CS11-711 Advanced NLP

Pragmatics

Daniel Fried and Saujas Vaduguru

Carnegie Mellon University

Language Technologies Institute

Site

https://cmu-anlp.github.io

With slides from Nick Tomlin

Using Language in Context

Text

Language Modeling, Semantics, QA…

How do I view whether a job was run in interactive mode in Slurm?

You can use the sacct command with the –o option to specify the fields you want to display, and the –j option to specify the job ID. For example…

This Lecture

Physical World

Grounding

“Take me to the airport”

Intents and Effects

Pragmatics

“My neck hurts”

Using Language in Context

Grounding

“Stop at the second car”

Pragmatics

“Stop at the car”

Language is Contextual

- Some knowledge seems easier to get with grounding:

Winograd schemas

*The large ball crashed right through the table because it was made of **steel**. What was made of steel?*

-> **ball**

*The large ball crashed right through the table because it was made of **styrofoam**. What was made of styrofoam?*

-> **table**

Winograd 1972; Levesque 2013; Wang et al. 2018

"blinking and breathing problem"

Word	*Teraword*	*Knext*	*Word*	*Teraword*	*Knext*
spoke	11,577,917	244,458	hugged	610,040	10,378
laughed	3,904,519	169,347	blinked	390,692	20,624
murdered	2,843,529	11,284	was late	368,922	31,168
inhaled	984,613	4,412	exhaled	168,985	3,490
breathed	725,034	34,912	was punctual	5,045	511

Table 1: Frequencies from [3] and the number of times Knext learns that *A person may* ⟨*x*⟩, including appropriate arguments, e.g., *A person may hug a person*. For *murder*, more frequently encountered in the passive, we include *be murdered*.

Gordon and Van Durme, 2013

Language is Contextual

- Some problems depend on grounding references to context
- Indexicals and *Deixis*: “pointing or indicating”(e.g. pronouns, “this”, “that”, “here”, “now”)
  - *I am speaking*
  - *We won* (a team I’m on, OR a team I support)
  - *He had rich taste* (walking through the Taj Mahal)

  - *I am here* (at CMU; in this classroom)
  - *We are here* (pointing to a map)

  - *I’m in a class now*
  - *I’m in a graduate program now*
  - *I’m not here right now* (voicemail greeting)

Language is Contextual

- Some problems depend on grounding into speaker intents or goals:
  - “Can you pass me the salt”

    -> please pass me the salt
  - “Do you have any kombucha?” // “I have tea”

    -> I don’t have any kombucha
  - “You’re fired!”

    -> *performative*, that changes the state of the world

Using Language in Context

Saying something will often… produce certain consequential effects upon the feelings, thoughts, or actions of the audience.

[*How to Do Things with Words.* Austin, 1962]

Our talk exchanges … are cooperative efforts… One of my avowed aims is to see talking as purposive, indeed rational, behavior.

[*Logic and Conversation.* Grice, 1975]

Language is an act people take to produce effects on others and the world!

Reference Games

Generation

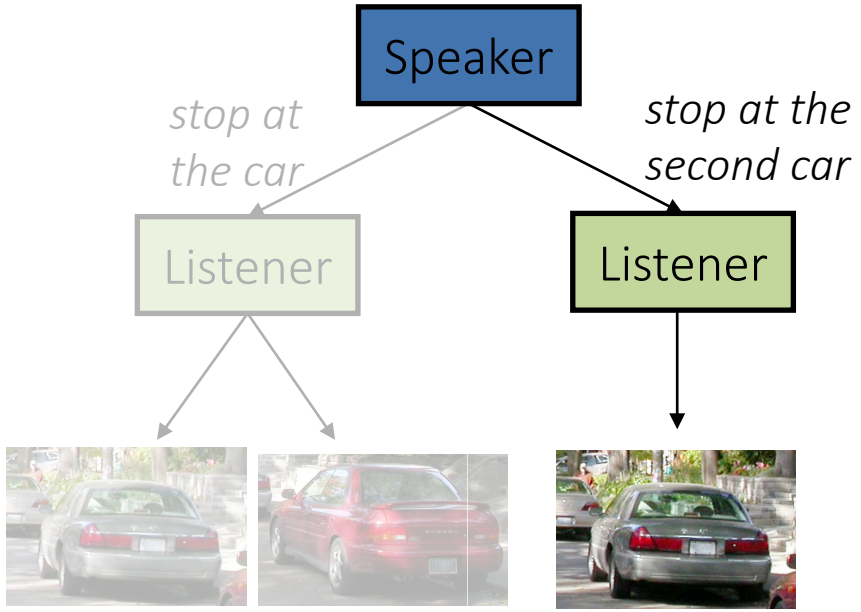

Interpretation

[e.g. Lewis 1969; Golland et al. 2010;
Frank and Goodman 2012; Degen et al. 2013]

Reference Games

Generation

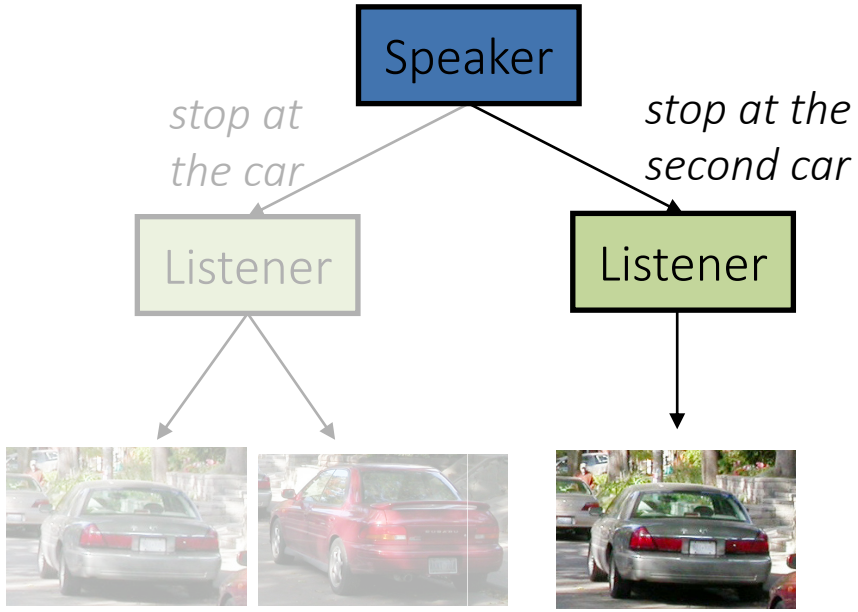

Interpretation

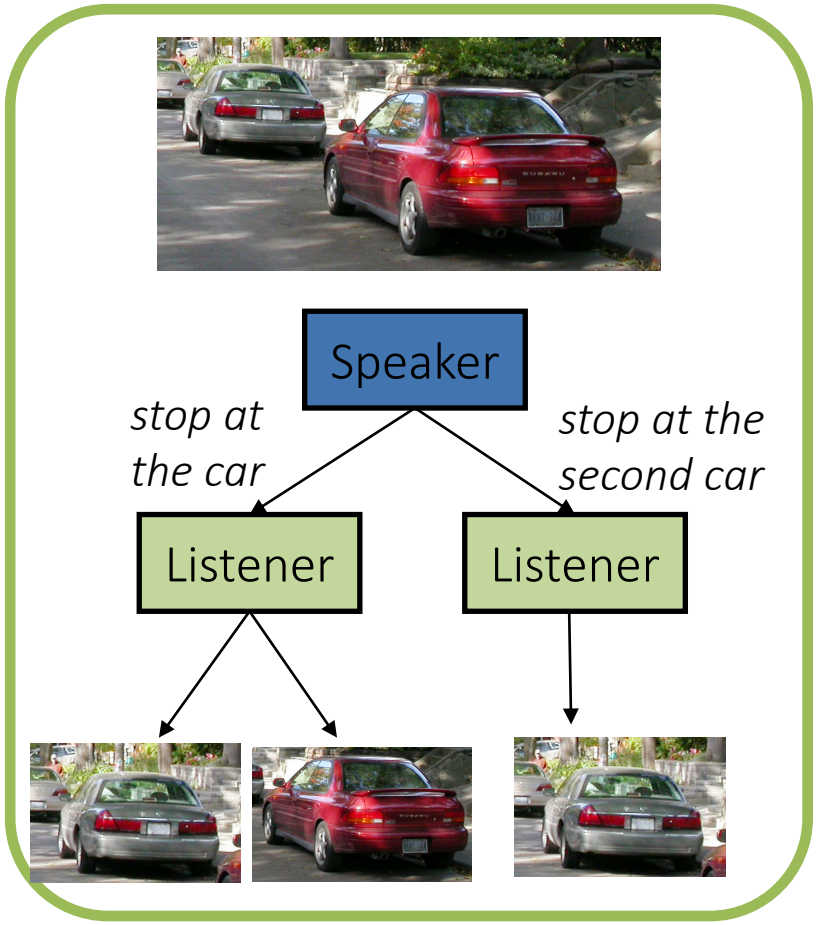

[e.g. Lewis 1969; Golland et al. 2010; **Frank and Goodman 2012; Degen et al. 2013**]

Reference Games

Generation

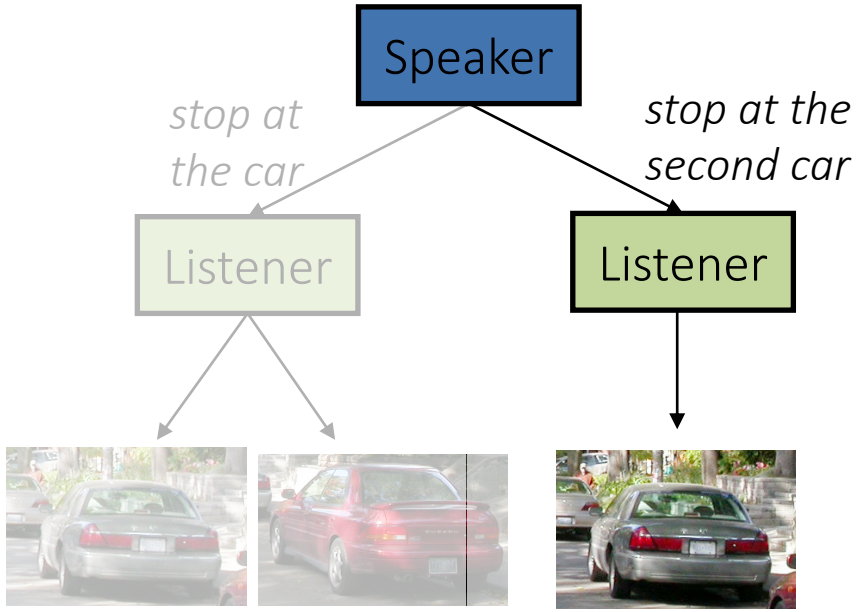

Interpretation

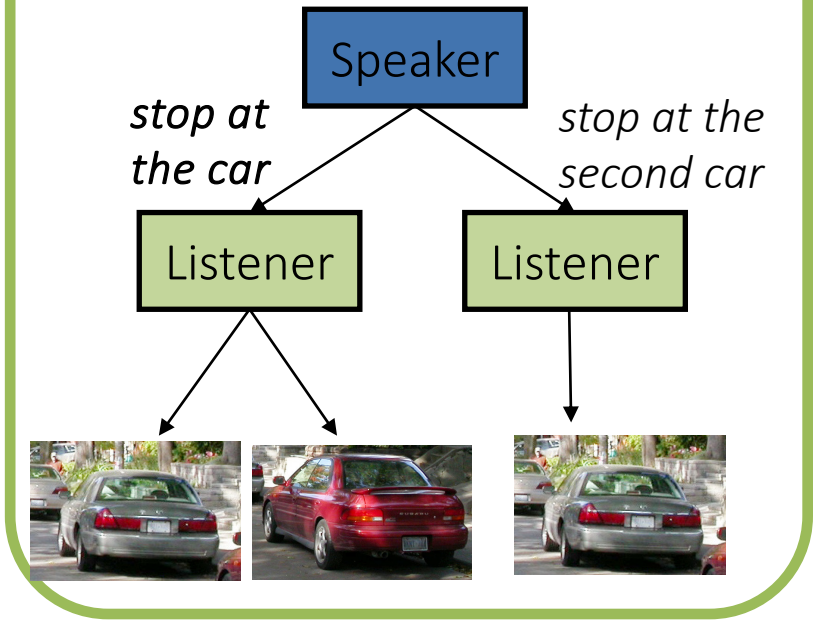

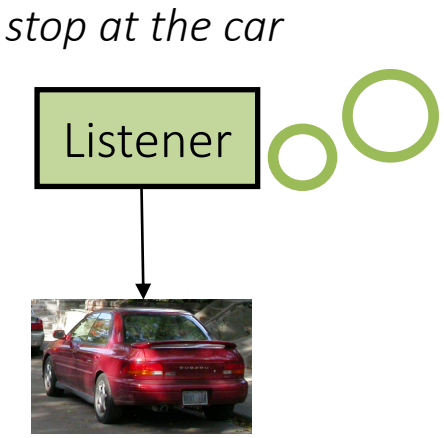

[e.g. Lewis 1969; Golland et al. 2010; Frank and Goodman 2012; Degen et al. 2013]

Reasoning About Alternatives

Core Idea:

Large chunks of linguistic understanding can be attributed to reasoning about alternatives. E.g., if a speaker says X but not Y, then perhaps Y isn't true, or the speaker doesn't want to talk about Y.

Reasoning About Alternatives

Core Idea:

Large chunks of linguistic understanding can be attributed to reasoning about alternatives. E.g., if a speaker says X but not Y, then perhaps Y isn't true, or the speaker doesn't want to talk about Y.

Example:

"I didn't steal your car."

Reasoning About Alternatives

Core Idea:

Large chunks of linguistic understanding can be attributed to reasoning about alternatives. E.g., if a speaker says X but not Y, then perhaps Y isn't true, or the speaker doesn't want to talk about Y.

Example:

"I didn't steal your car."

Conveyed meaning:

Someone stole your car, but it wasn't me.

Reasoning About Alternatives

Core Idea:

Large chunks of linguistic understanding can be attributed to reasoning about alternatives. E.g., if a speaker says X but not Y, then perhaps Y isn't true, or the speaker doesn't want to talk about Y.

Example:

"I didn't steal your car."

Conveyed meaning:

Contrary to what you think, I did not steal your car.

Reasoning About Alternatives

Core Idea:

Large chunks of linguistic understanding can be attributed to reasoning about alternatives. E.g., if a speaker says X but not Y, then perhaps Y isn't true, or the speaker doesn't want to talk about Y.

Example:

"I didn't steal your car."

Conveyed meaning:

I did something to your car, but not stealing it. E.g., I just borrowed it.

Reasoning About Alternatives

Core Idea:

Large chunks of linguistic understanding can be attributed to reasoning about alternatives. E.g., if a speaker says X but not Y, then perhaps Y isn't true, or the speaker doesn't want to talk about Y.

Example:

"I didn't steal <u>your</u> car."

Conveyed meaning:

I stole somebody else's car.

Reasoning About Alternatives

Core Idea:

Large chunks of linguistic understanding can be attributed to reasoning about alternatives. E.g., if a speaker says X but not Y, then perhaps Y isn't true, or the speaker doesn't want to talk about Y.

Example:

"I didn't steal your car."

Conveyed meaning:

I stole something you own, but not your car.

Implicatures

Implicatures

Q: Does *some* mean *not all*?

A: Not always:

- "Some of the students were late for class; in fact, they all were."
- "I'd be much happier if some grocery stores had eggs in stock."

We call this *implicature*. The implicature occurs because a rational listener might assume that the speaker would have said *all* if they meant to, since *all* is the more informative choice.

Implicatures

"The car was stolen."

- *The speaker doesn't know, or doesn't want to tell, who stole it.*

"Did you invite Alice and Bob?" // "I invited Alice."

- *The speaker didn't invite Bob.*

"I'm out of gas." // "There's a station round the corner."

- *You can get gas there (e.g. it's open).*

"He overslept and failed the test."

- *Those events happened in that order.*

What Are People's Goals in Conversation?

Grice (1975) claims that many of these phenomena are explained by the tensions between the following ***maxims***:

1. **Quantity** – be as informative as possible, give as much information as needed, but no more. (*"The car was stolen."*)
2. **Quality** - be truthful, and don't give information that is false or unsupported by evidence. (*"Did you invite A and B?" // "I invited B."*)
3. **Relation** – be relevant, and say things that are pertinent to the discussion. (*"I'm out of gas" // "There's a station round the corner."*)
4. **Manner** – be clear, brief, and orderly as possible; avoid unnecessary prolixity. (*"He overslept and failed the test."*)

The Cooperative Principle

The Cooperative Principle (Grice 1975):

"Make your contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged."

Language is a rational action in a cooperative game.

