Planning, Inference, and Pragmatics in Sequential Language Games

Fereshte Khani Noah Goodman Percy Liang

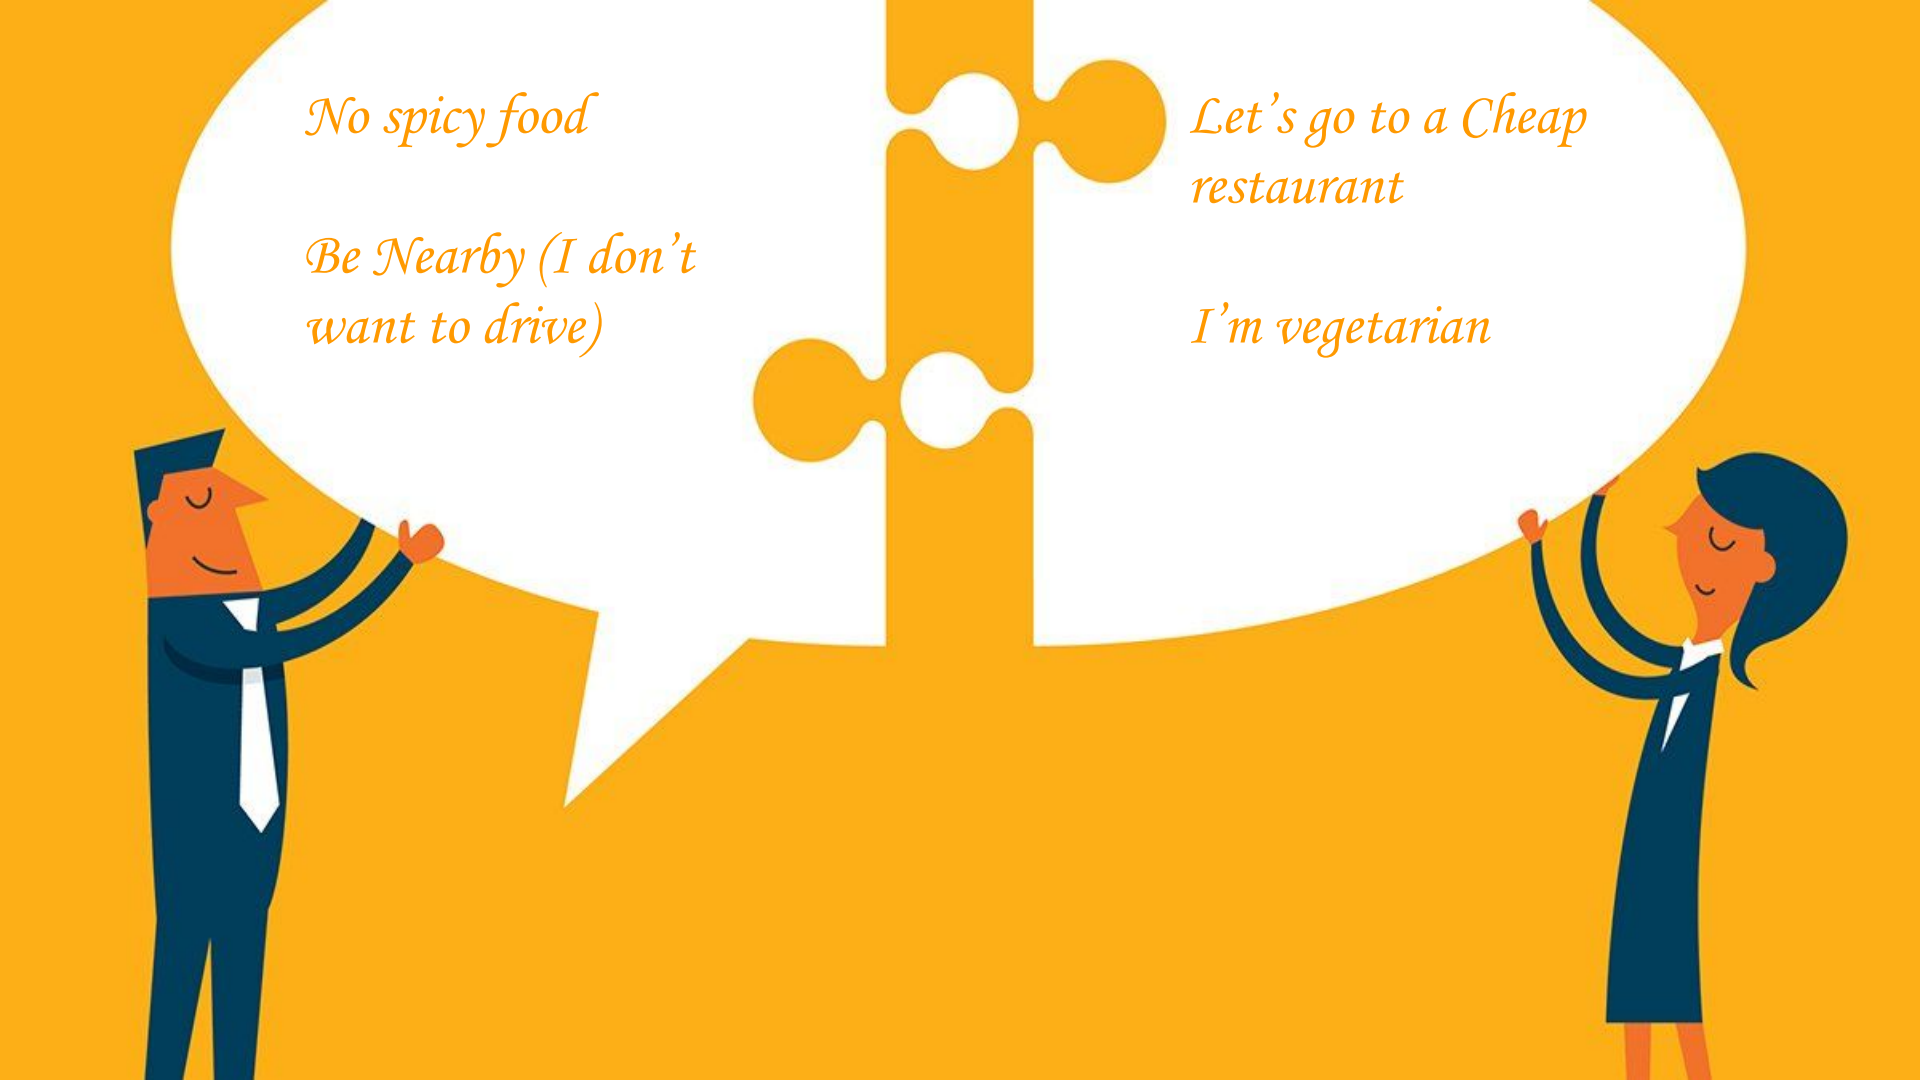
No spicy food
Be Nearby (I don't want to drive)
Let's go to a Cheap restaurant
I'm vegetarian

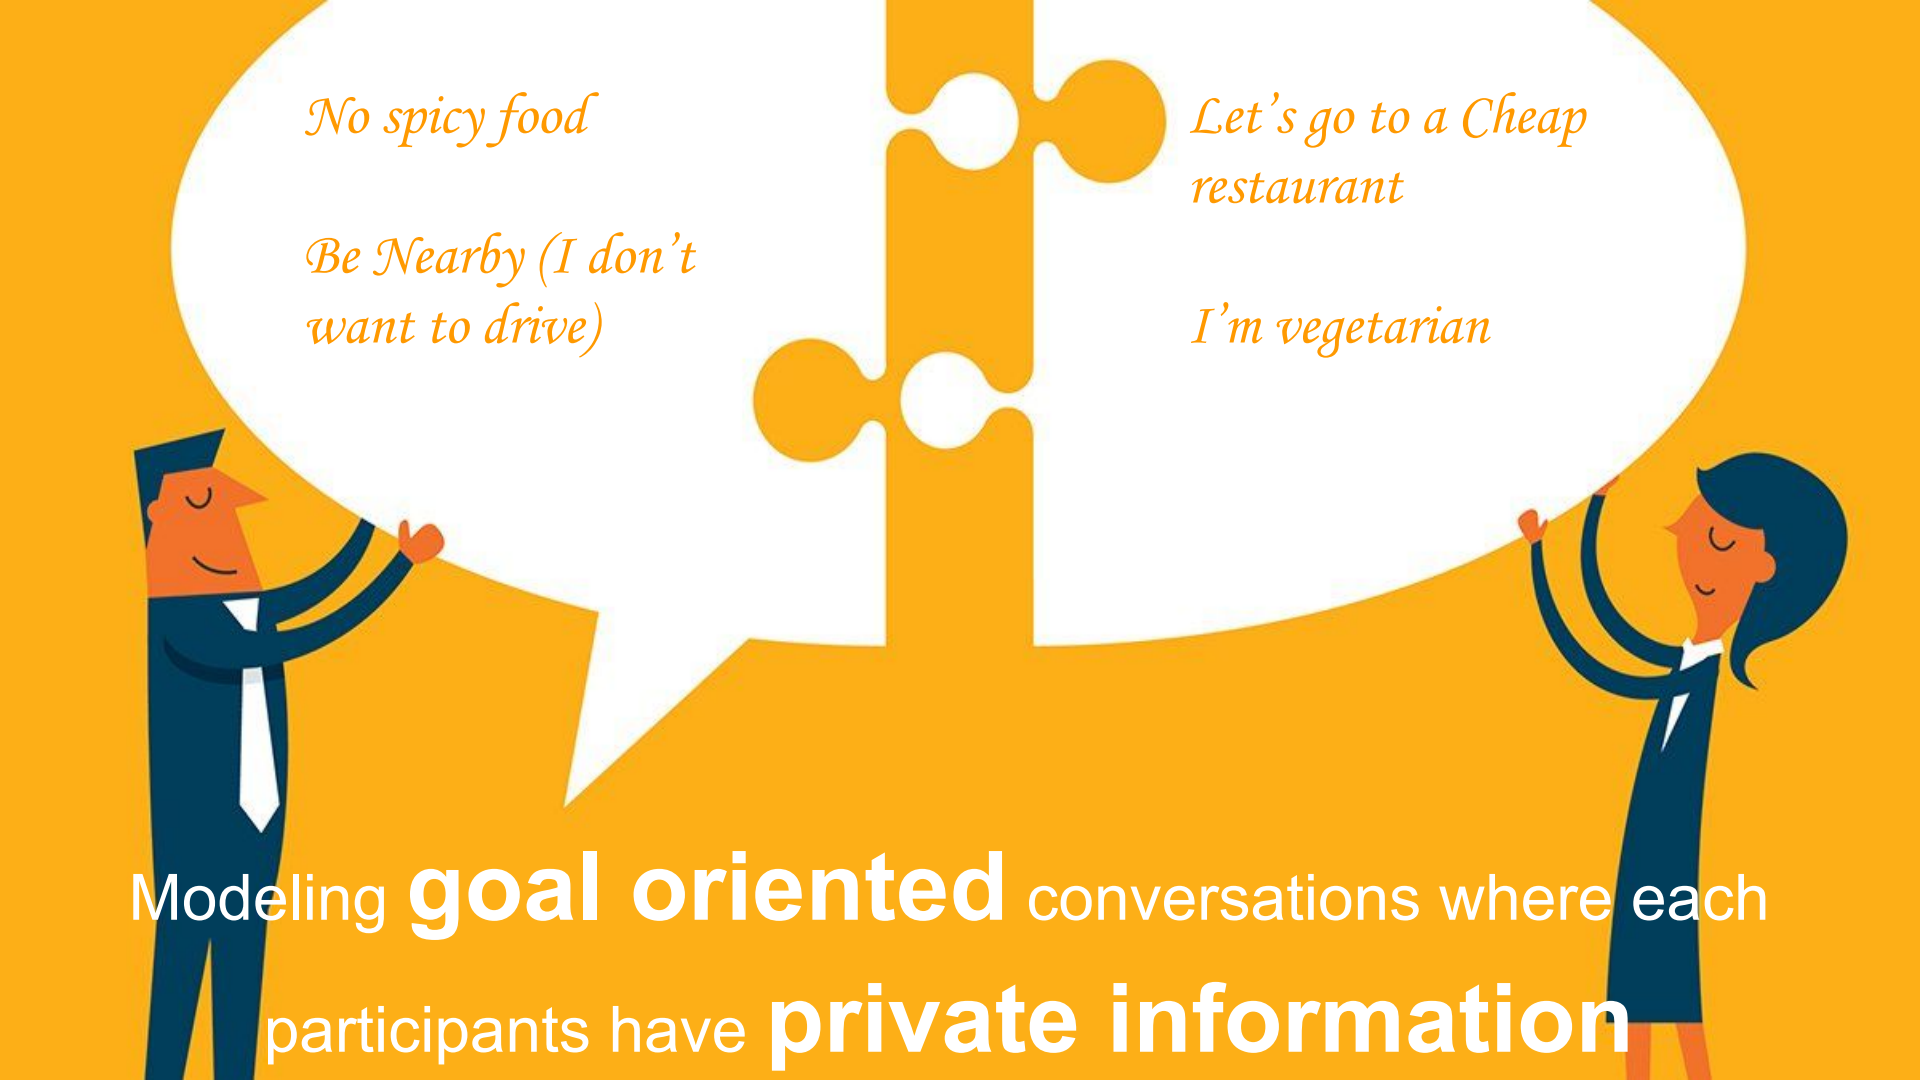

Modeling **goal oriented** conversations where each participants have **private information**

Inference

They figure out the state of the world based on what others say.

Planning

People routinely choose what to say based on their goals.

Pragmatics

They always take into account that others are strategizing agents too.

Inference

They figure out the state of the world based on what others say.

Planning

People routinely choose what to say based on their goals.

Pragmatics

They always take into account that others are strategizing agents too.

Inference

They figure out the state of the world based on what others say.

Planning

People routinely choose what to say based on their goals.

Pragmatics

They always take into account that others are strategizing agents too.

How to have all three in one model?

