Connecting Language, Perception and Interaction using Type Theory with Records

Staffan Larsson

Centre for Linguistic Theory and Studies in Probability (CLASP)
Dept. of Philosophy, Linguistics and Theory of Science
University of Gothenburg

Referential Semantics One Step Further, ESSLLI 2016
August 23, 2016

Outline

Introduction

- Questions
  - How is linguistic meaning related to perception?
  - How do we learn and agree on the meanings of our words?
- We are developing a formal *judgement-based semantics* where notions such as perception, classification, judgement, learning and dialogue coordination play a central role
  - See e.g. Cooper (2005), Cooper and Larsson (2009), Larsson (2011), Dobnik *et al.* (2011), Cooper (2012), Dobnik and Cooper (2013), Cooper *et al.* (2015a)
- Key idea:
  - modeling (perceptual) meanings as classifiers of real-valued (perceptual) data, and training these classifiers in interaction with the world and other agents
- This presentation based on Larsson (2011) and Larsson (2015)

Classification

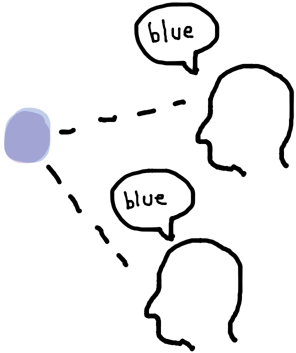

Classification is subjective?

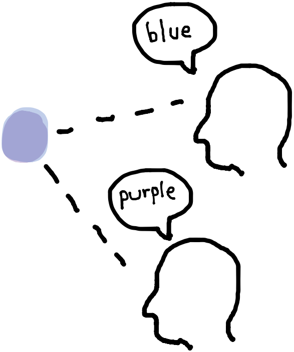

Coordination process

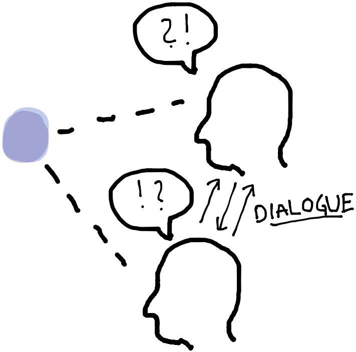

Classification is coordinated

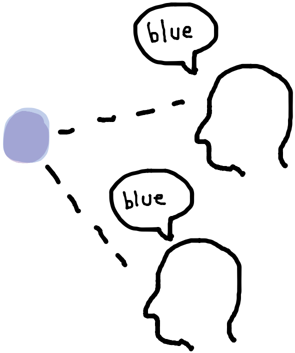

Classification is coordinated

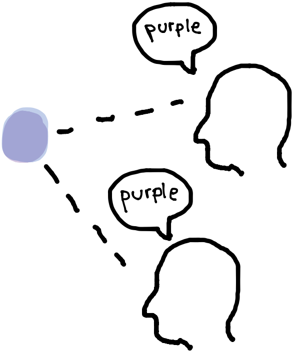

Coordination can be creative

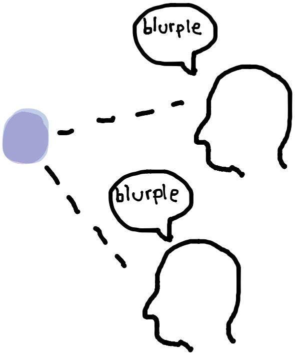

- What is meaning?
  - When a community is coordinated on the use of an expression, that expression has meaning in that community; it can be used for communicating
- Meaning is regarded as being acquired by an agent through its perception of, and interaction with, the world and other agents.
- This makes meaning agent-relative but essentially
  - *social* and *intersubjective*, in the sense of being coordinated in interaction between individuals
  - *dynamic*, in the sense of always being up for revision and negotiation as new perceptual and conversationally mediated information is encountered

Outline

Communicative grounding

- Utterances incrementally add to Common Ground
  - The collection of mutual knowledge, mutual beliefs, and mutual assumptions that is essential for communication between two people (Clark and Schaefer, 1989)
- "To ground a thing ... is to establish it as part of common ground well enough for current purposes."
- Making sure that the participants are perceiving, understanding, and accepting each other's utterances; dealing with miscommunication
  - See e.g. Clark and Schaefer (1989), Clark and Brennan (1990), Clark (1996)