Cooperative Principle via Game Theory

- **Best-response** [Franke 2009; Golland 2010; Jäger 2014]
- Recursive Bayesian agents
  - RSA [Frank and Goodman 2012, 2016]
  - Reward-rational implicit choice [Jeon et al. 2020]
- Other formalisms (info-theoretic):
  - Optimal transport of beliefs [Wang et al. 2020]
  - Rate distortion [Zaslavsky et al. 2020]

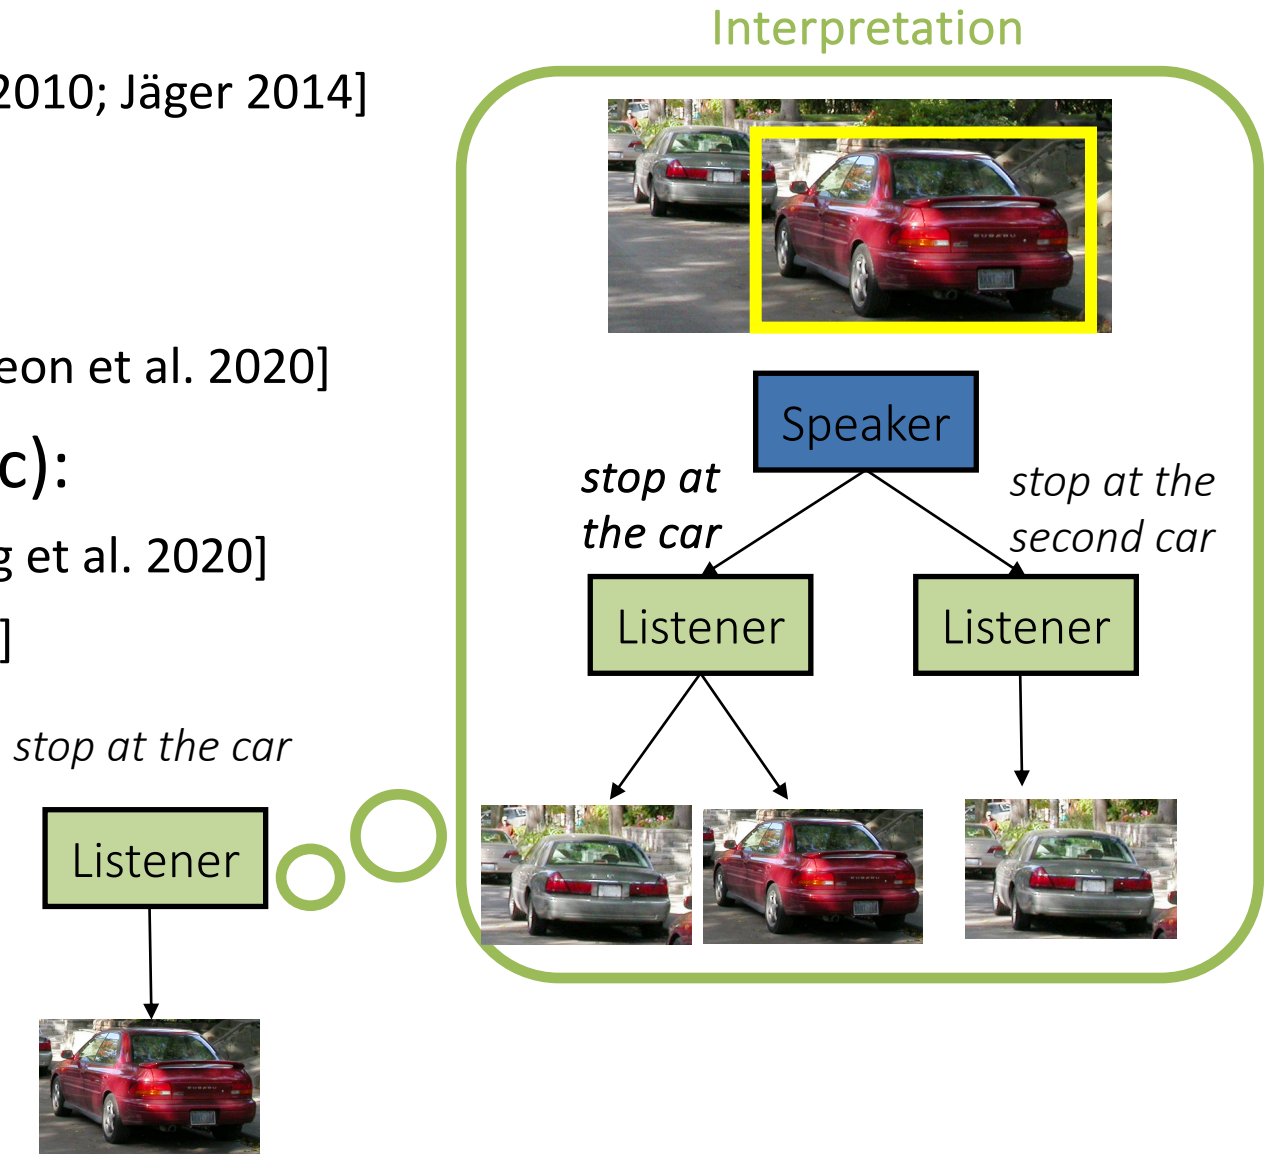

A Navigation Task

Human View:

Agent View:

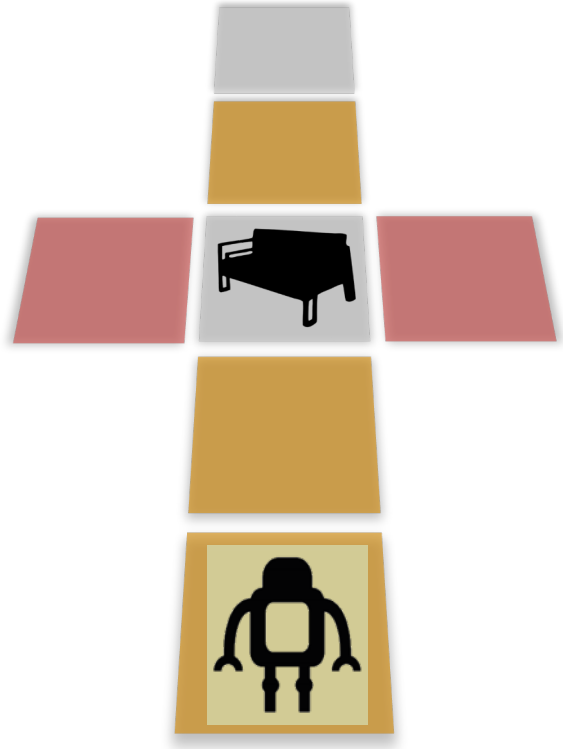

SAIL [MacMahon et al., 2006; Chen and Mooney, 2011]

Interpretation Task

Input instruction: *go forward to the grey hallway*

Output actions:

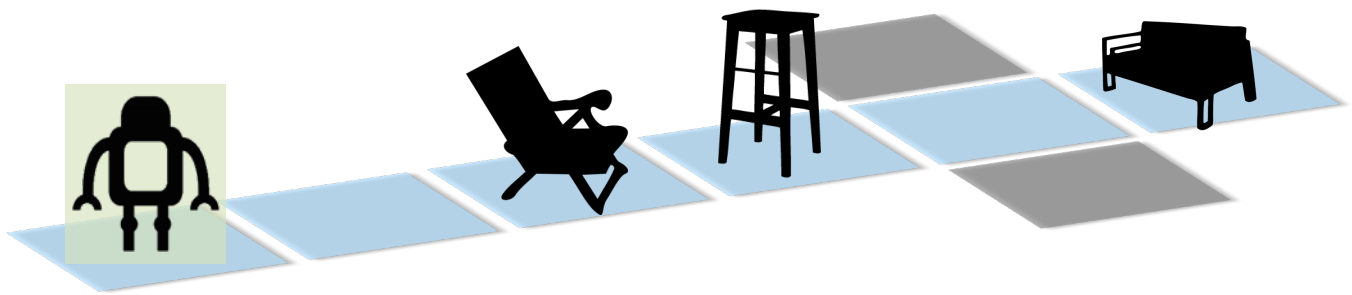

Interpretation Task

Input instruction: *go forward to the grey hallway*

Output actions:

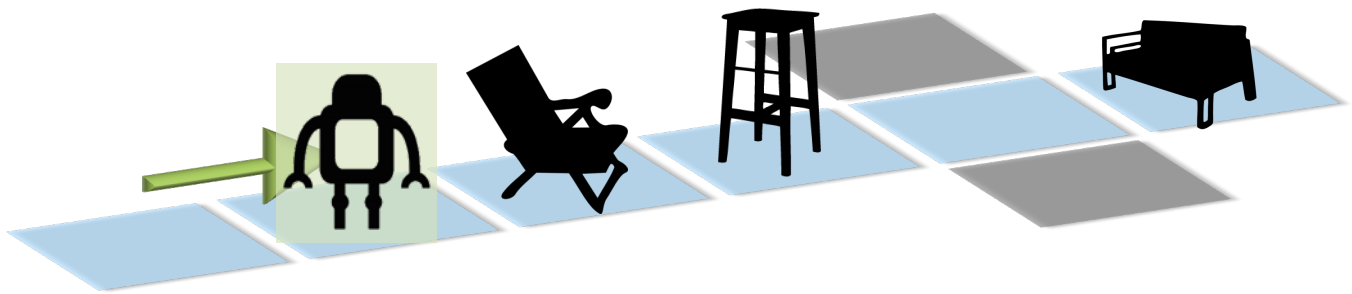

Interpretation Task

Input instruction: *go forward to the grey hallway*

Output actions:

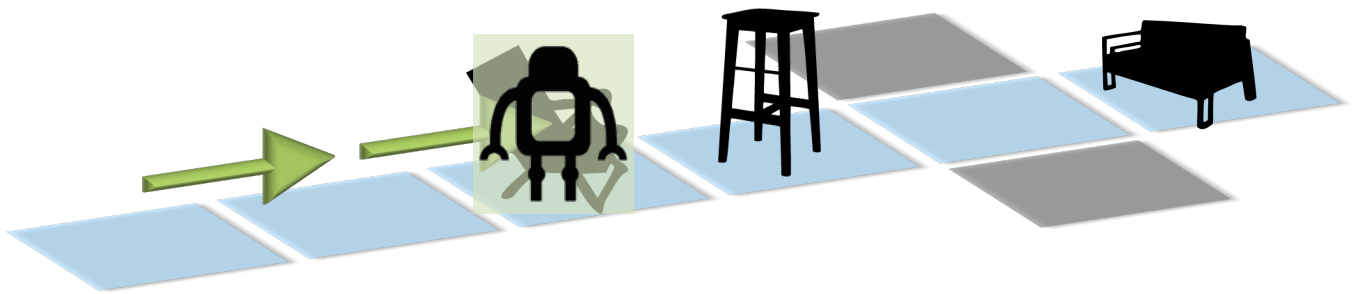

Interpretation Task

Input instruction: *go forward to the grey hallway*

Output actions:

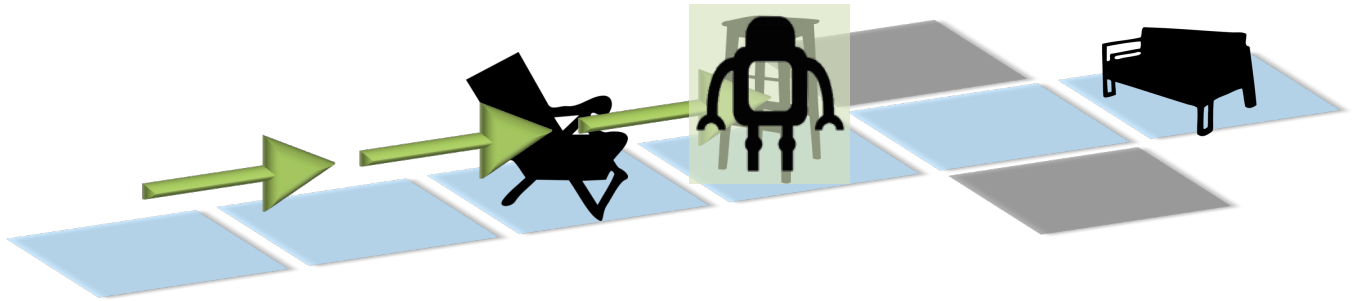

Interpretation Task

Input instruction: *go forward to the grey hallway*

Output actions:

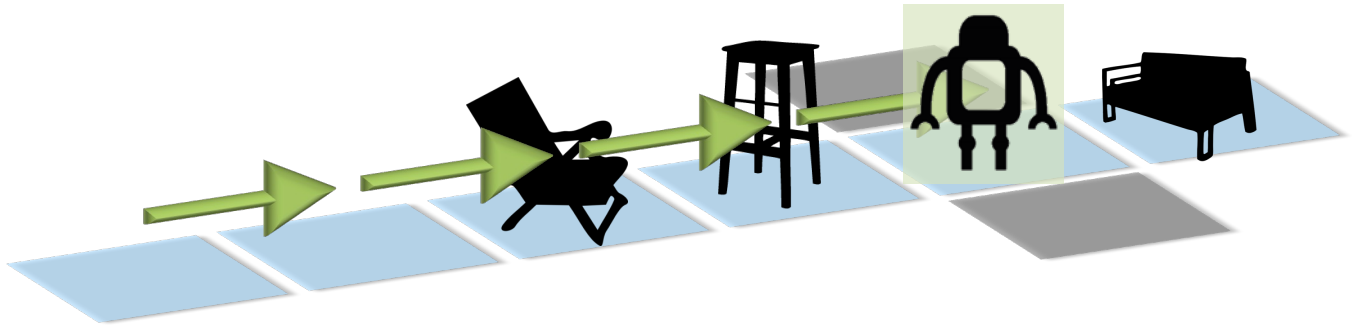

Interpretation Task

Input instruction: *go forward to the grey hallway*

Output actions:

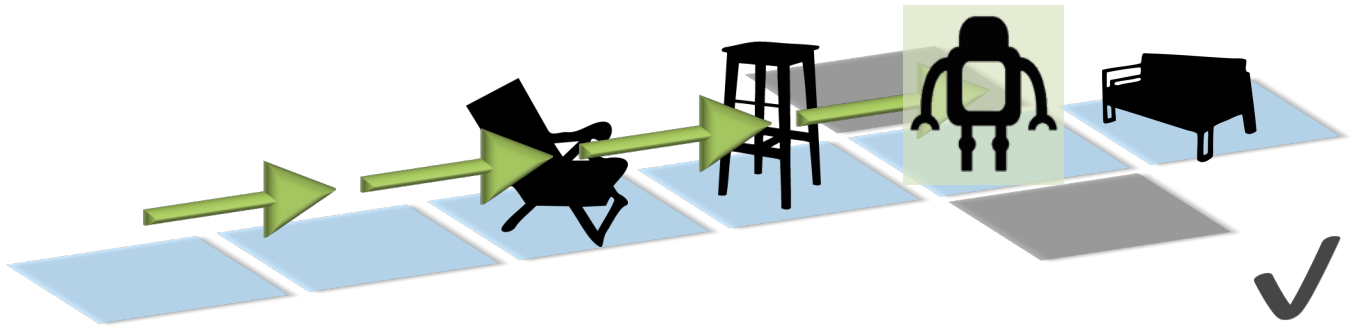

Generation Task

Input actions:

