Pragmatics – Rational Speech Act Theory

Prof. Dr. Vera Demberg

Saarland University

Elements of Data Science and Artificial Intelligence

UNIVERSITÄT
DES
SAARLANDES

Motivation

It's the taste!

It's the taste!

Motivation

It's the taste!
+> The tea tastes fantastic.

It's the taste!
+> The food tastes bad.

Plan for today

- Pragmatics – meaning beyond semantics
- Grice: Cooperative principle and conversational maxims, implicatures
- Pragmatics for language generation
- Modern incarnation: Rational speech act model
- Instruction giving as an application

Motivation

A: Can you tell me what time it is?
B: Well, the milkman has already been there.

Motivation

A: Can you tell me what time it is?
B: Well, the milkman has already been there.
$+>$ A does not know exactly what time it is. However, the information that the milkman has already been there may help B guess the current time, because it is later than the time the milkman normally comes.

Motivation

A: Can you tell me what time it is?
B: Well, the milkman has already been there.
$+>$ A does not know exactly what time it is. However, the information that the milkman has already been there may help B guess the current time, because it is later than the time the milkman normally comes.

A: Einige Studenten haben die Klausur bestanden.

Motivation

A: Can you tell me what time it is?
B: Well, the milkman has already been there.
+> A does not know exactly what time it is. However, the information that the milkman has already been there may help B guess the current time, because it is later than the time the milkman normally comes.

A: Einige Studenten haben die Klausur bestanden.
+> Nicht alle Studenten haben die Klausur bestanden.

Motivation

A: Can you tell me what time it is?
B: Well, the milkman has already been there.
+> A does not know exactly what time it is. However, the information that the milkman has already been there may help B guess the current time, because it is later than the time the milkman normally comes.

A: Einige Studenten haben die Klausur bestanden.
+> Nicht alle Studenten haben die Klausur bestanden.

John F. Kennedy: Ich bin ein Berliner.

Motivation

A: Can you tell me what time it is?
B: Well, the milkman has already been there.
+> A does not know exactly what time it is. However, the information that the milkman has already been there may help B guess the current time, because it is later than the time the milkman normally comes.

A: Einige Studenten haben die Klausur bestanden.
+> Nicht alle Studenten haben die Klausur bestanden.

John F. Kennedy: Ich bin ein Berliner.
+>Wir sind solidarisch mit Berlin (Luftbrücke).

Table of Contents

Meaning and Communication

Language philosopher H. Paul Grice (1975)

Communicated meaning includes

- sentence meaning (literal semantic content of a message)
- speaker meaning (ironic, metaphorical and implicit or indirect communicative content) = additional inferred information that the speaker intends that the listener recognizes as part of the intended communicated meaning.

Meaning and Communication

Language philosopher H. Paul Grice (1975)

Communicated meaning includes

- sentence meaning (literal semantic content of a message)
- speaker meaning (ironic, metaphorical and implicit or indirect communicative content) = additional inferred information that the speaker intends that the listener recognizes as part of the intended communicated meaning.

Sentence meaning and speaker meaning may be different.

Meaning and Communication

Language philosopher H. Paul Grice (1975)

Communicated meaning includes

- sentence meaning (literal semantic content of a message)
- speaker meaning (ironic, metaphorical and implicit or indirect communicative content) = additional inferred information that the speaker intends that the listener recognizes as part of the intended communicated meaning.

Sentence meaning and speaker meaning may be different.

Speaker meaning is

- more than the linguistic meaning (sentence meaning, utterance meaning)
- less than the total information that is inferable from the utterance.

Speaker meaning

Grice calls the intended inference "implicatures", verb is "implicate" in contrast to "imply":
Implicatures are inferences intended by the speaker, based both on the content of what has been said and on some particular assumptions about the cooperative nature of a normal verbal interaction.

Grice's Cooperative Principle and Conversational Maxims

Cooperative Principle: Make your conversational 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 ('Super-maxim').

Conversational Maxims: Principles of effective and efficient language use, as a basis for every conversation.

Conversational Maxims:

1. Quality: Make your contribution one that is true, that is to say:
   (i) Don't say what you believe to be false.
   (ii) Don't say that for which you lack adequate evidence.
2. Quantity:
   (i) Make your contribution as informative as is required (for the current purposes of the exchange).
   (ii) Do not make your contribution more informative than is required.
3. Relevance: Make your contribution relevant.
4. Manner: Be perspicuous (clear), that is to say:
   (i) Avoid obscurity of expression
   (ii) Avoid ambiguity
   (iii) Be brief (avoid unnecessary prolixity)
   (iv) Be orderly

Table of Contents

Conversational implicatures

Conversational Implicature (CI)
is a conclusion that

- goes beyond the semantic content of the uttered sentences, and
- is derived from the maxims
  - Attention / Observance of the maxims = Standard-CIs
    observing the maxims of conversational implicature = implicature_O, or
  - Disregard / Violation of maxims = non-standard CIs
    "flouting" the maxims of conversational implicature = implicature_F

Conversational Implicature – Definition

- If the speaker says that p, q is a conversational implicature of p iff:
  - (i) we can assume that the speaker obeyed the maxims or at least (in case of flouting the maxims) the principle of cooperation
  - (ii) Based on this assumption, the hearer must assume that the speaker believes that q
  - (iii) the speaker believes that the speaker and the hearer mutually know that the hearer can infer that q is necessary for maintain the assumption in (i)

Conversational Implicature – Definition

- If the speaker says that p, q is a conversational implicature of p iff:
  - (i) we can assume that the speaker obeyed the maxims or at least (in case of flouting the maxims) the principle of cooperation
  - (ii) Based on this assumption, the hearer must assume that the speaker believes that q
  - (iii) the speaker believes that the speaker and the hearer mutually know that the hearer can infer that q is necessary for maintain the assumption in (i)
- To calculate the implicature q, the hearer needs to know:
  - (i) the conventional / semantic content of the uttered sentence p
  - (ii) the principle of cooperation and its maxims
  - (iii) the context of p
  - (iv) certain background information / world knowledge
  - (v) that (i) - (iv) is common knowledge of the speaker and the hearer

Table of Contents

Implicature$_O$

If the hearer assumes that the speaker follows the maxims, and the speaker relies on the hearer to draw conclusions from what is said, then the resulting inferences are called **standard conversational implicatures** (Implicature$_O$ for "implicature observed").

A: (to a passer-by) I just ran out of gas.
B: Oh, there's a service station around the corner.