Data Collection

Sequential Language Game: InfoJigsaw

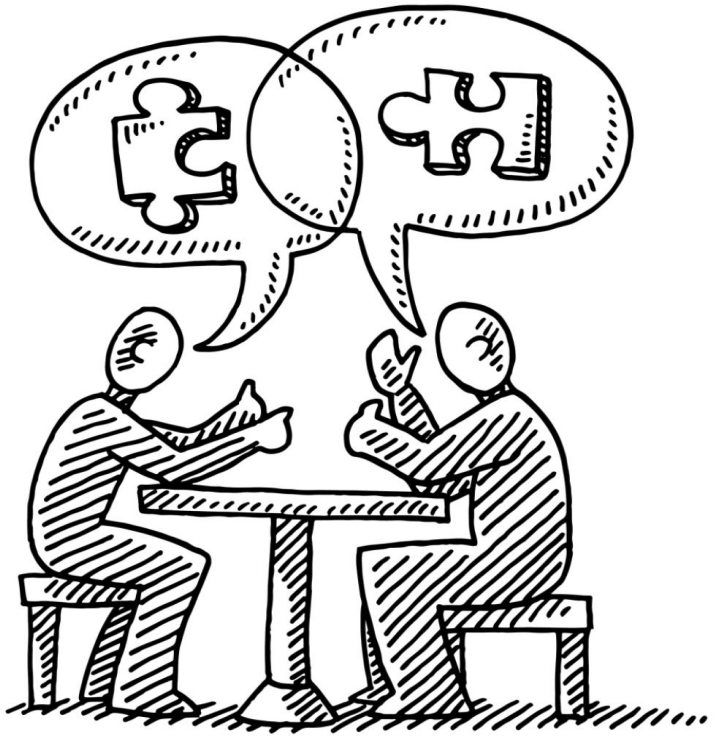

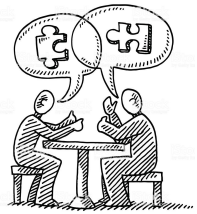

InfoJigsaw

Find A1

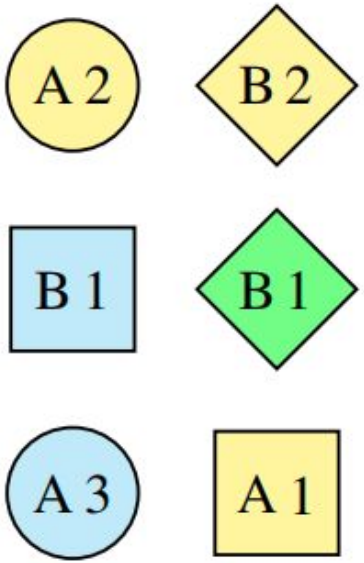

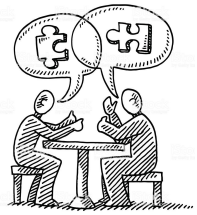

InfoJigsaw

Find A1

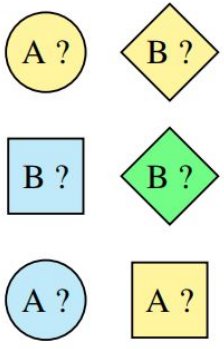

P_{letter} view

Find A1

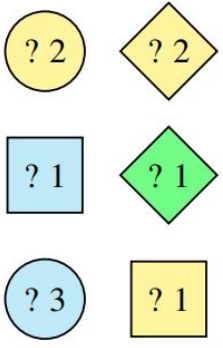

P_{digit} view

The goal is A 1

Round: 1/10

Your partner turn

Partner: square

You: right

Send

Correct Clicks: 0

Wrong Clicks: 0

Word Usage: 2

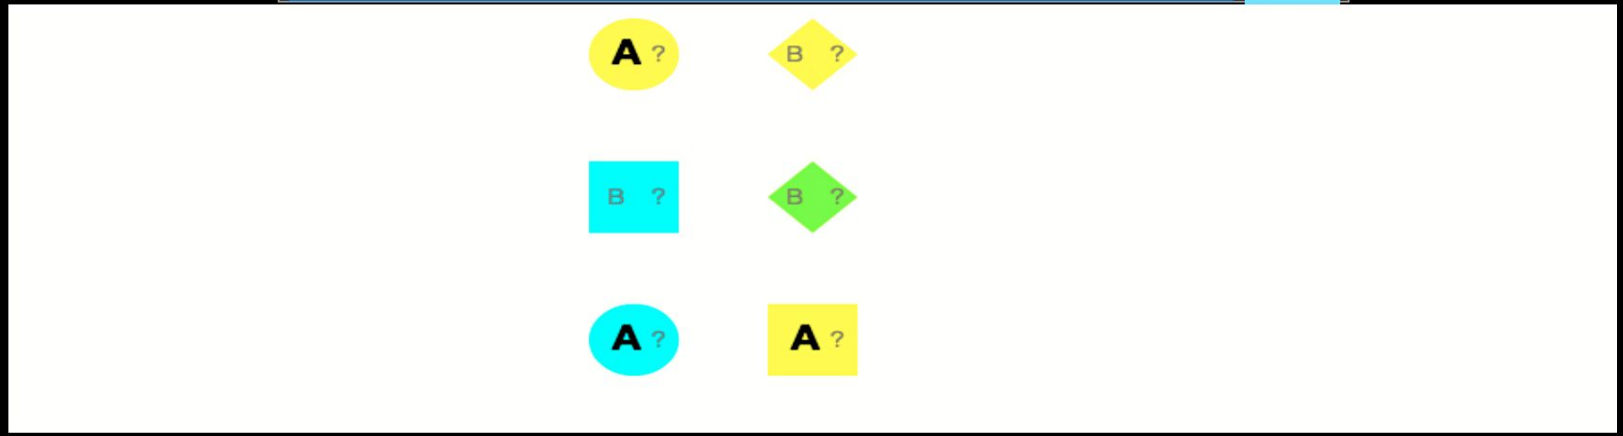

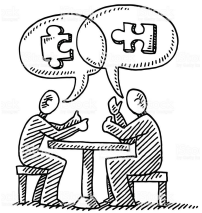

InfoJigsaw

- 10 InfoJigsaw scenarios (either 2*3 or 3*2)

- 200 pairs played 10 scenarios in random error

InfoJigsaw: Example

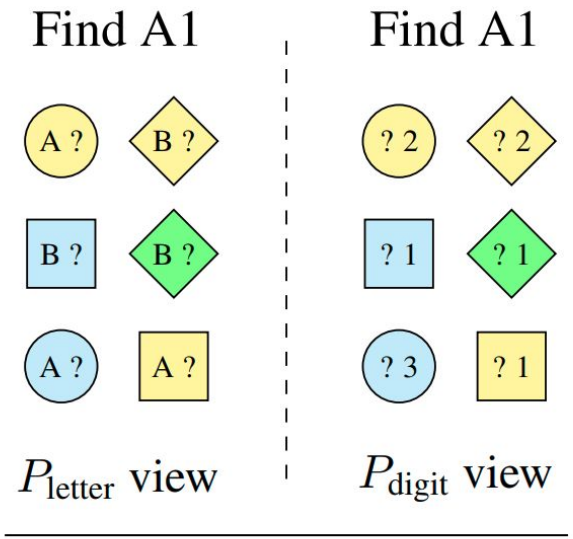

P_{digit}: `middle`

P_{letter}: `yellow circle`

P_{digit}: `bottom right`

P_{letter}: `click (1,2)`

InfoJigsaw: Example

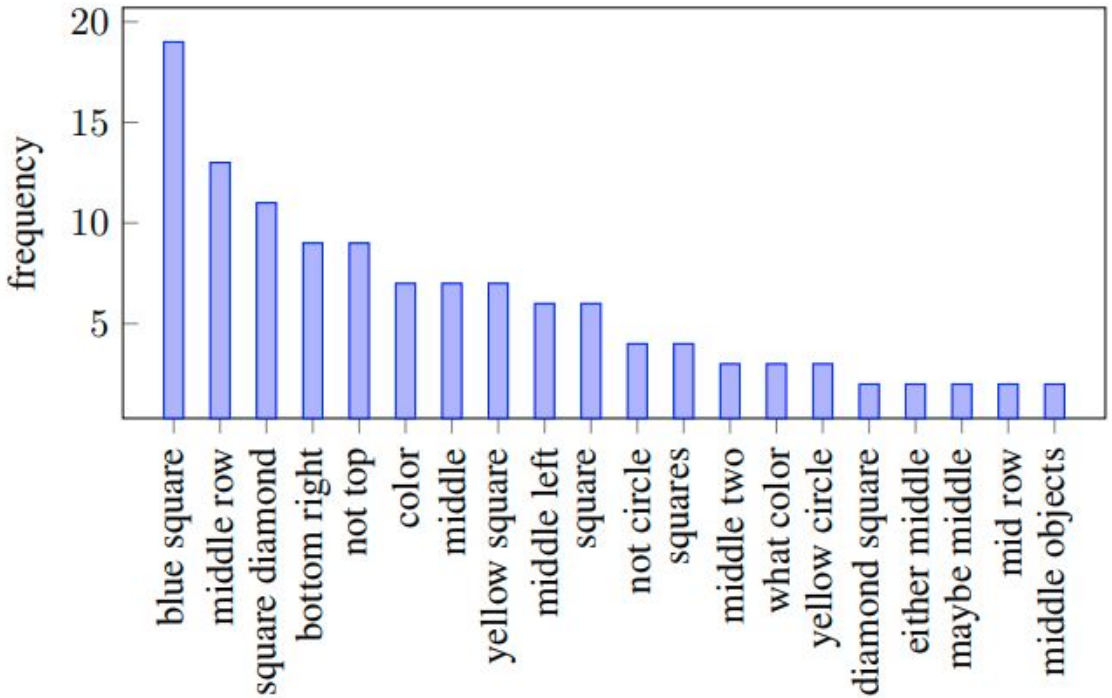

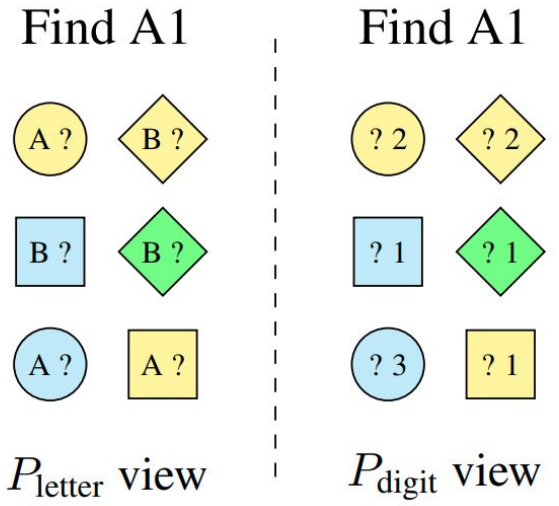

P_{digit}: `middle`

P_{letter}: `yellow circle`

P_{digit}: `bottom right`

P_{letter}: `click (1,2)`

InfoJigsaw statistics

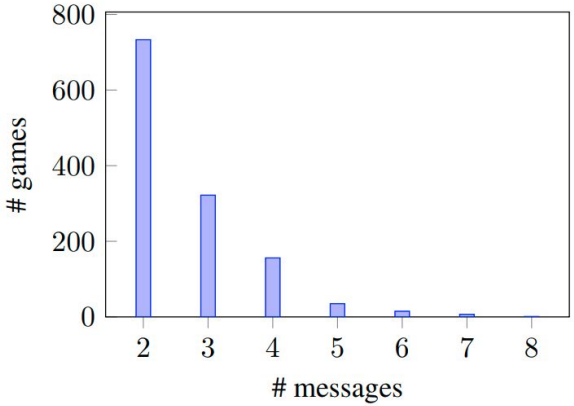

(a) Number of exchanged messages per game.

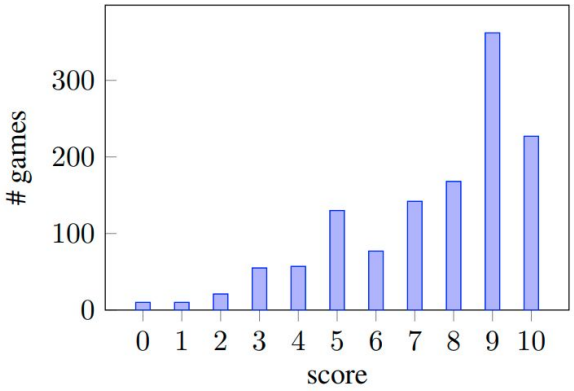

(b) Distribution of final game scores.

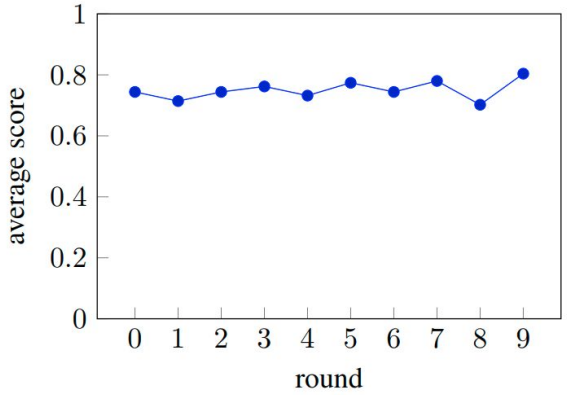

(c) Average score over multiple rounds of game play, which interestingly remains constant.

InfoJigsaw statistics

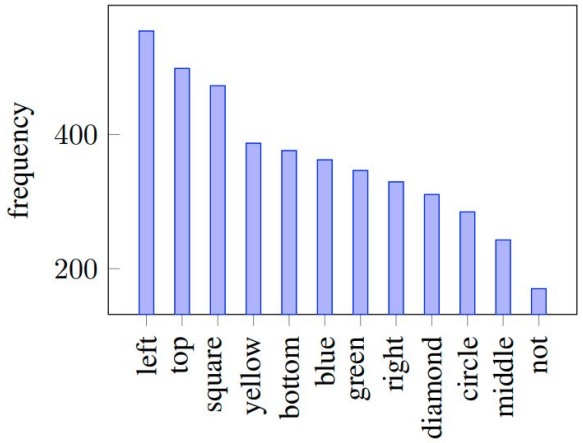

(d) 12 most frequent words, which make up 73% of all tokens.

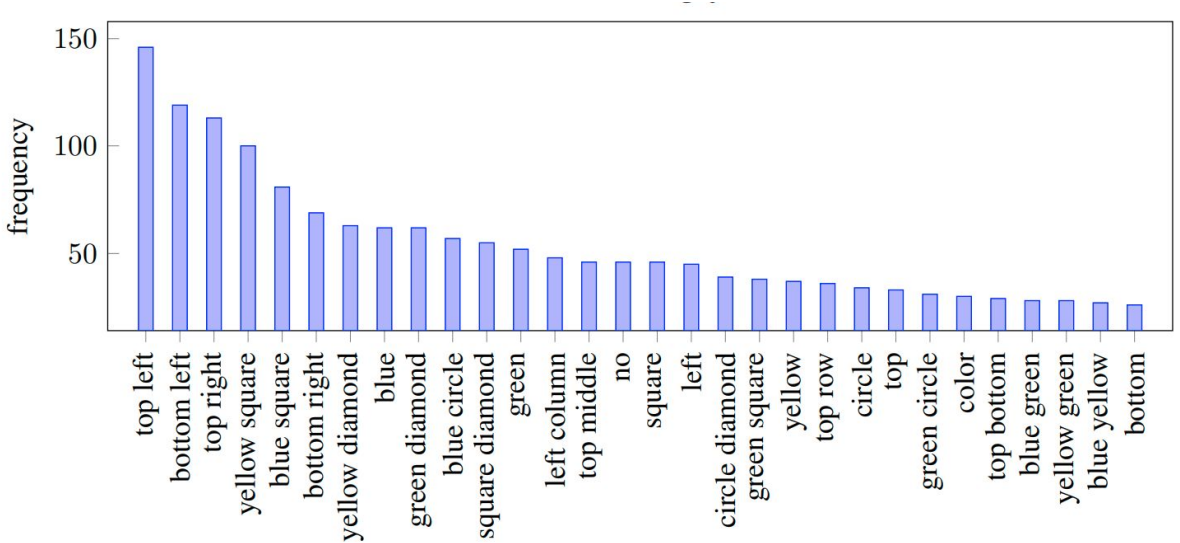

(e) 30 most frequent messages, which make up 49% of all messages.

Dataset statistics

1680	4987	7.5
Games	Messages	Average score

Dataset statistics

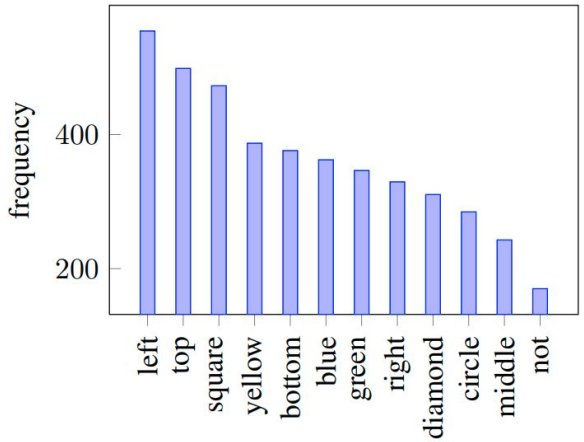

top - middle - bottom
yellow - blue - green
square - circle - diamond
left - right
not - yes - no

1680 Games	4987 Messages	7.5 Average score
1259	3358	7.48

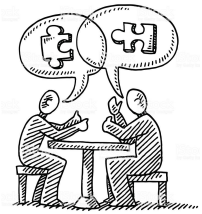

InfoJigsaw

top - middle - bottom
yellow - blue - green
square - circle - diamond
left - right
not - yes - no

146
Different messages

64
Different states

How to have all three in one model?

