575
27	Hold before answer/agreement	^h	I'm drawing a blank.	540	556
28	Reject	ar	Well, no	338	346
29	Negative non-no answers	ng	Uh, not a whole lot.	292	302
30	Signal-non-understanding	br	Excuse me?	288	298
31	Other answers	no	I don't know	279	286
32	Conventional-opening	fp	How are you?	220	225
33	Or-Clause	qrr	or is it more of a company?	207	209
34	Dispreferred answers	arp_nd	Well, not so much that.	205	207
35	3rd-party-talk	t3	My goodness, Diane, get down from there.	115	117
36	Offers, Options, Commits	oo_co_cc	I'll have to check that out	109	110
37	Self-talk	t1	What's the word I'm looking for	102	103
38	Downplayer	bd	That's all right.	100	103
39	Maybe/Accept-part	aap_am	Something like that	98	105
40	Tag-Question	^g	Right?	93	92
41	Declarative Wh-Question	qw^d	You are what kind of buff?	80	80
42	Apology	fa	I'm sorry.	76	79
43	Thanking	ft	Hey thanks a lot	67	78



## Switchboard Dialog Act Corpus with parsetrees

- My release of the SwDA includes the Treebank3 POS tags.
- It also includes the Treebank3 trees, but these are somewhat more challenging to work with:
  - Only 118,218 (53%) of utterances have trees.
  - The Treebank3 team merged some utterances into single trees.
  - Other utterances were split across trees.
  - The turn numbering was altered, often dramatically.
- On the bright side:
  - 82% of the utterances with trees correspond to a single tree.
  - With the exception of non-verbal (x) and tag-questions (^g), the distribution of tags in this subset is basically the same as the distribution for the whole corpus.
- Additional details: <http://compprag.christopherpotts.net/swda.html>

# Act-tag and syntactic category alignment

A quick experiment: to what extent are dialog act tags and clause-types aligned?

## 1 Request act

- a. Take these pills twice a day.
- b. You should take these twice a day.
- c. Could you please take these twice a day?

## 2 Question act

- a. Is today Tuesday?
- b. It's Tuesday, right?
- c. I need to confirm that it's Tuesday.

## 3 Imperative form

- a. Take these pills twice a day.
- b. Have a seat.
- c. Get well soon.

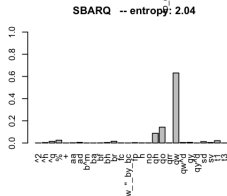
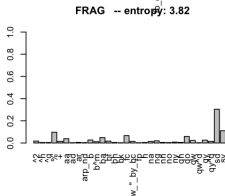
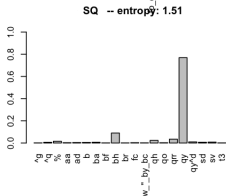
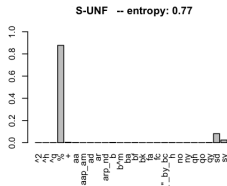
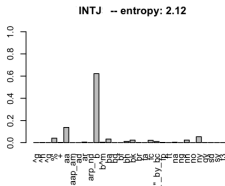
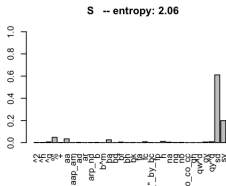
## 4 Interrogative

- a. Is today Tuesday?
- b. Is he ever tall!
- c. Can you pass the salt?

# Act-tag and syntactic category alignment

A quick experiment: to what extent are dialog act tags and clause-types aligned?

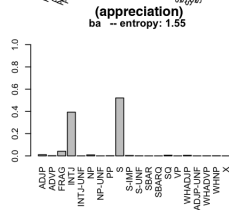
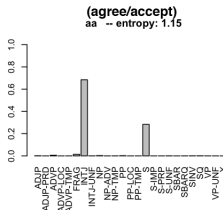
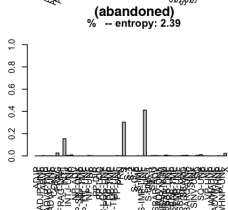
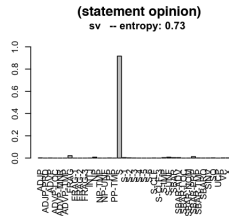
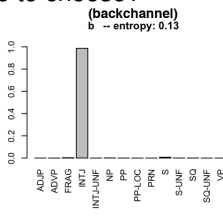
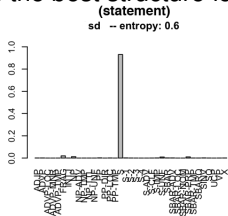
**The hearer's perspective:** given that I heard a syntactic structure with root label  $L$ , what are the speaker's possible intended dialog acts?



# Act-tag and syntactic category alignment

A quick experiment: to what extent are dialog act tags and clause-types aligned?

**The speaker's perspective:** given that I want to convey dialog act  $D$ , what is the best structure for me to choose?



## Modeling act sequences

- Modeling act sequences could be an important step towards realistic interpretation and production.
- Shriberg et al. (1998) and Stolcke et al. (2000) use acoustic features to predict general dialog act labels, using the SwDA. Their model is a decision-tree classifier.
- Other classifiers might also be appropriate; the natural assumption here is that the classifications decisions are made on a by-utterance basis, with no inspection of neighboring utterances (Bangalore et al. 2006; Kumar Rangarajan Sridhar et al. 2009).
- Dialog act prediction can also be viewed as a sequence modeling problem akin to POS tagging, and thus Hidden Markov Models and Conditional Random Fields models are often used. Such models incorporate earlier and/or later tags to make classification decisions.

# On the SwDA for dialogue research

## Advantages

- Richly annotated.
- Includes speech data.
- Includes sociolinguistic metadata.
- Long conversations, and lots of them.
- Participants did not typically know each other before the conversation, so most of their common ground is general knowledge.

# On the SwDA for dialogue research

## Advantages

- Richly annotated.
- Includes speech data.
- Includes sociolinguistic metadata.
- Long conversations, and lots of them.
- Participants did not typically know each other before the conversation, so most of their common ground is general knowledge.

## Disadvantages

- Open-domain, unfocussed (participants do not stick to their topics).
- Virtually no hope of modeling the context or grounding the language in the world or in action.

# PLOW: webpage structure as context

For the PLOW system, the context is the webpage:

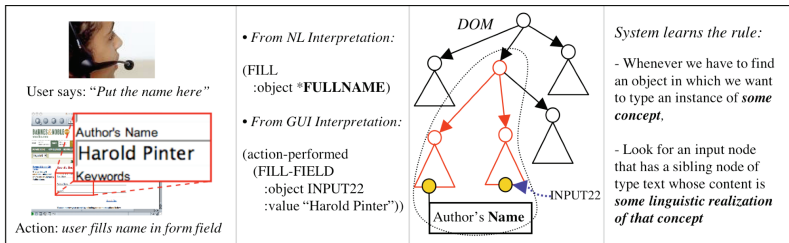


Figure 4: Learning to find and fill a text field

- Project homepage:

<http://www.cs.rochester.edu/research/cisd/projects/plow/>

- Language processing with the TRIPS parser:

<http://www.cs.rochester.edu/research/cisd/projects/trips/parser/cgi/web-parser-xml.cgi>

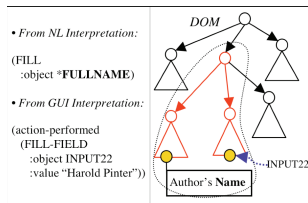


## Learning new rules and generalizations

- Learning rules of the form 'If A, then B, else C' is a challenge because the latent variable A is generally not observed. Rather, one sees only B or C.
- In an interactive, instructional setting, one needn't rely entirely on abduction or probabilistic inference: users generally state the needed rules during their interactions.

# Language-based principles

- 1 The user's actions ground the parsed language.



- 2 The DOM structure grounds the user's indexicals and other referential devices.