Implicature$_O$:
quality: B knows there's a service station around the corner
quantity: B does not know if you can get gas in the service station
relevance: B thinks you can get gas in the service station

An example of an observed conversational implicature

Quantity Ally looked me right in the eye and said, “I need to know how you feel about me.” I didn’t say anything for a good time... “I care deeply about you,” I said. “But you don’t love me?” “I don’t know.” She nodded. Tears streamed down her face (Peter David Marks, New York Times)

“I care deeply about you.”

An example of an observed conversational implicature

Quantity Ally looked me right in the eye and said, "I need to know how you feel about me." I didn't say anything for a good time... "I care deeply about you," I said. "But you don't love me?" "I don't know." She nodded. Tears streamed down her face (Peter David Marks, New York Times)

"I care deeply about you."
$+>$ the speaker does not love the addressee

Exercise

Children with pragmatic disorders may often fail to observe the maxims. In the following conversation, which maxim did the child violate?

Speech therapist: “So you like ice cream. What are your favourite flavours?”
Child: “Hamburger... fish and chips.”

Table of Contents

Non-standard-implicatures - Implicature$_F$ (flouted)

- If the speaker obviously and intentionally disregards the maxims, the hearer still assumes cooperation and draws the necessary conclusions to explain this violation. These are **non-standard implicatures** (implicature$_F$).

  A: Lass uns was für die Kinder kaufen.
  B: Ja, aber kein E.I.S.

Non-standard-implicatures - Implicature$_F$ (flouted)

- If the speaker obviously and intentionally disregards the maxims, the hearer still assumes cooperation and draws the necessary conclusions to explain this violation. These are **non-standard implicatures** (implicature$_F$).

  A: Lass uns was für die Kinder kaufen.
  B: Ja, aber kein E.I.S.

  Spelling out the word violates the maxim of manner.
  Implicature$_F$: B doesn't want the kids to hear the word ice cream

- Many traditional rhetorical figures including metaphor, irony, rhetorical questions etc. depend on flouting a conversational maxim.

Exercises

Quality Tony Blair is no longer the Prime Minister of Britain, he is the Foreign Minister of the United States (Nelson Mandela, quoted in Susie Dent, Language Report 2003:62)

Exercises

Quality Tony Blair is no longer the Prime Minister of Britain, he is the Foreign Minister of the United States (Nelson Mandela, quoted in Susie Dent, Language Report 2003:62)
+> Tony Blair has followed the American foreign policies too closely.

Exercises

Quality Tony Blair is no longer the Prime Minister of Britain, he is the Foreign Minister of the United States (Nelson Mandela, quoted in Susie Dent, Language Report 2003:62)
+> Tony Blair has followed the American foreign policies too closely.

Relevance Could you hand me the butter?

Exercises

Quality Tony Blair is no longer the Prime Minister of Britain, he is the Foreign Minister of the United States (Nelson Mandela, quoted in Susie Dent, Language Report 2003:62)
+> Tony Blair has followed the American foreign policies too closely.

Relevance Could you hand me the butter?
+> Please hand me the butter.

Exercises

Quality Tony Blair is no longer the Prime Minister of Britain, he is the Foreign Minister of the United States (Nelson Mandela, quoted in Susie Dent, Language Report 2003:62)
+> Tony Blair has followed the American foreign policies too closely.

Relevance Could you hand me the butter?
+> Please hand me the butter.

Manner The corners of John's lips turned slightly upward.

Exercises

Quality Tony Blair is no longer the Prime Minister of Britain, he is the Foreign Minister of the United States (Nelson Mandela, quoted in Susie Dent, Language Report 2003:62)
+> Tony Blair has followed the American foreign policies too closely.

Relevance Could you hand me the butter?
+> Please hand me the butter.

Manner The corners of John's lips turned slightly upward.
+> John did not exactly smile.

Placing conversational implicatures in pragmatics

Nonliteral / additional meaning beyond semantics

Presupposition:
“the king of France is bald”
-> there is a king of France

Inference, e.g., entailment:
“Peter slept well last night”
-> Peter slept last night

Implicature

Conversational Implicature

Generalized Convers. Implicature
“Mary has three children”
-> Mary has not more than three children

Particularized Convers. Implicature
“Where is the steak?”
“The dog looks happy.”
-> the dog ate the steak

Conventional Implicature
“Even Peter was on time today.”
-> I did not expect that Peter would be punctual.
other words that give rise to conventional implicatures:
“but”, “only”, “even”, “although”

Summary so far

- Grice's theory of meaning
  - Communicative meaning is a complex intention which is achieved by being recognized by the hearer.
  - There are guidelines for effective and rational language usage
    - cooperation principle
    - conversational maxims

    On the assumption that the principle of cooperation is observed (and follows the maxims), inferences are generated: conversational implicatures.
- Conversational implicatures allow us to infer different meanings from the same utterance in different situations.

Summary so far

- Grice's theory of meaning
  - Communicative meaning is a complex intention which is achieved by being recognized by the hearer.
  - There are guidelines for effective and rational language usage
    - cooperation principle
    - conversational maxims

  On the assumption that the principle of cooperation is observed (and follows the maxims), inferences are generated: conversational implicatures.
- Conversational implicatures allow us to infer different meanings from the same utterance in different situations.
- This brings advantages for the formulation of semantics.

Summary so far

- Grice's theory of meaning
  - Communicative meaning is a complex intention which is achieved by being recognized by the hearer.
  - There are guidelines for effective and rational language usage
    - cooperation principle
    - conversational maxims

    On the assumption that the principle of cooperation is observed (and follows the maxims), inferences are generated: conversational implicatures.
- Conversational implicatures allow us to infer different meanings from the same utterance in different situations.
- This brings advantages for the formulation of semantics.

Let's now look at some more concrete application examples.

Table of Contents

Haushaltsroboter

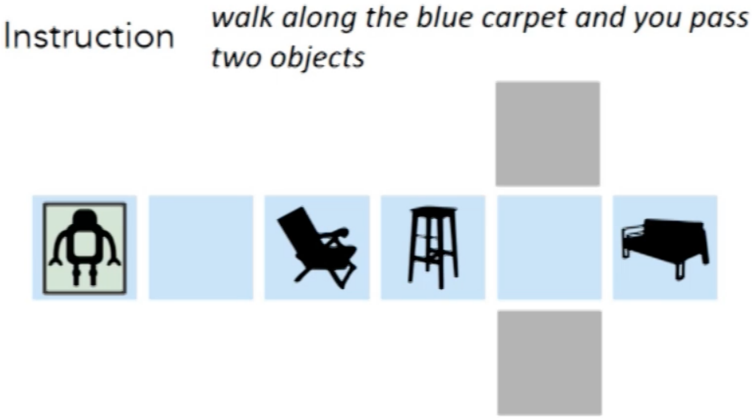

Imagine you were a robot and received this instruction. What would you do?