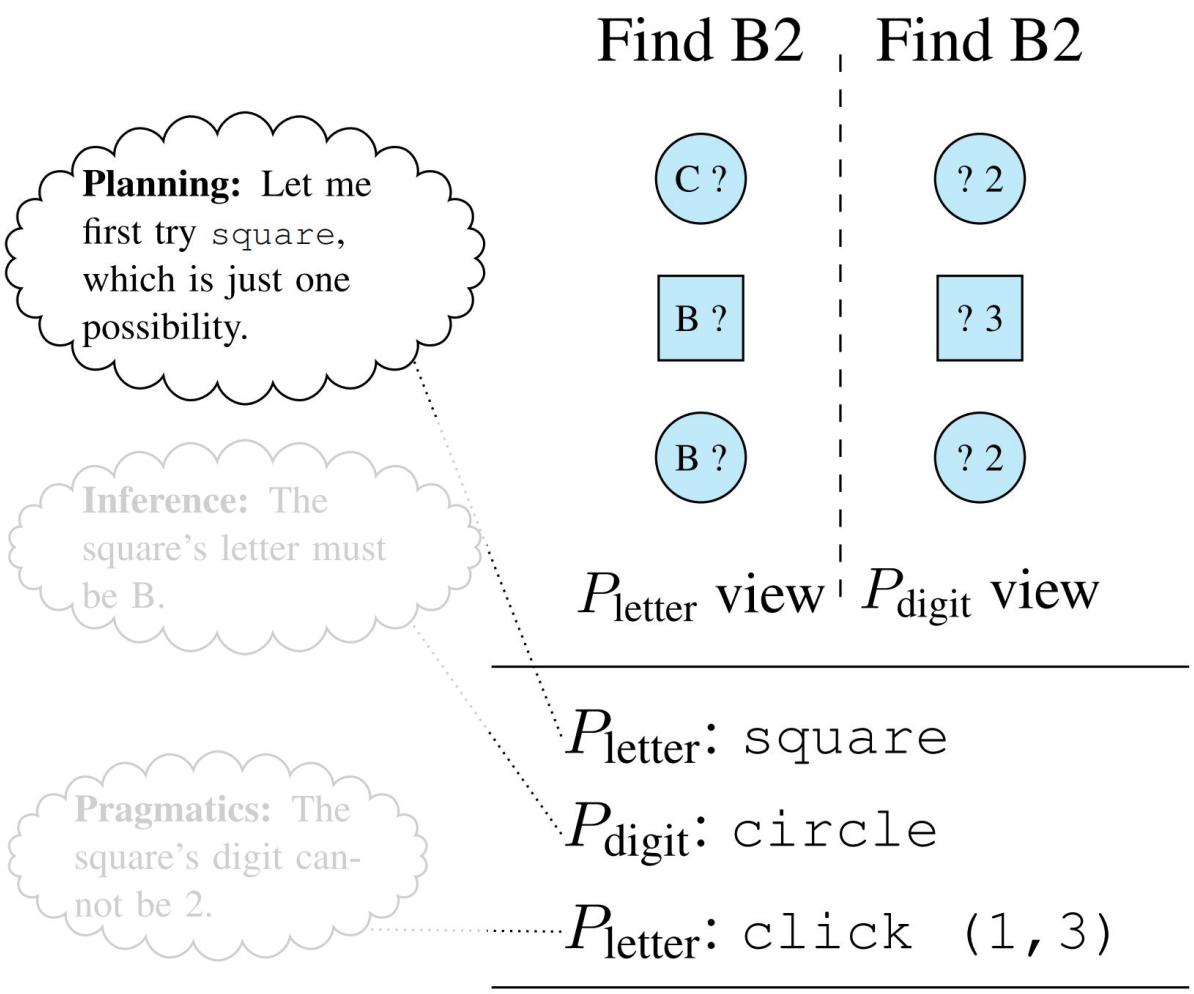
Find B2
Find B2
Planning: Let me first try square, which is just one possibility.
C ?
? 2
B ?
? 3
B ?
? 2
Inference: The square's letter must be B.
P_{letter} view
P_{digit} view
Pragmatics: The square's digit cannot be 2.
P_{letter}: square
P_{digit}: circle
P_{letter}: click (1,3)

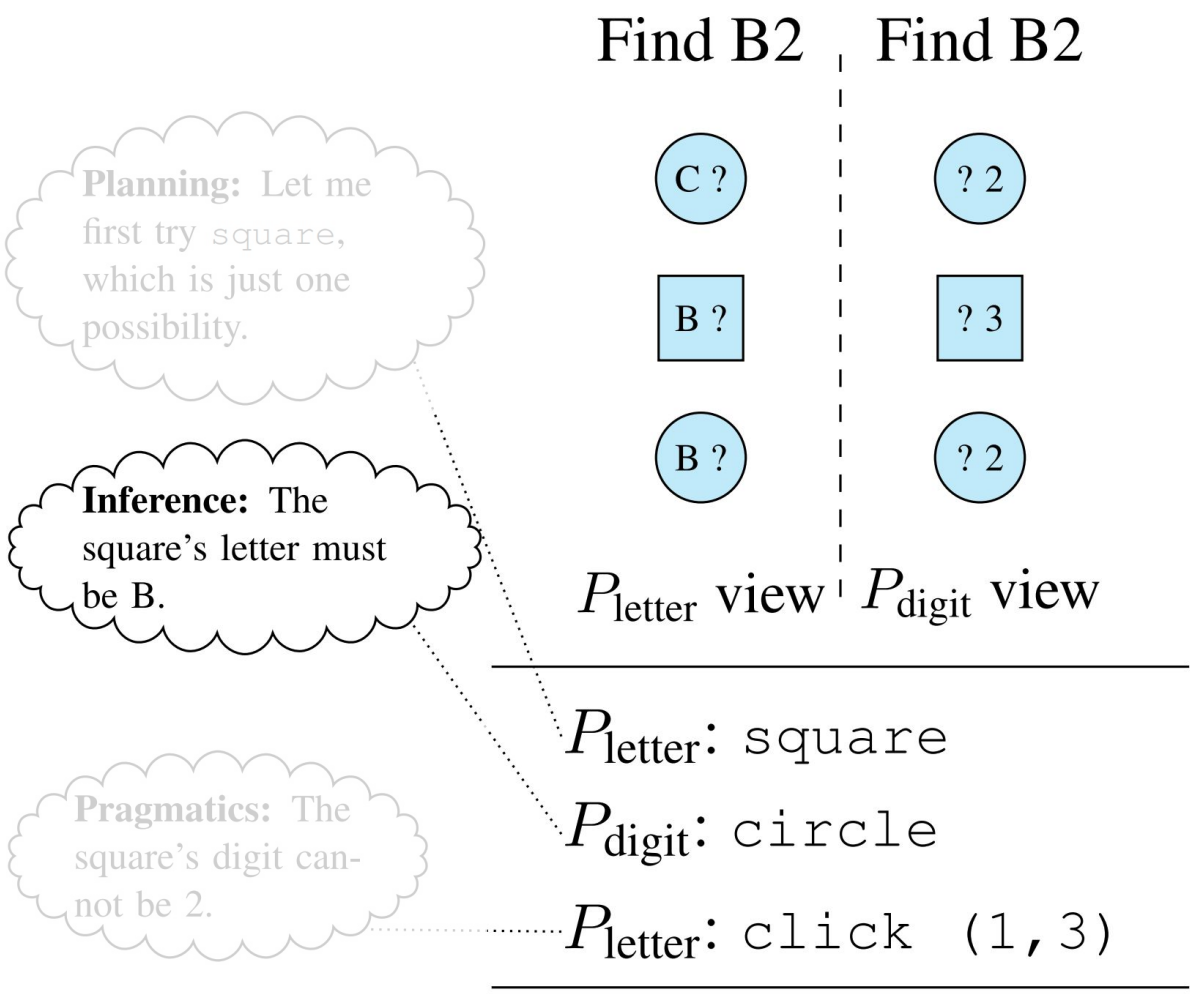
Find B2
Find B2
C ?
? 2
B ?
? 3
B ?
? 2
Planning: Let me first try square, which is just one possibility.
Inference: The square's letter must be B.
Pragmatics: The square's digit cannot be 2.
P_{letter} view
P_{digit} view
P_{letter}: square
P_{digit}: circle
P_{letter}: click (1,3)

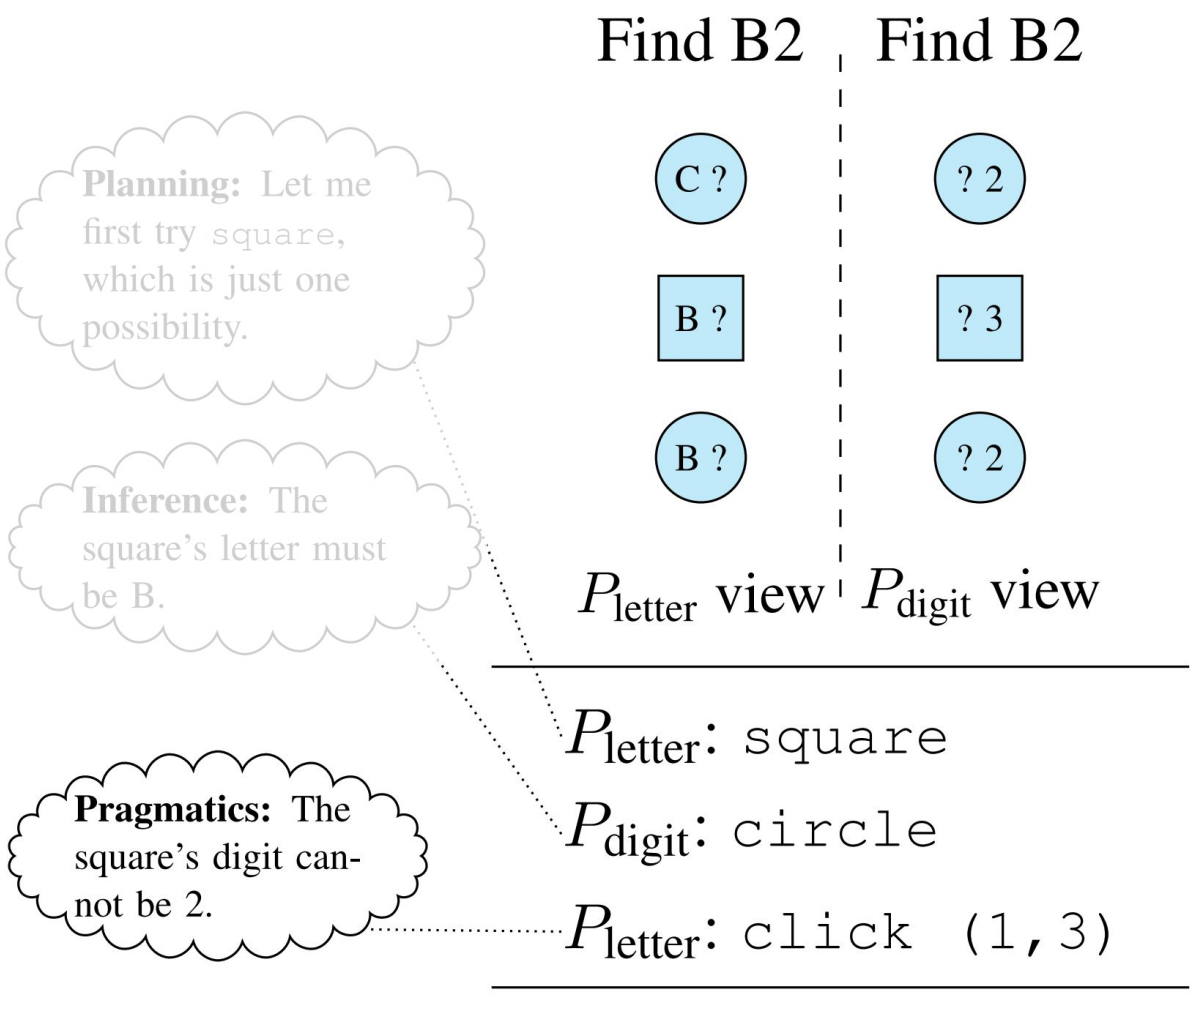
Find B2
Find B2
C ?
? 2
B ?
? 3
B ?
? 2
P_{letter} view
P_{digit} view
Planning: Let me first try square, which is just one possibility.
Inference: The square's letter must be B.
Pragmatics: The square's digit cannot be 2.
P_{letter}: square
P_{digit}: circle
P_{letter}: click (1,3)

Pragmatics

Literal semantics

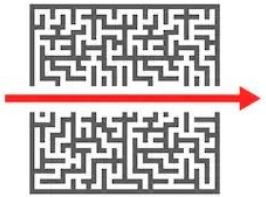

All together

Planning

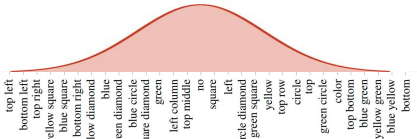

Inference

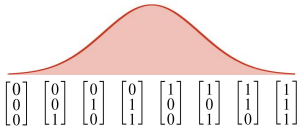

Outline

- Literal semantics
- Inference
- Planning
- Pragmatics
- All three together

Literal Semantics

Purpose: A mapping from each word to its literal semantics,

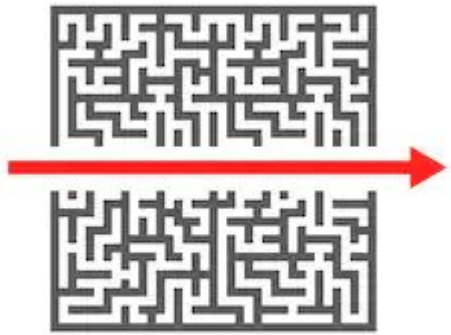

Literal Semantics

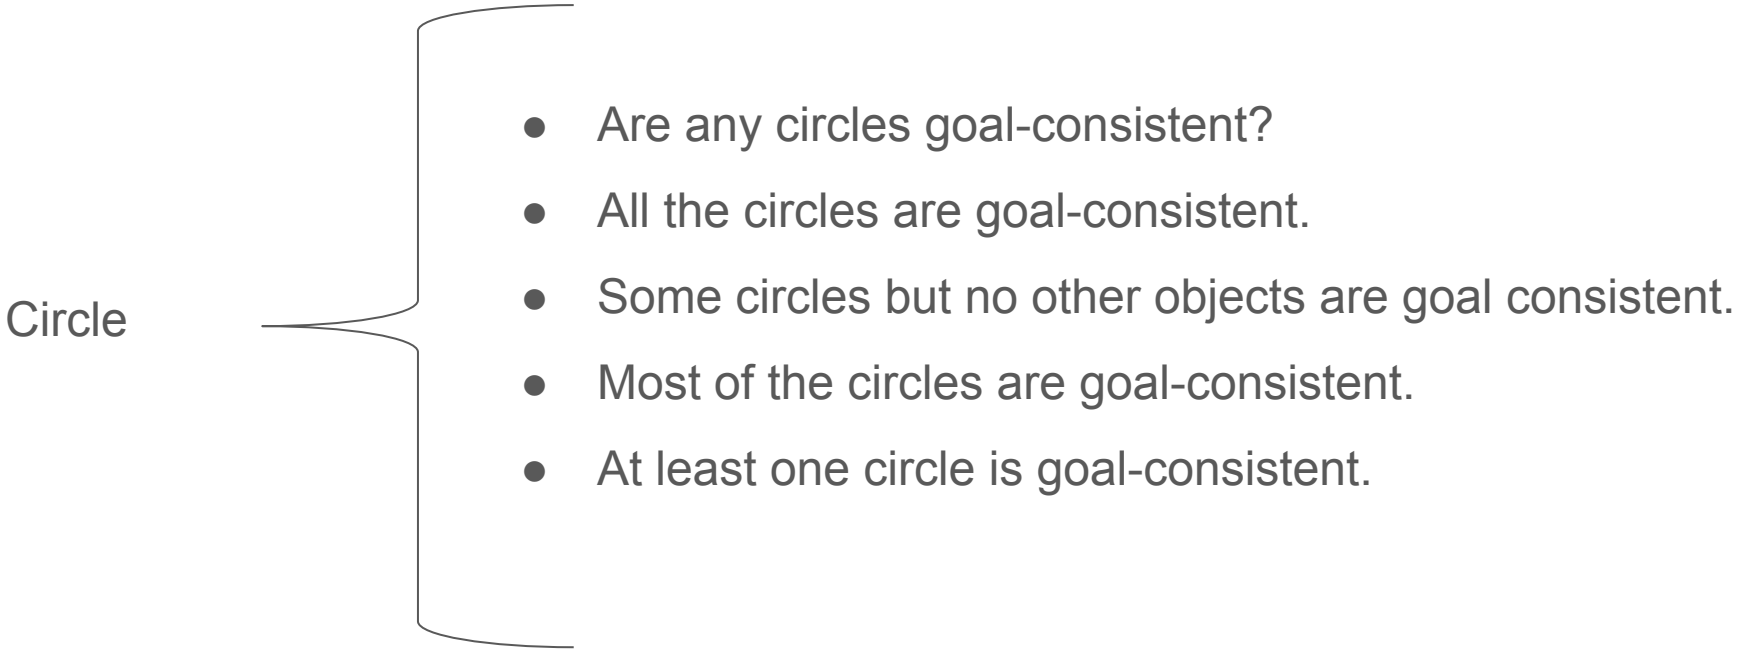

Private states

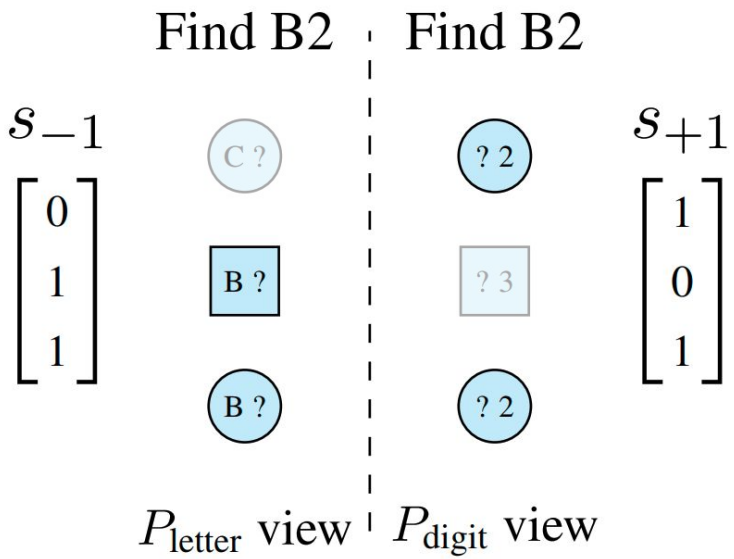

Literal Semantics: informative message

Informative messages describe constraints on the speaker's private state.