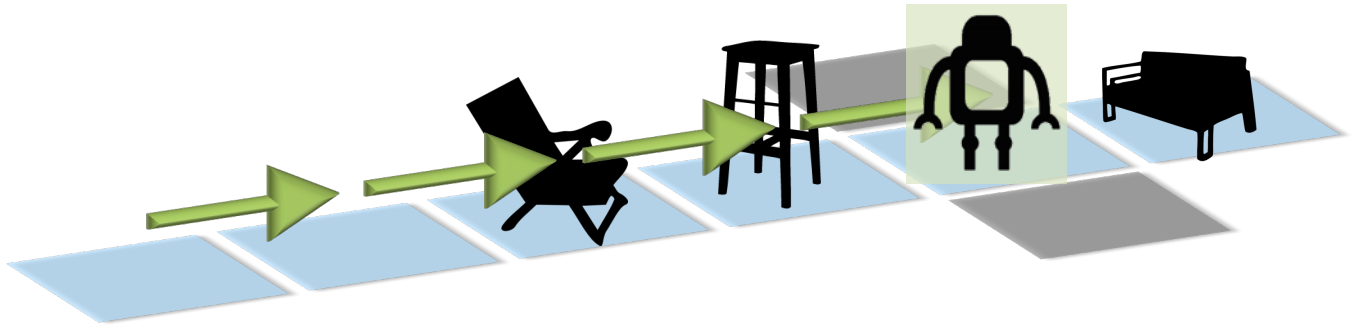

Output Instruction: *go forward to the grey hallway*

Models of Listeners and Speakers

Inputs

Outputs

go forward to the grey hallway

Instruction, i

Listener

$P_L(a \mid i)$

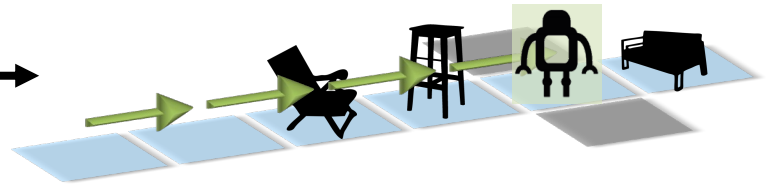

Actions, a

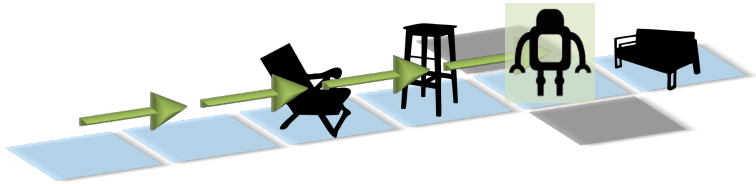

Actions, a

Speaker

$P_S(i \mid a)$

go forward to the grey hallway

Instruction, i

Base Models

Base Listener

$$P_L(a \mid i) = \prod_t P_L(a_t | a_{1:t-1}, i)$$

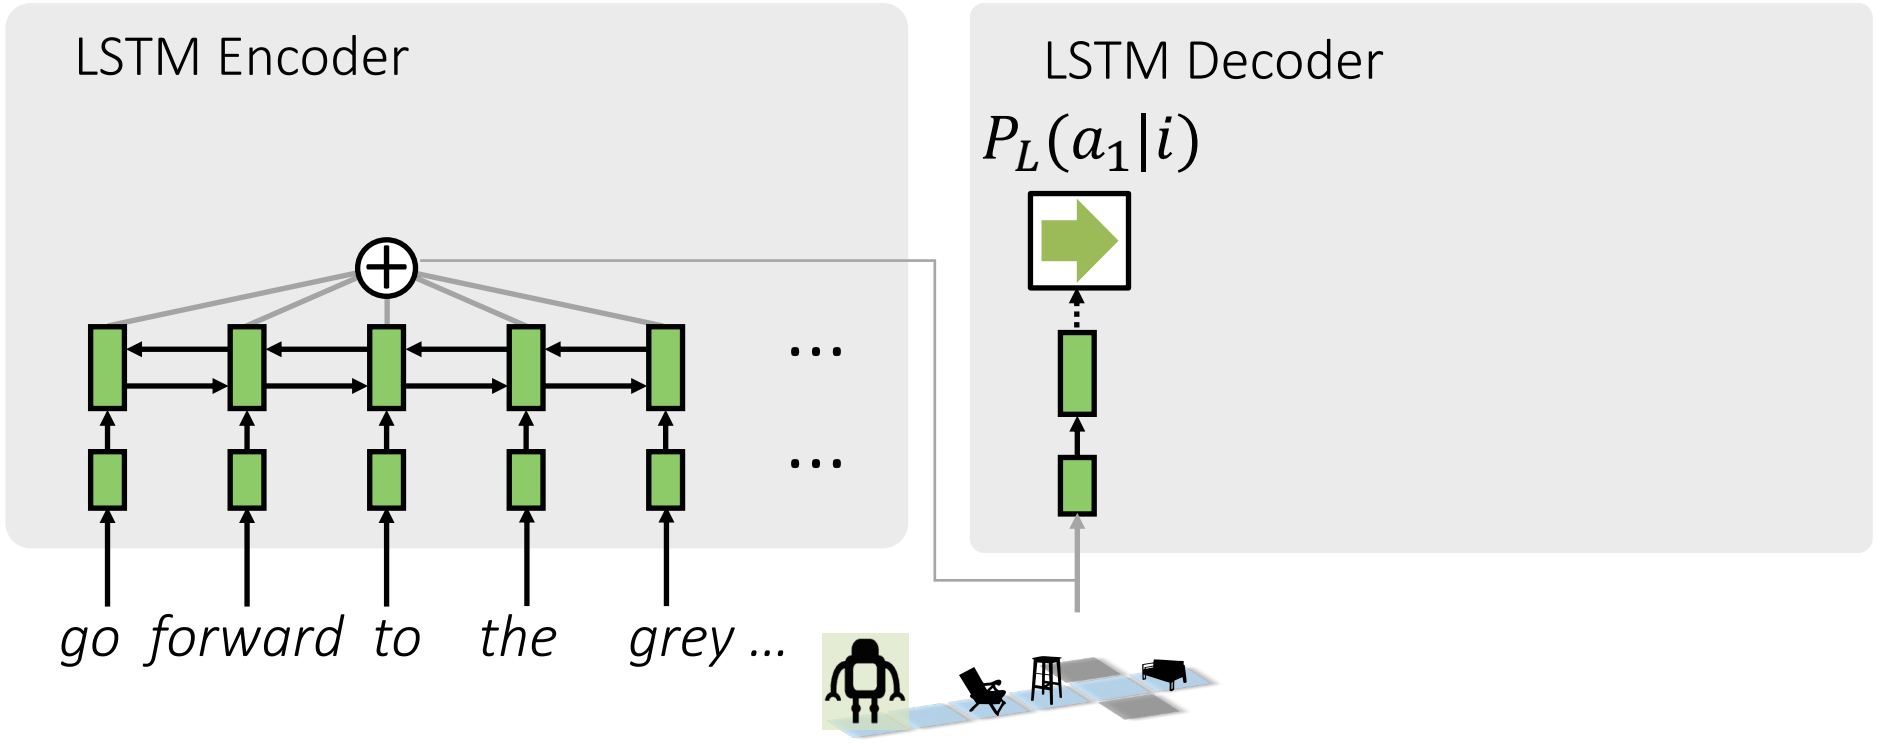

Based on Mei et al. [2016]

Base Models

Base Listener

$$P_L(a \mid i) = \prod_t P_L(a_t | a_{1:t-1}, i)$$

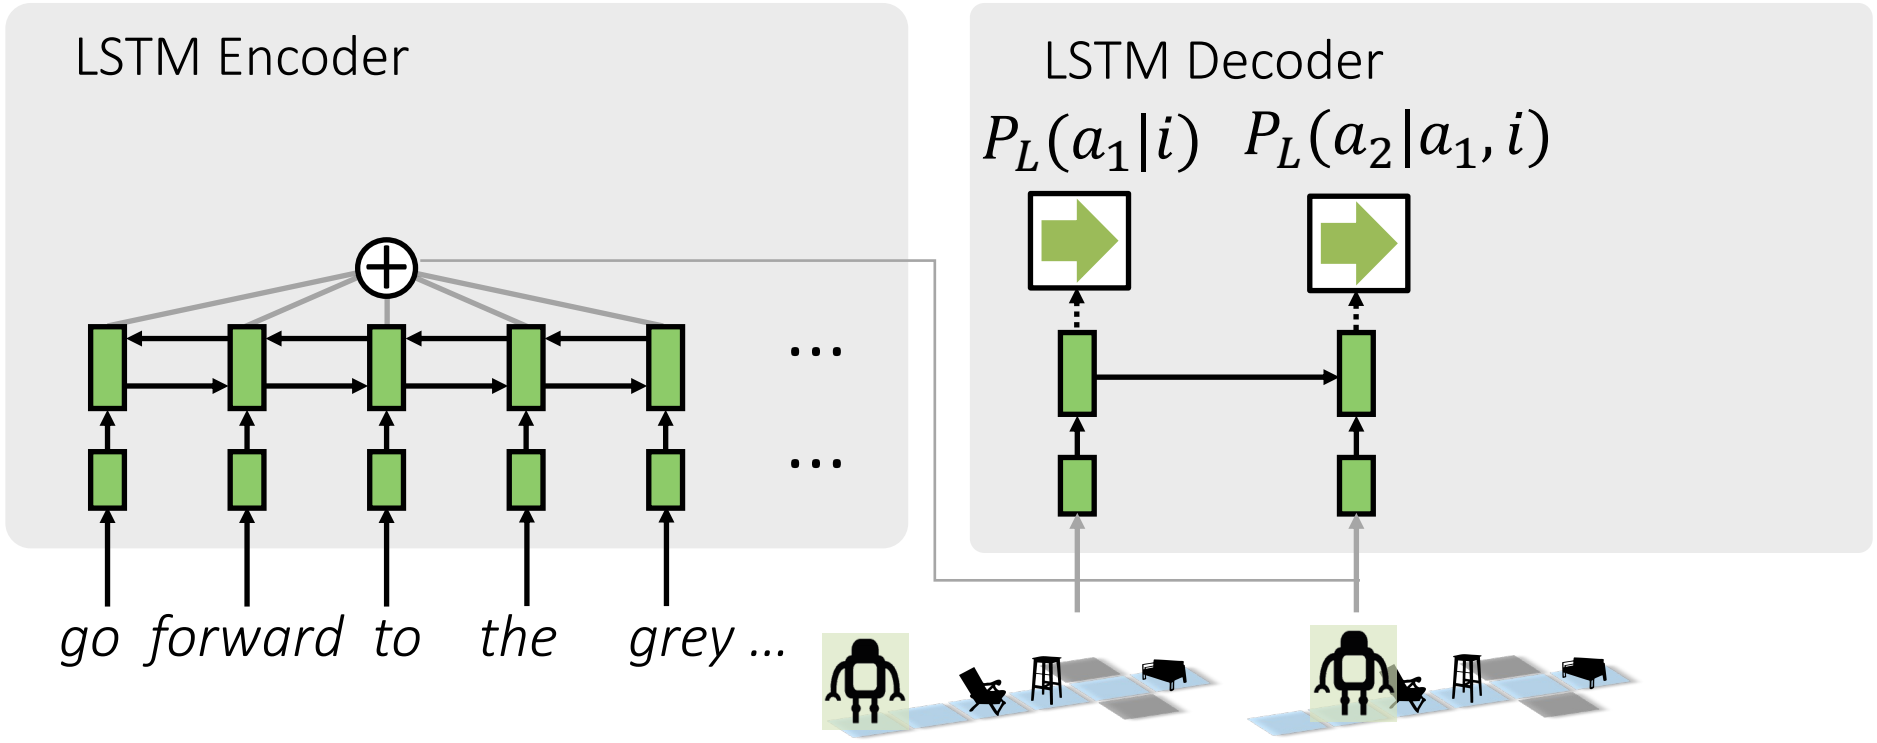

Based on Mei et al. [2016]

Base Models

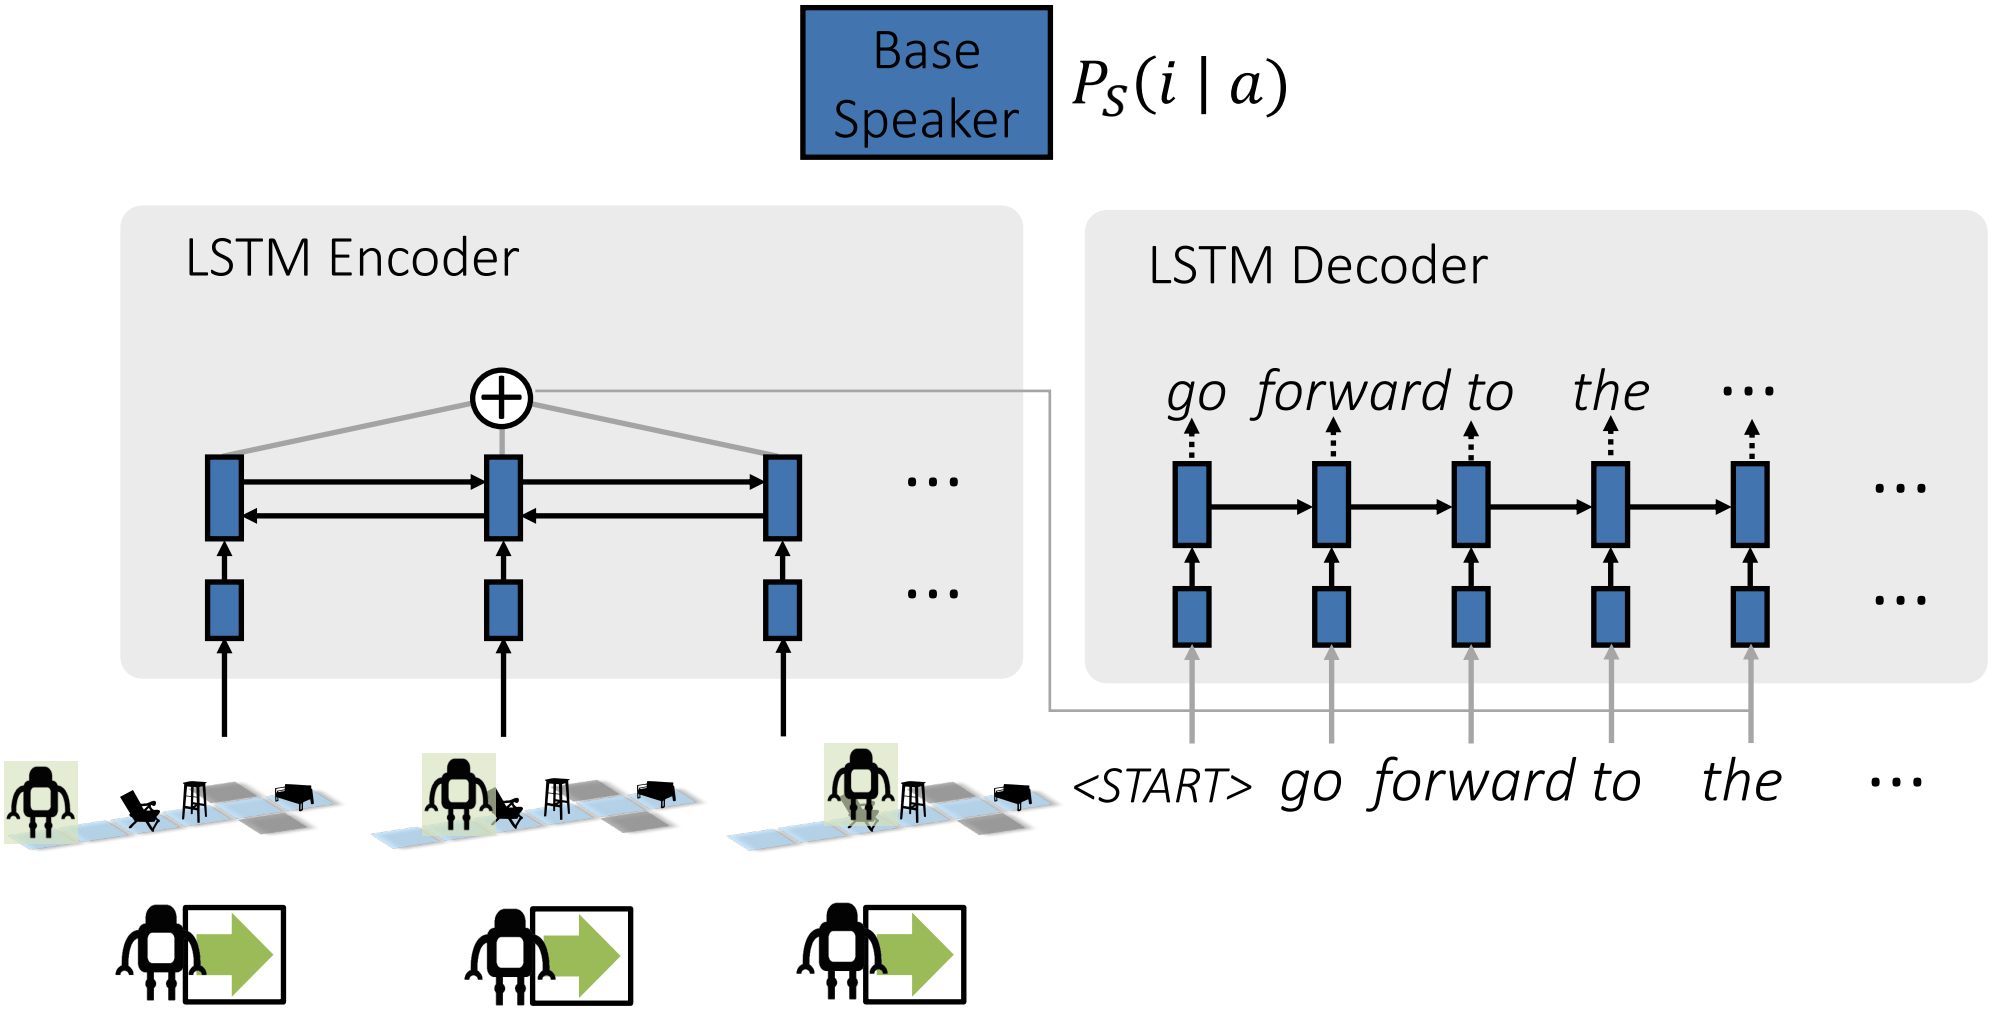
Base Speaker
$P_S(i \mid a)$
LSTM Encoder
LSTM Decoder
go forward to the ...
<START> go forward to the ...