Semantic coordination

- Research on alignment shows that agents negotiate domain-specific microlanguages for the purposes of discussing the particular domain at hand
  - See e.g. Clark and Gerrig (1983), Clark and Wilkes-Gibbs (1986), Garrod and Anderson (1987), Pickering and Garrod (2004), Brennan and Clark (1996), Healey (1997), Larsson (2007)
- Two agents do not need to share exactly the same linguistic resources (grammar, lexicon etc.) in order to be able to communicate
- An agent's linguistic resources can change during the course of a dialogue when she is confronted with a (for her) innovative use
- *Semantic coordination*: the process of interactively coordinating the meanings of linguistic expressions

Communicative grounding and semantic coordination

- Two kinds of coordination in dialogue:
  - Information coordination: agreeing on information (facts, what is true, what the relevant questions are, etc.)
  - Language coordination: agreeing on how to talk; incl. semantic coordination

Semantic coordination

- Semantic coordination can occur as a side-effect of information coordination, e.g.
  - Acknowledgements
  - Clarification requests
  - Repair
  - Accommodation/deference: "silent" coordination where a DP observes the language use of another and adapts to it
- There are also dialogue strategies whose primary purpouse is to aid semantic coordination, e.g.
  - Word meaning negotiation / litigation (Myrendal, 2015; Ludlow, 2014)
  - Corrective feedback
  - Clarification requests

Examples of semantic coordination strategies in 1LA

- "non-repair" indirect offer:
  - D (1;8.2, having his shoes put on; points at some ants on the floor): Ant. Ant.
  - Father (indicating a small beetle nearby): And that's a bug.
  - D: bug.
- offers-in-repairs
  - explicit
    - explicit replace ("That's not an X, that's a Y")
    - clarification question ("You mean Y?")
  - implicit/embedded (reformulation, corrective feedback)

Examples of semantic coordination strategies in 1LA, cont'd

(examples from Eve Clark et. al., most from CHILDES corpus)

- Example 1: "In-repair"
  - Abe: I'm trying to tip this over, can you tip it over? Can you tip it over?
  - Mother: Okay I'll turn it over for you.
- Example 2: Clarification request
  - Adam: Mommy, where my plate?
  - Mother: You mean your saucer?
- Example 3: "Explicit replace"
  - Naomi: Birdie birdie.
  - Mother: Not a birdie, a seal.
- Example 4: "Bare" correction
  - Naomi: mittens.
  - Father: gloves.

Outline

The Symbol Grounding Problem

- If a speaker of English is unable to distinguish gloves from mittens, most people would probably agree that something is missing in this person's knowledge of the meaning of "glove".
- Similarly, if we tell A to find some nice pictures of dogs chasing cats, and A comes back happily with an assortment of pictures displaying lions chasing zebras, we would question whether A really knows the full meaning of the words "dog" and "cat"

Perception and meaning

- Part of learning a language is learning to identify individuals and situations that are in the extension of the phrases and sentences of the language
- For many concrete expressions, this identification relies crucially on the ability to
  - perceive the world
  - use perceptual information to classify individuals and situations as falling under a given linguistic description or not
- This view was put forward by Harnad (1990) as a way of addressing the "symbol grounding problem" in artificial intelligence:

> *How can the meanings of the meaningless symbol tokens, manipulated solely on the basis of their (arbitrary) shapes, be grounded in anything but other meaningless symbols?"*

(Harnad, 1990)

How to solve the symbol grounding problem

- Harnad's own sketch of a solution to the symbol grounding problem:
  - A **hybrid** system encompassing both symbolic and non-symbolic representations, the latter such that they "can pick out the objects to which they refer, via connectionist networks that extract the invariant features of their analog sensory projections"
  - **Learning** non-symbolic representations from interaction; "a connectionist network that learns to identify icons correctly from the sample of confusable alternatives it has encountered by dynamically adjusting the weights of the features"
  - **Compositionality**, where complex constructions "will all inherit the intrinsic grounding of [the grounded set of elementary symbols]"
- All these components are needed for a solution to the symbol grounding problem
- We follow these ideas, specify them further and formalize them

Statistical classifiers

- Harnad proposed using connectionist networks to ground symbols
  - This was also followed by Steels and Belpaeme (2005)
- Connectionist networks are one kind of *(statistical) classifier*, a computational device determining what class an item belongs to, based on various properties of the item.
- Crucially, these properties need not be encoded in some high-level representation language (such as logic or natural language)
- Instead, it may consist entirely of numeric data encoding more or less "low-level" information about the item in question, for example perceptual data.