Haushaltsroboter

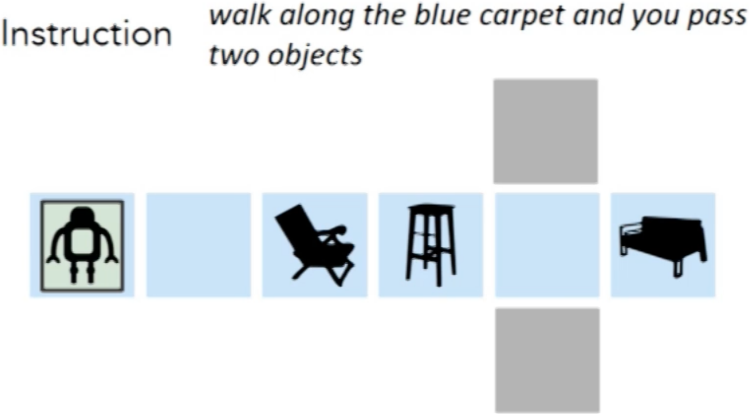

Imagine you were a robot and received this instruction. What would you do?
Can you explain this in terms of Grice's pragmatics (conversational implicature)?
How?

Communication in dialog systems

Need to distinguish understanding from production; both of them are affected by pragmatics.

- Understanding
  Need to draw correct pragmatic inferences, just like a human would draw them.
- Production
  Need to be understood by humans. Therefore need to take into account what a human would understand (potentially even including the pragmatic inferences they might draw).

Let's start out with a look at language production.

Choices in language production

- What content should be included/omitted?
- How should that content be organised to be coherent?
- Which syntactic construtions should be used?
- How should entities be referred to?
- Which words should be chosen?

Choices in language production

- What content should be included/omitted?
- How should that content be organised to be coherent?
- Which syntactic construtions should be used?
- How should entities be referred to?
- Which words should be chosen?

We here focus on a sub-task of NLG: generation of “referring expressions”

Referring expression generation

Dale and Reiter 1995: “Computational Interpretations of the Gricean Maxims in the Generation of Referring Expressions”

Referring expression generation

Dale and Reiter 1995: “Computational Interpretations of the Gricean Maxims in the Generation of Referring Expressions”

Dale and Reiter focus on referring expressions that:

- are realised as definite NPs
- refer to physical objects
- their communicative goal is solely to identify a target object

Examples

the black dog, the woman with the glasses, the upside-down cup

Referring expression generation

Dale and Reiter 1995: "Computational Interpretations of the Gricean Maxims in the Generation of Referring Expressions"

Dale and Reiter focus on referring expressions that:

- are realised as definite NPs
- refer to physical objects
- their communicative goal is solely to identify a target object

Examples

the black dog, the woman with the glasses, the upside-down cup

Dale and Reiter focus on three criteria an algorithm for GRE should satisfy:

1. it should produce expressions that satisfy the communicative goal: that allow the hearer to identify the intended object
2. it should produce expressions that do not lead the hearer to derive false implicatures
3. it should be computationally efficient / similar to how humans refer

Example for choice of category

"Look at the dog!" vs. "Look at the pitbull!"

When there is just a single dog, using the more specific term "pitbull" may lead to unintended pragmatic inferences (e.g., warning of a danger).
When there are two dogs, a dalmatian and a pitbull, then using the more specific term "pitbull" is necessary to refer successfully.

Another example

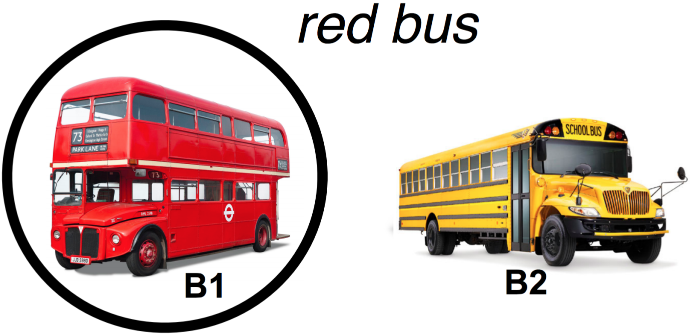

(Cohn-Gordon et al., 2018)

- "red bus" is a good referring expression in this context, as it allows to identify the target B1 and is not overly specific
- "old red English bus" refers to same object, but hearer may draw additional inferences (i.e., what is important about this being an old English bus).

Another example

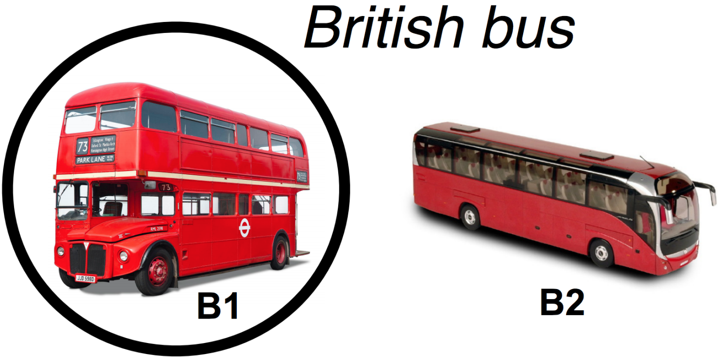

(Cohn-Gordon et al., 2018)

Another example

old British bus

(Cohn-Gordon et al., 2018)

Table of Contents

What does it mean to satisfy a communicative goal?

A referring definite description satisfies its communicative goal if it is a distinguishing description.

- Let D be the set of entities that are in the focus of attention of speaker and hearer (the context set).
- Each entity in D is characterised by means of a set of properties or attribute-value pairs such as $\langle \texttt{colour}, \texttt{red} \rangle$ or `colour=red`.
- If a property p does *not* apply to an entity $d \in D$, we say that p rules out d.
- Let $r \in D$ be the target referent, and C the contrast set: the set of all elements in D except r.

A set L of properties is a distinguishing description if the following conditions hold:

C1. Every property in L applies to r.

C2. For every $c \in C$, there is at least one property in L that rules out c.