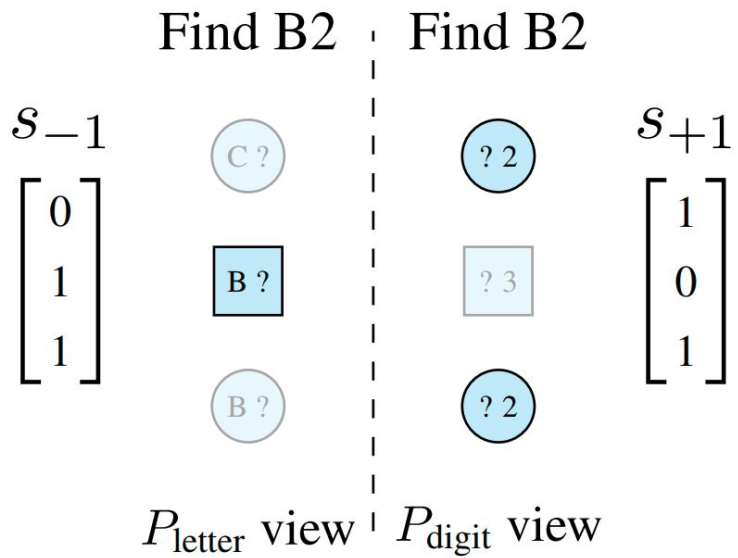

Literal Semantics: informative message

$$[\![\texttt{square}]\!] = \begin{bmatrix}0\\0\\0\end{bmatrix} \begin{bmatrix}0\\0\\1\end{bmatrix} \begin{bmatrix}0\\1\\0\end{bmatrix} \begin{bmatrix}0\\1\\1\end{bmatrix} \begin{bmatrix}1\\0\\0\end{bmatrix} \begin{bmatrix}1\\0\\1\end{bmatrix} \begin{bmatrix}1\\1\\0\end{bmatrix} \begin{bmatrix}1\\1\\1\end{bmatrix}$$

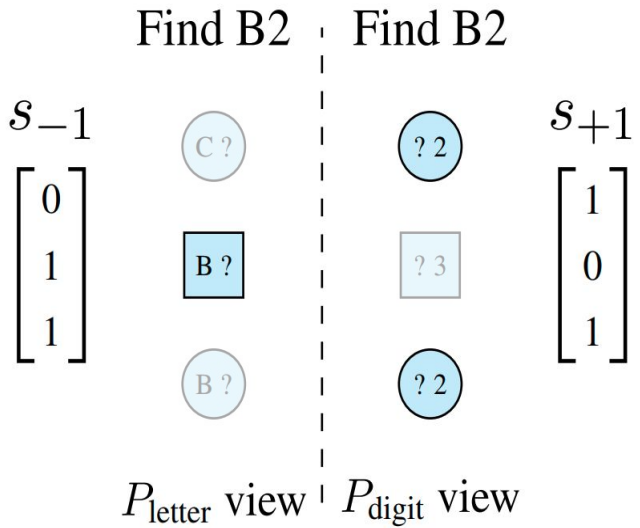

Literal Semantics: informative message

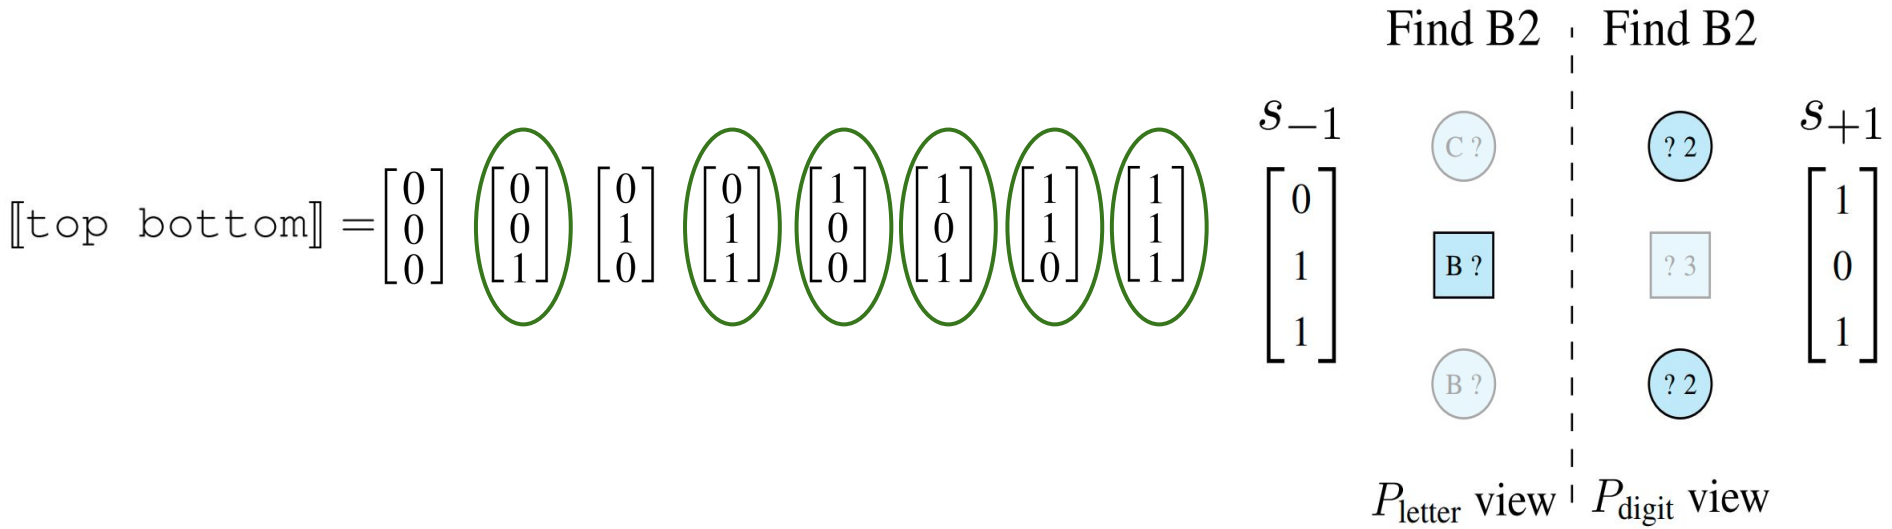

Literal semantics: action sequence

$$P_{\text{digit}} : \text{circle}$$
$$P_{\text{letter}} : \text{square} \implies s \in [\![circle]\!], s \notin [\![square]\!]$$
$$P_{\text{digit}} : \text{no}$$

Literal semantics: action sequence

P_{digit} : circle
P_{letter} : square
P_{digit} : no

$$\begin{bmatrix}0\\0\\0\end{bmatrix} \begin{bmatrix}0\\0\\1\end{bmatrix} \begin{bmatrix}0\\1\\0\end{bmatrix} \begin{bmatrix}0\\1\\1\end{bmatrix} \begin{bmatrix}1\\0\\0\end{bmatrix} \begin{bmatrix}1\\0\\1\end{bmatrix} \begin{bmatrix}1\\1\\0\end{bmatrix} \begin{bmatrix}1\\1\\1\end{bmatrix}$$

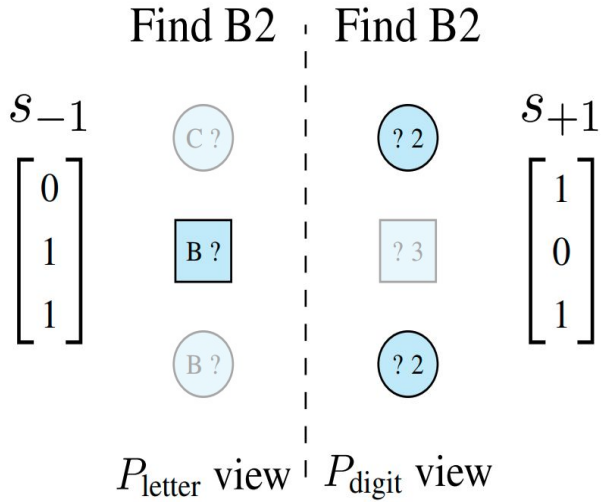

$$\Longrightarrow \; s \in [\![circle]\!], s \notin [\![square]\!]$$

Outline

Literal semantics

- Inference
- Planning
- Pragmatics
- All three together

Inference

Purpose: computing a distribution over the partner's private state

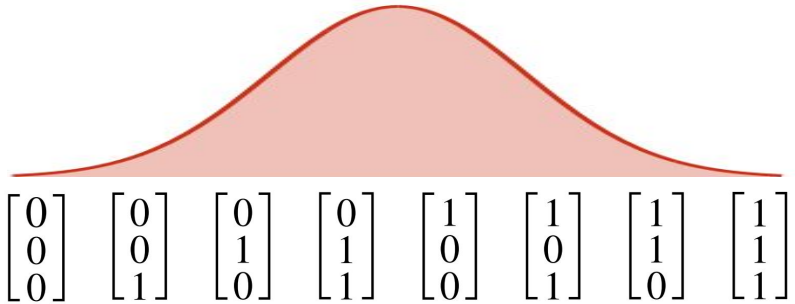

Inference

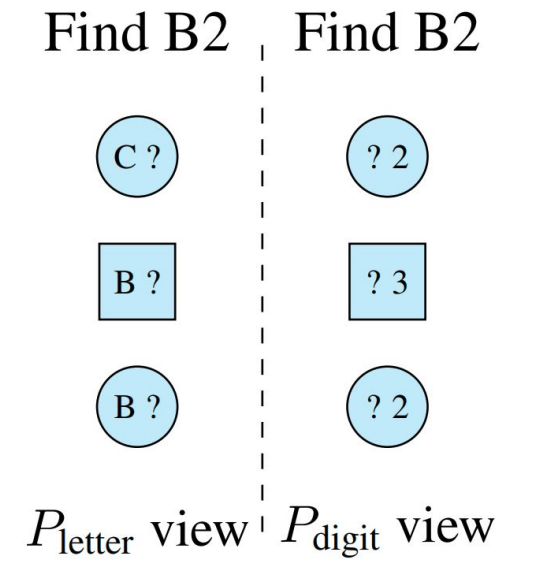

P_{letter}: `square`

P_{digit}: `circle`

What is belief of P_{letter} about P_{digit} state?

Inference

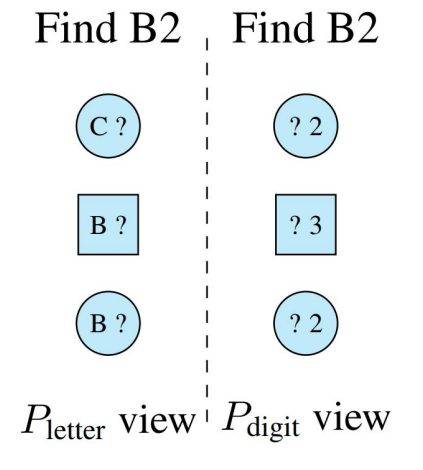

P_{letter}: `square`

P_{digit}: `circle`

What is belief of P_{letter} about P_{digit} state?

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \begin{cases} 1 & s_{-j} \text{ consistent with } a_{1:t} \\ 0 & o.w. \end{cases}$$

Inference

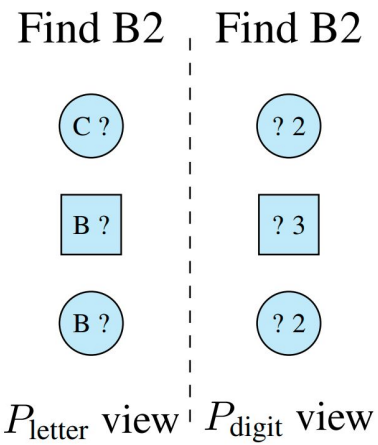

P_{letter}: `square`

P_{digit}: `circle`

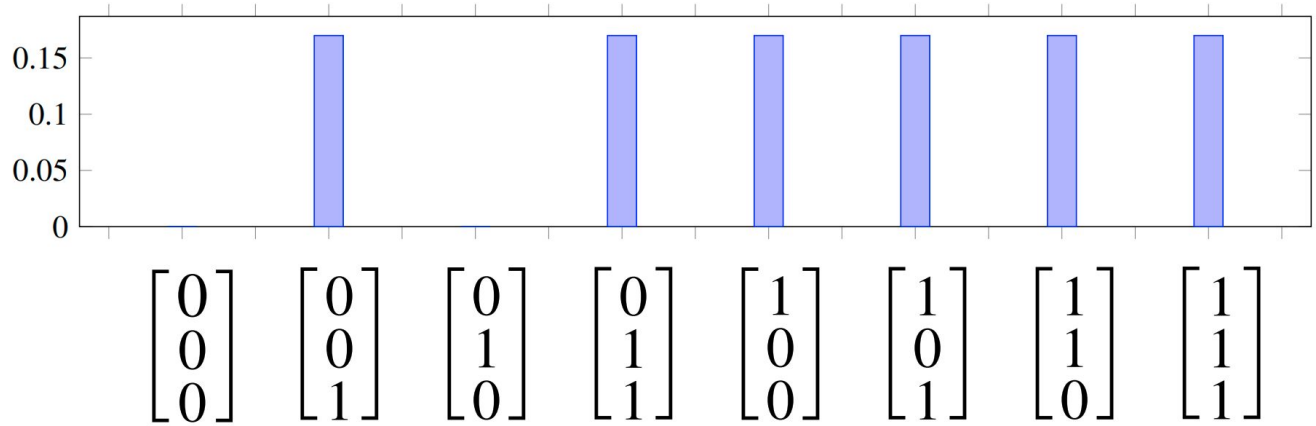

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \begin{cases} 1 & s_{-j} \text{ consistent with } a_{1:t} \\ 0 & o.w. \end{cases}$$

Outline


Literal semantics


Inference

- Planning
- Pragmatics
- All three together

Planning

Purpose: computing a policy, which specifies a distribution over a player actions

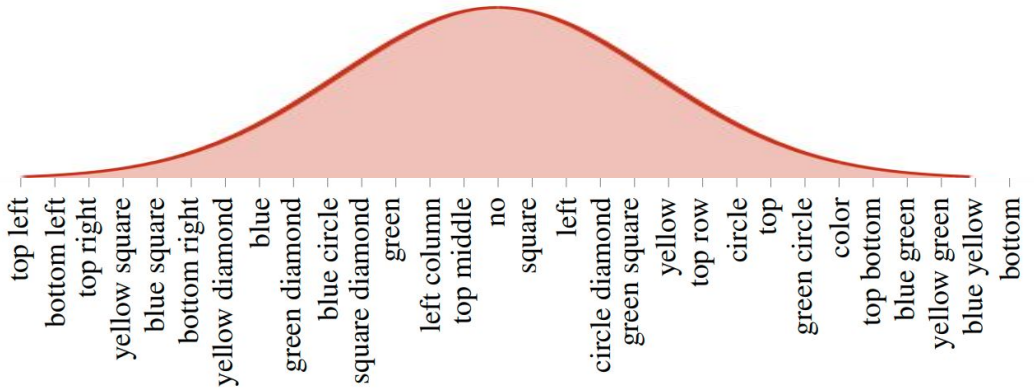

Planning

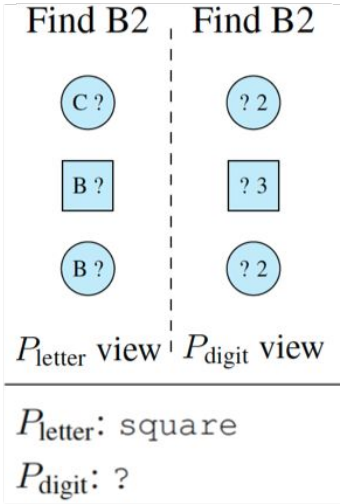
Find B2
Find B2
C ?
? 2
B ?
? 3
B ?
? 2
P_{letter} view
P_{digit} view
P_{letter}: square
P_{digit}: ?

Planning

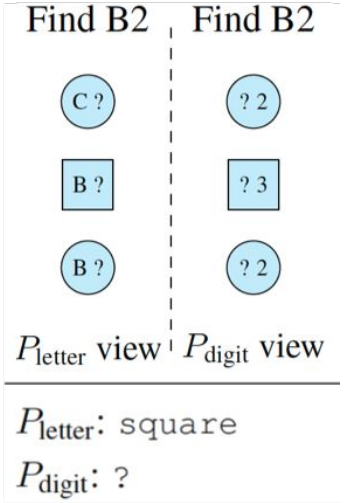

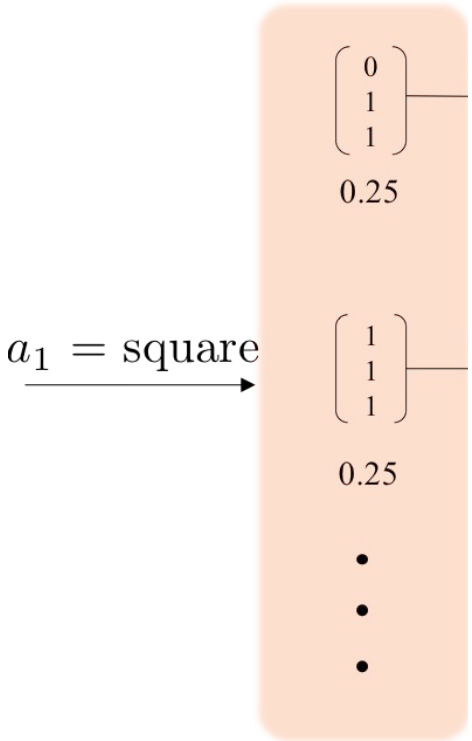

(a)

(a) Infer other player private state from utterances so far.