- Put the name here. (user clicks on the DOM element)
- This is the ISBN number. (user highlights some text)
- Find another tab. (user has selected a tab)

- 3 Indefinites mark new information; definites refer to established information:

- *A man walked in. He/ The man looked tired.*
- *an address* ⇒ new input parameter
- *the address* ⇒ existing input parameter

## Interaction and error correction

- PLOW is tested with human users in real scenarios. (It has been used by the US Military Health System to set up doctor's appointments.)
- Thus, PLOW tries to immediately apply the rules it infers, so that the user will correct it. This helps with:
  - finding the right level of generalization; and
  - overcoming noise in the context (from poor HTML mark-up)

# Evaluation

16 independent evaluators trained on PLOW and three other systems.

## Phase 1

- 1 The evaluators taught the systems some predefined tasks.
- 2 The system then performed those tasks with different input parameters.

## Phase 2

- 1 The evaluators used the systems to teach some of the tasks at right.
- 2 PLOW received the highest average score of all systems.
- 3 Evaluators had free choice of which system to use. 13 chose PLOW for at least one task, and PLOW was chosen for 30 of the 55 evaluation tasks.

1. What <businesses> are within <distance> of <address>?
2. Get directions for <integer> number of restaurants within <distance> of <address>.
3. Find articles related to <topic> written for <project>.
4. Which <project> had the greatest travel expenses between <start date> and <end date>?
5. What is the most expensive purchase approved between <start date> and <end date>?
6. For what reason did <person> travel for <project> between <start date> and <end date>?
7. Find <ground-transport, parking> information for <airport>.
8. Who should have been notified that <person> was out of the office between <start date> and <end date>?
9. Summarize all travel and purchase costs for <project> between <date> and <date> by expense category
10. Which projects exceeded the current government maximum allowable expense for travel costs?

Figure 1: Previously unseen tasks used in the evaluation

# Markov Decision Processes (MDPs)

- The agent has complete knowledge of the environment and its own current state, but the effects of its actions are non-deterministic.
- MDPs were developed starting in the 1950s by Richard Bellman (1957), Ronald Howard (1960), Karl Åström (1965), Edward Sondik (1971), Richard Sutton (1988), and others. Most of this work concerns efficiently finding the agent's optimal action.
- Howard (1978) describes one of the earliest applications: programming the Sears, Roebuck, and Co.'s giant Addressograph mechanical computer to optimize the process of choosing which customers to send which catalogues (late 1950s): "The optimum policy was confirmed by applying it to [...] a selected set of customers whose purchases were very carefully monitored. When the policy was later implemented on the full customer set, the results closely confirmed the model predictions" (p. 100).

# Defined

## Definition (MDP)

- ①  $S$  is a finite set of states.
- ②  $A$  is a finite set of actions.
- ③  $R : (S \times A) \mapsto \mathbb{R}$  is the reward function.
- ④  $T : (S \times A \times S) \mapsto [0, 1]$  is the state transition function.

## Example

Cab driver Ron serves towns A and B. He has two actions: **cruise** for fares or wait at a cab **stand**.

<b>cruise</b>	A	B
A	0.9	0.1
B	0.1	0.9

(a)  $T$  for cruising around

<b>stand</b>	A	B
A	0.4	0.6
B	0.6	0.4

(b)  $T$  for the cab stand

	A	B
<b>cruise</b>	\$8	\$20
<b>stand</b>	\$5	\$22

(c)  $R$

**Table:** Optimizing Ron's plans based on his data.

# Optimization

## Definition (Bellman operator for MDPs)

Define  $\mathcal{B}_0(s) = 0$  for all  $s \in \mathcal{S}$ . Then for all  $t > 0$ :

$$\mathcal{B}_t(s, a) = R(s, a) + \gamma \sum_{s' \in \mathcal{S}} T(s, a, s') \mathcal{B}_{t-1}(s')$$

where  $0 < \gamma \leq 1$  is a discounting term (a dollar today is worth more than a dollar tomorrow).

VALUEITERATION( $\mathcal{S}, A, R, T, \gamma, \varepsilon$ )

```

1   $V(s) = 0, V'(s) = 0$  for all  $s \in \mathcal{S}$ 
2  while True
3      for  $s \in \mathcal{S}$ 
4          # argmax for policy too:
5           $V'(s) = \max_{a \in A} [R(s, a) + \gamma \sum_{s' \in \mathcal{S}} T(s, a, s') V(s')]$ 
6          if  $|V'(s) - V(s)| < \varepsilon$  for all  $s \in \mathcal{S}$ 
7              return  $V'$ 
8          else  $V = V'$ 

```

# Optimal planning under uncertainty

## Example

Cab driver Ron serves towns A and B. He has two actions: **cruise** for fares or wait at a cab **stand**.

<b>cruise</b>	A	B
A	0.9	0.1
B	0.1	0.9

(a)  $T$  for cruising around

<b>stand</b>	A	B
A	0.4	0.6
B	0.6	0.4

(b)  $T$  for the cab stand

	A	B
<b>cruise</b>	\$8	\$20
<b>stand</b>	\$5	\$22

(c)  $R$

A  $\mapsto$  **stand**      B  $\mapsto$  **cruise**

(d) Optimal policy

**Table:** Optimizing Ron's plans based on his data.



# A simple robot controller (Russell and Norvig 2003:§17)

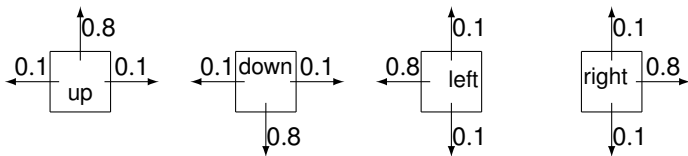


Figure: Action-specific state transitions

→	→	→	+1
↑		↑	-1
↑	←	←	←

(a) Optimal policy when the reward (penalty) for being in a blank square is  $-0.04$ .

→	→	→	+1
↑		↑	-1
↑	→	↑	←

(b) Optimal policy when the reward (penalty) for being in a blank square is  $-0.3$ .

Figure: Optimality for different reward functions.

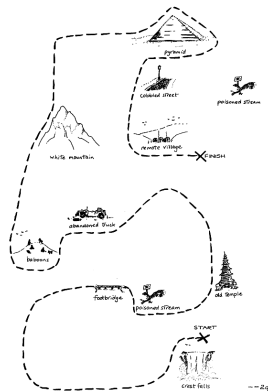
## Vogel and Jurafsky (2010)

- Agents that learn to follow navigational instructions on maps.
- MDP formulation with online reinforcement learning.
- **Inspiring idea:** feature functions  $\phi(s, a)$  and associated learned weights, to process unknown utterances, landmarks, etc.
- **Inspiring idea:** learning probabilistic word meanings from the interaction of language, the world, and the rewards.
- Limitations begin to show us the need for more complex agents.

# The Edinburgh Map Corpus (Thompson et al. 1993)

One participant tells the other how to reproduce a path through a map.