An example case

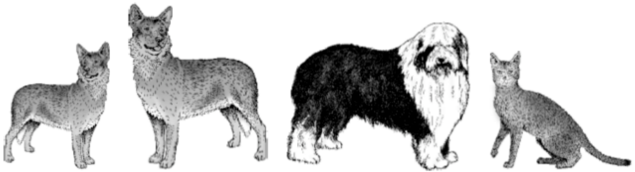

d_1 d_2 d_3 d_4

$D = \{d_1, d_2, d_3, d_4\}$
$r = d_1$
$C = \{d_2, d_3, d_4\}$

Knowledge base representing the scene:

d_1 : `type=dog, size=small, color=brown`
d_2 : `type=dog, size=large, color=brown`
d_3 : `type=dog, size=large, color=black+white`
d_4 : `type=cat, size=small, color=brown`

Some examples of possible descriptions in this scenario:

content determination	possible realisation	distinguishing
$L = \{$type=dog,size=small$\}$	*'the small dog'*	✓
$L = \{$type=dog,colour=brown$\}$	*'the brown dog'*	×
$L = \{$type=dog,size=small,colour=brown$\}$	*'the small brown dog'*	✓

Are all distinguishing descriptions equally felicitous or appropriate?

Full brevity algorithm

The full brevity algorithm calculates all possible referring expressions and then selects the shortest one that describes the target object unambiguously.

Full brevity algorithm

The full brevity algorithm calculates all possible referring expressions and then selects the shortest one that describes the target object unambiguously.

Plausibility of this algorithm for human language production?

Full brevity algorithm

The full brevity algorithm calculates all possible referring expressions and then selects the shortest one that describes the target object unambiguously.

Plausibility of this algorithm for human language production?

Observation 1: Human speakers in many cases include unnecessary modifiers in the referring expressions they construct.

Observation 2: Human speakers can begin to utter a referring expression before they have finished scanning the set of distractors.

Incremental algorithm

The incremental algorithm simply sequentially iterates through a (task-dependent) list of attributes, adding an attribute to the description being constructed if it rules out any distracters that have not already been ruled out, and terminates when a distinguishing description has been constructed. (Dale and Reiter 1995)

Incremental algorithm

Example

Object1: <type, Chihuahua>, <size, small> , <colour, black>
Object2: <type, Chihuahua>, <size, large>, <colour, white>
Object3: <type, Siamese-cat> , <size, small > , <colour, black>

target referent = Object1;
Distractors = (Object2, Object3);
Attribute preference P = type, colour, size,. . . (a list ordered in terms of how much people like using it in their referring expressions)

Algorithm:

- first attribute: type; use base expression: “dog”; check what is ruled out (Object 3).
- try more specific type: “chihuahua”; however this doesn’t rule out additional objects, so stay with “dog”.
- take second attribute: color; “black” rules out Object 2.
- all distractors have been ruled out. So terminate and say “black dog”.

Table of Contents

The Rational Speech Act (RSA) model

The rational speech act (RSA) theory is a framework for pragmatic reasoning, that implements on core ideas regarding language communication as **rational, goal-directed behaviour going back to Grice**, but builds on more recent developments in **game theory and probabilistic modeling**.

The Rational Speech Act (RSA) model

The rational speech act (RSA) theory is a framework for pragmatic reasoning, that implements on core ideas regarding language communication as **rational, goal-directed behaviour going back to Grice**, but builds on more recent developments in **game theory and probabilistic modeling**.

Rational speech act (RSA) models provide **a quantitative framework** to capture intuitions about pragmatic reasoning in language understanding.

The Rational Speech Act (RSA) model

The rational speech act (RSA) theory is a framework for pragmatic reasoning, that implements on core ideas regarding language communication as **rational, goal-directed behaviour going back to Grice**, but builds on more recent developments in **game theory and probabilistic modeling**.

Rational speech act (RSA) models provide **a quantitative framework** to capture intuitions about pragmatic reasoning in language understanding.

Extensions to RSA that allow for reasoning about the speaker (for instance, her goals and word usage) can capture many otherwise puzzling phenomena, including vagueness, embedded implicatures, hyperbole, irony, and metaphor.

(Frank & Goodman, 2012; Goodman & Frank, 2016)

The Rational Speech Act (RSA) model

Literal Listener

When you hear someone say "blue" to refer to one of the objects, what do you think they mean?

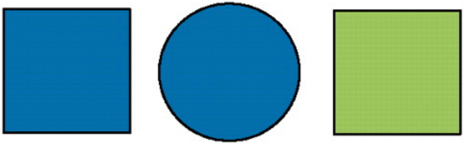

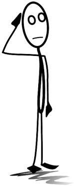

Literal listener L_0

(Frank and Goodman, 2012)

The Rational Speech Act (RSA) model

Literal Listener

When you hear someone say "blue" to refer to one of the objects, what do you think they mean?

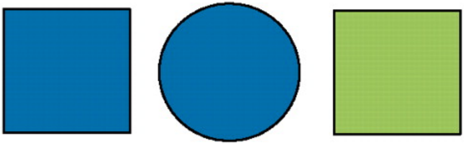

	blue square	blue circle	green square
"blue"	true	true	false
"circle"	false	true	false

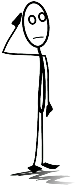

Literal listener L_0

(Frank and Goodman, 2012)

The Rational Speech Act (RSA) model

Literal Listener

When you hear someone say "blue" to refer to one of the objects, what do you think they mean?

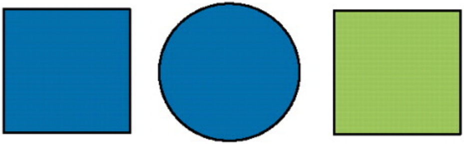

	blue square	blue circle	green square
"blue"	true	true	false
"circle"	false	true	false

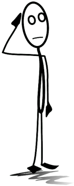

Literal listener L_0

(Frank and Goodman, 2012)

- "Blue" is **literally true** for both the blue square and the blue circle.

The Rational Speech Act (RSA) model

Pragmatic speaker

Imagine you are talking to someone, which word will you use to refer to **the object in the middle**, “blue” or **“circle”**?

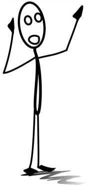

Pragmatic speaker S_1

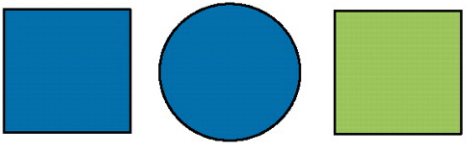

The Rational Speech Act (RSA) model

Pragmatic speaker

Imagine you are talking to someone, which word will you use to refer to **the object in the middle**, "blue" or **"circle"**?

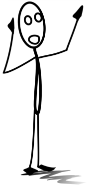

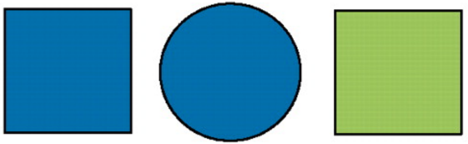

Pragmatic speaker S_1

	blue square	blue circle	green square
"blue"	true	true	false
"circle"	false	true	false

(Frank and Goodman, 2012)

The Rational Speech Act (RSA) model

Pragmatic speaker

Imagine you are talking to someone, which word will you use to refer to **the object in the middle**, "blue" or **"circle"**?