Training Models on Human Instructions

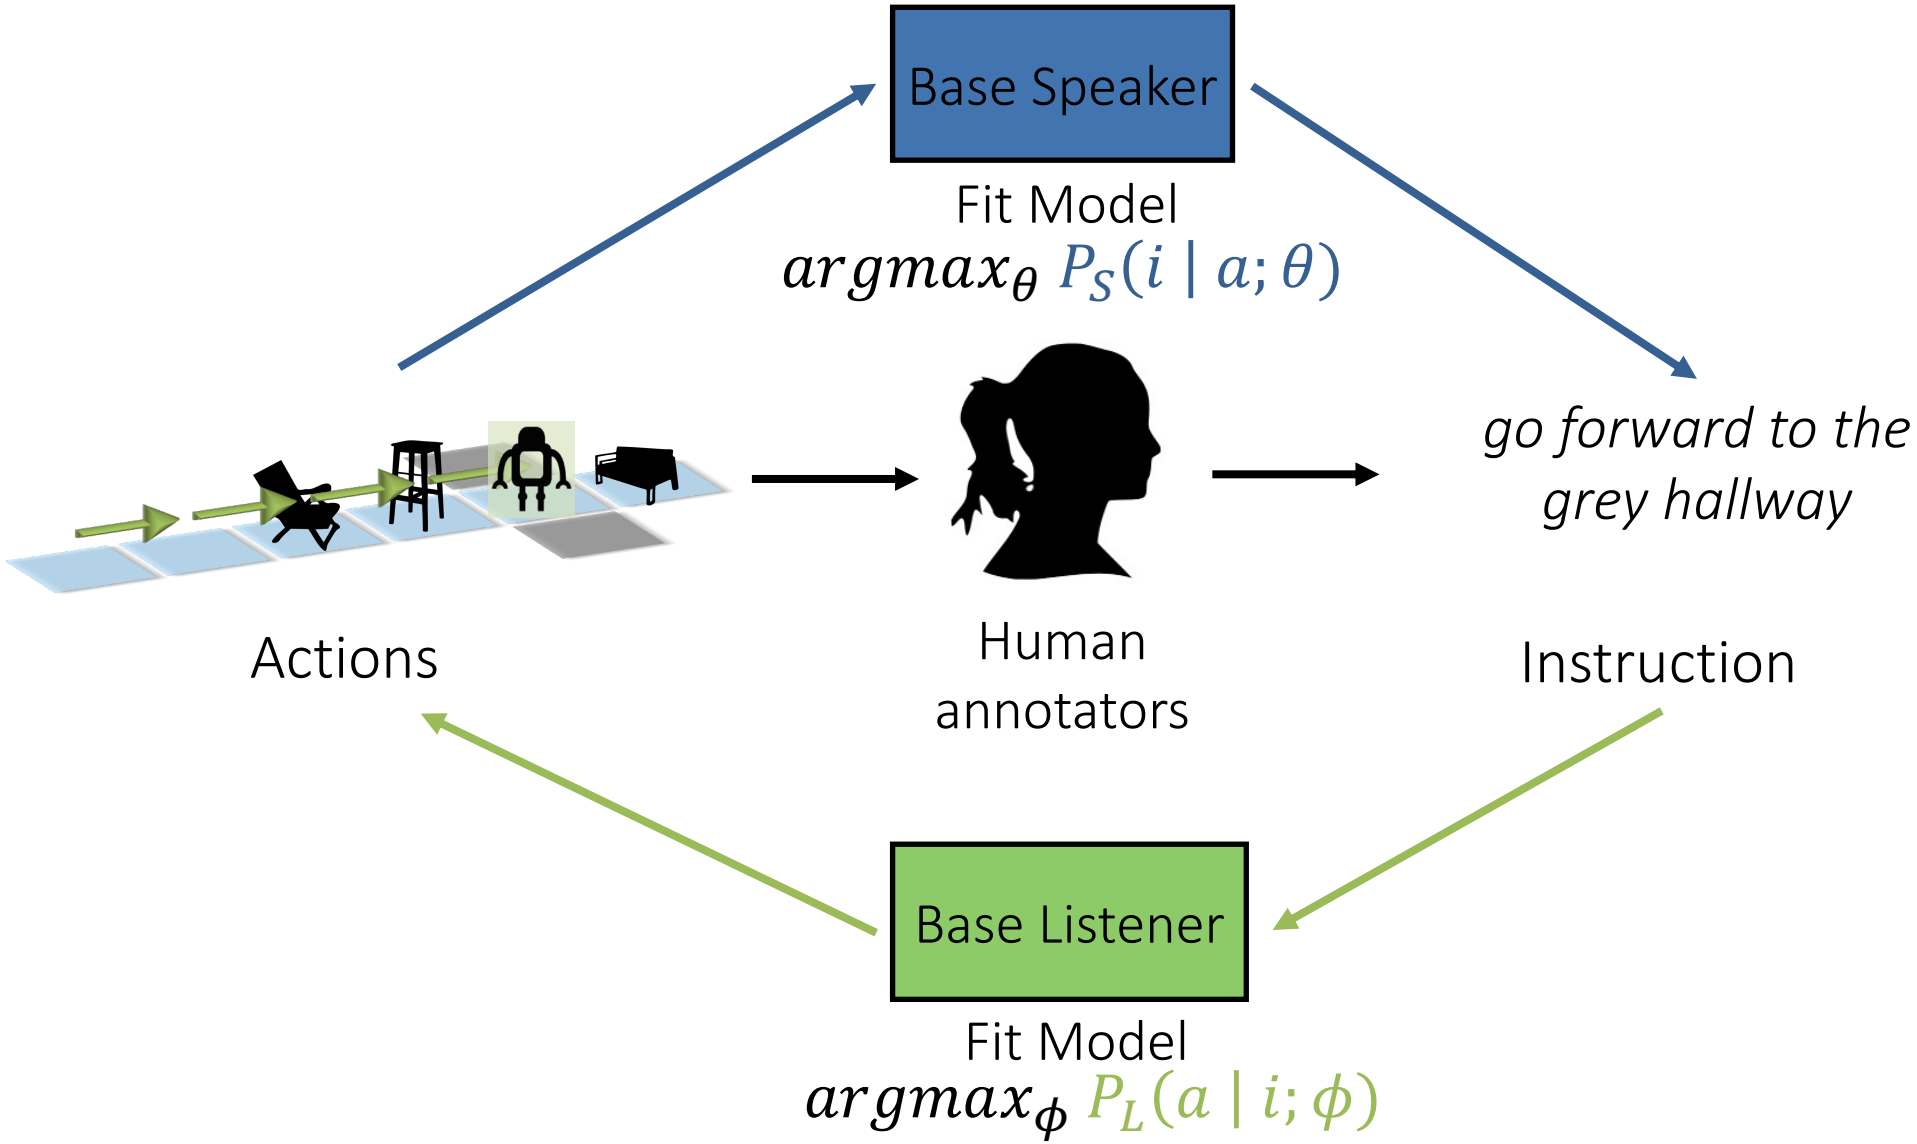

Speaker Tasks and Evaluation

Speaker produces an instruction

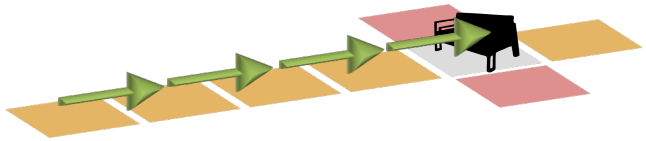

Speaker

walk along the wood path to the chair

Humans try to interpret it

walk along the wood path to the chair

Human direction followers (MTurk)

SAIL navigation [MacMahon et al., 2006; Chen and Mooney, 2011]

Speaker Tasks and Evaluation

Alchemy

1. remove all the purple chemical from the beaker on the far left
2. do the same with one unit of green chemical
3. …

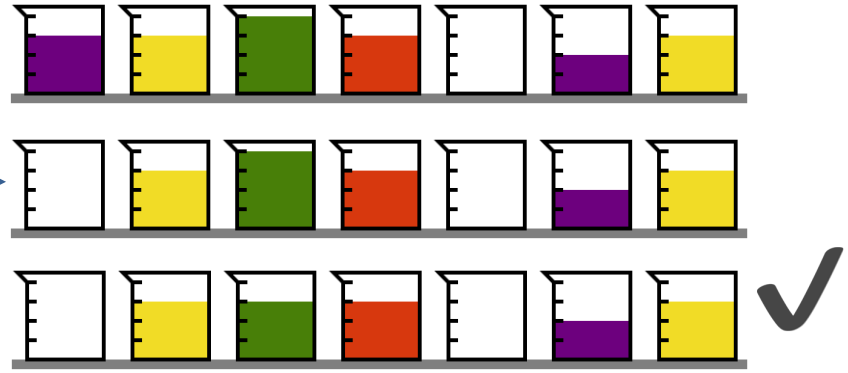

Human direction followers (MTurk)

Tangrams

1. remove first figure
2. add it back into middle spot
3. …

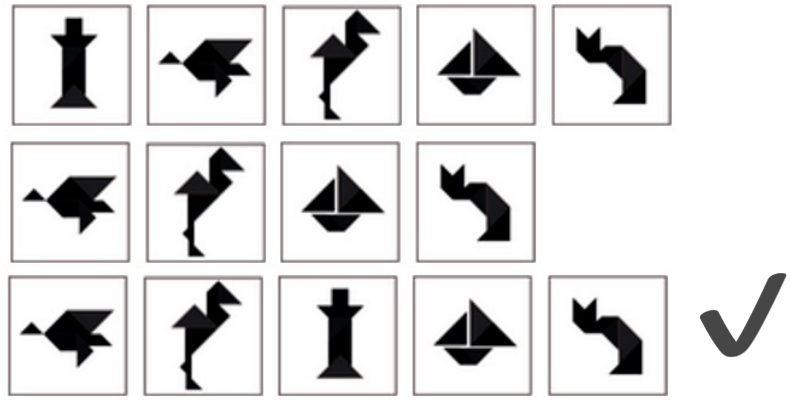

SCONE contextual instruction following [Long et al. 2016]

Generation is Hard to Imitate!

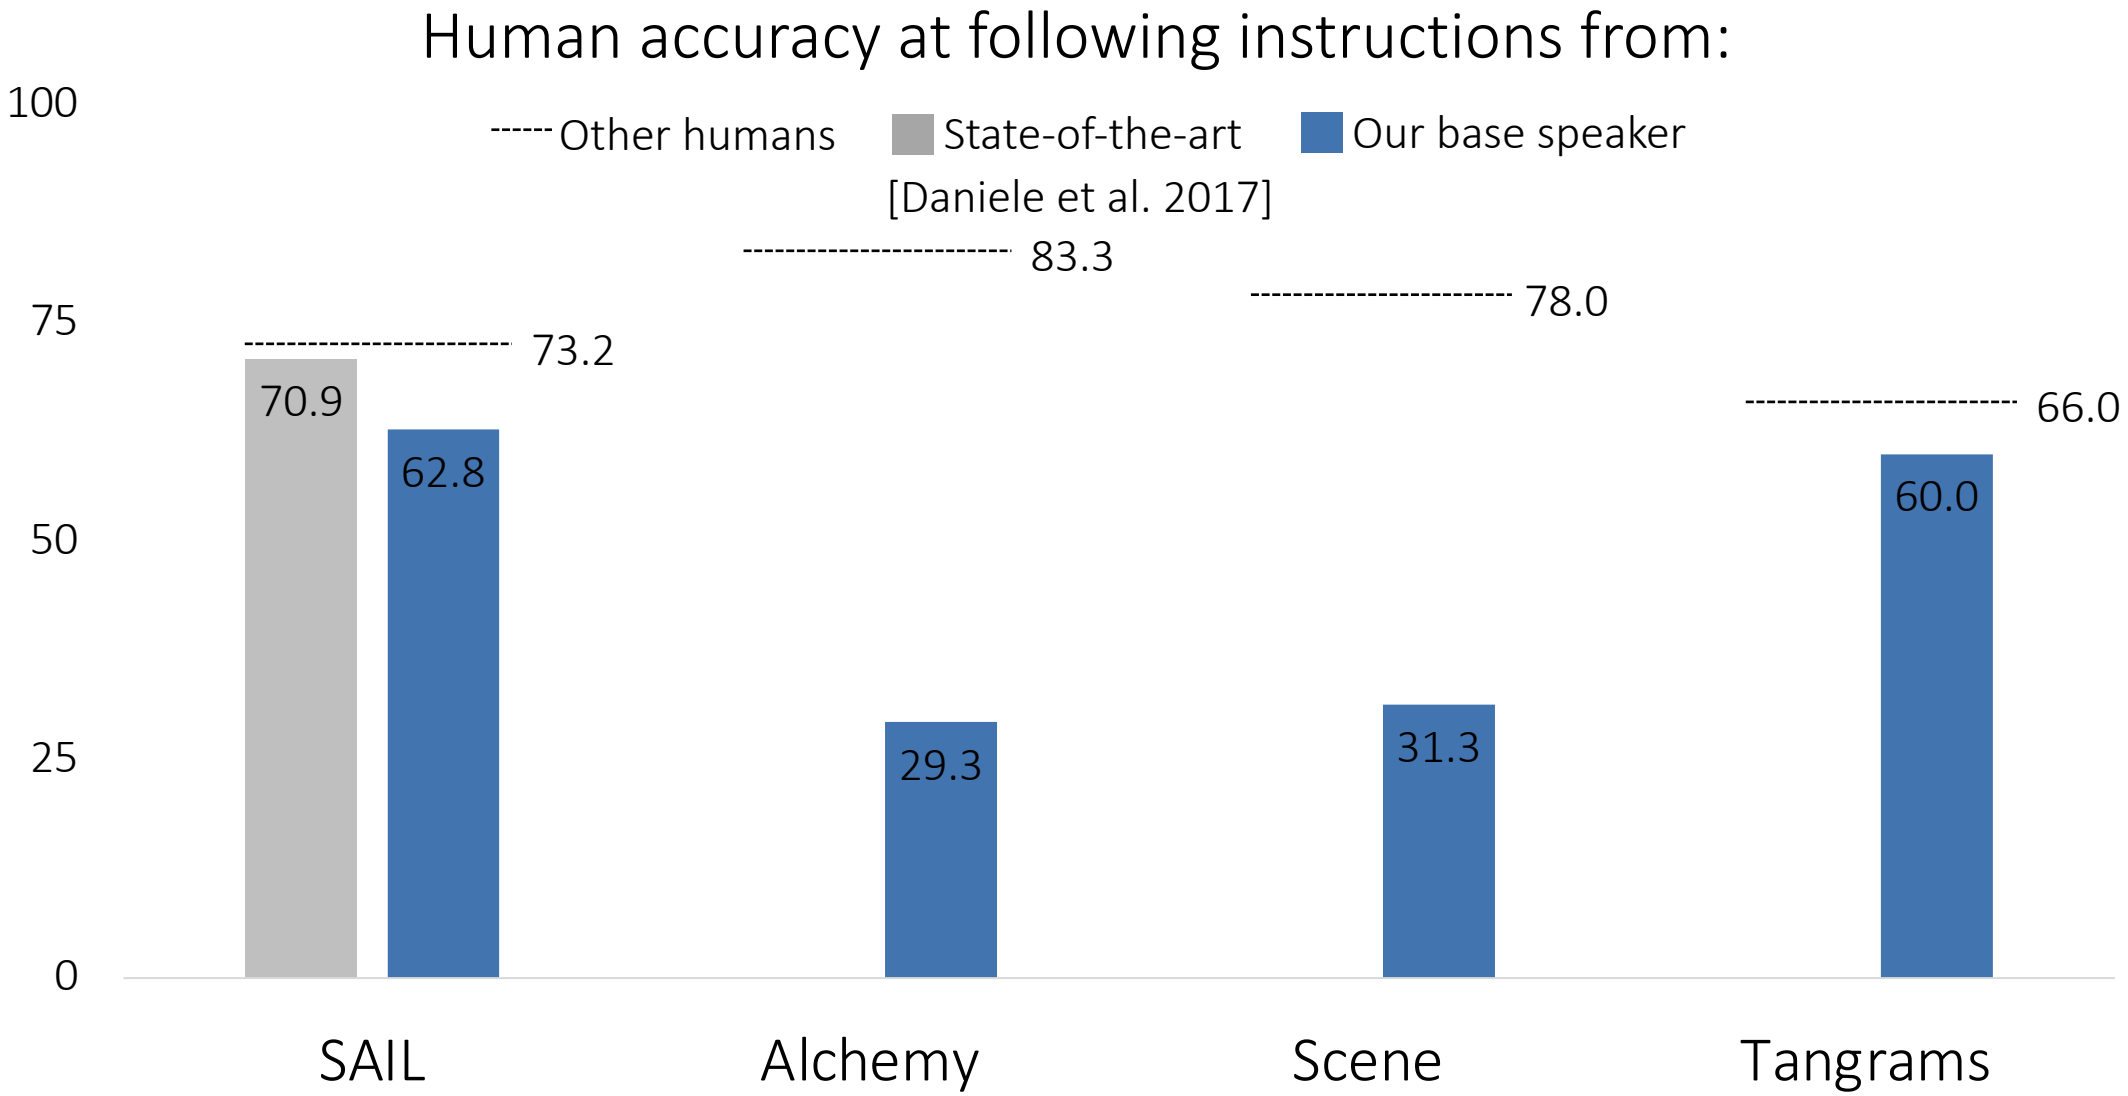

A Failure Mode: Underspecification

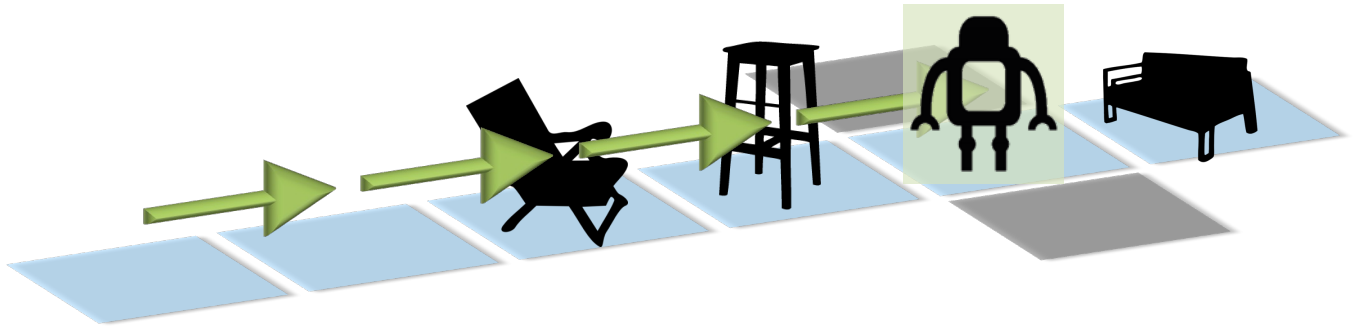

Base Speaker

go forward past the stool ?