- g right it starts directly above the crest falls if you go  
to the left of your page just to the edge of the crest  
falls
- f mmhmm
- g come south due south to the bottom of the page
- f mmhmm
- g go to the left of the page to about an inch from the  
end
- f over the banana tree
- g i suppose so yeah eh
- f mmhmm
- g go north to the level of the footbridge
- f mmhmm
- g go up and go across the footbridge and stop  
exact— right at the end edge of the footbridge
- f above the footbridge
- g o— over the footbridge
- f mm
- g and stop right at the end of it
- g there is a poisoned stream on mine but which you  
don't have

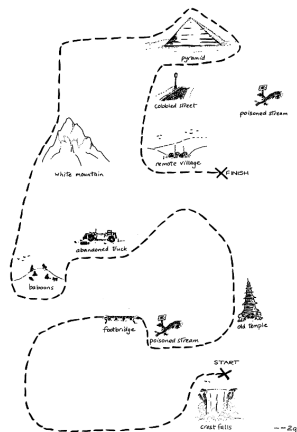


Transcripts, audio, maps, etc.: <http://groups.inf.ed.ac.uk/maptask/>

# MDP formulation and learning

- 1  $S$ : a set of  $s = (u, l, c)$  triples:

  - A set of utterances  $u$
  - A set of landmarks  $l$
  - $c \in \{\text{North, South, East, West}\}$
- 2  $A$ :  $(l, c)$ , meaning pass  $l$  on side  $c$
- 3  $R((u, l, c), (l', c')) = \begin{bmatrix} \mathbb{I}[l = l'] & + \\ \mathbb{I}[c = c'] & + \\ \text{sim}(u, l') & \end{bmatrix}$
- 4  $T(s, a) = s'$
- 5  $\phi(s, a) \in \mathbb{R}^n$  capturing world and linguistic information



# MDP formulation and learning

- 1  $S$ : a set of  $s = (u, l, c)$  triples:
  - A set of utterances  $u$
  - A set of landmarks  $l$
  - $c \in \{\text{North, South, East, West}\}$
- 2  $A$ :  $(l, c)$ , meaning pass  $l$  on side  $c$
- 3  $R((u, l, c), (l', c')) = \begin{bmatrix} \mathbb{I}[l = l'] & + \\ \mathbb{I}[c = c'] & + \\ \text{sim}(u, l') & \end{bmatrix}$
- 4  $T(s, a) = s'$
- 5  $\phi(s, a) \in \mathbb{R}^n$  capturing world and linguistic information

**Input:** Dialog set  $D$   
 Reward function  $R$   
 Feature function  $\phi$   
 Transition function  $T$   
 Learning rate  $\alpha_t$

**Output:** Feature weights  $\theta$

```

1 Initialize  $\theta$  to small random values
2 until  $\theta$  converges do
3   foreach Dialog  $d \in D$  do
4     Initialize  $s_0 = (l_1, u_1, \theta)$ ,
       $a_0 \sim \text{Pr}(a_0 | s_0; \theta)$ 
5     for  $t = 0$ ;  $s_t$  non-terminal;  $t++$  do
6       Act:  $s_{t+1} = T(s_t, a_t)$ 
7       Decide:  $a_{t+1} \sim \text{Pr}(a_{t+1} | s_{t+1}; \theta)$ 
8       Update:
9          $\Delta \leftarrow R(s_t, a_t) + \theta^T \phi(s_{t+1}, a_{t+1})$ 
10         $\quad - \theta^T \phi(s_t, a_t)$ 
11         $\theta \leftarrow \theta + \alpha_t \phi(s_t, a_t) \Delta$ 
12     end
13   end
14 end
15 return  $\theta$ 
  
```

**Algorithm 1:** The SARSA learning algorithm.

# Learned paths

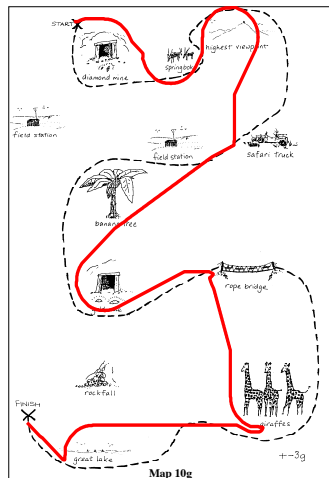
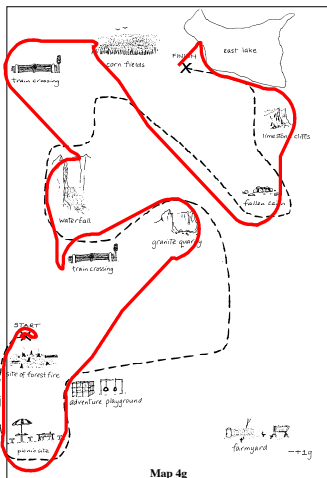


Figure 4: Sample output from the SARSA policy. The dashed black line is the reference path and the solid red line is the path the system follows.

## Learned meanings

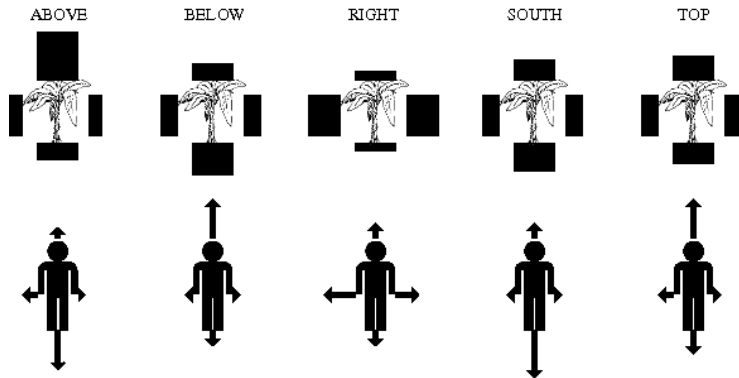


Figure 5: This figure shows the relative weights of spatial features organized by spatial word. The top row shows the weights of allocentric (landmark-centered) features. For example, the top left figure shows that when the word *above* occurs, our policy prefers to go to the north of the target landmark. The bottom row shows the weights of egocentric (absolute) spatial features. The bottom left figure shows that given the word *above*, our policy prefers to move in a southerly cardinal direction.

# The Cards Corpus

<http://CardsCorpus.christopherpotts.net/>

## Included

- The transcripts in CSV format
- Python classes for working with the transcripts
- Examples of the Python classes in action
- R code for reading in the corpus as a data frame
- **All the annotations used in the work described here**

## By the numbers

- 1,266 transcripts
- Game length mean: 373.21 actions (median 305, sd 215.20)
- Card pickup: 19,157
- Card drop: 12,325
- Move: 371,811
- Utterance: 45,805 (260,788 words,  $\approx$ 4,000 word vocab)



# Amazon Mechanical Turk HIT (Human Intelligence Task)

- **Title:** Collaborative Search Game with Chat
- **Description:** Two-player collaborative video game involving dialogue/chat with other Turkers.
- **Payment:** \$1.00, and up to \$0.50 cents for rich, collaborative problem-solving using meaningful dialogue.
- **Restrictions:** US IP addresses; at least 95%. approval rating
- **Timing:** mid-week, 7:00 am – 3:00 pm Pacific time
- **Turker Nation:** posting on Turker Nation about our HIT and its goals, responding to Turkers' questions and concerns, and learning from Turkers' about what life is like for them.

# Amazon Mechanical Turk HIT (Human Intelligence Task)

## Instructions

1. You'll first be presented with an experimental consent form; the transcripts from these games will be used to study collaborative problem solving.
2. Once you click through, you'll go into "Wait for Participant" mode until another player shows up. **It might take a while for another player to show up! We suggest opening a new window and doing other things until someone arrives**
3. When play begins, you'll be presented with a task description: you and your partner are tasked with finding six consecutive cards of the same suit.
4. When you've solved the game, click the "Task Complete" button to register that the task is complete. You can continue playing and chatting at that point, or you can click **Complete HIT and return to MTurk** at the top of the screen.
5. You will receive a bonus of up to \$0.50 for good collaboration with meaningful communication.
6. Notes:
  - Once the game begins, you can move around with the arrow keys or the on-screen buttons.
  - The chat window is near the top of the screen.
  - You can't see your partner, and your partner can't see you!
  - The yellow boxes mark cards in your "line of sight".
7. **Caution: Players who do not complete the task properly will not be paid. It is vital that you read and understand the task description and follow the instructions.**
8. Feel free to play multiple times — especially useful if you get good at the game.