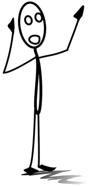

Pragmatic speaker S_1

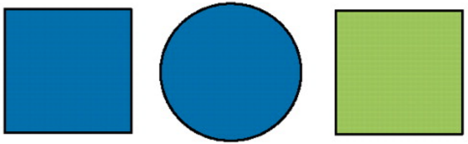

	blue square	blue circle	green square
"blue"	true	true	false
"circle"	false	true	false

(Frank and Goodman, 2012)

- "Blue" is **literally true** for the blue circle in middle.
  - But it is better to say **"circle"**. Why?
  - Because if you say "blue", the listener may think you meant the blue square.

The Rational Speech Act (RSA) model

Pragmatic speaker

Imagine you are talking to someone, which word will you use to refer to **the object in the middle**, “blue” or **“circle”**?

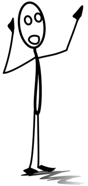

Pragmatic speaker S_1

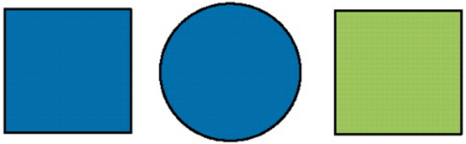

	blue square	blue circle	green square
“blue“	true	true	false
“circle“	false	true	false

(Frank and Goodman, 2012)

- “Blue” is **literally true** for the blue circle in middle.
  - But it is better to say **“circle“**. Why?
  - Because if you say “blue”, the listener may think you meant the blue square.
- So the **pragmatic** speaker should say **“circle”**, because it is **less ambiguous for the listener**.

The Rational Speech Act (RSA) model

Pragmatic Listener

When you hear someone say "blue" to refer to one of the objects, which do you think it is, the blue square or the blue circle?

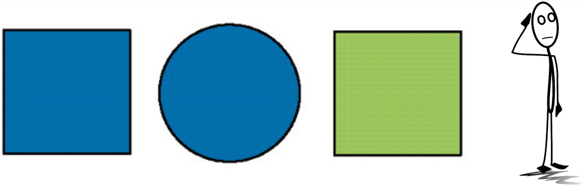

Pragmatic Listener
L_1
(Frank and Goodman, 2012)

The Rational Speech Act (RSA) model

Pragmatic Listener

When you hear someone say "blue" to refer to one of the objects, which do you think it is, the blue square or the blue circle?

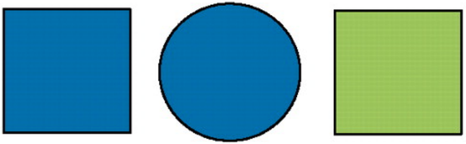

	blue square	blue circle	green square
"blue"	true	true	false
"circle"	false	true	false

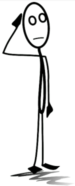

Pragmatic Listener L_1

(Frank and Goodman, 2012)

- "Blue" is **literally true** for both the blue square and the blue circle.
  - But it is more likely that the speaker means the **blue square**. Why?
  - Because if he had meant the blue circle, he should have said the less ambiguous **"circle"**

The Rational Speech Act (RSA) model

Pragmatic Listener

When you hear someone say “blue” to refer to one of the objects, which do you think it is, the blue square or the blue circle?

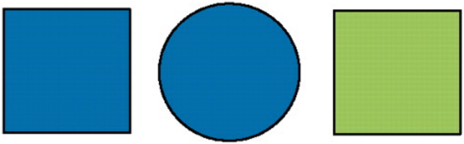

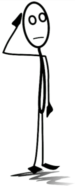

	blue square	blue circle	green square
“blue“	true	true	false
“circle“	false	true	false

(Frank and Goodman, 2012)

- “Blue” is **literally true** for both the blue square and the blue circle.
  - But it is more likely that the speaker means the **blue square**. Why?
  - Because if he had meant the blue circle, he should have said the less ambiguous **“circle”**
- So the **pragmatic** listener would interpret **“blue”** as the **blue square**.

The Rational Speech Act (RSA) model

According to RSA, speakers and listeners reason about each other, but the reasoning is **grounded on the literal meaning of the utterances**.

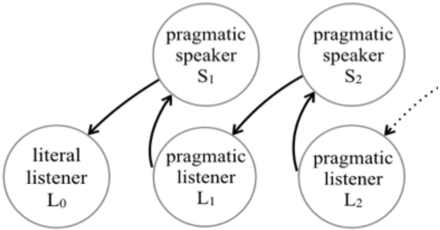

The **speaker** reasons about '**what the listener may misunderstand** if I say this'.

In turn, **listeners** also reason about '**what the speaker could have said** instead of what she actually said, so as to avoid my misunderstanding'.