Planning

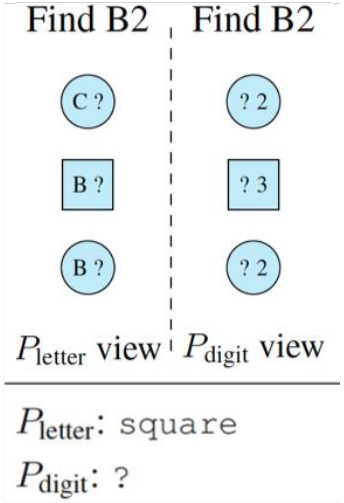

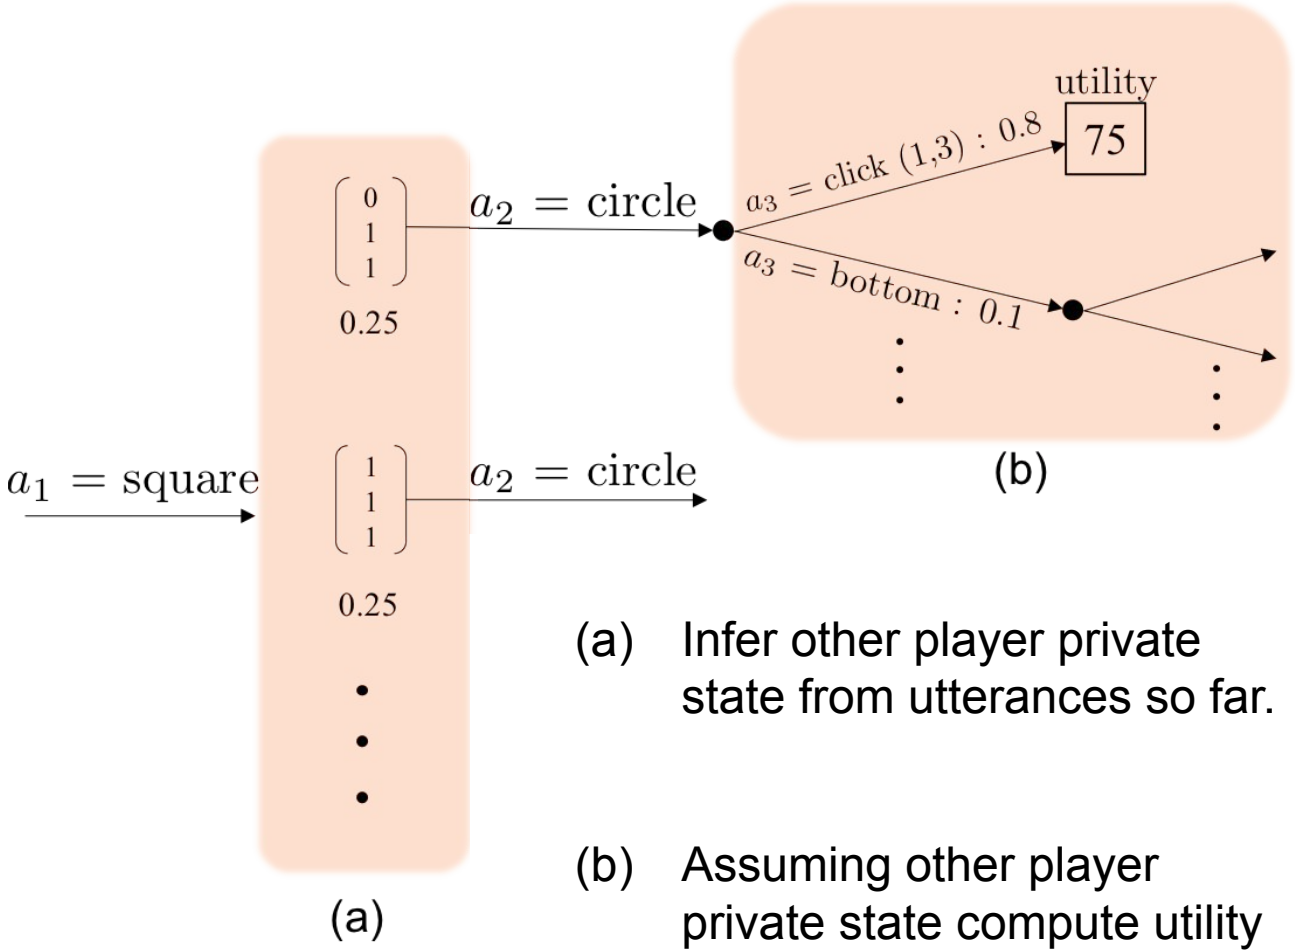

(a) Infer other player private state from utterances so far.

(b) Assuming other player private state compute utility of saying some utterances.

Planning

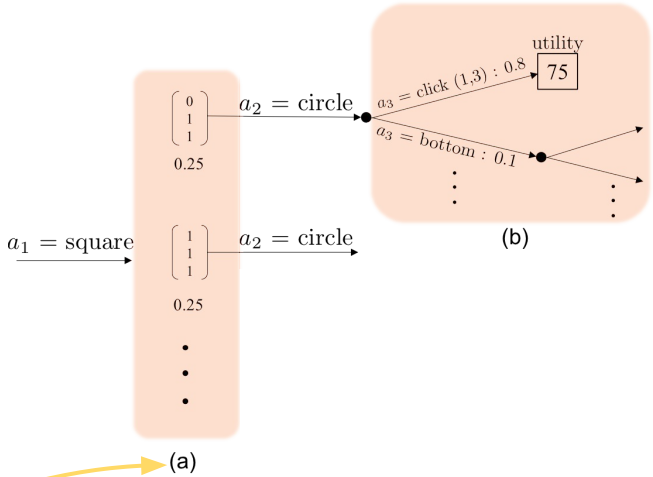

$$\pi(a_t \mid s_j, a_{1:t}) = \sum_{s_{-j}} p(s_{-j} \mid s_j, a_{1:t-1}) V(s_{-j}, s_j, a_{1:t})$$

Inferring other player's private state

Planning

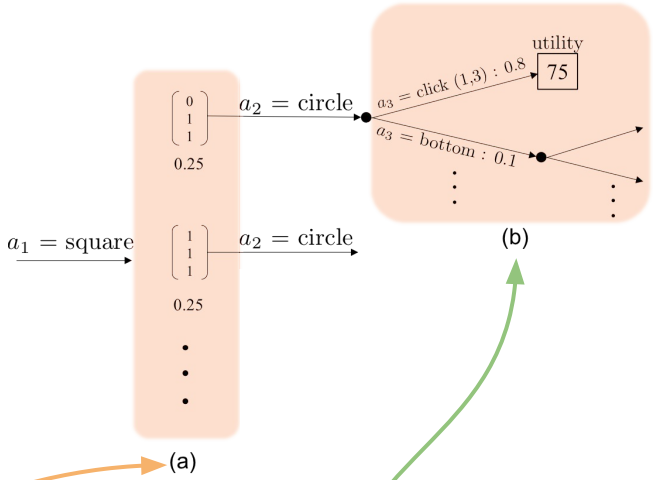

$$\pi(a_t \mid s_j, a_{1:t}) = \sum_{s_{-j}} p(s_{-j} \mid s_j, a_{1:t-1}) V(s_{-j}, s_j, a_{1:t})$$

Inferring other player's private state

Utility of the game given states and actions

Expected Utility

If the game is over

- A penalty for not reaching the goal
- A reward for reached the goal
- Penalty for each action

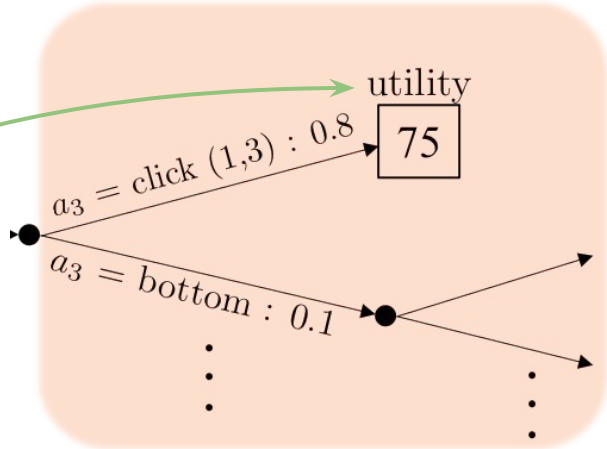

Expected Utility

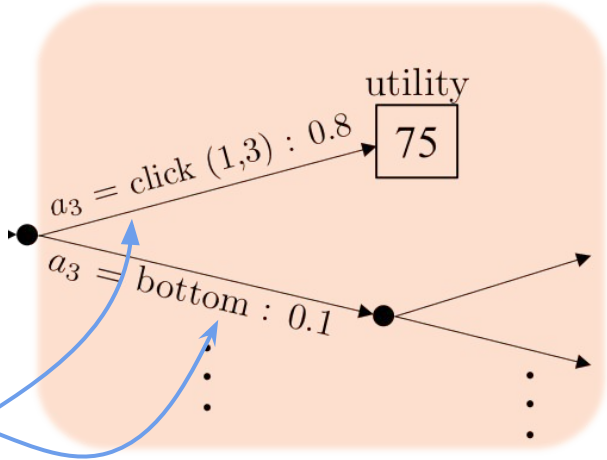

$$V(s_j, s_{-j}, a_{1:t}) = \sum_{a_{t+1}} \pi(a_{t+1} \mid s_j, a_{1:t}) V(s_{-j}, s_j, a_{1:t+1})$$

Probability of choosing next action

Expected Utility

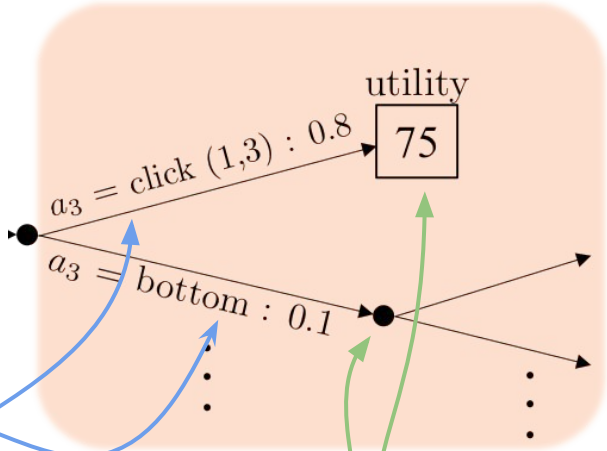

$$V(s_j, s_{-j}, a_{1:t}) = \sum_{a_{t+1}} \pi(a_{t+1} \mid s_j, a_{1:t}) V(s_{-j}, s_j, a_{1:t+1})$$

Probability of choosing next action

Utility of the games given all actions

Outline


Literal semantics


Inference


Planning

- Pragmatics
- All three together

Pragmatics

Purpose: Taking into account the partner's strategizing.

Inference: Recap

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \begin{cases} 1 & s_{-j} \text{ consistent with } a_{1:t} \\ 0 & o.w. \end{cases}$$

Pragmatics

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \begin{cases} 1 & s_{-j} \text{ consistent with } a_{1:t} \\ 0 & o.w. \end{cases}$$

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \text{ Probability of choosing } a_{1:t} \text{ in } s_{-j}$$

Pragmatics

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \text{Probability of choosing } a_{1:t} \text{ in } s_{-j}$$

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \pi(a_t \mid s_{-j}, a_{1:t-1}) p(s_{-j} \mid s_j, a_{1:t-2})$$

Pragmatics

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \text{Probability of choosing } a_{1:t} \text{ in } s_{-j}$$

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \pi(a_t \mid s_{-j}, a_{1:t-1}) p(s_{-j} \mid s_j, a_{1:t-2})$$

Probability of choosing the last action in that state

Pragmatics

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \text{Probability of choosing } a_{1:t} \text{ in } s_{-j}$$

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \pi(a_t \mid s_{-j}, a_{1:t-1}) p(s_{-j} \mid s_j, a_{1:t-2})$$

Probability of choosing the last action in that state

Probability of being on that stat given the rest of actions

Outline


Literal semantics


Inference


Planning


Pragmatics

- All three together

All three together

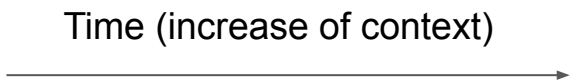

utility

planning

Inference

All three together: computing utility

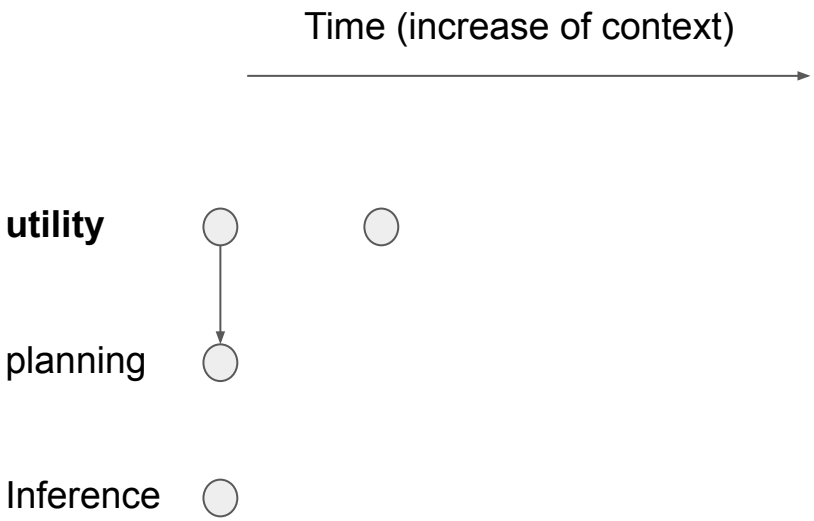

$$V(s_j, s_{-j}, a_{1:t}) = \sum_{a_{t+1}} \pi(a_{t+1} \mid s_j, a_{1:t}) V(s_{-j}, s_j, a_{1:t+1})$$

All three together: computing utility

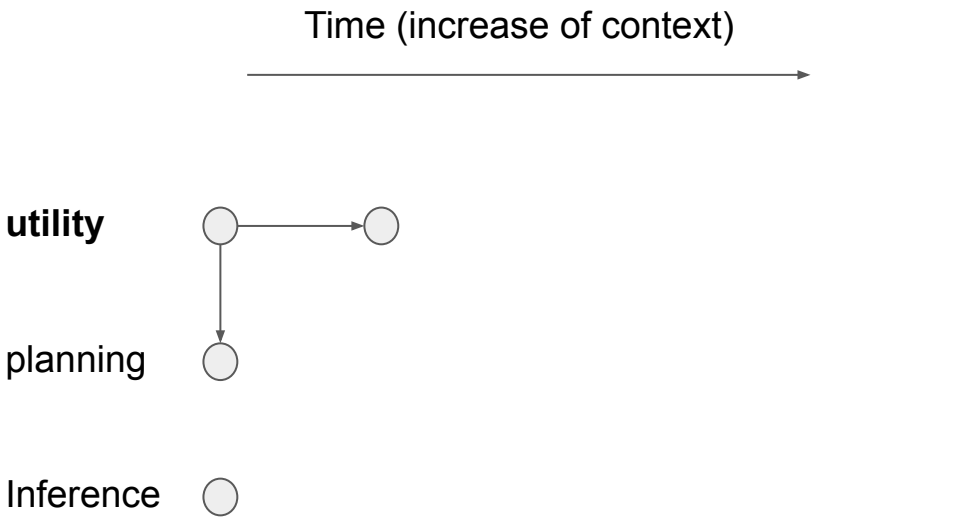

$$V(s_j, s_{-j}, a_{1:t}) = \sum_{a_{t+1}} \pi(a_{t+1} \mid s_j, a_{1:t}) V(s_{-j}, s_j, a_{1:t+1})$$