## Annotated game screen

# Amazon Mechanical Turk HIT (Human Intelligence Task)

**TYPE HERE**

**Yellow boxes mark cards in your line of sight.**

**You are on 2D**

**Task description: Six consecutive cards of the same suit**

Received: hi  
Sent: I have the JH  
Received: I have the 8H

Type text here:  
Disable Sound

I'm on 2D, which isn't too useful. There are cards to my right and below, though. I'll check them out.

Gather six consecutive cards of a particular suit (decide which suit together). Each of you can hold only three cards at a time, so you'll have to coordinate your efforts. You can talk all you

P1 turns remaining: 546  
P2 turns remaining: 599

Indicate Task Complete

up  
Click a card to pick it up:  
2D

left  
Click a card to drop it from your hand:  
JH  
right

down

**The cards you are holding**

**Move with the arrow keys or these buttons.**

# Amazon Mechanical Turk HIT (Human Intelligence Task)

*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.*

# Amazon Mechanical Turk HIT (Human Intelligence Task)

*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?



Which suit should we pursue?



Which sequence should we pursue?



Where is card X?

# Transcripts: environment metadata

Agent	Time	Action type	Contents
Server	0	COLLECTION_SITE	Amazon Mechanical Turk
Server	0	TASK_COMPLETED	2010-06-17 10:10:53 EDT
Server	0	PLAYER_1	A00048
Server	0	PLAYER_2	A00069
Server	2	P1_MAX_LINEOFSIGHT	3
Server	2	P2_MAX_LINEOFSIGHT	3
Server	2	P1_MAX_CARDS	3
Server	2	P2_MAX_CARDS	3
Server	2	P1_MAX_TURNS	200
Server	2	P2_MAX_TURNS	200
Server	2	GOAL_DESCRIPTION	Gather six consecutive cards ...
Server	2	CREATE_ENVIRONMENT	[ASCII representation]
Player 1	2092	PLAYER_INITIAL_LOCATION	16,15
Player 2	2732	PLAYER_INITIAL_LOCATION	9,10

# Transcripts: environment metadata

```

-----;
-                                     -;
- ----- - - - - - - - - - - -;
- - - - - - - - - - - - - - - -;
- --- ----- - - - - - - - - -;
-      - - - - - - - - - - -;
- b - - - - - - - - - - -;
- - - - - - - - - - - - - - -;
- - - - - - - - - - - - - - -;
-      - - - - - - - - - - -;
- - - - - - - - - - - - - - -;
- - - - - - - - - - - - - - -;
- - - - - - - - - - - - - - -;
- - - - - - - - - - - - - - -;
-      - - - - - - - - - - -;
- - - - - - - - - - - - - - -;
-      b-----;
-----;

```

## NEW\_SECTION

```

1,2:2D;1,7:KH;1,7:9S;1,11:6C;1,13:QC;1,14:QS;
2,18:3H;2,18:9H;
3,19:4H;4,8:AC;4,19:3D;
4,19:KD;
5,14:QH;5,15:5S;5,15:2S;5,16:4D;5,16:10C;5,18:4S;
6,11:KC;6,15:9C;
7,11:2H;7,13:7S;
8,2:QD;8,4:AD;8,11:JC;8,20:8S;
9,9:10S;9,9:6H;9,9:8C;9,10:7H;9,14:JS;
10,1:2C;10,10:8D;11,14:6D;11,14:10H;
11,18:4C;11,18:9D;
12,10:3S;12,12:6S;12,16:5H;12,16:JD;12,20:3C;
13,4:5C;13,4:JH;13,15:KS;
14,2:5D;14,20:10D;15,2:AH;
15,13:7D;15,15:8H;15,17:AS;15,20:7C;

```

# Transcripts: game play

Agent	Time	Action type	Contents
Player 1	566650	PLAYER_MOVE	7,11
Player 2	567771	CHAT_MESSAGE_PREFIX	which c's do you have again?
Player 1	576500	CHAT_MESSAGE_PREFIX	i have a 5c and an 8c
Player 2	577907	CHAT_MESSAGE_PREFIX	i jsut found a 4 of clubs
Player 1	581474	PLAYER_PICKUP_CARD	7,11:8C
Player 1	586098	PLAYER_MOVE	7,10



## Novice strategy

Player 1: Hello. Are you here?  
 Player 2: yes  
 Player 2: do you see any cards  
 Player 1: Yes. I see a yellow spot. Those are our cards. We'll only be able to see the ones that are in our view until we move with our arrows.  
 Player 2: i see 3 of them  
 Player 1: We only have a certain number of moves, so we should decide how we're going to do this before we use them, do you think?  
 Player 2: sure  
 Player 1: Ok. So, we have to pick up six cards of the same suit, in a row...  
 Player 1: each of us can hold three, so...  
 Player 1: I think I should get my three, then you should get your three or vice versa  
 Player 2: ok  
 Player 2: you go ahead  
 Player 1: What suit should we do?  
 Player 1: And which six cards do you want to try for?  
 Player 2: whatever you want  
 Player 1: I'm Courtney, by the way- nice to meet you.  
 Player 2: i'm becky....nice to meet you too  
 Player 1: Hi Becky. How about we go for hearts? And take 234567  
 [...]

# Journeyman strategy

These players have explored and are now forming a strategy:

Player 1 I have 9 clubs and K clubs

Player 1 want to look for clubs?

Player 2 ok

[...]

The players then find various clubs, checking with each other frequently, until they gain an implicit understanding of which specific sequences to try for (either 8C-KC or 9C-AC):

Player 1 so you are holding Jc and Kc now?

Player 2 i now have 10d JC and KC

Player 2 yes

Player 1 drop 10d and look for either 8c or Ace of clubs

# Expert strategy

Player 2: hi

Player 1: hi--which side r u on?

Player 2: right side

Player 2: u?

Player 1: left/middle

Player 1: ok i gathered everything in my area

Player 2: i think i have all of them also

Player 1: how bout 5C - 10C?

Player 2: ok

Player 1: i have 5C, 8C, 9C, and you should have 6C, 7C, 10C

Player 2: got them

## Asymmetric play

Player 1: very limited number of moves but infinite line-of-sight; Player 2: large number of moves but very limited line of sight.

Player 1: Hi  
 Player 2: hi where are you  
 Player 1: near the upper right  
 Player 2: ok any cards that way  
 Player 1: lots of cards near me to the upper right corner  
 Player 2: did you get that  
 Player 1: get wjat ?  
 Player 2: the drop in the top right  
 Player 1: I have not gone there yet  
 Player 2: ok I'll wait  
 Player 2: we have the 4 8 j h  
 Player 2: 3 k c  
 Player 1: ok  
 Player 1: the cards are pretty scattered  
 Player 1: did you check the entire right column?

⋮

# Language in context

Each transcript is a data structure that is intuitively a list of temporally-ordered states

$$\left( \text{context, event} \right)$$

The context includes

- local information (the state of play at that point)
- historical information (the events up to that point)
- global information (limitations of the game, the task, etc.)

When the event is an utterance, we can interpret it *in context*.

This is what pragmatics is all about, but it is very rare to have a dataset that truly lets you do it.

## 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 <b>in context</b>

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.

## Papers using the Cards corpus

- Djalali et al. (2012): anapora and domain restriction
- Djalali et al. (2011): presuppositions
- Potts (2012): goal-orientation of underspecified locative expressions
- Vogel et al. (2013a): emergent Gricean behavior with Dec-POMDPs
- Vogel et al. (2013b): conversational implicature with Dec-POMDPs

# POMDPs and approximate Dec-POMDPs

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 locative advice from the other player.
- Give locative advice to the other player.