Model predictions vs. experimental data

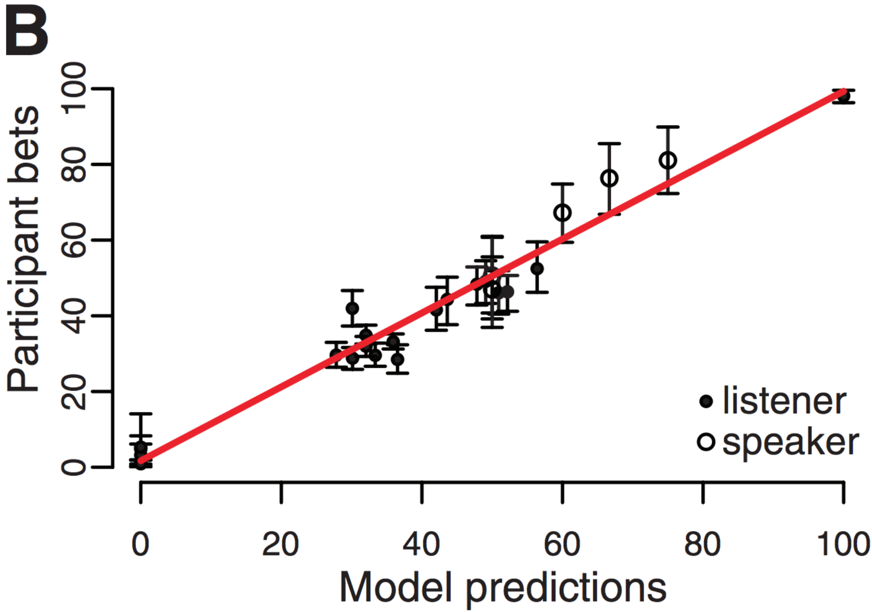

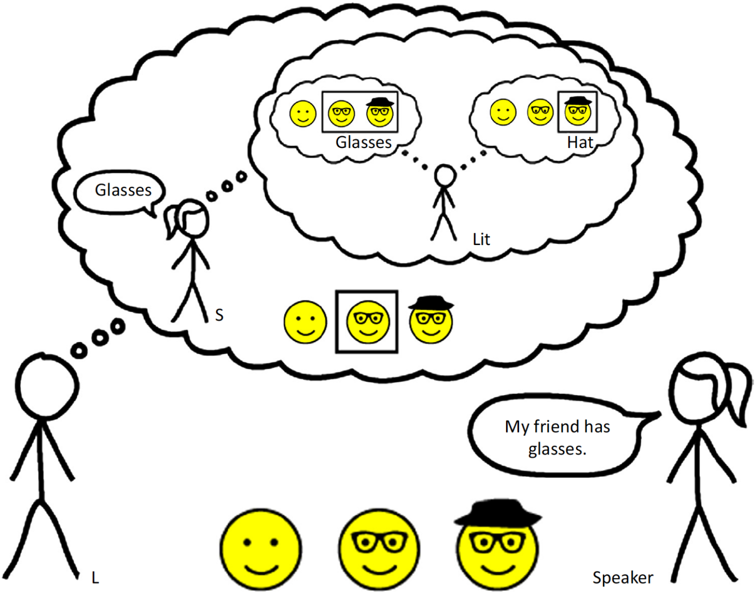

Trends in Cognitive Sciences

RSA model

These concepts of RSA are formally defined by the following probabilities:

- $P_{Speaker}(words|referent, context)$
- $P_{Listener}(referent|words, context)$

speaker model

- the speaker is approximately rational; that is, she chooses her utterances in proportion to the utility she expects to gain:

$$P_S(w|r_S, c) \propto \exp\left(\alpha Utility(w; r_S, c)\right)$$
$$Utility(w; r_S, c) = \log P_L(r|w, c) - cost(w)$$

(w: words; r: referent; c:context)

- $P_L(r|w, c)$ is the literal listener which assigns uniform probability to all literally true referring expressions.

Formal Model

- $P_{L_0}(r|w,c)$ is the literal listener which assigns uniform probability to all literally true referring expressions.
- $P_S(w|r_S,c)$ is the level-1 pragmatic speaker

$$P_S(w|r_S,c) \propto \exp\left(\alpha Utility(w;r_S,c)\right)$$
$$Utility(w;r_S,c) = \log P_{L_0}(r_S|w,c) - cost(w)$$

- $P_{L_1}(r|w,c)$ is the first-level pragmatic listener.

$$P_{L_1}(r|w,c) \propto P_S(w|r,c) * P(r)$$

Table of Contents

Computational challenges

For the pragmatic listener, we need to estimate:

$$P_{L_1}(r|w,c) \propto P_S(w|r,c) * P(r)$$

To be able to estimate the probability $P(r)$ correctly, we need to know the likelihood of r (this can be a referent, or an event that is being talked about). To be able to do this automatically for any problem, one would need to have extensive "world knowledge", i.e., we need to know exactly which events normally follow one another.

Computational challenges

For the pragmatic listener, we need to estimate:

$$P_{L_1}(r|w,c) \propto P_S(w|r,c) * P(r)$$

To be able to estimate the probability $P(r)$ correctly, we need to know the likelihood of r (this can be a referent, or an event that is being talked about). To be able to do this automatically for any problem, one would need to have extensive "world knowledge", i.e., we need to know exactly which events normally follow one another.

In order to estimate the probability $P(w|r,c)$, we have to compute which alternative formulations w are appropriate in context c for referent r.

Instruction giving (Listener)

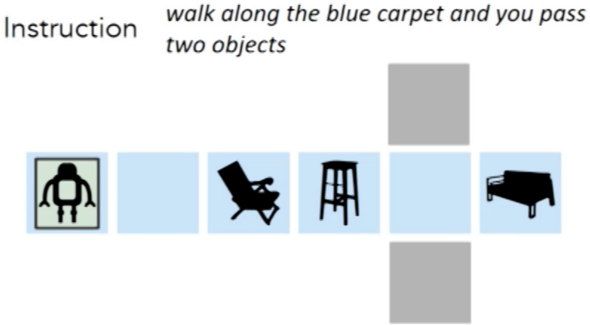

(Fried et al., 2018)

Instruction giving (Listener)

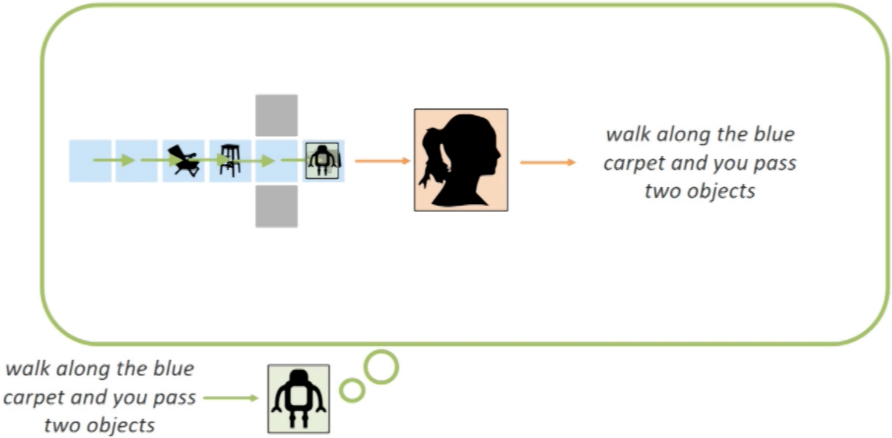

The robot has to consider whether the given instruction is a good way to express a particular intention. If he had to get to the sofa, there would have been another, more appropriate instruction that the speaker could have said.

(Fried et al., 2018)

Instruction giving: (Speaker)

Now let's change the perspective and think about what the speaker would have to say for the listener to get it right.