All three together: computing planning

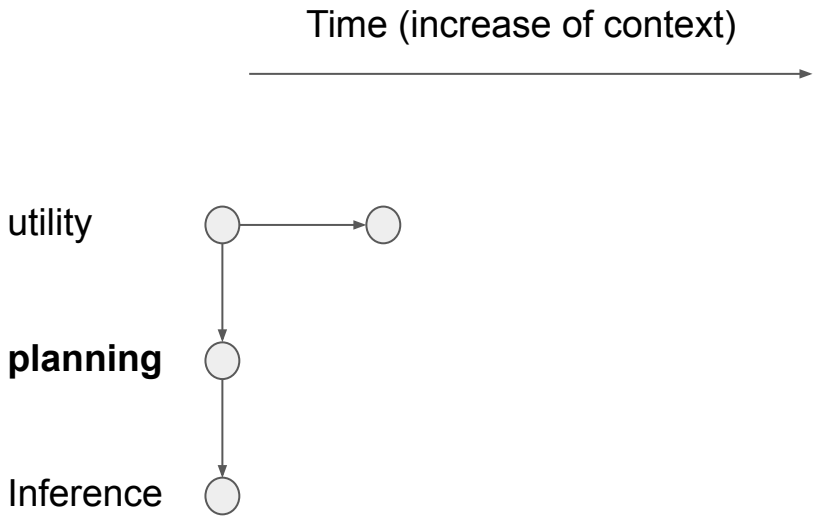

$$\pi(a_t \mid s_j, a_{1:t}) = \sum_{s_{-j}} p(s_{-j} \mid s_j, a_{1:t-1}) V(s_{-j}, s_j, a_{1:t})$$

All three together: computing planning

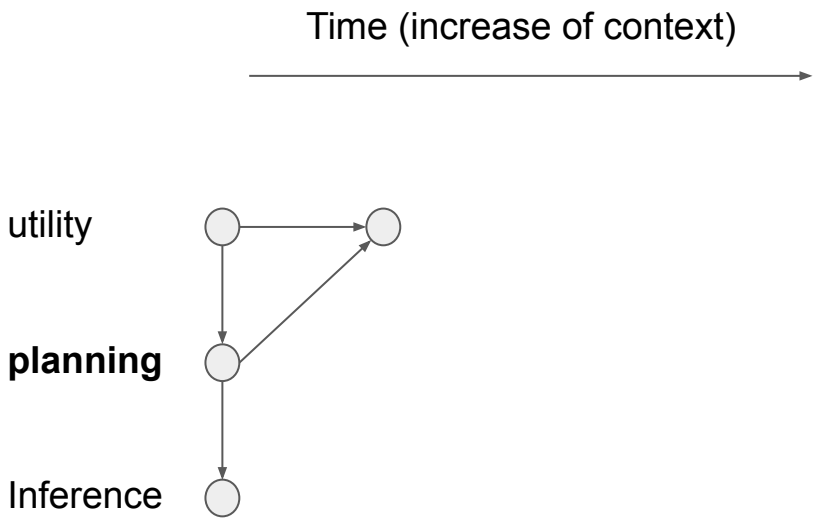

$$\pi(a_t \mid s_j, a_{1:t}) = \sum_{s_{-j}} p(s_{-j} \mid s_j, a_{1:t-1}) V(s_{-j}, s_j, a_{1:t})$$

All three together

Time (increase of context)

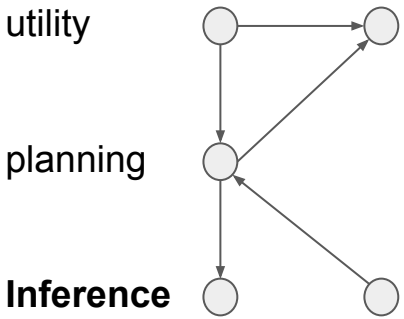

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \pi(a_t \mid s_{-j}, a_{1:t-1}) p(s_{-j} \mid s_j, a_{1:t-2})$$

All three together

Time (increase of context)

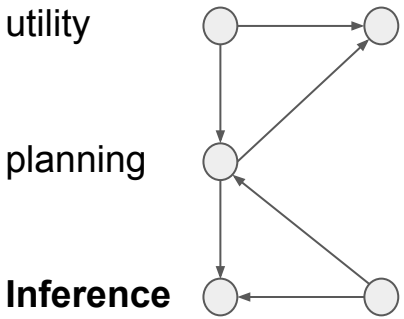

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \pi(a_t \mid s_{-j}, a_{1:t-1}) p(s_{-j} \mid s_j, a_{1:t-2})$$

All three together

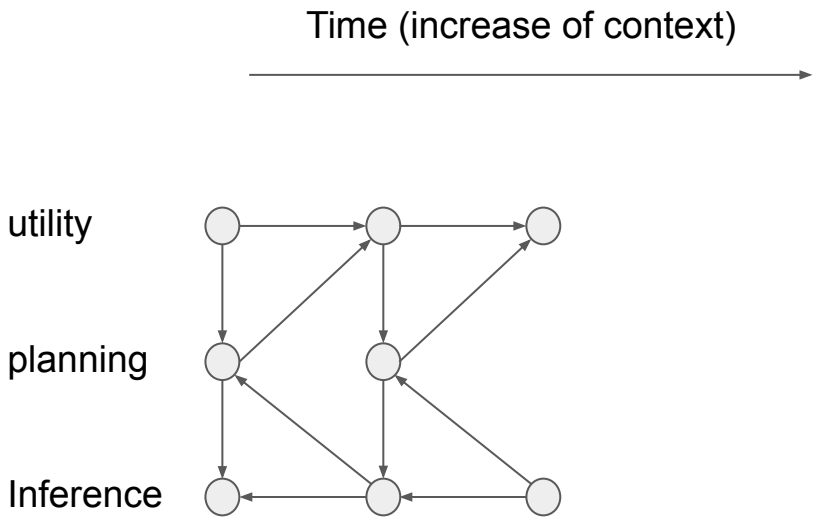

All three together

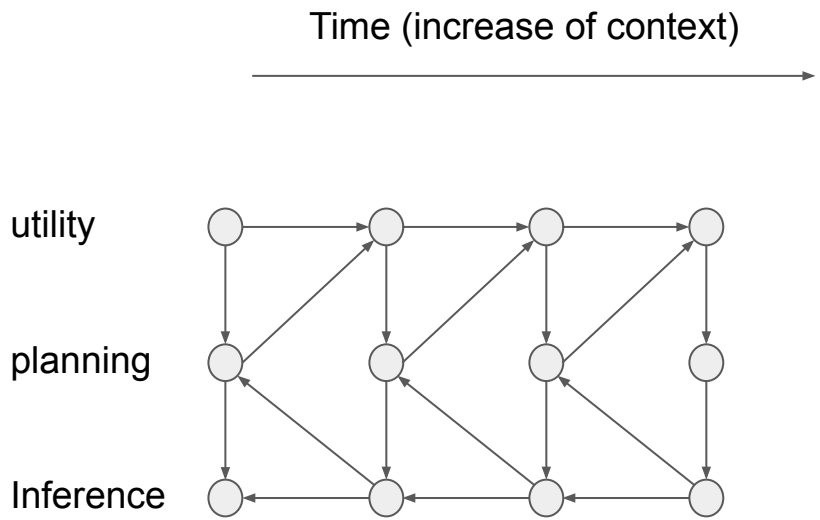

All three together

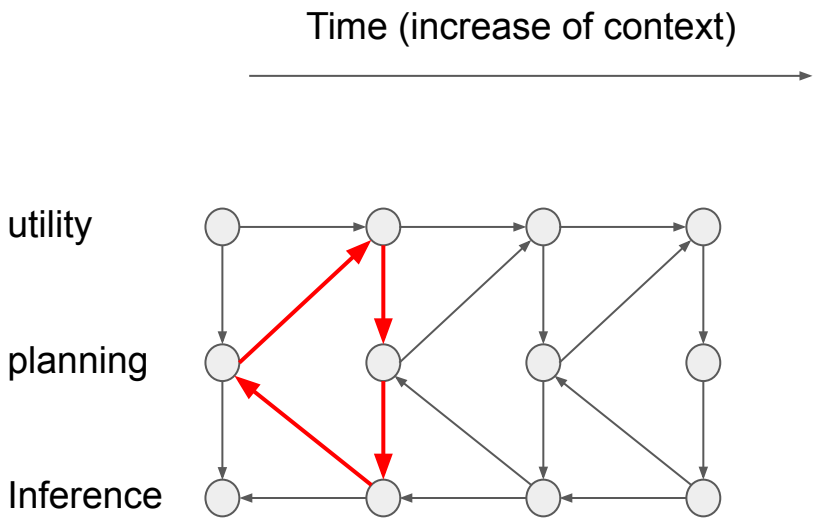

All three together

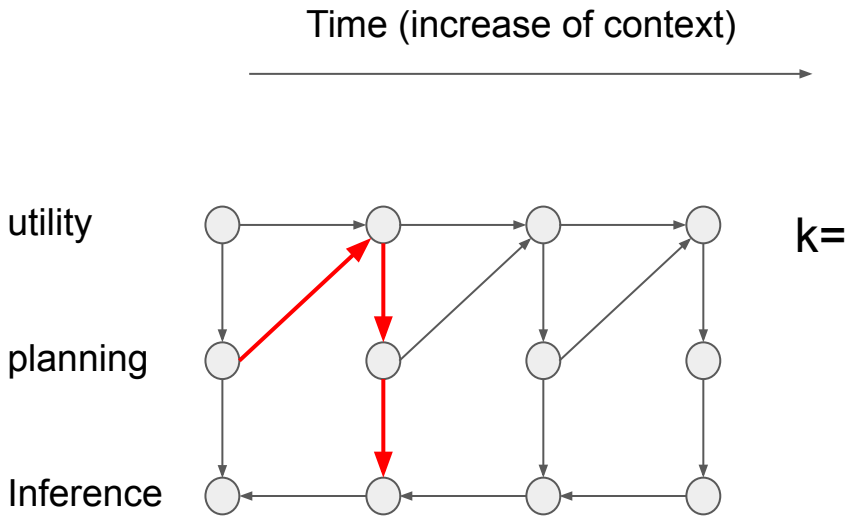

k=0

Basic inference

$$p(s_{-j} \mid s_j, a_{1:t}) \propto \begin{cases} 1 & s_{-j} \text{ consistent with } a_{1:t} \\ 0 & o.w. \end{cases}$$

All three together

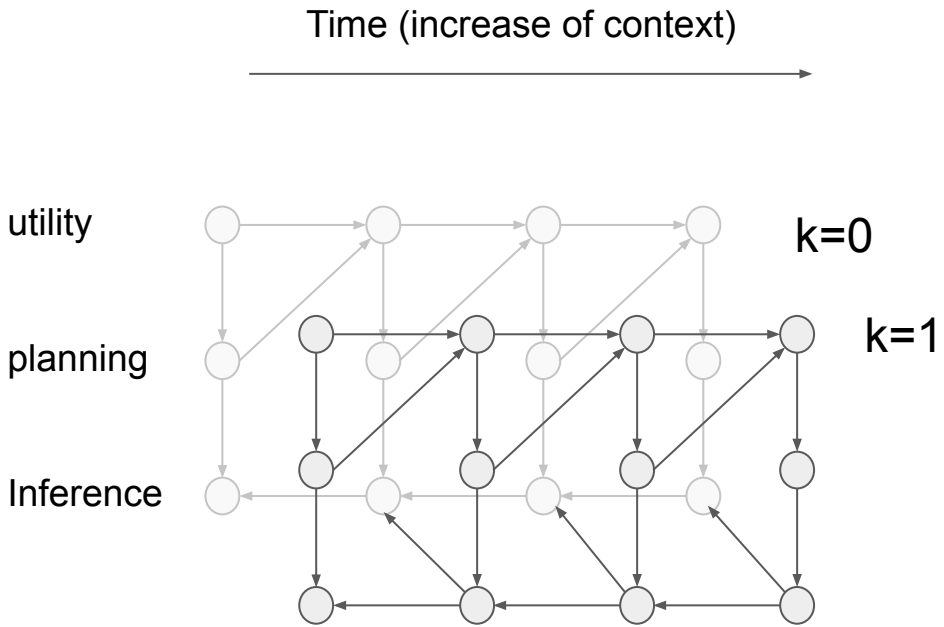

All three together

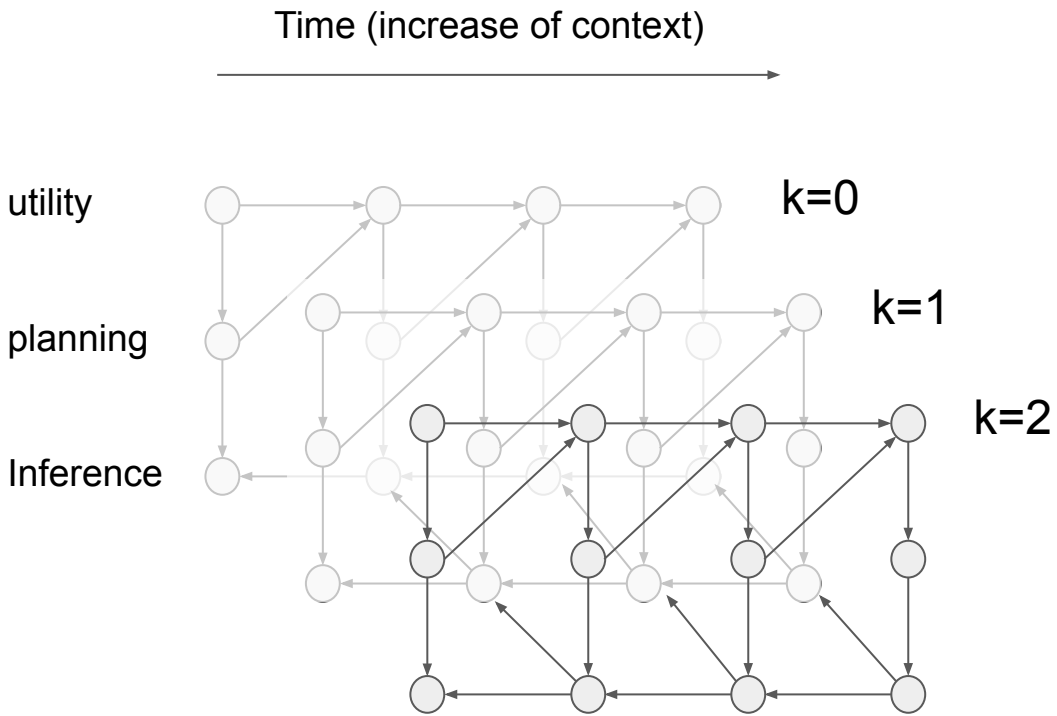

All three together

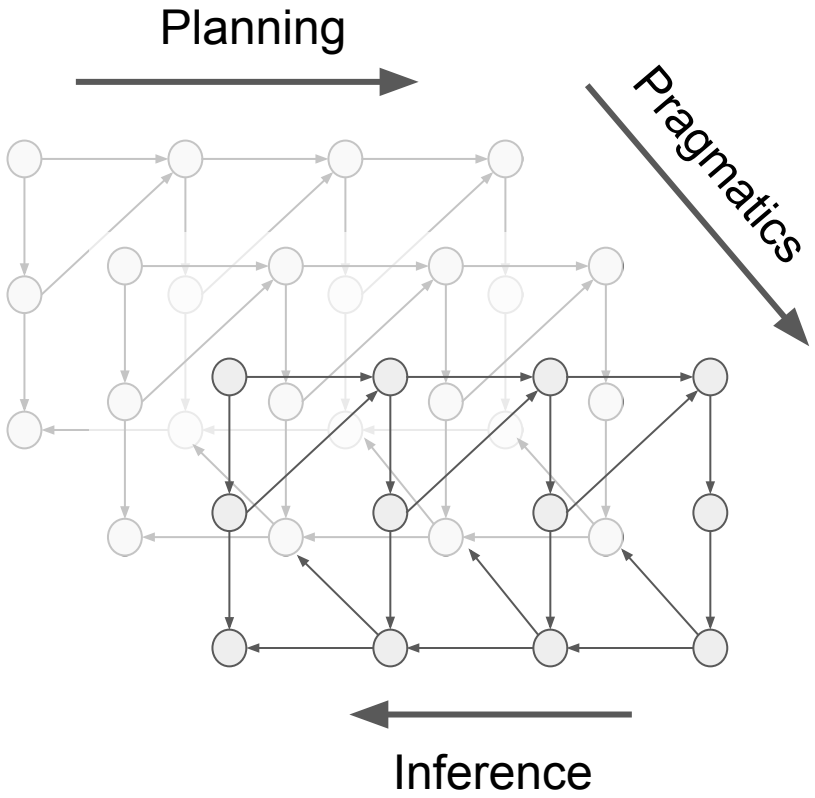

Pragmatics: example

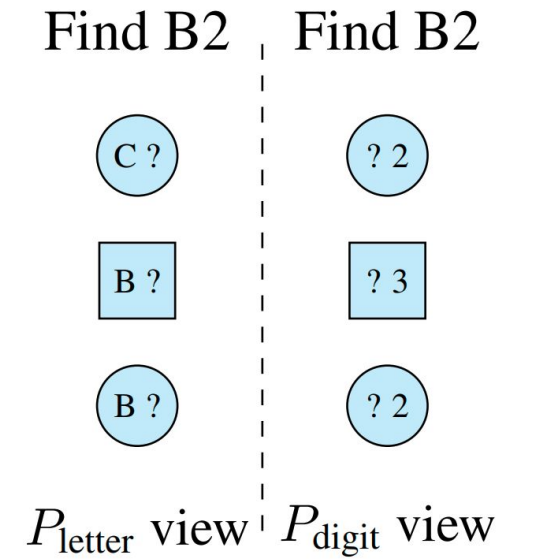

What is belief of P_{letter} about P_{digit} state?