Classifiers, intensions and extensions

- Classifiers can be defined formally as mathematical functions.
- Typically, the domain of a classifier function is numerical (e.g. real-valued, integer or binary) vectors and the range is a set of categories
- When making use of classifiers in formal semantics we will regard them as (parts of) representations of (agents' takes on) intensions of linguistic expressions.
- Classifiers (as intensions) produce judgements whether some perceived thing or situation falls within the extension of a linguistic expression

Perceptual meaning

- Perceptual meaning is an important aspect of the meaning of linguistic expressions referring to physical objects (such as concrete nouns or noun phrases).
- Knowing the perceptual meaning of an expression allows an agent to identify perceived objects and situations falling under the meaning of the expression.
- For example, knowing the perceptual meaning of “blue” would allow an agent to correctly identify blue objects.
- Similarly, an agent which is able to compute the perceptual meaning of “a boy hugs a dog” will be able to correctly classify situations where a boy hugs a dog.

Using classifiers to represent perceptual meanings

- Steels & Belpaeme (2005): Robots coordinating on colour terms through a simple language game of pointing and guessing; meanings of colour terms are captured in (weight vectors describing) neural networks; utterances describe single objects
- This can be seen as a further specification implementation of Harnad's ideas, adding interaction to the mix
- We follow Steels & Belpaeme in representing (takes on) meanings using classifiers, and training these classifiers based on dialogue interaction
- We add a connection to formal semantics as well as an account of compositionality

Formal semantics for perceptual meanings

- We want to integrate perceptual meanings and low-level perceptual data into formal semantics
- This means mixing low-level (perceptual) and high-level (logical-inferential) meaning in a single framework
  - A hybrid system, as proposed by Harnad
- To enable learning and coordination, we need a framework where intensions
  1) are represented independently of extensions, and
  2) are structured objects which can be modified (updated)
  3) can be modeled as classifiers of perceptual data
- (Possible worlds semantics does not represent intensions independently of extensions)

Type Theory with Records

- We will be using Type Theory with Records, or TTR (Cooper, 2012)
- TTR starts from the idea that information and meaning is founded on our ability to perceive and classify the world
- Based on the notion of *judgements* of entities and situations being of certain *types*
- TTR integrates logical techniques such as binding and the lambda-calculus into feature-structure like objects called *record types*
- Feature structure-like properties are important for the straightforward definition of meaning modifications
- Logical aspects are important for relating our semantics to the model- and proof-theoretic tradition associated with compositional semantics

Related work

- Perceptual aspects of meanings have been explored in previous research, e.g. Barsalou *et al.* (2003),Roy (2005),Steels and Belpaeme (2005),Kelleher *et al.* (2005),Skočaj *et al.* (2010).
  - However, the connection to logical-inferential meaning and compositionality as traditionally studied in formal semantics has not been a focus of this body of work.
- There have also been attempts to extend semantic formalisms to cover embodied meaning, e.g. Feldman (2010)
  - However, this line of work has tended to concentrate on abstract (high-level) representations and has generally not paid attention to low-level perceptual aspects of context.
- More recently, there has been computational work which is more in line with the approach taken here, e.g. Kennington and Schlangen (2015)
  - We propose a way of connecting this line of work to formal semantics, to enable combining it with the successes of formal semantics (compositionality, quantification, etc.)

The Perceptron

- The general account is intended to work for all kinds of classifiers
- As a simple example of how perceptual classifiers can be integrated in formal semantics, we will use the perceptron (Rosenblatt, 1958)
- Classification of perceptual input can be regarded as a mapping of sensor readings (corresponding to situations) to types
- The perceptron is a very simple neuron-like object with several inputs and one output.

$$o(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{w} \cdot \mathbf{x} > \text{t} \\ 0 & \text{otherwise} \end{cases}$$

  where $\mathbf{w} \cdot \mathbf{x} = \sum_{i=1}^{n} w_i x_i = w_1 x_1 + w_2 x_2 + \ldots + w_n x_n$

- Limited to learning problems which are linearly separable; the distinction between left and right is one such problem.