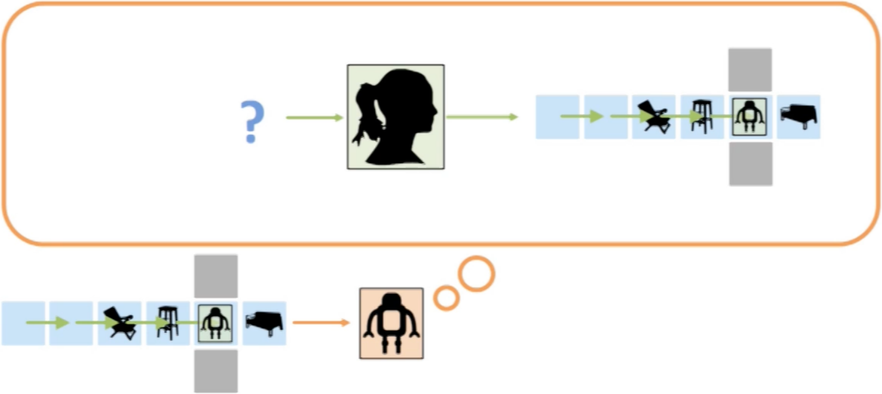

(Fried et al., 2018)

Instruction giving: (Speaker)

Now let's change the perspective and think about what the speaker would have to say for the listener to get it right.

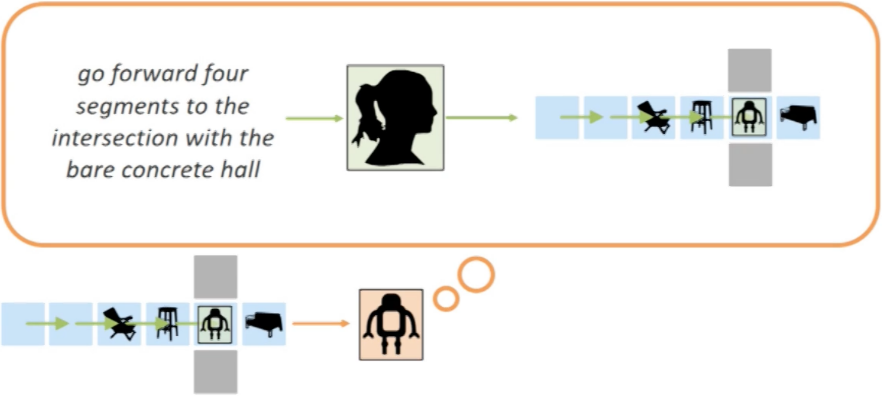

(Fried et al., 2018)

Instruction giving: Model

First, we need a model for the literal interpretation of the instruction, and a model for how an intended instruction can be formulated.

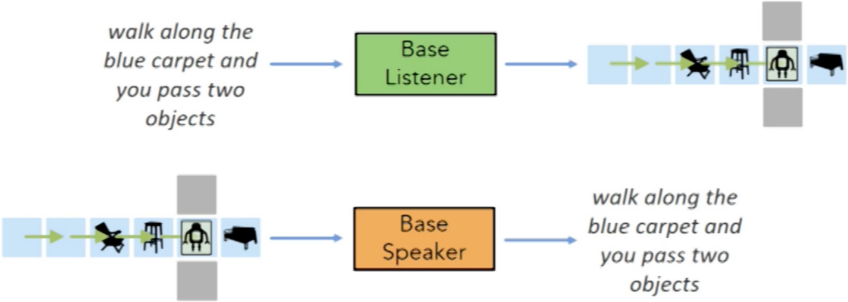

These models help us to estimate the probability of an instruction give an intention, and for estimating an interpretation given an instruction.