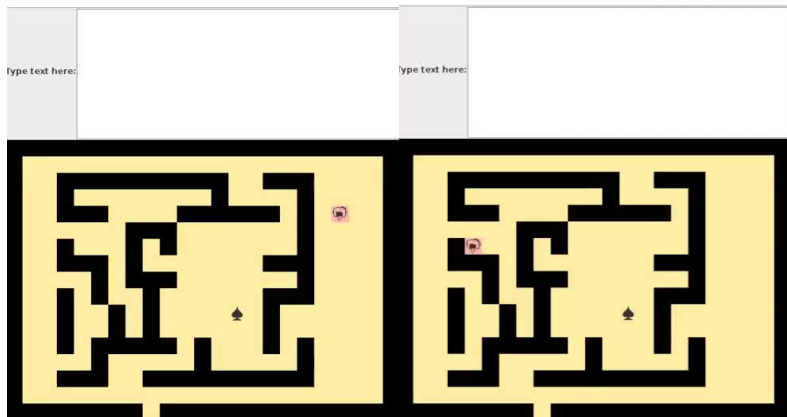
Modeling the problem as a POMDP allows us to train agents that have these properties.



# Simplified cards scenario

Both players must find the ace of spades.

DialogBot:

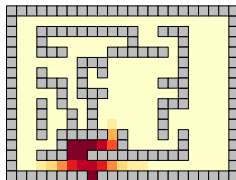


# Grounded language interpretation

“in the bottom you see the opening on the bottom row”



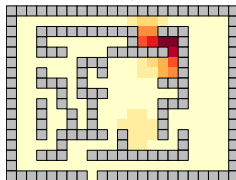
BOARD(entrance & bottom);  $H: 5.48$



“in the top right of the middle part of the board”



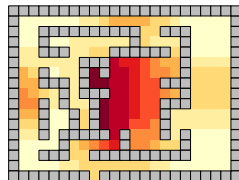
middle(top & right);  $H: 5.27$



“i'm in the center”



BOARD(middle);  $H: 7.37$



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

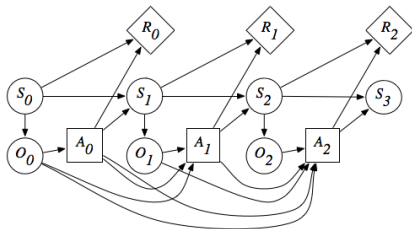
# POMDPs

The agent has only probabilistic information about its current state (and the effects of its actions are non-deterministic, as in MDPs).

## Definition (POMDP)

A POMDP is a structure  $(S, A, R, T, \Omega, O)$ :

- $(S, A, R, T)$  is an MDP.
- $\Omega$  is a finite set of observations.
- $O : (A \times S \times \Omega) \mapsto [0, 1]$  is the observation function.



## ListenerBot (a POMDP agent)

- $S$ : all combinations of the player's region and the card's region
- $b_0$ : initial belief state (distribution over  $S$ )
- $A$ : *travel* actions for each region, and a single *search* action
- $\Omega$ : {AS seen, AS not seen}
- $\Sigma$ : a set of messages, treated as observations; each message  $\sigma$  denotes a distribution  $P(s | \sigma)$  over states  $s$ . We apply Bayes rule to incorporate these into the POMDP observations.
- $T$ : distributions  $P(s' | s, a)$ , except *travel* actions fail between nonadjacent regions
- $O$ : distributions  $P(o | s, a)$ ; *travel* actions never return positive observations; *search* actions return positive observations only if the player's current region contains the AS
- $R$ : small negative for not being on the card, large positive for being on it. No sensitivity to the other player.

# Optimization

A belief state for  $(S, A, R, T, \Omega, O)$  is a probability distribution  $b$  over  $S$ .

$$P(s, a, o, b) = O(s, a, o) \sum_{s' \in S} T(s', a, s) b(s') \quad (1)$$

$$b_o^a(s) = \frac{P(s, a, o, b)}{\sum_{s' \in S} P(s', a, o, b)} \quad (2)$$

## Definition (Bellman operator for POMDPs)

Let  $b$  be a belief state for  $(S, A, R, T, \Omega, O)$ . Set  $\mathcal{P}_0(b') = 0$  for all belief states  $b'$ . Then for all  $t > 0$ :

$$\mathcal{P}_t(b, a) = \left( \sum_{s \in S} b(s) R(s, a) \right) + \gamma \sum_{o \in \Omega} \left( \sum_{s \in S} P(s, a, o, b) \right) \mathcal{P}_{t-1}(b_o^a)$$

where  $0 < \gamma \leq 1$  is a discounting term.

## Approximate solutions take us (only) part of the way

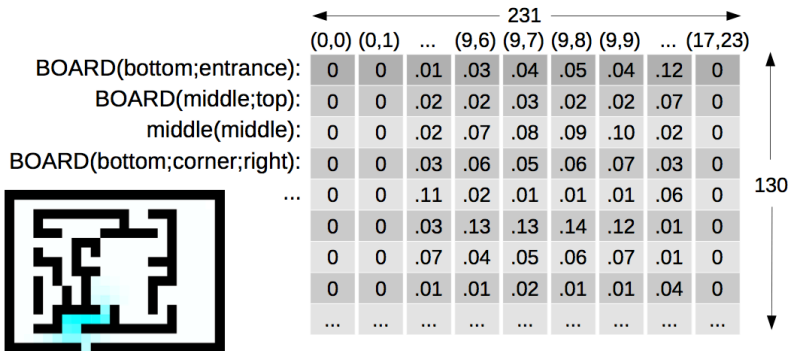
- An exact solution specifies the value of every action at any reachable belief state.
- In practice, only approximate solutions are tractable. We used the PERSEUS solution algorithm.
- Even approximate solutions are generally only possible for problems with  $< 10K$  states.

Card location	Agent location	Partner location	Partner's card beliefs	
231	×	231	×	231
		≈ 50K		≈ 12M
				≈ 3B

**Table:** Size of the state-space for the one-card game.

# Language as a representation for planning

- Divide the board up into  $n$  regions, for some tractable  $n$
- Generate this partition using our locative phrase distributions.
- $k$ -means clustering in locative phrase space.



# Clusters induced

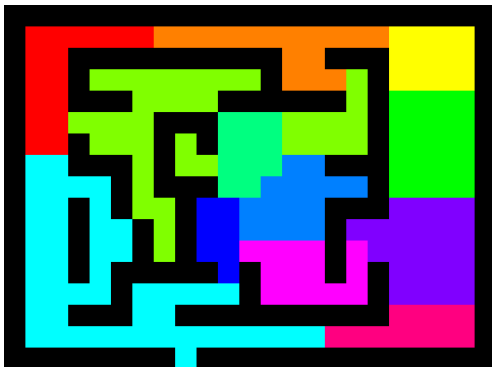


Figure: 12-cell clustering.



# Clusters induced

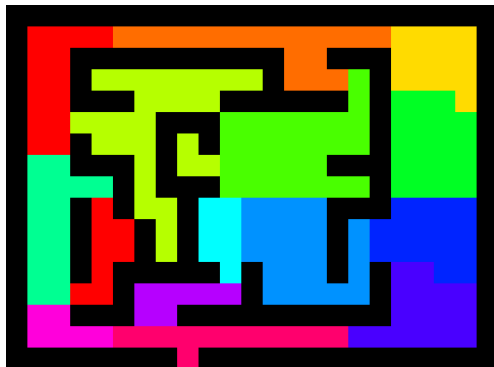


Figure: 14-cell clustering.

# Clusters induced

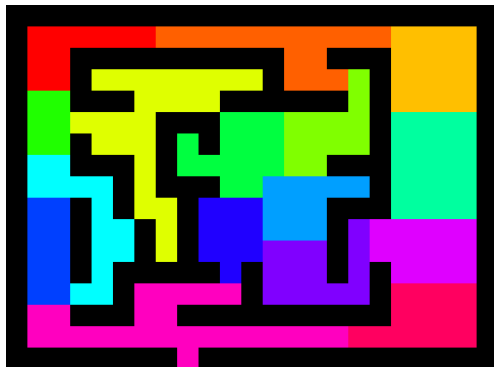


Figure: 16-cell clustering.