P_{letter}: `square`

P_{digit}: `circle`

Pragmatics: example

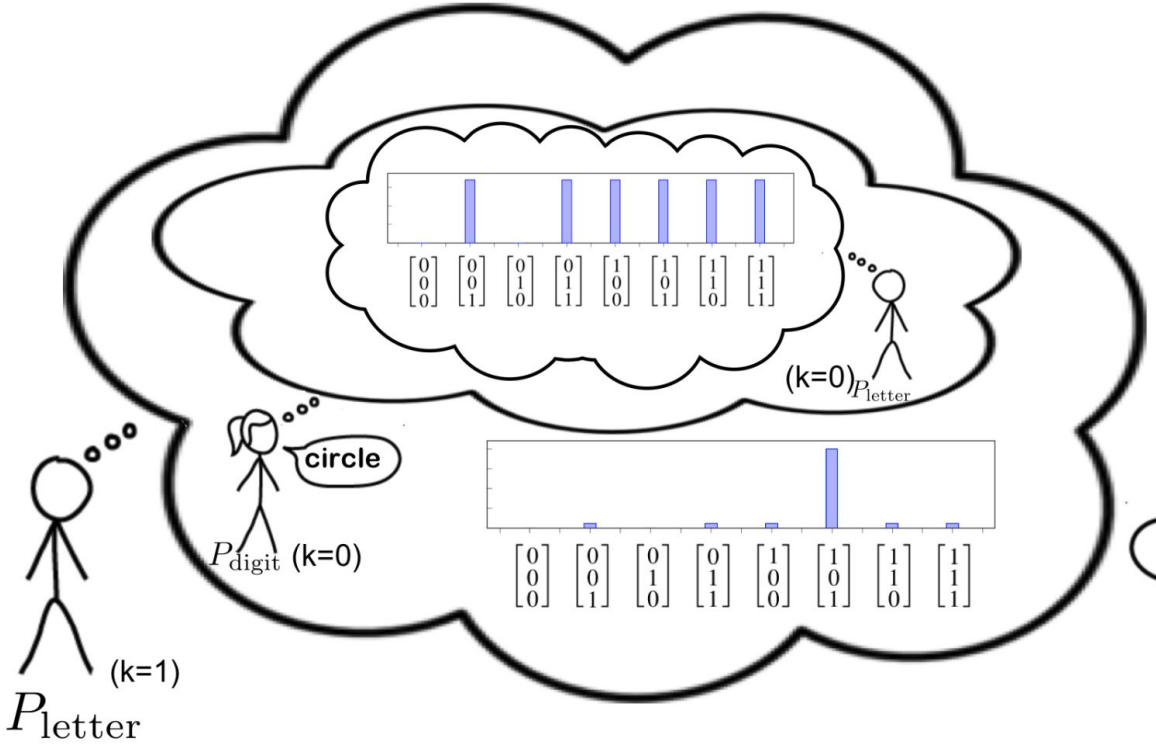

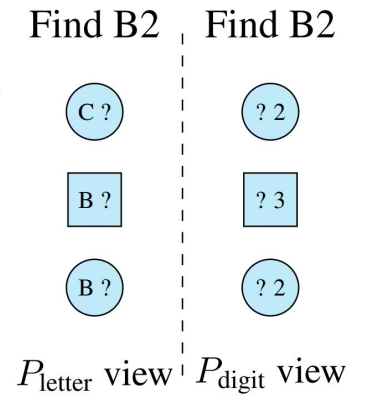

P_{letter}: `square`

P_{digit}: `circle`

Pragmatics: example

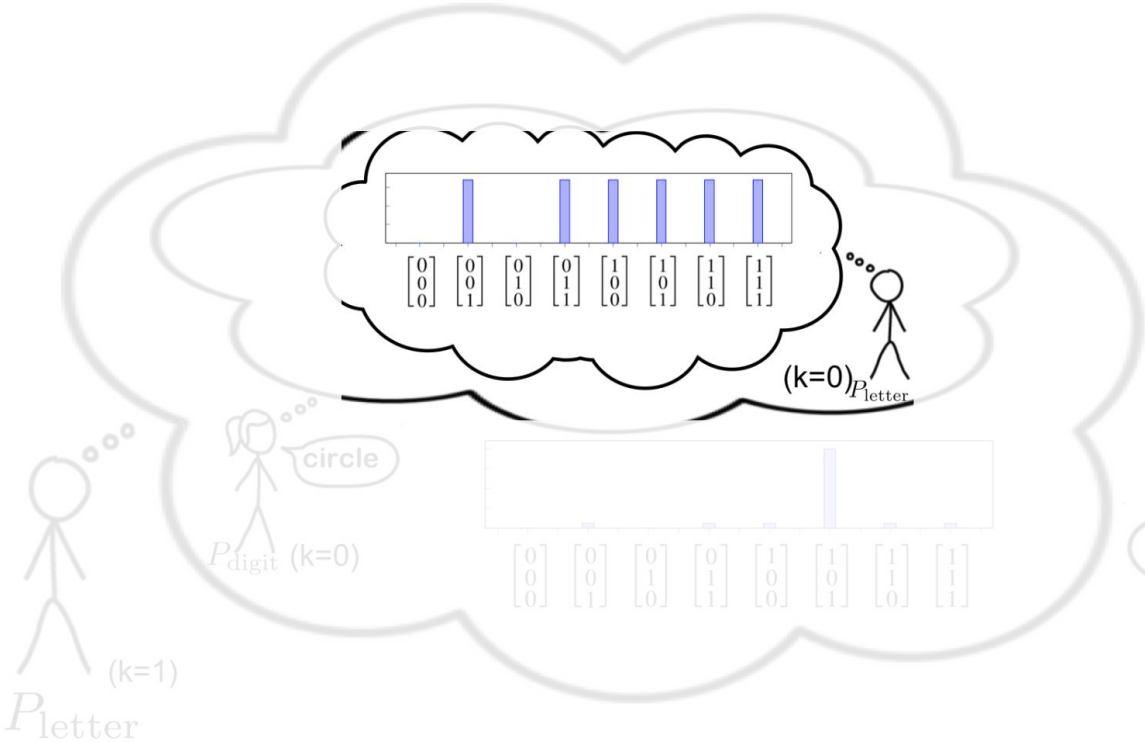

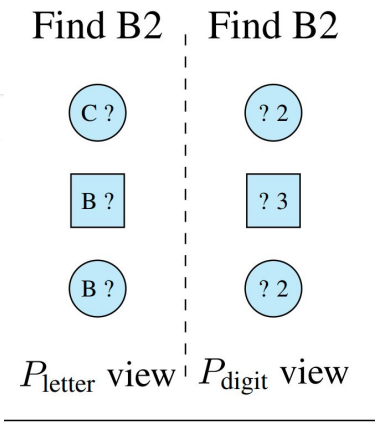

P_{letter}: `square`

P_{digit}: `circle`

Pragmatics: example

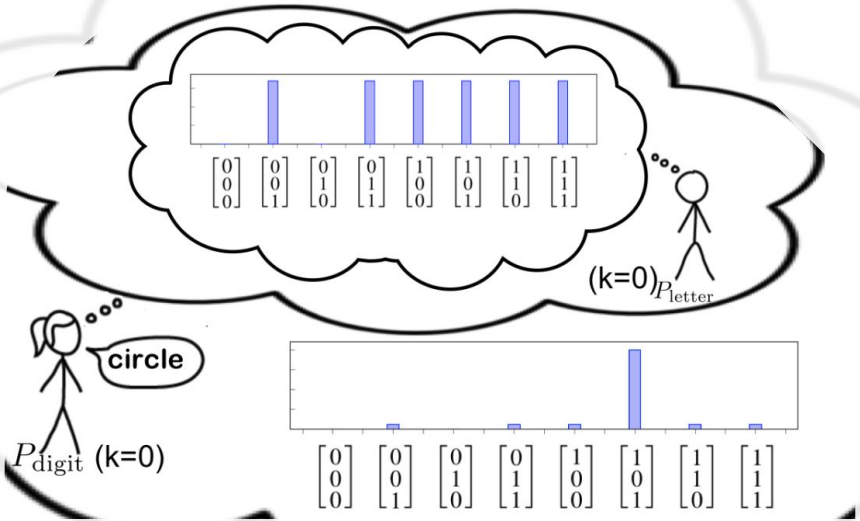

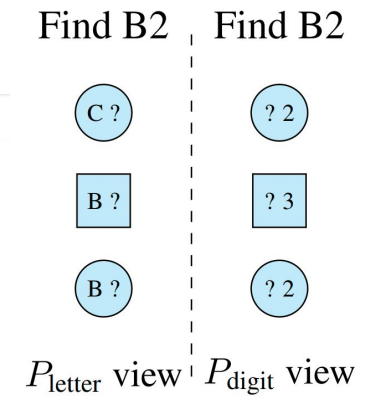

P_{letter}: `square`

P_{digit}: `circle`

Pragmatics: example

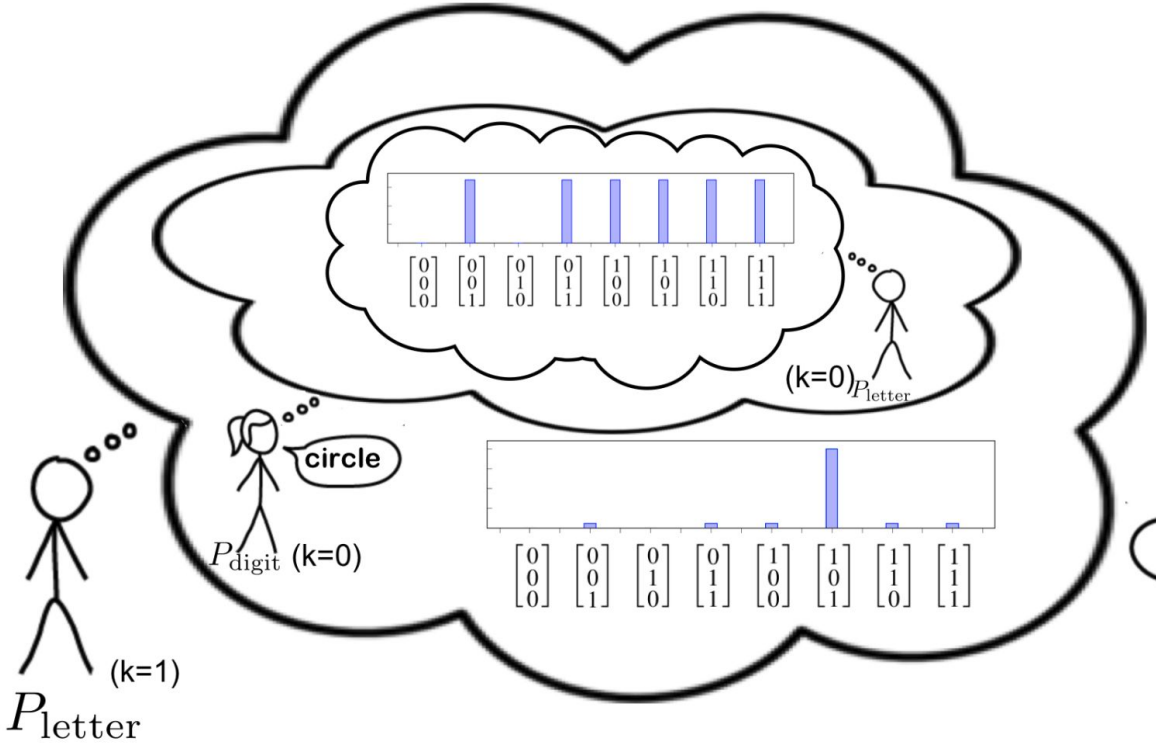

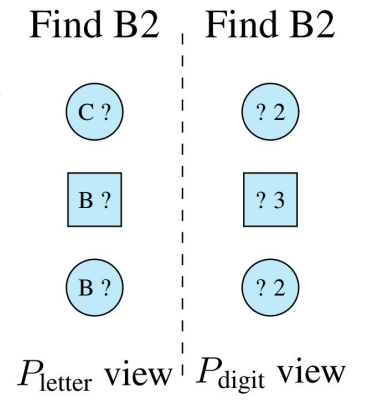

P_{letter}: `square`

P_{digit}: `circle`

Pragmatics: example

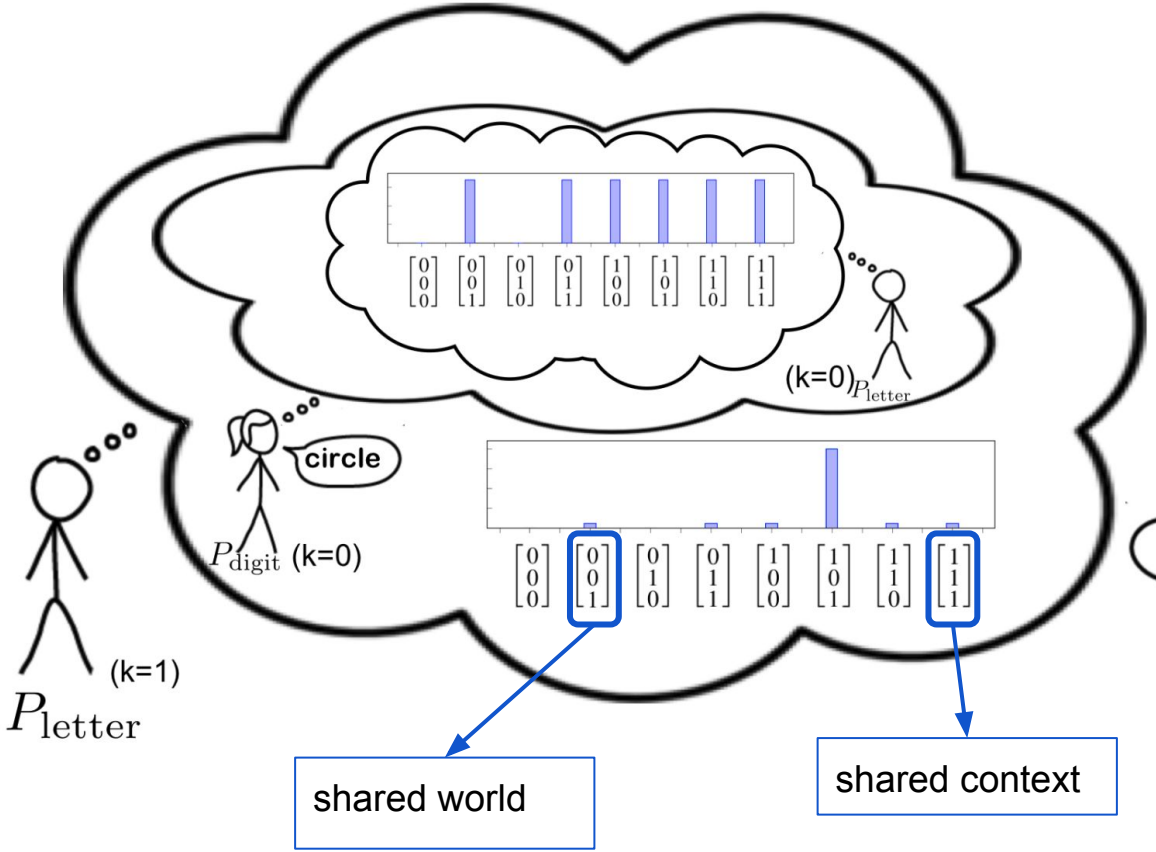

Find B2 | Find B2

C ? | ? 2

B ? | ? 3

B ? | ? 2

P_{letter} view | P_{digit} view

P_{letter}: `square`

P_{digit}: `circle`

PIP infers rich meaning

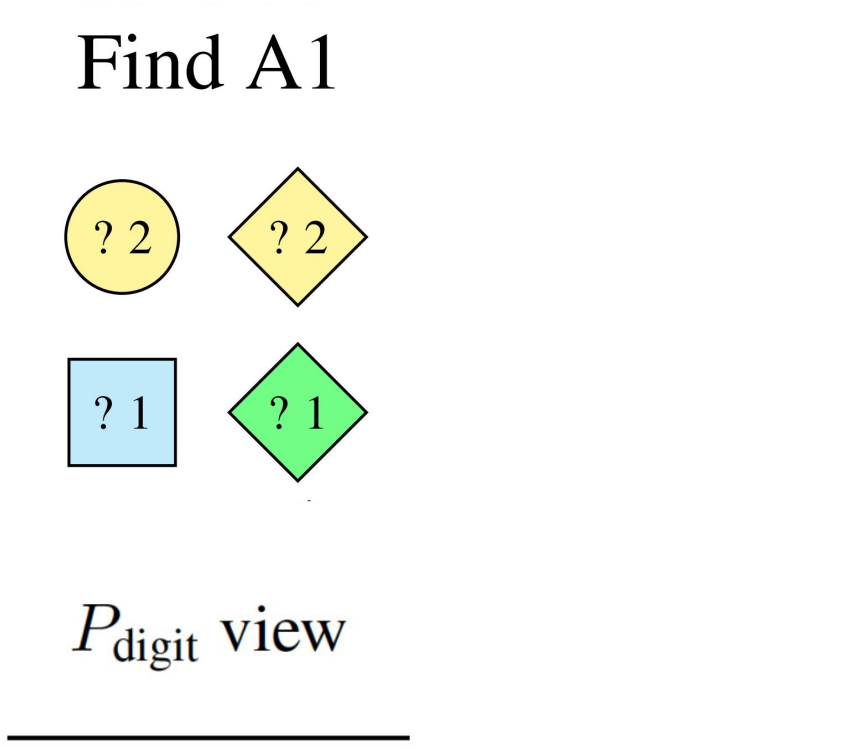

PIP infers rich meaning

Find A1

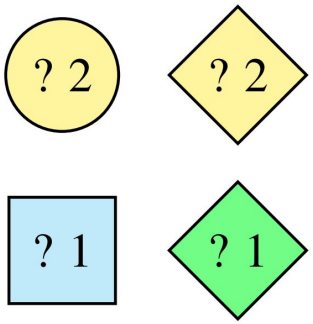

P_{digit} view

What does P_{digit} think about P_{letter} state?