Instruction giving: Model for pragmatic listener

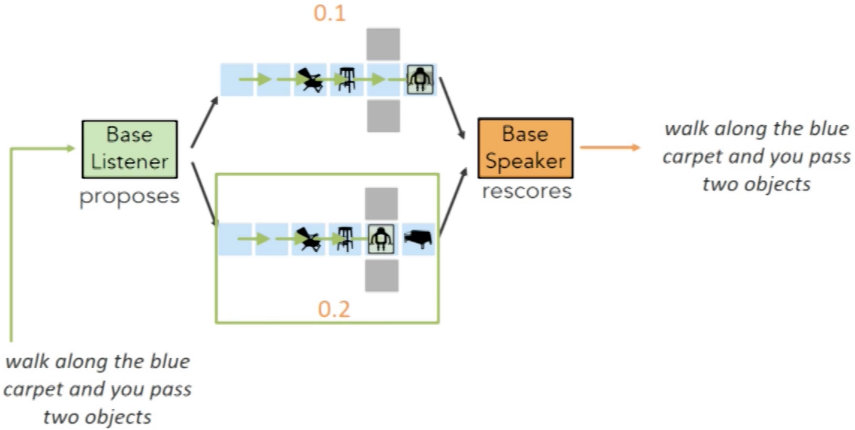

Instruction giving: Model

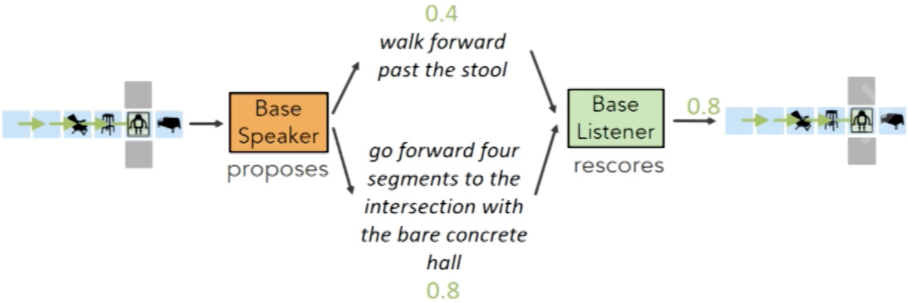

Instruction giving: Model Implementation

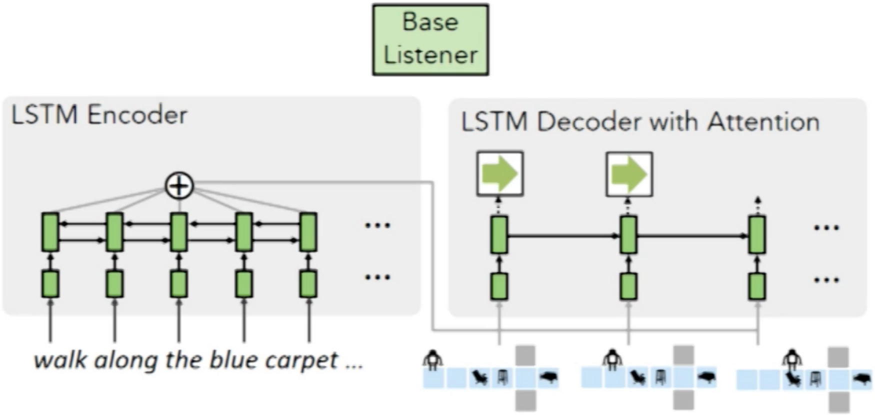

Instruction giving: Model Implementation

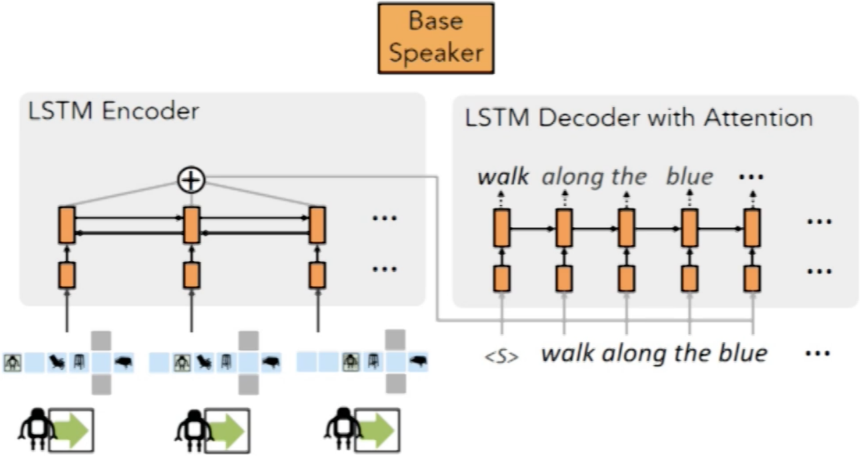

Instruction giving: Model Performance

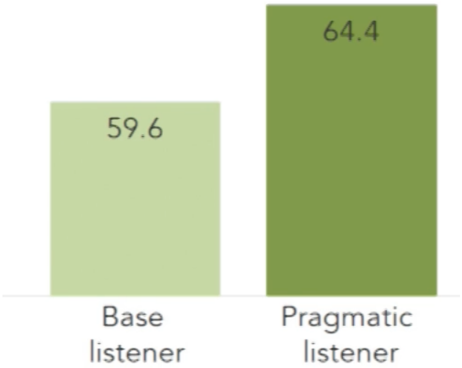

Instruction giving: Summary

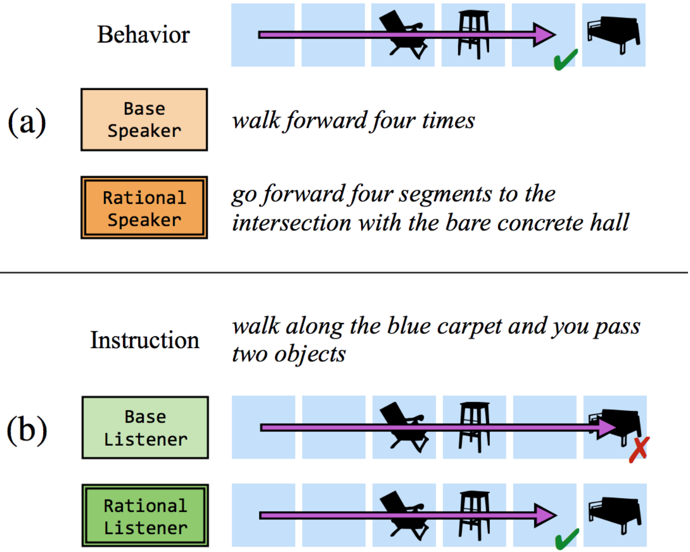

(Fried et al., 2018)

Instruction giving: Model

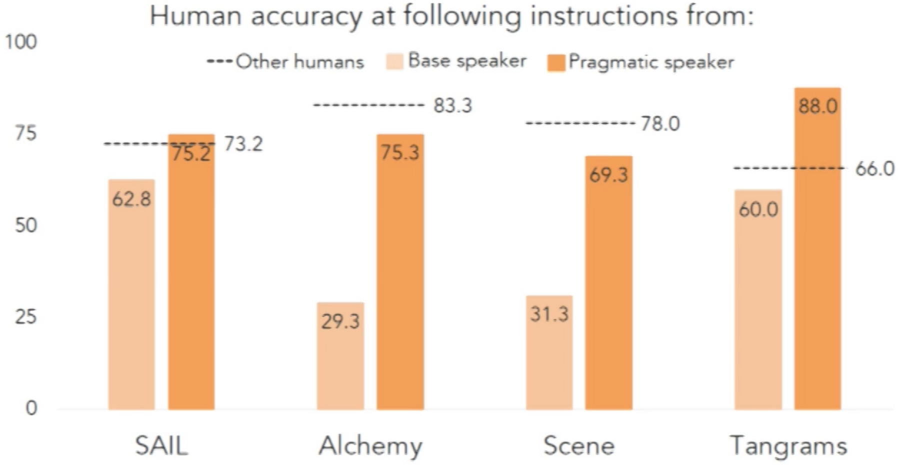

Instruction giving: Model

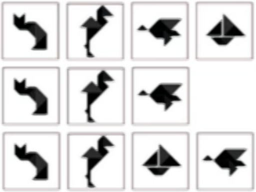

Base Speaker	*remove the last figure* *add it back*	✗
Pragmatic Speaker	*remove the last figure* *add it back in the 3rd position*	✓
Human	*take away the last item* *undo the last step*	✗

Instruction giving: Model

A more realistic domain (here, we also need image recognition, there are a lot more possible referents and more competitors).

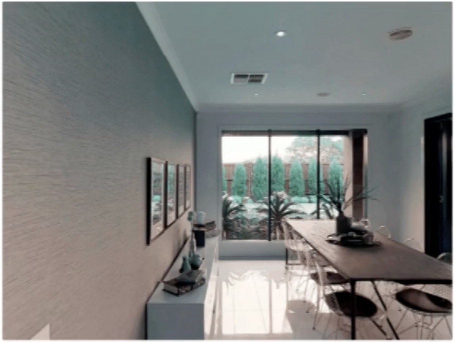

human description:
walk through the kitchen. go right into the living room and stop by the rug.

base speaker:
walk past the dining room table and chairs and wait there .

pragmatic speaker:
walk past the dining room table and chairs and take a right into the living room . stop once you are on the rug .

Summary

- In addition to literal meaning, humans frequently infer additional "pragmatic" meaning given the utterance context.
- Paul Grice proposed conversational maxims that allow one to derive additional meanings in a principled way.
- The strict implementation of Gricean maxims in a referring expression algorithm calculates non-redundant referring expressions, but it does not fit well with human data (humans often overspecify), and it is computationally expensive.
- The incremental algorithm is a lot more efficient and fits human data better
- The rational speech act model (RSA) combines ideas based on Grice (reasoning about speakers and listeners) with game theoretic models and probability theory.
- RSA allows us to make quantitative predictions of pragmatic inferences (not just qualitative ones, as we can do based on Gricean maxims).
- Recent instruction giving models can reason rationally following the RSA model, even for more complex environments.