# Clusters induced

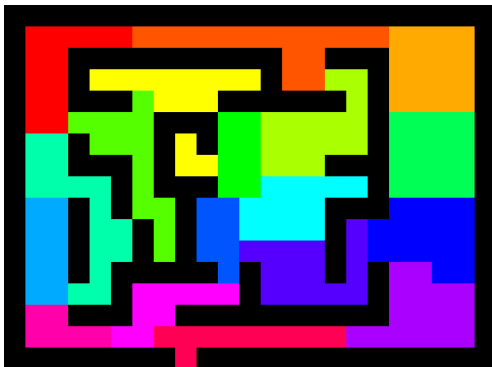
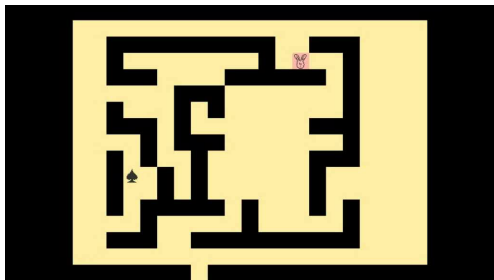


Figure: 18-cell clustering.

# ListenerBot example



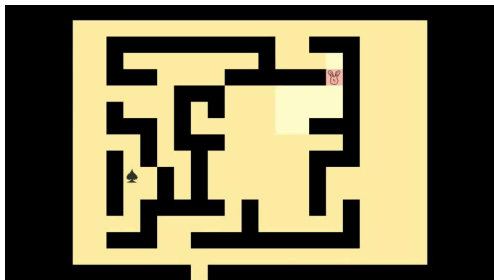
ListenerBot:



# ListenerBot example



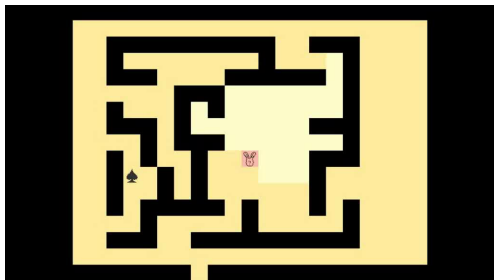
ListenerBot:



# ListenerBot example



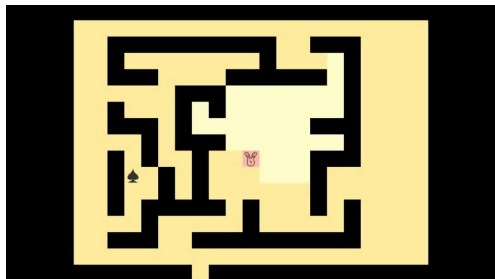
ListenerBot:



# ListenerBot example



ListenerBot:



“it’s on the left side”



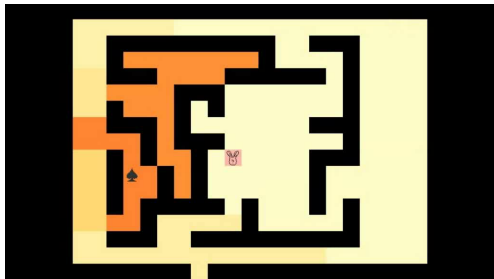
**BOARD(left)**



# ListenerBot example



ListenerBot:



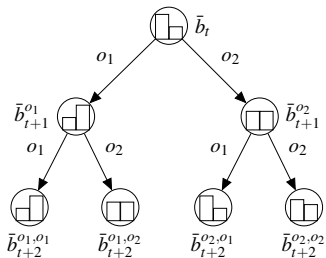


## DialogBot (an approximate Dec-POMDP)

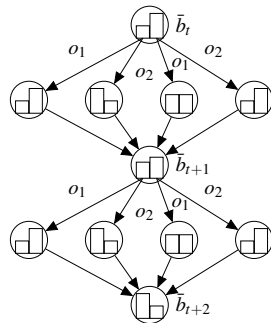
DialogBot is a strict extension of ListenerBot:

- 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 dialog actions.
- (The player assumes that) a dialog action  $U$  alters the other player's beliefs in the same way that  $U$  would impact his own beliefs.
- Same basic reward structure as for Listenerbot, except now also sensitive to whether the other player has found the card.

# Belief-state approximation

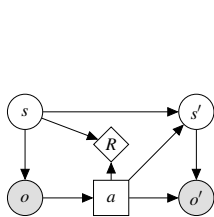


(a) Exact multi-agent belief tracking

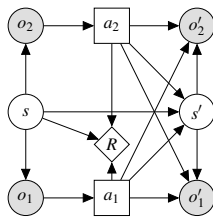


(b) Approximate multi-agent belief tracking

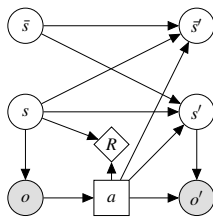
## How the agents relate to each other



(a) ListenerBot POMDP



(b) Full Dec-POMDP

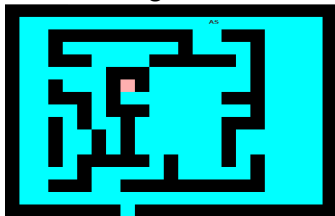


(c) DialogBot POMDP

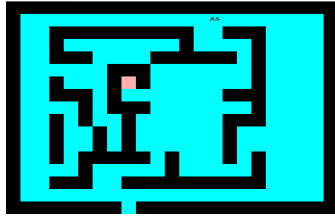
**Figure:** In the full Dec-POMDP (b), both agents receive individual observations and choose actions independently. Optimal decision making requires tracking all possible histories of beliefs of the other agent. DialogBot approximates the full Dec-POMDP as single-agent POMDP. At each time step, DialogBot marginalizes out the possible observations  $\bar{o}$  that ListenerBot received, yielding an *expected belief state*  $\bar{s}$ .

# DialogBot and ListenerBot play together

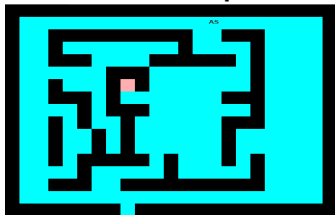
**DialogBot beliefs**



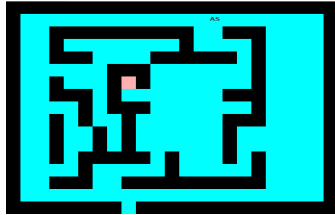
**ListenerBot beliefs**



**DialogBot beliefs:  
ListenerBot's position**

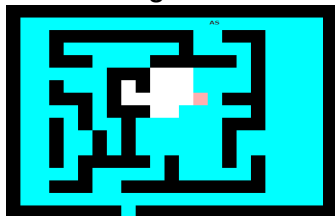


**DialogBot beliefs:  
ListenerBot's beliefs**



# DialogBot and ListenerBot play together

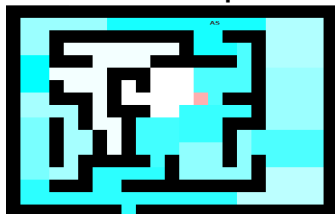
**DialogBot beliefs**



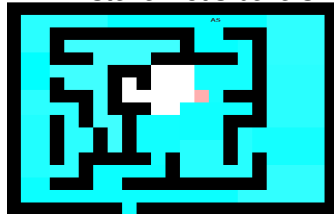
**ListenerBot beliefs**



**DialogBot beliefs:  
ListenerBot's position**



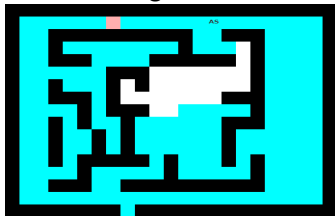
**DialogBot beliefs:  
ListenerBot's beliefs**