A Failure Mode: Contextual Ambiguity

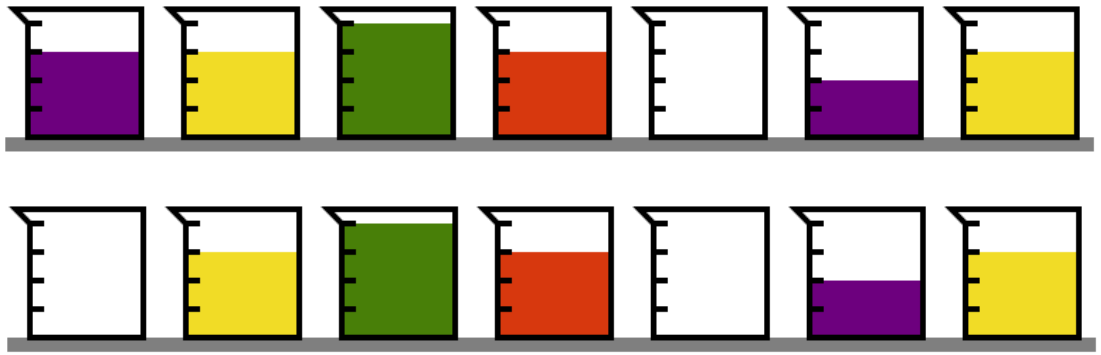

Base Speaker

throw out the purple chemical ✗

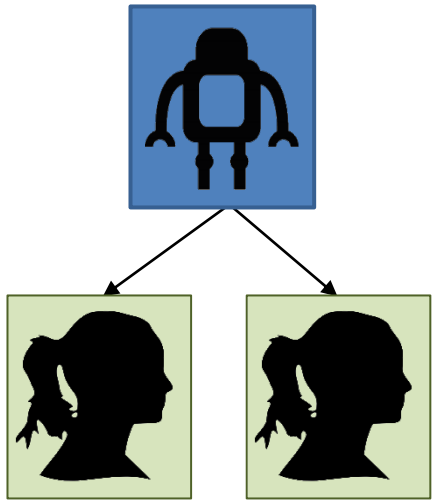

Making Text Informative with Pragmatic Speakers

Pragmatic Speakers Simulate Interpretation

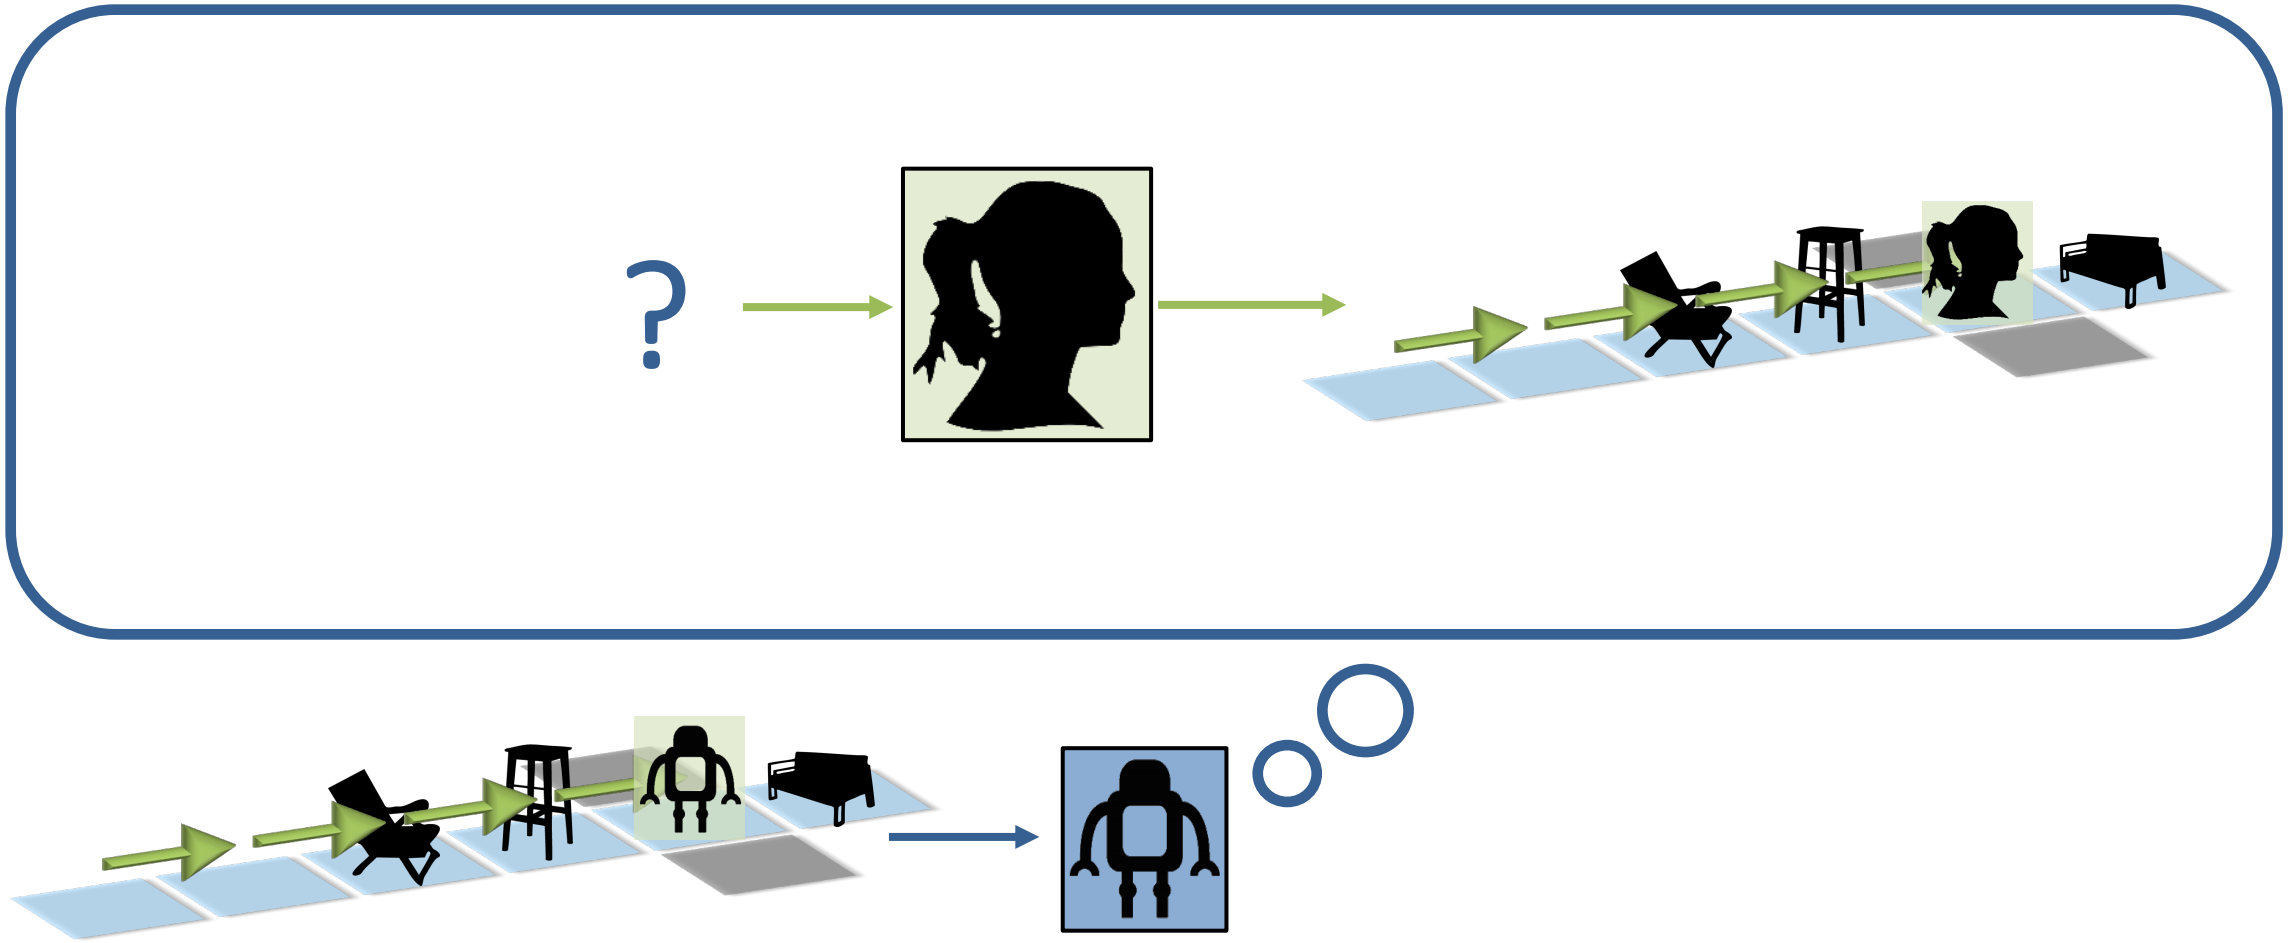

Pragmatic Speakers Simulate Interpretation

go forward
past the stool

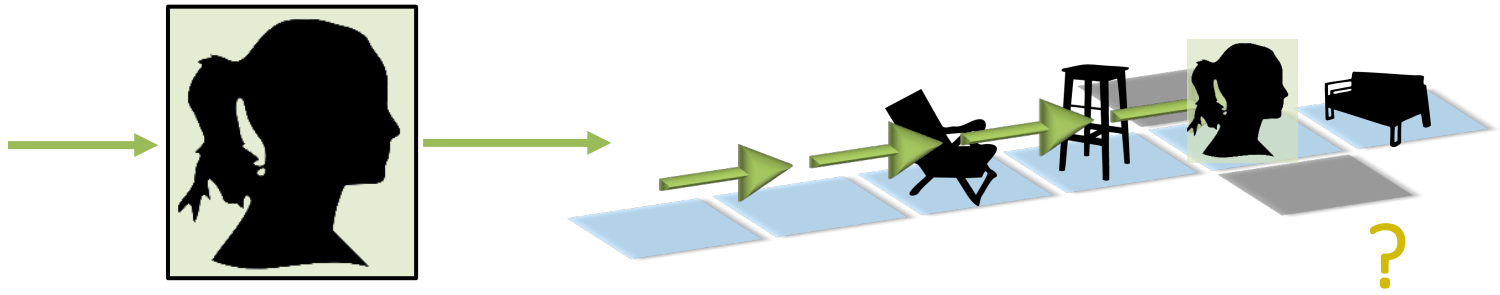

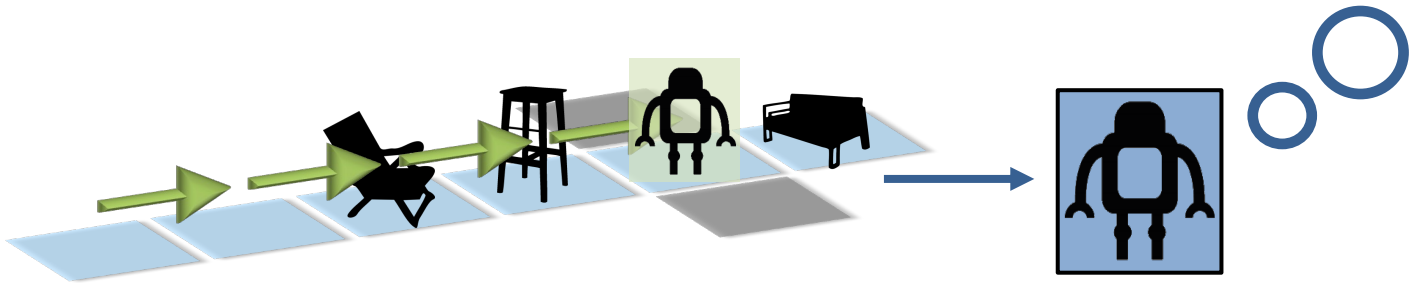

Pragmatic Speakers Simulate Interpretation

go forward four segments to the intersection with the bare concrete hall

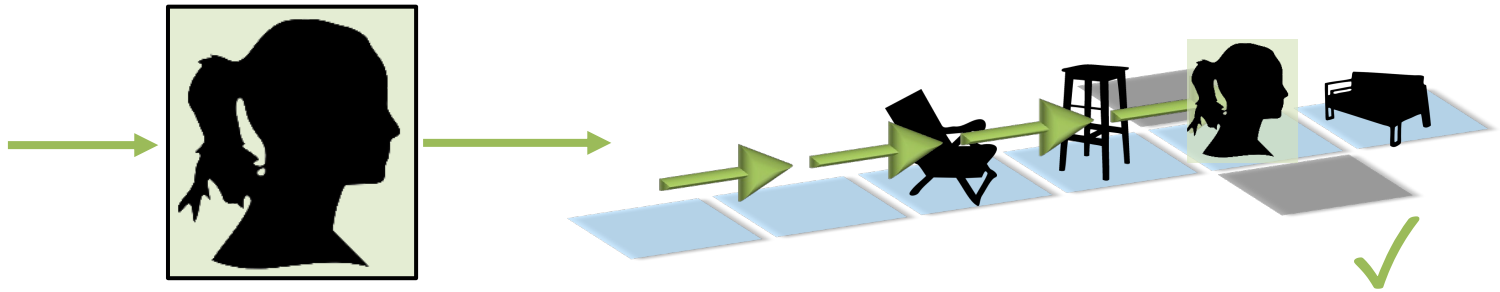

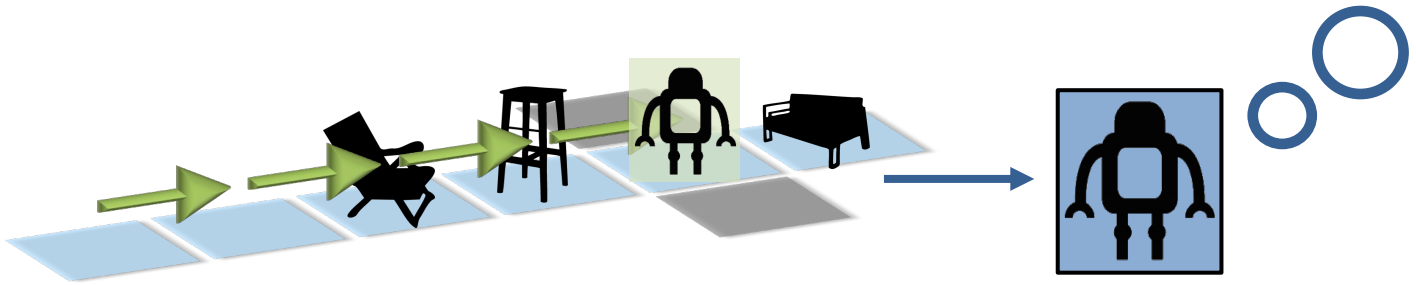

Building a Pragmatic Speaker

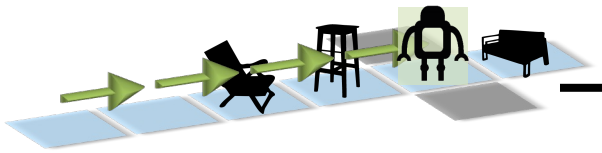

Base Speaker

proposes

0.4

walk forward past the stool

go forward four segments to the intersection with the bare concrete hall

Base Listener

rescores

$P_L(a \mid i)$

0.4

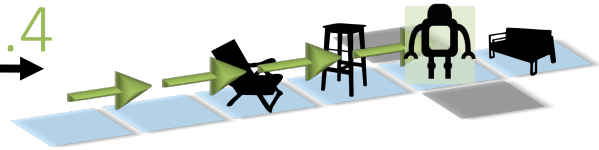

Building a Pragmatic Speaker

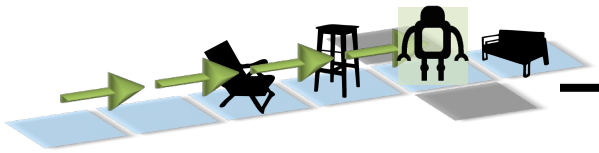

Base Speaker