PIP infers rich meaning

$$B(s_1, a_{1:t}) = \sum_{s_{-1}} p(s_{-1} \mid s_1, a_{1:t}) s_{-1}$$

PIP infers rich meaning

A matrix with 1 in any goal-consistent element

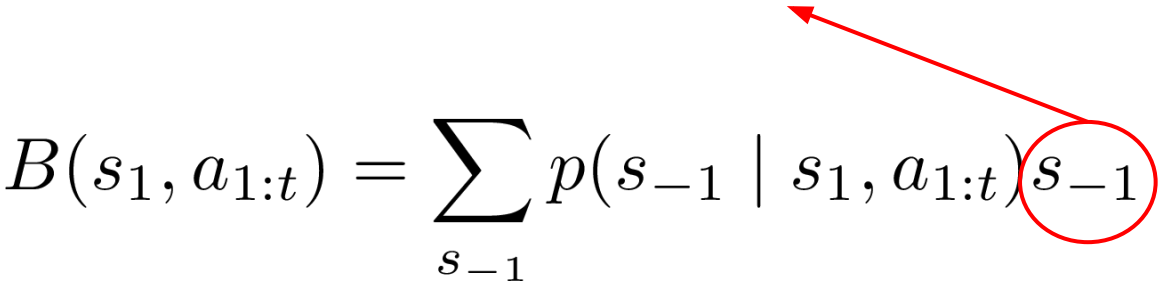

PIP infers rich meaning

A matrix with 1 in any goal-consistent element

$$B(s_1, a_{1:t}) = \sum_{s_{-1}} p(s_{-1} \mid s_1, a_{1:t}) s_{-1}$$

The probability of this entry having the goal letter

P_{digit} estimation of P_{letter} state

$$\begin{bmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{bmatrix}$$

PIP infers rich meaning

Find A1

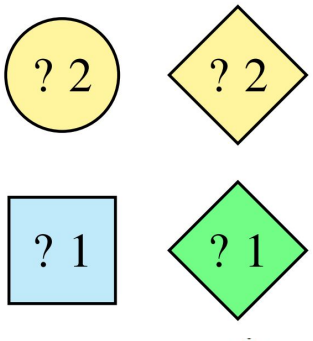

P_{digit} view

P_{digit} estimation of P_{letter} state

P_{letter}: `right` (k=0)

$$\begin{bmatrix} 0.500 & 0.667 \\ 0.500 & 0.667 \end{bmatrix}$$

P_{letter}: `right` (k=1)

$$\begin{bmatrix} 0.422 & 0.805 \\ 0.422 & 0.805 \end{bmatrix}$$

PIP infers rich meaning

Find A1

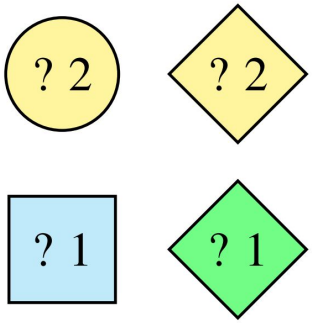

P_{digit} view

P_{digit} estimation of P_{letter} state

P_{digit}: `bottom`
P_{letter}: `right`
(k=0)

$$\begin{bmatrix} 0.500 & 0.667 \\ 0.500 & 0.667 \end{bmatrix}$$

P_{digit}: `bottom`
P_{letter}: `right`
(k=1)

$$\begin{bmatrix} 0.424 & 0.769 \\ 0.423 & 0.940 \end{bmatrix}$$

Outline

- ✓ Literal semantics
- ✓ Inference
- ✓ Planning
- ✓ Pragmatics
- ✓ All three together

Outline

- ✓ Literal semantics
- ✓ Inference
- ✓ Planning
- ✓ Pragmatics
- ✓ All three together

Experiments: setup

Random

Randomly chooses one of the semantically valid actions

Experiments: setup

Random

Randomly chooses one of the semantically valid actions

Greedy

Assigns higher probability to the actions that convey more information

Experiments: setup

Random

Randomly chooses one of the semantically valid actions

Greedy

Assigns higher probability to the actions that convey more information

PIP

Computing probability of actions using PIP

Experiments: results

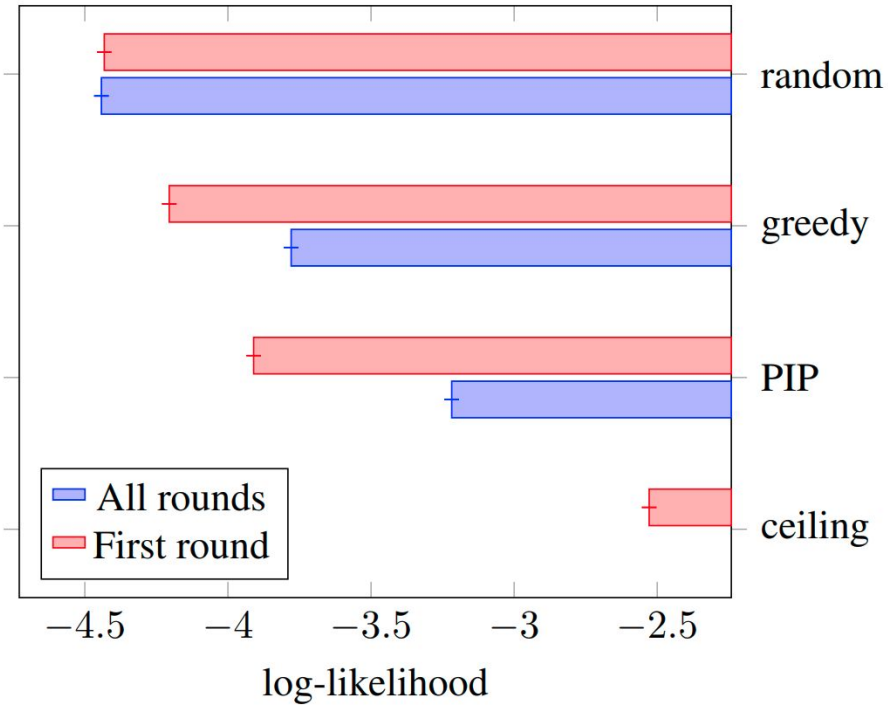

Experiments: results

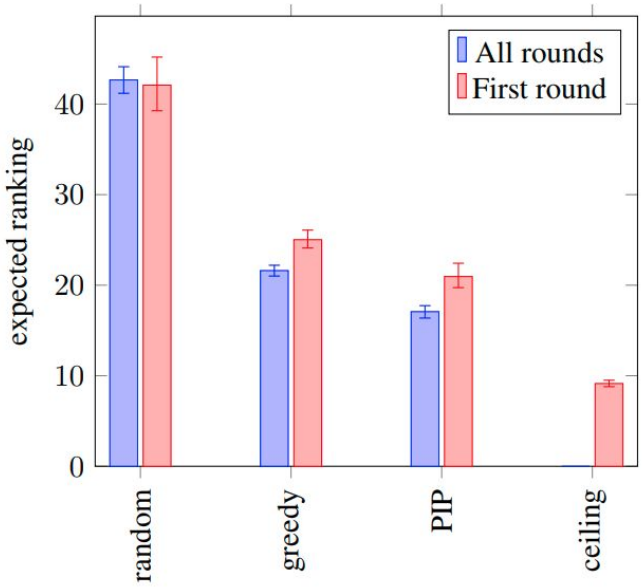

Expected ranking of human messages

Experiments: setup

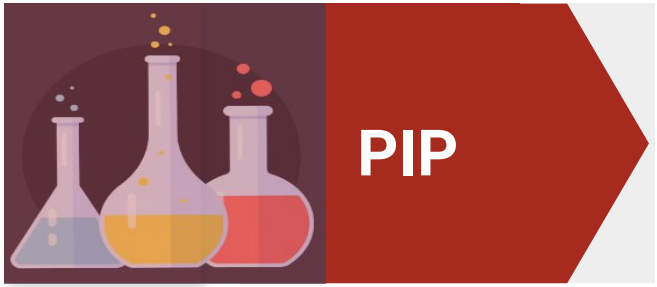

	PIP	PIP$_{\text{-prag}}$	PIP$_{\text{-plan}}$	PIP$_{\text{-infer}}$
k (pragmatics)	1	0	1	1
f (planning)	2	2	1	2
b (inference)	∞	∞	∞	1

Experiments: results

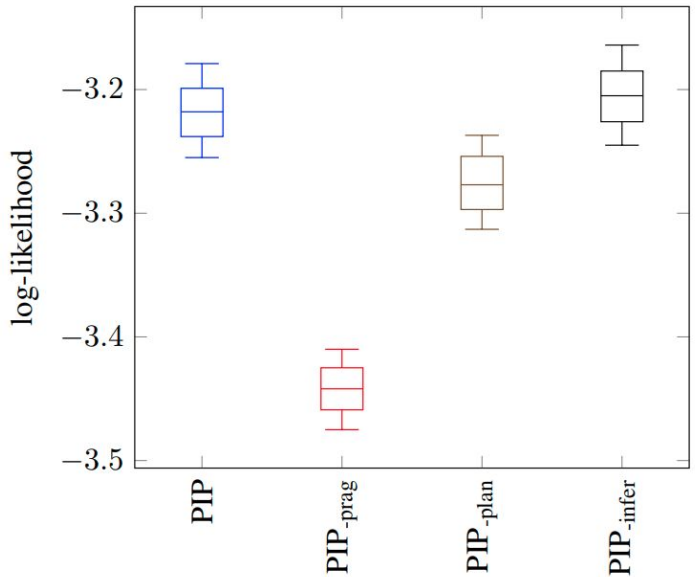

(a) Performance over all games and all rounds.

Experiments: results

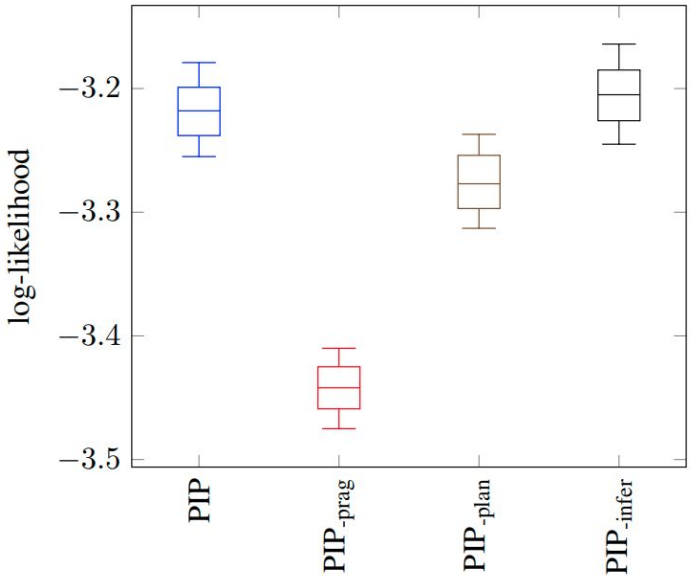

(a) Performance over all games and all rounds.

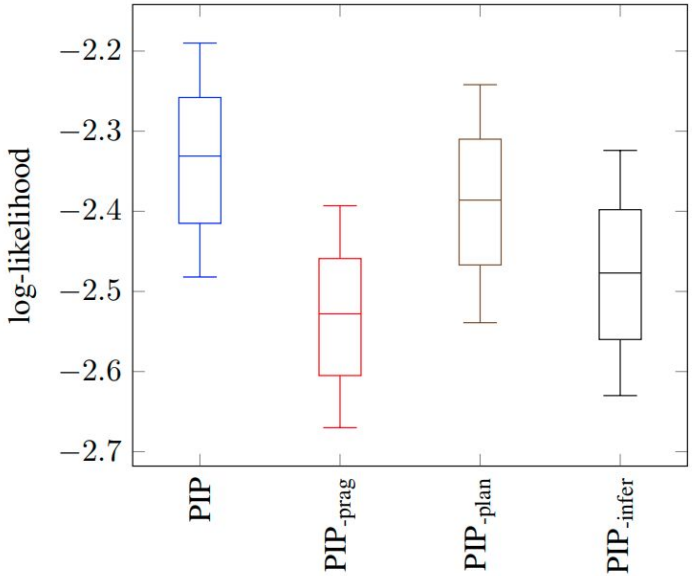

(b) Performance over messages after round 3.

Conclusion

Dataset We collect a dataset containing 1.7K games and 5k messages

Conclusion

Dataset

We collect a dataset containing 1.7K games and 5k messages

Model

1) Having all three (PIP) in one single unified model

2) Supports multiple round of communications

Conclusion

Dataset

We collect a dataset containing 1.7K games and 5k messages

Model

1) Having all three (PIP) in one single unified model
2) Supports multiple round of communications

Results

1) PIP is able to capture human behavior in InfoJigsaw
2) A very simple, context-independent literal semantics can give rise via PIP to rich phenomena.

Planning

Markov Decision Process and their extension (e.g., Vogel et al. (2013))

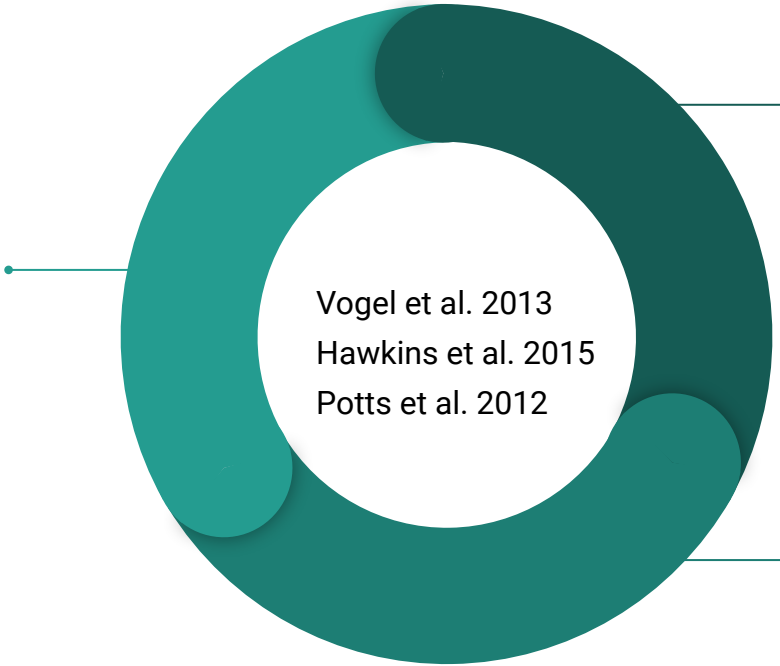

Inference

Model theoretic semantics (Montague, 1973, Matuszek et al., 2012)

Pragmatics

Cooperative principles of Grice (1975) can be realized (e.g., Franke (2009), Frank and Goodman (2012))