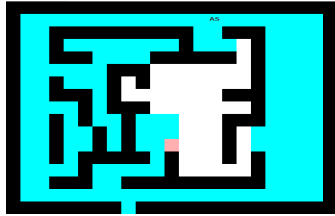
# DialogBot and ListenerBot play together

Dialogbot: "Top"

**DialogBot beliefs**



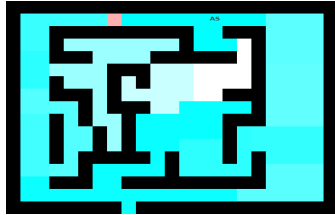
**ListenerBot beliefs**



**DialogBot beliefs:  
ListenerBot's position**



**DialogBot beliefs:  
ListenerBot's beliefs**



# DialogBot and ListenerBot play together

Dialogbot: "Top"

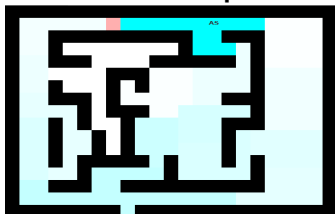
DialogBot beliefs



ListenerBot beliefs



DialogBot beliefs:  
ListenerBot's position

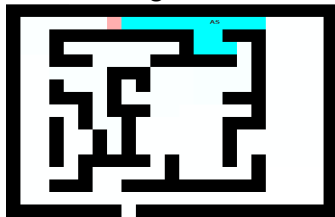


DialogBot beliefs:  
ListenerBot's beliefs

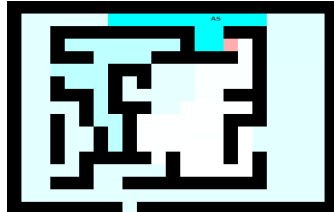


## DialogBot and ListenerBot play together

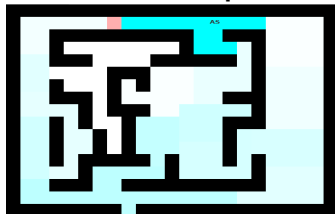
**DialogBot beliefs**



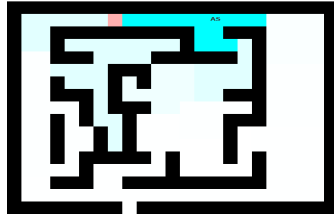
**ListenerBot beliefs**



**DialogBot beliefs:  
ListenerBot's position**



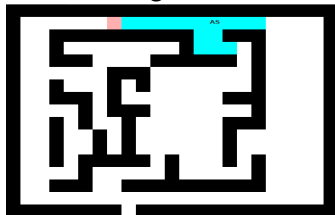
**DialogBot beliefs:  
ListenerBot's beliefs**





## DialogBot and ListenerBot play together

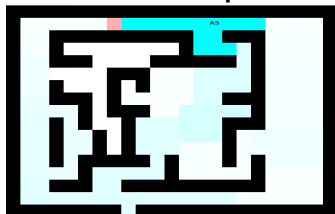
**DialogBot beliefs**



**ListenerBot beliefs**



**DialogBot beliefs:  
ListenerBot's position**

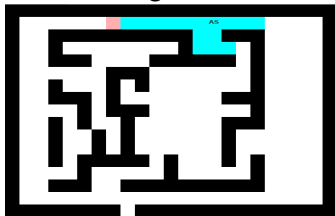


**DialogBot beliefs:  
ListenerBot's beliefs**

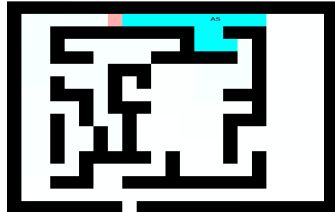


## DialogBot and ListenerBot play together

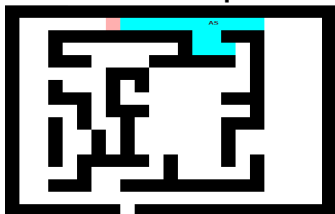
**DialogBot beliefs**



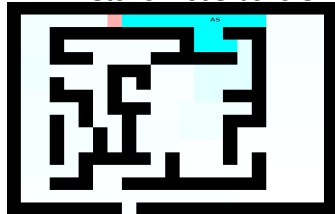
**ListenerBot beliefs**



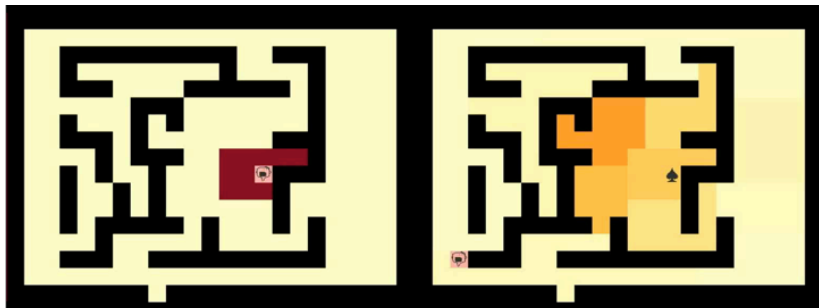
**DialogBot beliefs:  
ListenerBot's position**



**DialogBot beliefs:  
ListenerBot's beliefs**

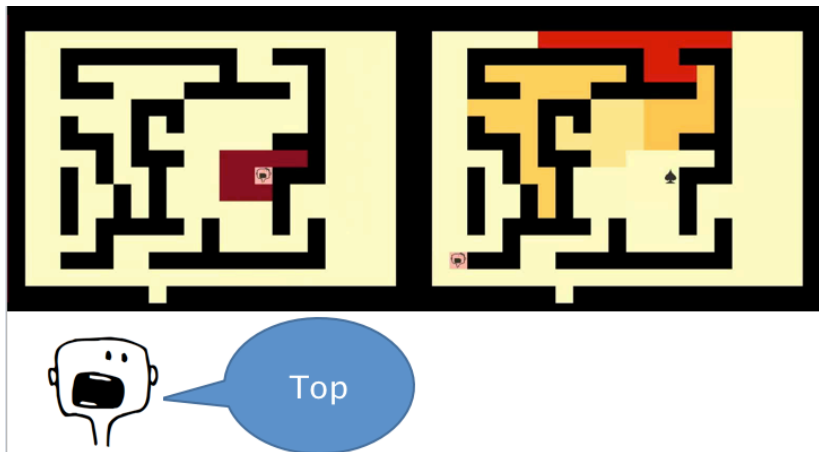


## Grown-up DialogBots (a week of policy exploration)



Middle  
of the  
board

## 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:** The evaluation for each combination of agents. 500 random initial states per agent combination. It pays to model other minds!