proposes

0.4

walk forward past the stool

go forward four segments to the intersection with the bare concrete hall

0.8

Base Listener

rescores

$P_L(a \mid i)$

0.8

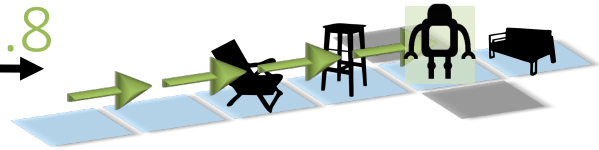

Building a Pragmatic Speaker

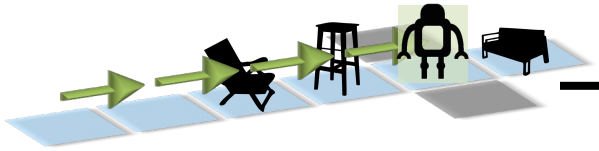

Base Speaker

proposes

0.4

walk forward past the stool

go forward four segments to the intersection with the bare concrete hall

0.8

Base Listener

rescores

$P_L(a \mid i)$

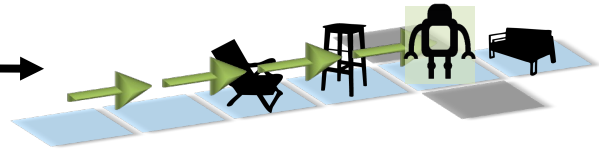

Speaker Results

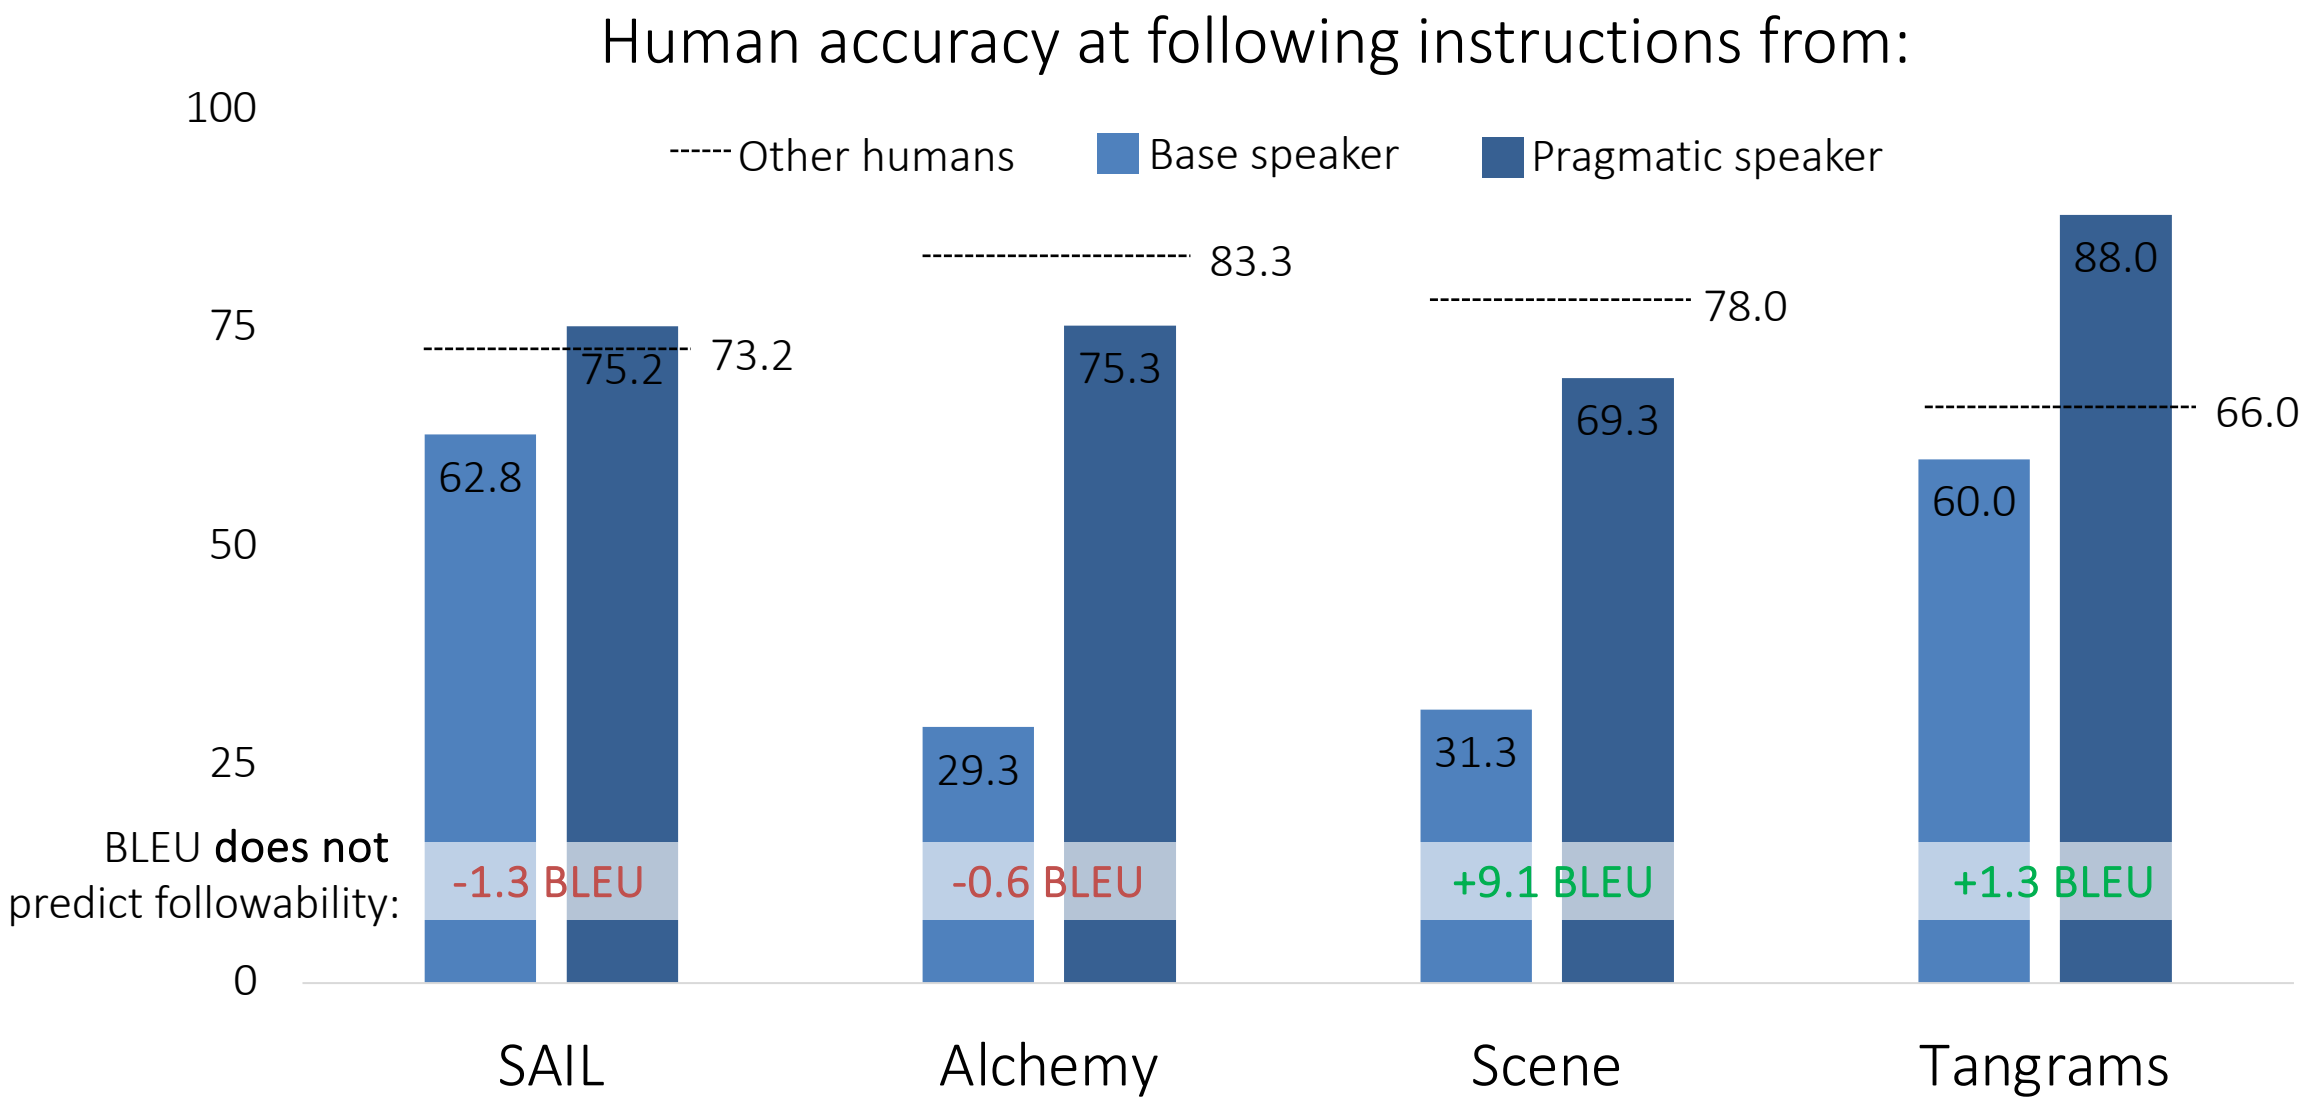

[Fried, Andreas, and Klein. NAACL 2018]

Pragmatics and Communicative Success

Base Speaker	*throw out the purple chemical*	✗
Pragmatic Speaker	*throw out the first purple chemical*	✓
Human	*remove all the purple chemical from the beaker on the far left*	✓

Instruction Quality: Alchemy

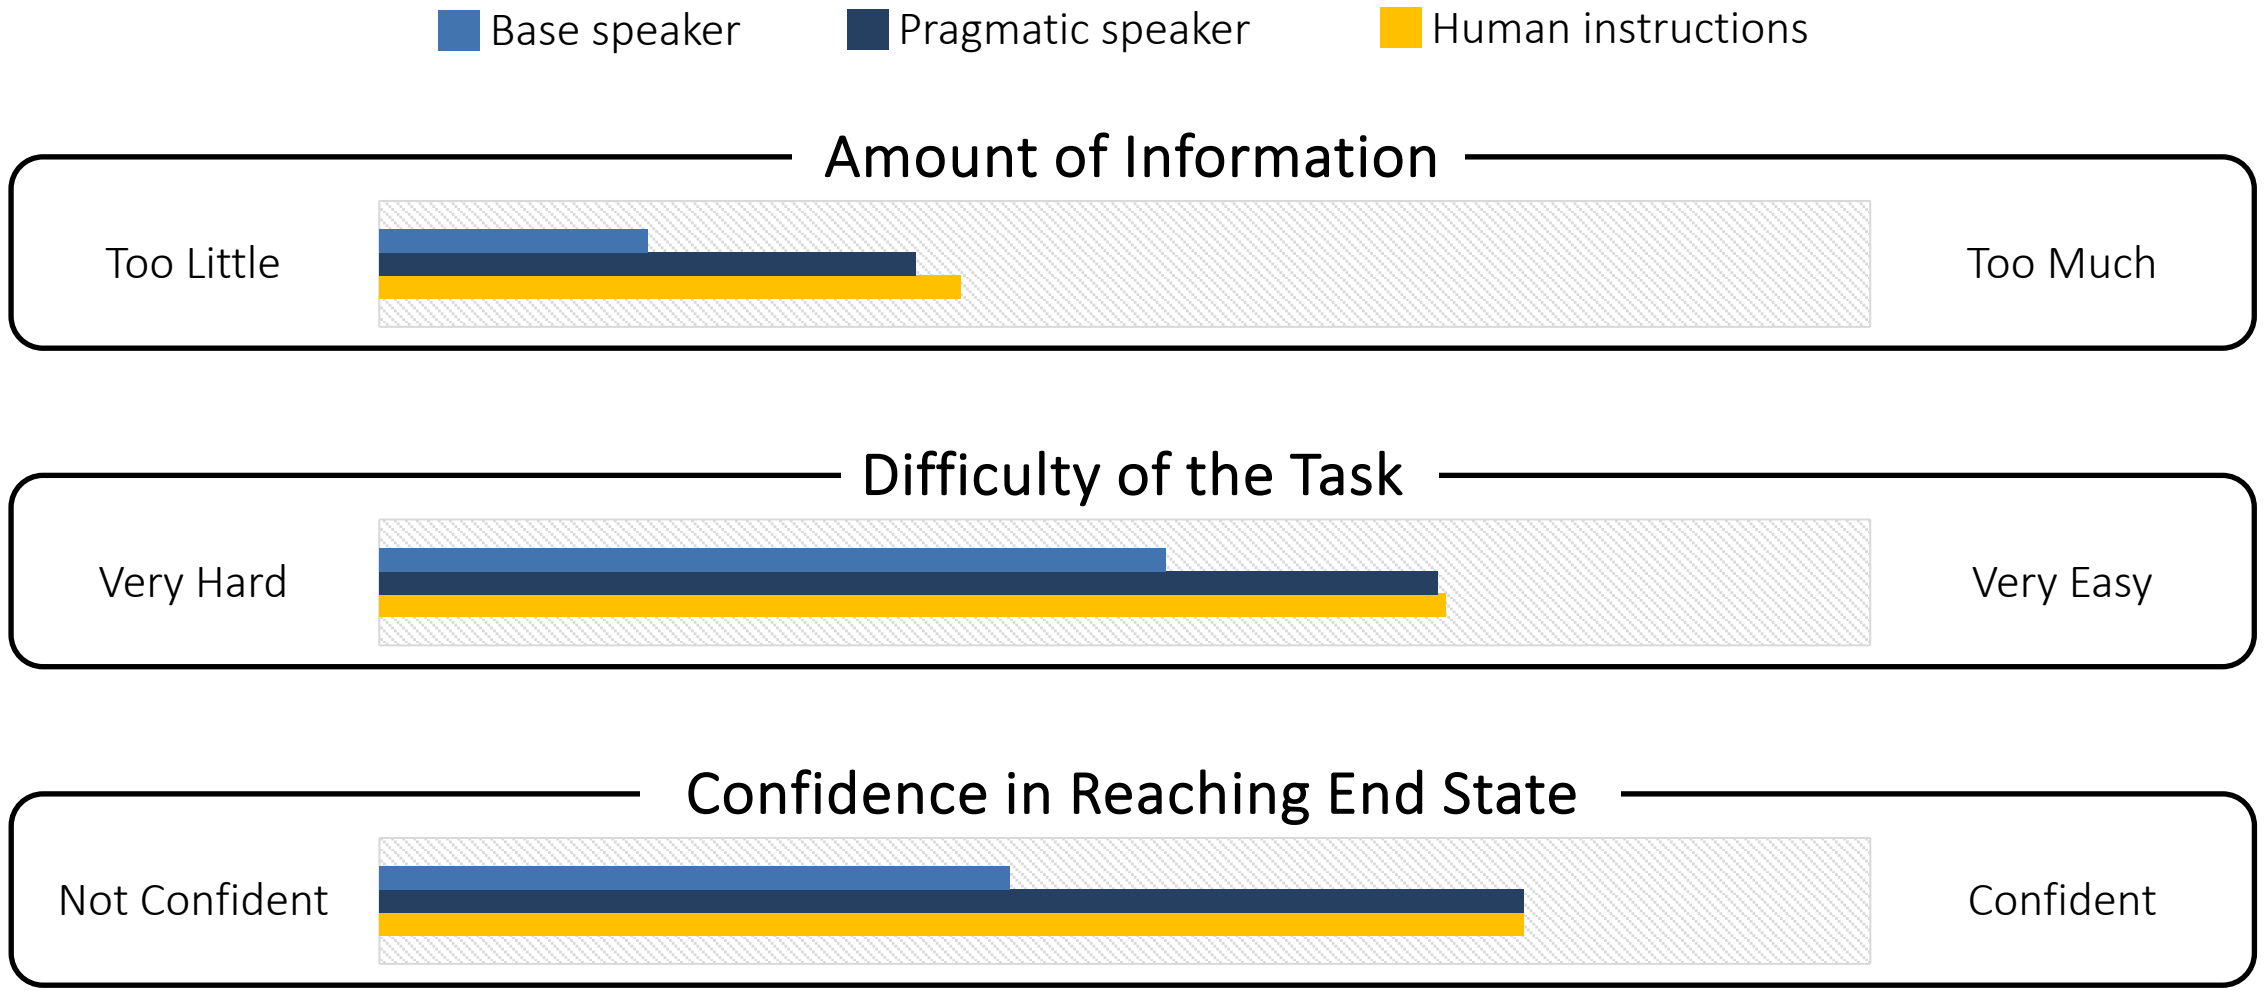

Averaged from 3 or 5 point Likert scales [Daniele et al. 2017]. Differences between base and pragmatic all statistically significant by χ^2 on counts.