# 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.

# References I

- Allen, James F. 1991. *Reasoning About Plans*. San Francisco: Morgan Kaufmann.
- Allen, James F.; Nathanael Chambers; George Ferguson; Lucian Galescu; Hyuckchul Jung; Mary Swift; and William Taysom. 2007. PLOW: A collaborative task learning agent. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*, 1514–1519. Vancouver, British Columbia, Canada: AAAI Press.
- Åström, Karl J. 1965. Optimal control of Markov decision processes with incomplete state estimation. *Journal of Mathematical Analysis and Applications* 10(1):174–205.
- Bangalore, Srinivas; Giuseppe Di Fabbrizio; and Amanda Stent. 2006. Learning the structure of task-driven human-human dialogs. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, 201–208. Sydney, Australia: Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P/P06/P06-1026>.
- Beaver, David I. 2002. Pragmatics, and that's an order. In David Barker-Plummer; David I. Beaver; Johan van Benthem; and Patrick Scotti di Luzio, eds., *Logic, Language, and Visual Information*, 192–215. Stanford, CA: CSLI.
- Bellman, Richard. 1957. A Markovian decision process. *Journal of Mathematics and Mechanics* 6(5):679–684.
- Benz, Anton; Gerhard Jäger; and Robert van Rooij, eds. 2005. *Game Theory and Pragmatics*. Basingstoke, Hampshire: Palgrave MacMillan.
- Blutner, Reinhard. 1998. Lexical pragmatics. *Journal of Semantics* 15(2):115–162.
- Büring, Daniel. 1999. Topic. In Peter Bosch and Rob van der Sandt, eds., *Focus — Linguistic, Cognitive, and Computational Perspectives*, 142–165. Cambridge: Cambridge University Press.
- Calhoun, Sasha; Jean Carletta; Jason Brenier; Neil Mayo; Daniel Jurafsky; Mark Steedman; and David I. Beaver. 2010. The NXT-format Switchboard corpus: A rich resource for investigating the syntax, semantics, pragmatics and prosody of dialogue. *Language Resources and Evaluation* 44(4):387–419.
- Champollion, Lucas. 2008. The influence of goals on ambiguities in certain donkey sentences. Presentation at the 4th Formal Semantics in Moscow Workshop.
- Clark, Herbert H. 1996. *Using Language*. Cambridge: Cambridge University Press.
- Djalali, Alex; David Clausen; Sven Lauer; Karl Schultz; and Christopher Potts. 2011. Modeling expert effects and common ground using Questions Under Discussion. In *Proceedings of the AAAI Workshop on Building Representations of Common Ground with Intelligent Agents*. Washington, DC: Association for the Advancement of Artificial Intelligence.
- Djalali, Alex; Sven Lauer; and Christopher Potts. 2012. Corpus evidence for preference-driven interpretation. In Maria Aloni; Vadim Kimmelman; Floris Roelofsen; Galit Weidman Sassoon; Katrin Schulz; and Matthijs Westera, eds., *Proceedings of the 18th Amsterdam Colloquium: Revised Selected Papers*, 150–159. Berlin: Springer.
- Franke, Michael. 2009. *Signal to Act: Game Theory in Pragmatics*. ILLC Dissertation Series. Institute for Logic, Language and Computation, University of Amsterdam.
- Ginzburg, Jonathan. 1996a. Dynamics and the semantics of dialogue. In Jerry Seligman, ed., *Language, Logic, and Computation*, volume 1, 221–237. Stanford, CA: CSLI.
- Ginzburg, Jonathan. 1996b. Interrogatives: Questions, facts, and dialogue. In Shalom Lappin, ed., *The Handbook of Contemporary Semantic Theory*, 385–422. Oxford: Blackwell Publishers.
- Graff, Delia. 2000. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics* 28(1):45–81.

# References II

- Groenendijk, Jeroen and Martin Stokhof. 1984. *Studies in the Semantics of Questions and the Pragmatics of Answers*. Ph.D. thesis, University of Amsterdam.
- Hobbs, Jerry R.; Mark Stickel; Douglas Appelt; and Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence* 63(1-2):69-142.
- Howard, Ronald Arthur. 1960. *Dynamic Programming and Markov Processes*. Cambridge, MA: Wiley.
- Howard, Ronald Arthur. 1978. Comments on the origin and application of Markov Decision Processes. In Martin L. Puterman, ed., *Dynamic Programming and its Applications*, 201-205. Orlando, FL: Academic Press. Republished with a new addendum, 2002, *Operations Research* 50(1):100-102.
- Kumar Rangarajan Sridhar, Vivek; Srinivas Bangalore; and Shrikanth Narayanan. 2009. Combining lexical, syntactic and prosodic cues for improved online dialog act tagging. *Computer Speech and Language* 23(4):407-422. doi:\bibinfo{doi}[10.1016/j.csl.2008.12.001]. URL <http://www.sciencedirect.com/science/article/pii/S0885230808000569>.
- Lewis, David. 1969. *Convention*. Cambridge, MA: Harvard University Press. Reprinted 2002 by Blackwell.
- Lewis, David. 1988. Relevant implication. *Theoria* 54(3):161-174.
- Malamud, Sophia. 2006. (Non)-maximality and distributivity: A decision theory approach. Paper presented at SALT 16, Tokyo, Japan.
- McCarthy, John. 1980. Circumscription — a form of non-monotonic reasoning. *Artificial Intelligence* 13(1):27-39.
- Merin, Arthur. 1997. If all our arguments had to be conclusive, there would be few of them. Arbeitspapiere SFB 340 101, University of Stuttgart, Stuttgart.
- Parikh, Prashant. 2001. *The Use of Language*. Stanford, CA: CSLI.
- Perrault, C. Raymond and James F. Allen. 1980. A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics* 6(3-4):167-182.
- Potts, Christopher. 2012. Goal-driven answers in the Cards dialogue corpus. In Nathan Arnett and Ryan Bennett, eds., *Proceedings of the 30th West Coast Conference on Formal Linguistics*, 1-20. Somerville, MA: Cascadilla Press.
- Roberts, Craige. 1996. Information structure: Towards an integrated formal theory of pragmatics. In Jae Hak Yoon and Andreas Kathol, eds., *OSU Working Papers in Linguistics*, volume 49: Papers in Semantics, 91-136. Columbus, OH: The Ohio State University Department of Linguistics. Revised 1998.
- Rooth, Mats. 1985. *Association with Focus*. Ph.D. thesis, UMass Amherst.
- van Rooij, Robert. 2003. Questioning to resolve decision problems. *Linguistics and Philosophy* 26(6):727-763.
- Russell, Stuart and Peter Norvig. 2003. *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ: Prentice-Hall, 2nd edition.
- Shriberg, Elizabeth; Rebecca Bates; Paul Taylor; Andreas Stolcke; Daniel Jurafsky; Klaus Ries; Noah Coccaro; Rachel Martin; Marie Meteer; and Carol Van Ess-Dykema. 1998. Can prosody aid the automatic classification of dialog acts in conversational speech? *Language and Speech* 41(3-4):439-487.
- Sondik, Edward J. 1971. *The Optimal Control of Partially Observable Markov Processes*. Ph.D. thesis, Stanford University, Stanford, CA.
- Stolcke, Andreas; Klaus Ries; Noah Coccaro; Elizabeth Shriberg; Rebecca Bates; Daniel Jurafsky; Paul Taylor; Rachel Martin; Marie Meteer; and Carol Van Ess-Dykema. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational Linguistics* 26(3):339-371.



# References III

- Stone, Matthew; Richmond Thomason; and David DeVault. 2007. Enlightened update: A computational architecture for presupposition and other pragmatic phenomena. To appear in Donna K. Byron; Craige Roberts; and Scott Schwenter, *Presupposition Accommodation*.
- Sutton, Richard S. 1988. Learning to predict by the methods of temporal differences. *Machine Learning* 3(1):9–44.
- Thompson, Henry S.; Anne Anderson; Ellen Gurman Bard; Gwyneth Doherty-Sneddon; Alison Newlands; and Cathy Sotillo. 1993. The HCRC map task corpus: Natural dialogue for speech recognition. In *HLT '93: Proceedings of the workshop on Human Language Technology*, 25–30. Morristown, NJ: ACL.
- Vogel, Adam; Max Bodoia; Christopher Potts; and 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 and Daniel Jurafsky. 2010. Learning to follow navigational directions. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, 806–814. Uppsala, Sweden: Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P10-1083>.
- Vogel, Adam; Christopher Potts; and 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.
- Weizenbaum, Joseph. 1966. ELIZA — a computer program for the study of natural language communication between man and machine. *Communications of the ACM* 9(1):36–45.
- Winograd, Terry. 1972. Understanding natural language. *Cognitive Psychology* 3(1):1–191.
- Winograd, Terry. 1986. A procedural model of language understanding. In Barbara J. Grosz; Karen Sparck-Jones; and Bonnie Lynn Webber, eds., *Readings in Natural Language Processing*, 249–266. San Francisco: Morgan Kaufmann Publishers Inc.