Pragmatics and Communicative Success

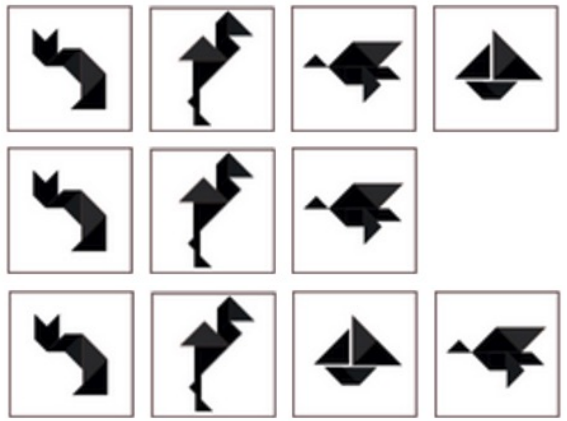

Base Speaker	*remove the last figure* *add it back*	✗
Pragmatic Speaker	*remove the last figure* *add it back in the 3rd position*	✓
Human	*take away the last item* *undo the last step*	✗

Outperforming Training Data (Toy Example)

Training Data

Context → "Language"

AX → x
AX → x
BX → x
BX → x

True, but under-informative

AX → ax
BX → bx

True and informative

Base Speaker

$P_S(\ x \mid AX) = 2/3$
$P_S(ax \mid AX) = 1/3$

Base Listener

$P_L(AX \mid x) = 1/2$
$P_L(AX \mid ax) = 1$

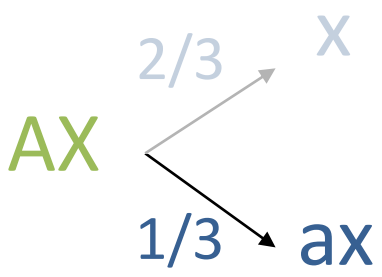

Pragmatics as best response [Franke 2009; Jäger 2014]
Other formalisms:

Recursive Bayesian agents [Frank and Goodman 2012; Jeon et al. 2020]
Optimal transport of beliefs [Wang et al. 2020]
Rate-distortion communication [Zaslavsky et al. 2020]

Speakers and listeners in machine translation

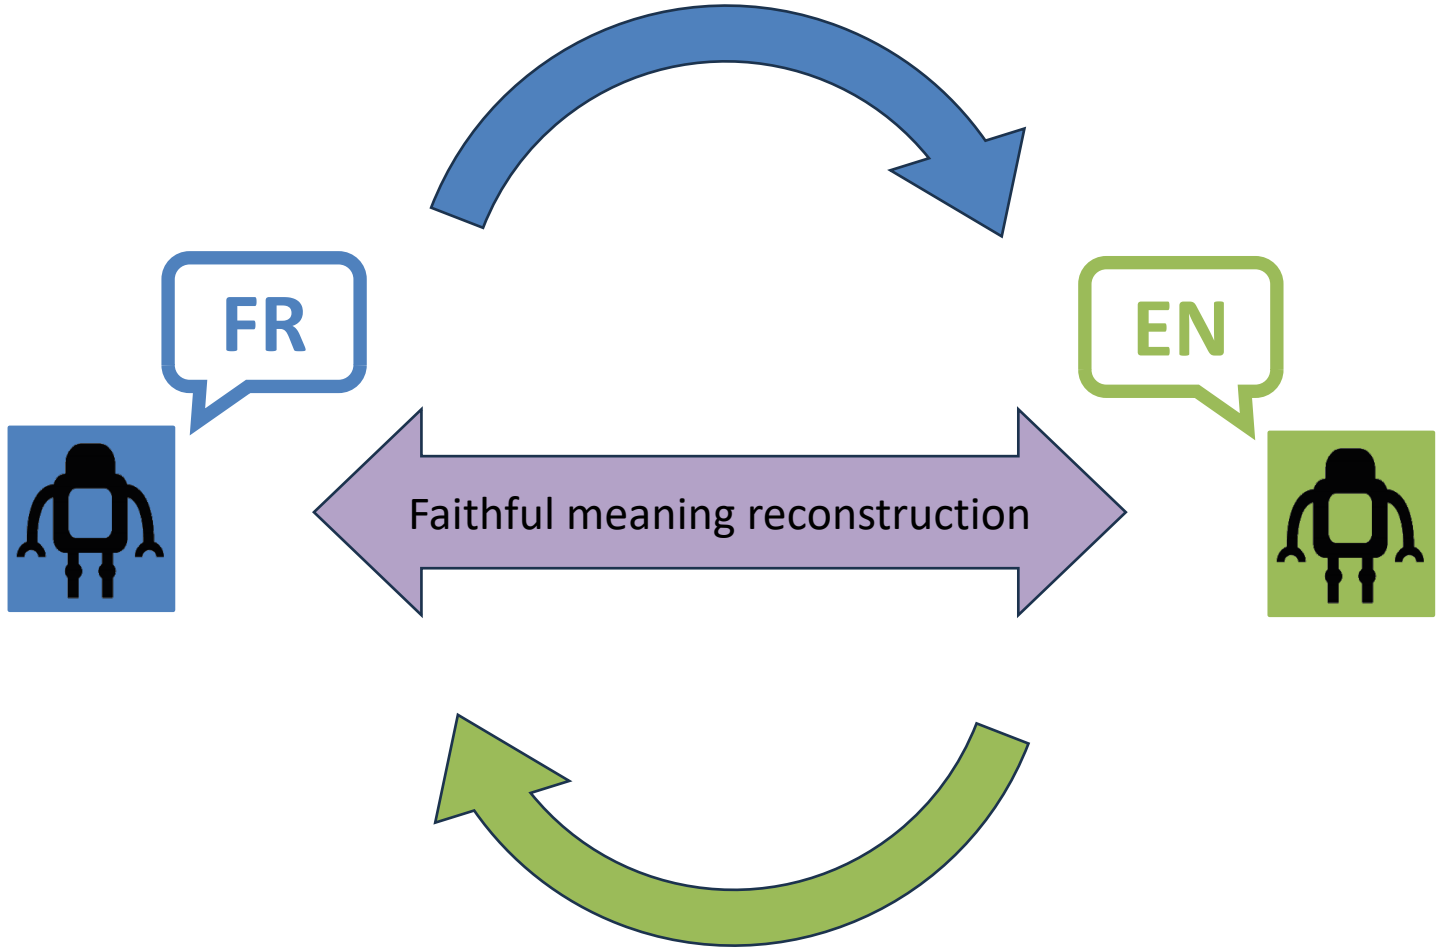

[Cohn-Gordon and Goodman 2019]

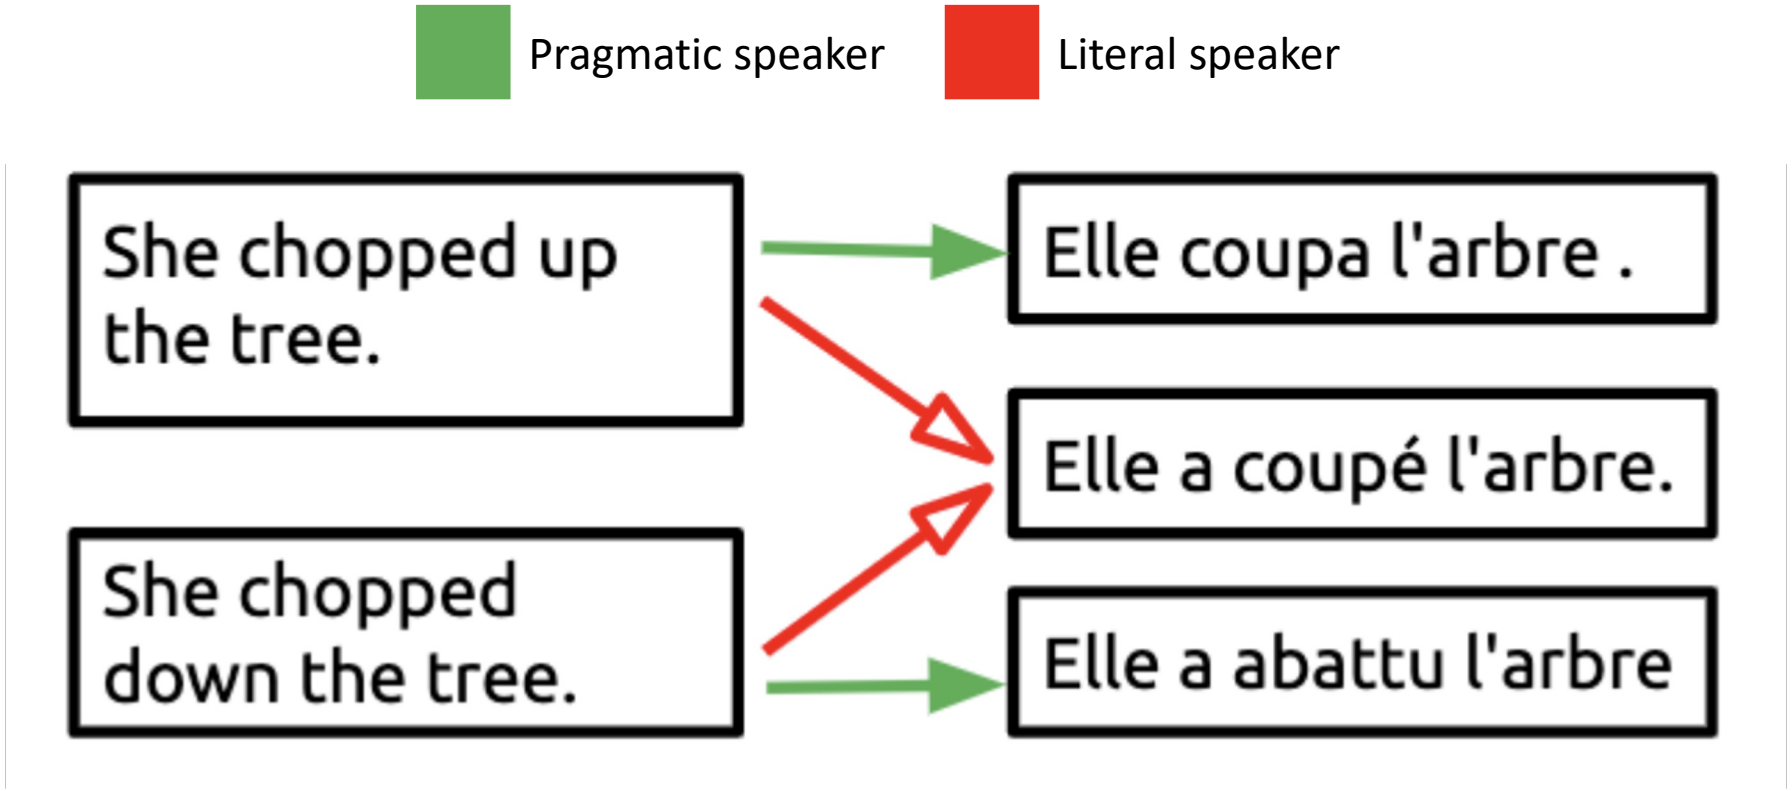
Pragmatic speaker
Literal speaker
She chopped up the tree.
She chopped down the tree.
Elle coupa l'arbre .
Elle a coupé l'arbre.
Elle a abattu l'arbre

Incremental generation with a listener

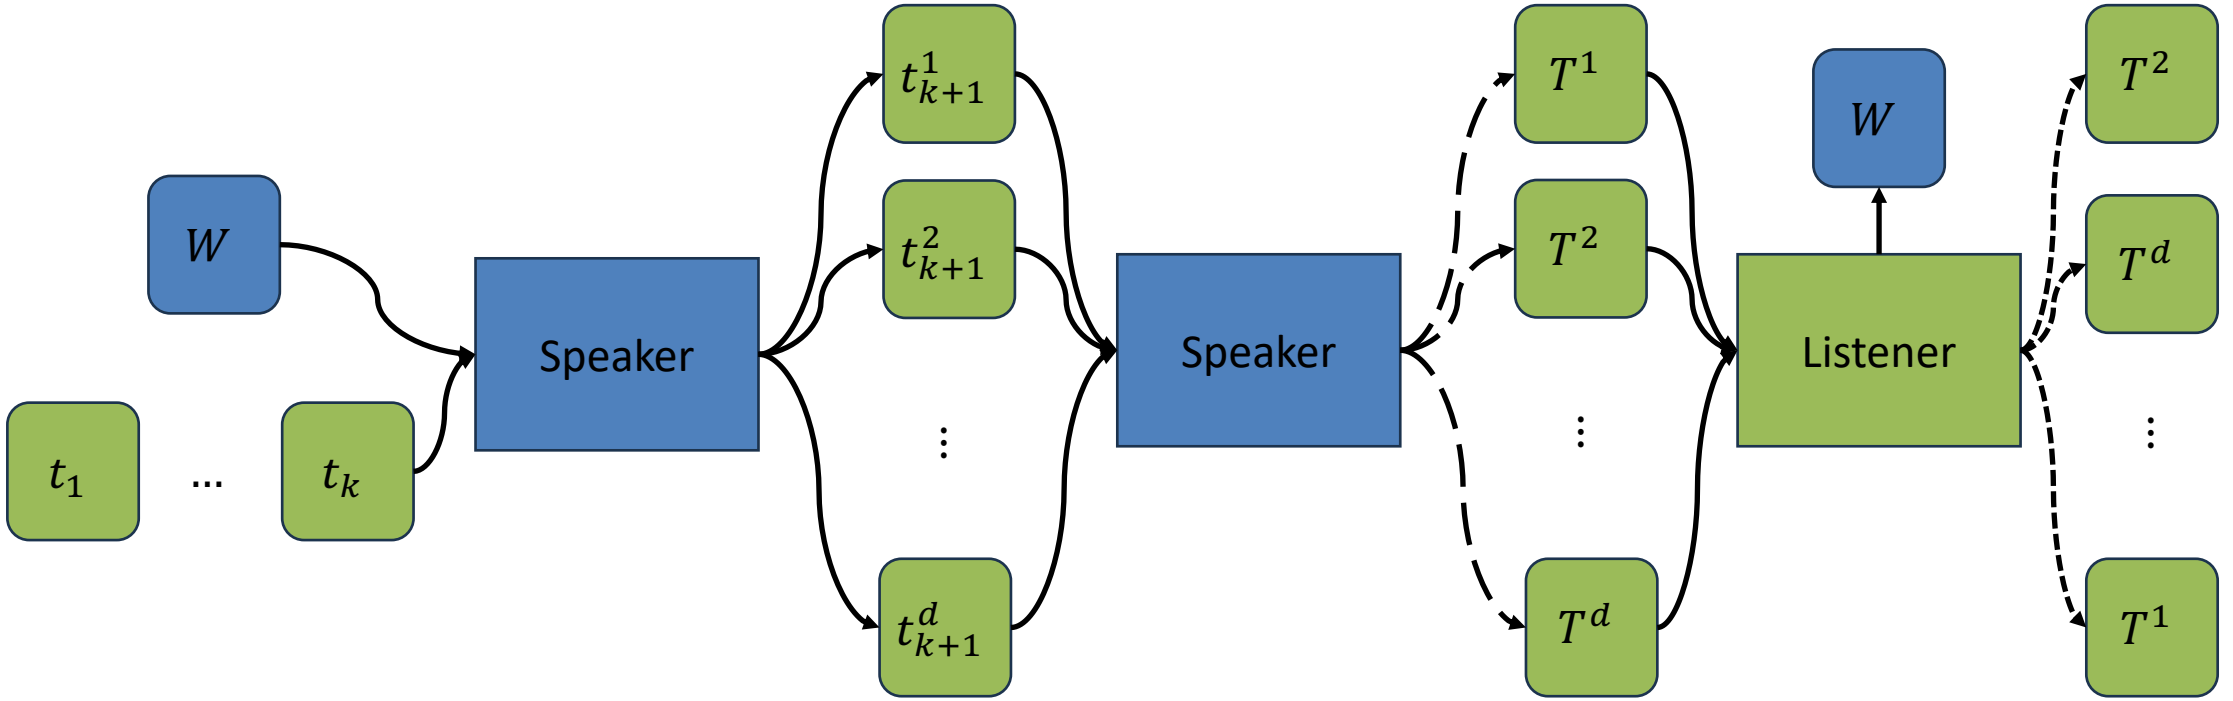

Speakers and listeners in code generation

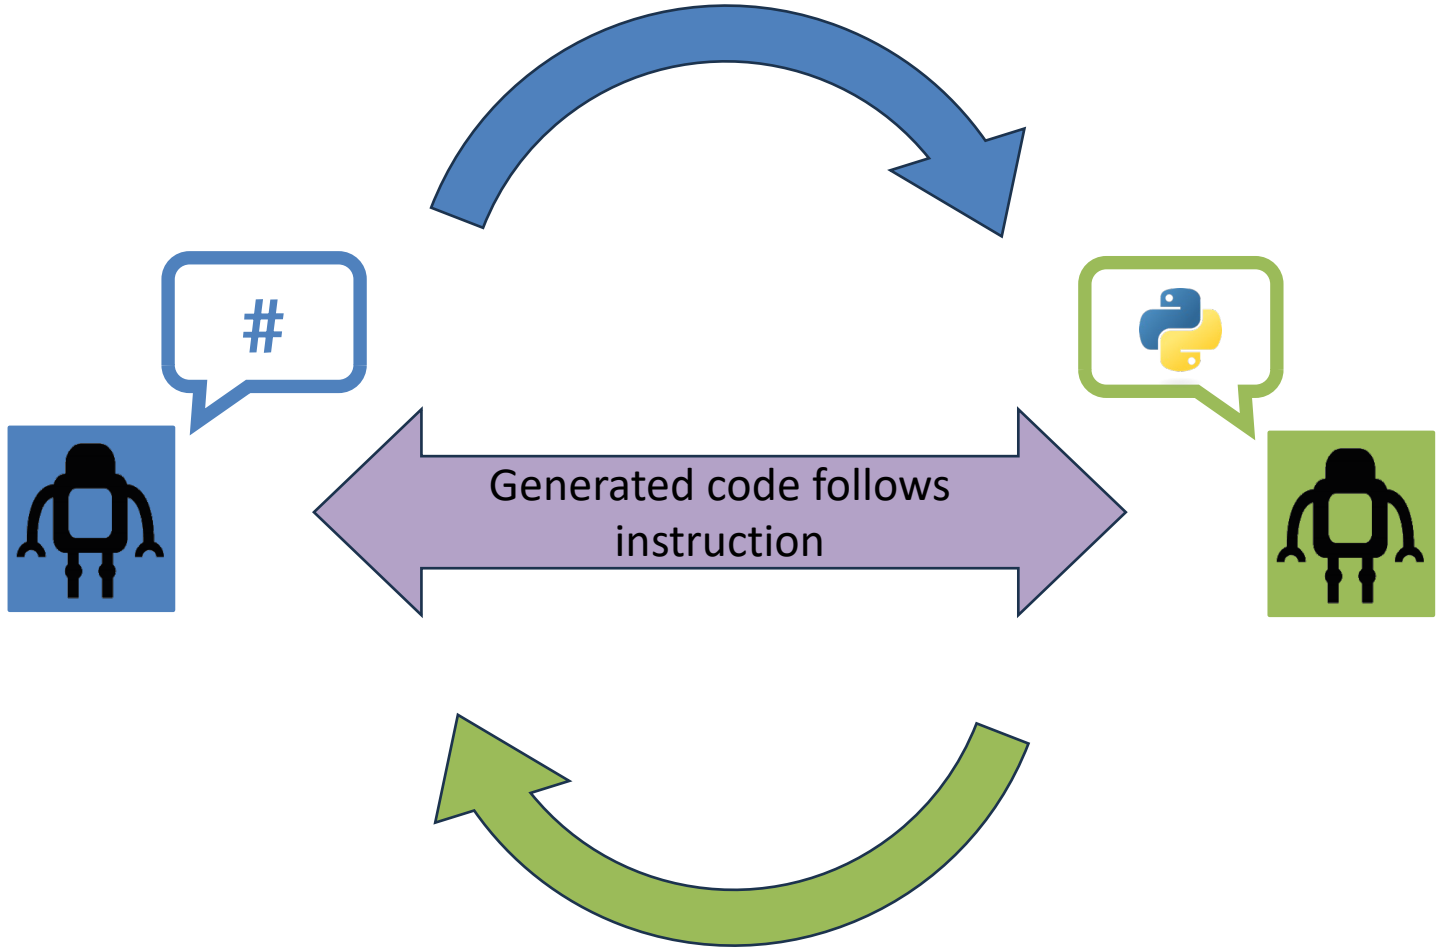

[Zhang et al. 2023]

```
import math
def get_decimal(num:float):
        """ return the decimal part of
        the output number
        """
```

Coder: sample from $P(y|x)$

```
frac, whole = math.modf(num)
return frac
```

Reviewer: sample from $P(x|y)$

```
import math
def get_decimal(num:float):
        frac, whole = math.modf(num)
        return frac
# write a docstring for the above function
def f(num:float):
        """ return the decimal part of
        the output number
        """
```

Rerank based on $P(y|x)P(x|y)$

[Zhang et al. 2023]

Speakers and listeners in image captioning

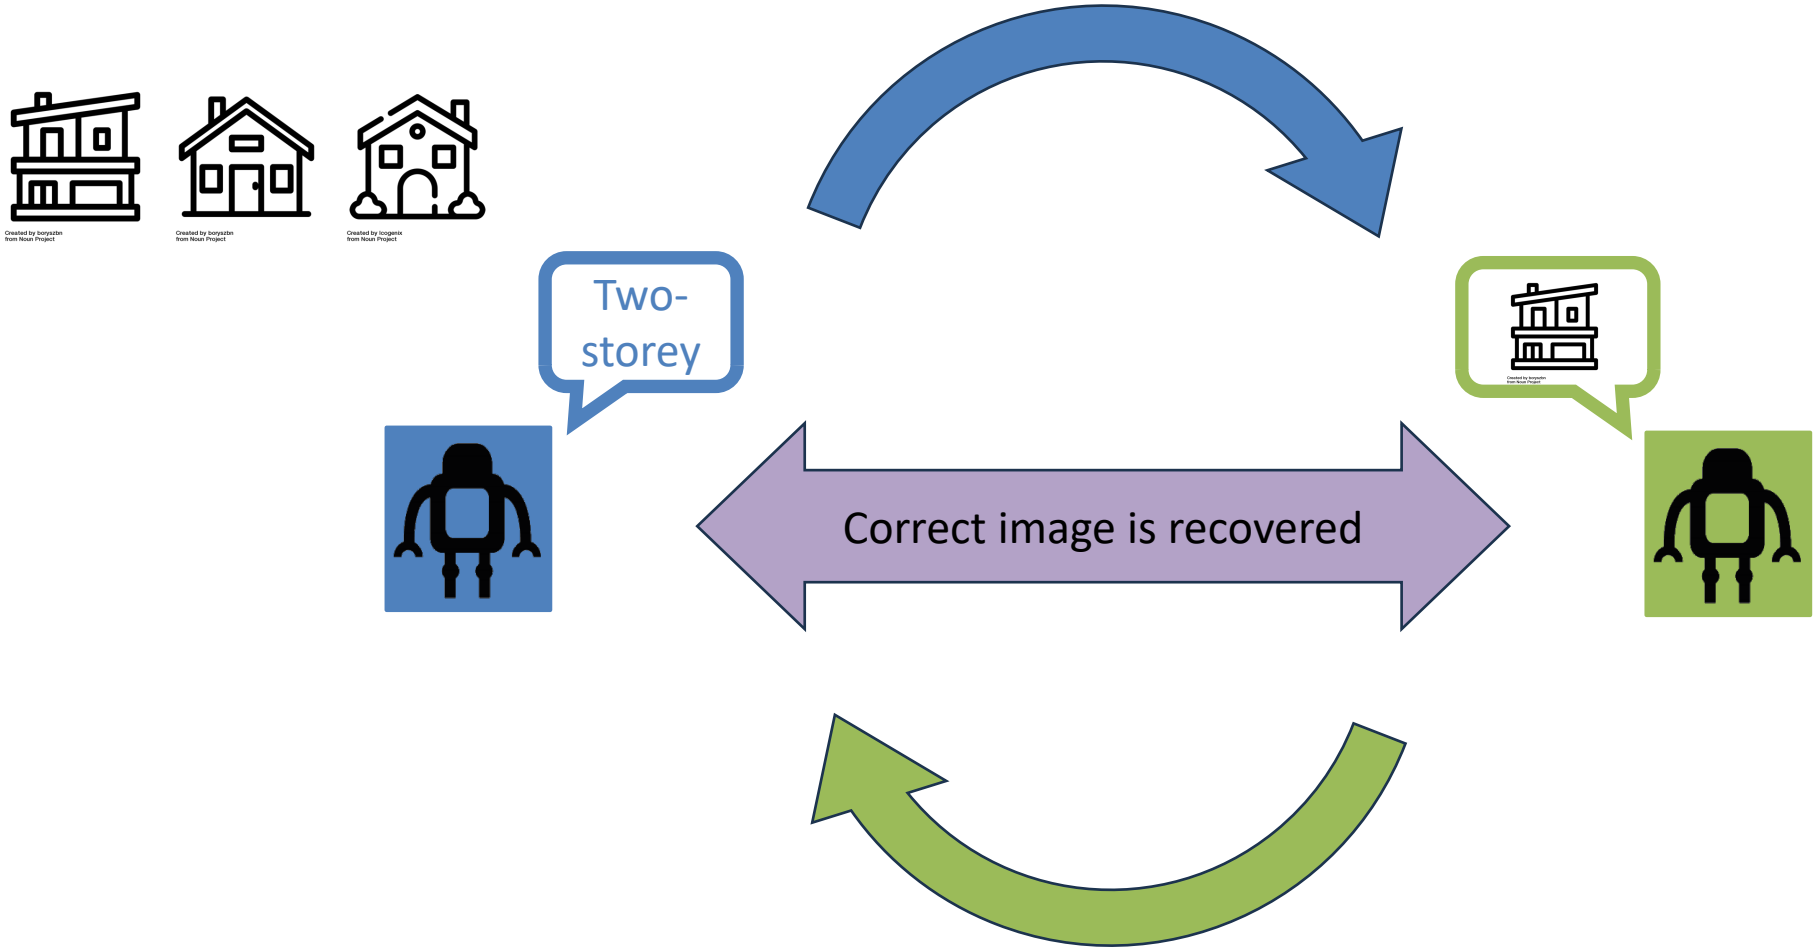

[Vedantam et al. 2017, Ou et al. 2023]

S: A train traveling down tracks next to a forest

IS: A red train is on the tracks in the woods

S: A giraffe standing next to a tree in a forest

IS: A giraffe standing next to a wooden fence

[Vedantam et al. 2017]

Emitter-suppressor beam search

Standard image captioning

$$c^* = \underset{c}{\text{argmax}} \log P(c|I) = \underset{c}{\text{argmax}} \sum \log P(c_i|c_1^{i-1}, I)$$

- Approximate argmax with beam search

Discriminative image captioning

$$c^* = \underset{c}{\text{argmax}} \sum \log \frac{P(c_i|c_1^{i-1}, I_t)}{P(c_i|c_1^{i-1}, I_d)^{1-\lambda}}$$

[Vedantam et al., 2017]

Repeated reference games

$t = 1$

all the ap[p]liances and cupboards are white except for the stove which is stainless [s]teel. the wall is white to the left and behind the appliances it is brown wood

$t = 2$

white appliances, wooden back wall, white wall to the left

…

$t = 6$

white ap[p]liances, brown back wall

[Hawkins et al., 2020]

Common ground

Set of contextual information shared between communicative partners

- task goals and collaborative actions
- social and communicative norms
  - including knowledge of how implicature is used!
- discourse context
- common knowledge

[Fried et al., 2023]

Meta-learning for continual adaptation

$$\theta_0 \leftarrow \Theta$$

$$\mathcal{L}(\theta_t) = \log P_{\theta_t}(\boldsymbol{u}|o_t) + \log P_{\theta_t}(o_t|\boldsymbol{u}) - \text{regularization}$$

$$\theta_t \leftarrow \theta_t + \beta \nabla \mathcal{L}(\theta_t)$$

[Hawkins et al., 2020]

Meta-learning for continual adaptation

$$\theta_0 \leftarrow \Theta$$

$$\mathcal{L}(\theta_t) = \log P_{\theta_t}(\boldsymbol{u}|o_t) + \log P_{\theta_t}(o_t|\boldsymbol{u}) - \text{regularization}$$

$$\theta_t \leftarrow \theta_t + \beta \nabla \mathcal{L}(\theta_t)$$

- Common knowledge
- Discourse context
- Task knowledge
- Cooperative objective
- Update discourse context

[Hawkins et al., 2020]

Assumed common ground: Presupposition

Assumptions and beliefs that are shared and taken for granted by discourse participants without an explicit mention in the discourse context by a particular utterance

Everyone enjoyed learning about pragmatics

presupposes *Everyone learned about pragmatics*

Presupposition in QA

Who is the current monarch of the UK?: King Charles.

Who is the current monarch of the France?: ???

- Google answers *Which linguist invented the lightbulb?* with Thomas Edison
- Bing answers the question *When did Marie Curie discover Uranium?* with 1896.

[Kim et al., 2021; example retrieved Jan 2021]

Handling false presuppositions

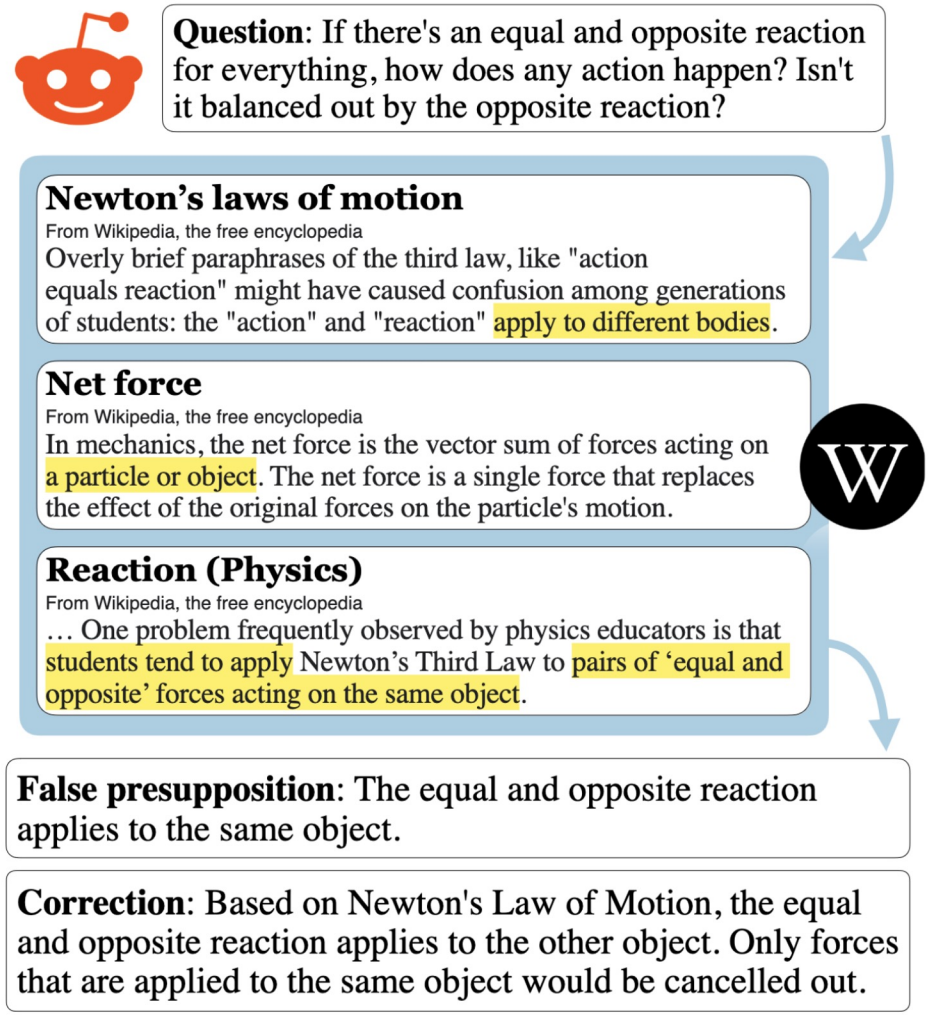

[Yu et al., 2023]

Handling false presuppositions in QA

1. Generate presuppositions
   - Use syntactic triggers
   - Train a model to generate the presupposition
2. Verify presuppositions
   - Train a classifier to predict whether presupposition is verified in source
     - Natural language inference
3. Respond to false presuppositions
   - The question is unanswerable because of the false presupposition
   - Issue a correction to the presupposition

[Kim et al., 2021; Yu et al., 2023]

Takeaways

- Language is used as an action
  - taken in context
  - with a purpose
- Reasoning explicitly about the context and goals of language can sometimes help model it better
  - Who is reading/listening to generated language?
  - What is the person who produced language trying to say?
  - How will the listener interpret it?