Classifying objects as being to the left or to the right

- Suppose we have a square surface, and object are placed on the surface
- To classify objects as being to the right or not:
    - Direct a sensor (e.g. a camera) towards the surface
    - Get a sensor reading (a picture from the camera)
    - Apply an algorithm which returns a vector of the coordinates of the object on the surface (assuming there is only one); this is a slightly higher-level rendering of our initial sensor reading
    -

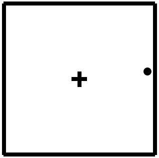

Classifying objects as being to the left or to the right

- Suppose we have a square surface, and object are placed on the surface
- To classify objects as being to the right or not:
  - Direct a sensor (e.g. a camera) towards the surface
  - Get a sensor reading (a picture from the camera)
  - Apply an algorithm which returns a vector of the coordinates of the object on the surface (assuming there is only one); this is a slightly higher-level rendering of our initial sensor reading
  - Apply a perceptron classifier to the coordinate vector and returns 1 or 0

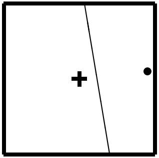

$$\Rightarrow 1$$

The TTR perceptron cont'd

A TTR perceptron classifier can be represented as a record:

$$\kappa = \begin{bmatrix} \mathrm{w} & = & \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} \\ \mathrm{t} & = & 0.090 \\ \mathrm{f} & = & \lambda v : \mathrm{RealVector}(\begin{cases} 1 & \text{if } v \cdot \mathrm{w} > \mathrm{t} \\ 0 & \text{otherwise} \end{cases}) \end{bmatrix}$$

Where $\kappa.\mathrm{f}$ will evaluate to

$$\lambda v : \mathrm{RealVector}\ (\begin{cases} 1 & \text{if } v \cdot \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} > 0.090 \\ 0 & \text{otherwise} \end{cases})$$

- This representation allows modifying w and t by updating the record

The TTR perceptron

- The basic perceptron returns a real-valued number (1 or 0) but when we use a perceptron as a classifier of situations we want it to instead return a type.
- Typically, such types will be built from a predicate and some number of arguments; a type of proof, or a "proposition".

A TTR classifier perceptron for a type P can be represented as a record:

$$\kappa = \begin{bmatrix} \mathrm{w} & = & \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} \\ \mathrm{t} & = & 0.090 \\ \mathrm{f} & = & \lambda v : \mathrm{RealVector}(\begin{cases} P & \text{if } v \cdot \mathrm{w} > \mathrm{t} \\ \neg P & \text{otherwise} \end{cases}) \end{bmatrix}$$

The meaning of "(that is to the) right" in TTR

Uses a TTR classifier perceptron to represent a agent's take on perceptual meaning:

$[\text{right}]^{Agt} =$

$$\left[\begin{array}{l} \text{w} = \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} \\ \text{t} = 0.090 \\ \text{bg} = \left[\begin{array}{lcl} \text{sr}_{\text{pos}} & : & \text{RealVector} \\ \text{foo} & : & \text{Ind} \\ \text{spkr} & : & \text{Ind} \end{array}\right] \\ \text{f} = \lambda r : \text{bg}(\left[\text{c}^{\text{perc}}_{\text{right}} = \begin{bmatrix} \text{foo} = r.\text{foo} \\ \text{sr}_{\text{pos}} = r.\text{sr}_{\text{pos}} \end{bmatrix} : \begin{cases} \text{right}(r.\text{foo}) & \text{if } r.\text{sr}_{\text{pos}} \cdot \text{w} > \text{t} \\ \neg\ \text{right}(r.\text{foo}) & \text{otherwise} \end{cases} \right]) \end{array}\right]$$

(Note how this representation combines low-level real-valued information and high-level logical/inferential information.)

Classifying objects as being to the right or not, TTR style

- Representation of current situation s
  - Coordinates of object in focus of attention
  - Label for object (obj_{45})

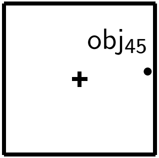

$$s = \begin{bmatrix} \mathrm{sr}_{\mathrm{pos}} = \begin{bmatrix} 0.900 & 0.100 \end{bmatrix} & : & \mathrm{RealVector} \\ \mathrm{foo} = \mathrm{obj}_{45} & : & \mathrm{Ind} \\ \mathrm{spkr} = \mathrm{A} & : & \mathrm{Ind} \end{bmatrix}$$

- Apply [right].f to s:

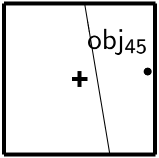

$$\Rightarrow \mathrm{right}(\mathrm{obj}_{45})$$

Outline

Communicative grounding and semantic coordination (reprise)

- Semantic coordination can occur as a side-effect of information coordination, e.g.
  - Accommodation/deference
  - Acknowledgements
  - Clarification requests
  - Repair
- There are also dialogue strategies whose primary purpouse is to aid semantic coordination, e.g.
  - Word meaning negotiation / litigation
  - Corrective feedback
  - Clarification requests
- **How are perceptual meanings learnt/updated based on dialogue interaction?**

The left-or-right game

- A and B are facing a framed surface on a wall, and A has a bag of objects which can be attached to the framed surface.
- A round of the game is played as follows:
  1. A places an object in the frame
  2. B orients to the new object, assigns it a unique individual marker and labels it "foo" in B's take on the situation
  3. A says either "left" or "right"
  4. B interprets A's utterance based on B's take on the situation. Interpretation includes determining whether B's understanding of A's utterance is consistent with B's take on the situation.
  5. If an inconsistency results from interpretation, B assumes A is right, says "aha", and learns from this exchange; otherwise, B says "okay"

- The left-or-right game can be regarded as a considerably pared-down version of the "guessing game" in Steels and Belpaeme (2005), where perceptually grounded colour terms are learnt from interaction.
- The kinds of meanings learnt in the left-or-right game may be considered trivial.
- However, at the moment we are mainly interested in the basic principles of combining formal dynamic semantics with learning of perceptual meaning from dialogue
- The hope is that these can be formulated in a general way which can later be used in more interesting settings.

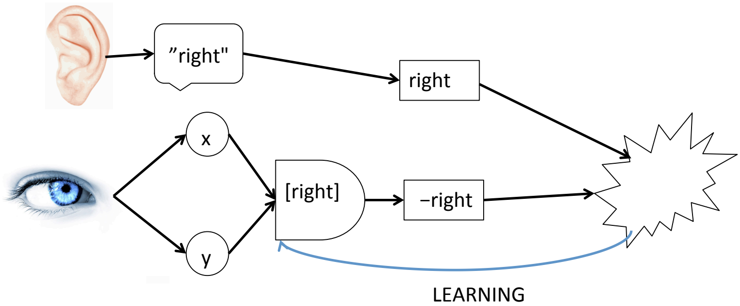

”right"
right
x
y
[right]
−right
LEARNING

Updating perceptual meaning

Perceptrons are updated using the *perceptron training rule*:

$$w_i \leftarrow w_i + \Delta w_i$$

where

$$\Delta w_i = \eta(o_t - o)x_i$$

where o_t is the target output, o is the actual output, and w_i is associated with input x_i.

- Note that if $o_t = o$, there is no learning.
- This rule can be formulated as a TTR update function (see Larsson, 2015)
- In the LoR-game, training results in moving the line dividing "(to the) right" from "not (to the) right"

Agent B's initial take on the meaning of "right":

$[\text{right}]^B =$

$$\left[\begin{array}{l} \text{w} = \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} \\ \text{t} = 0.090 \\ \text{bg} = \begin{bmatrix} \text{sr}_{\text{pos}} & : & \text{RealVector} \\ \text{foo} & : & \text{Ind} \\ \text{spkr} & : & \text{Ind} \end{bmatrix} \\ \text{f} = \lambda r : \text{bg}(\left[\text{c}^{\text{perc}}_{\text{right}} = \begin{bmatrix} \text{foo} = r.\text{foo} \\ \text{sr}_{\text{pos}} = r.\text{sr}_{\text{pos}} \end{bmatrix} : \begin{cases} \text{right}(r.\text{foo}) & \text{if } r.\text{sr}_{\text{pos}} \cdot \text{w} > \text{t} \\ \neg\ \text{right}(r.\text{foo}) & \text{otherwise} \end{cases} \right]) \end{array}\right]$$

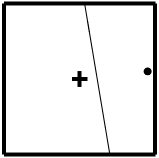

A: “right”

B: “okay”

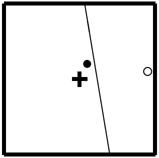

A: “right”

- B's classifier applied to this situation yields that the object is not to the right
- B applies the perceptron training rule to adjust the classifier

Agent B's revised on the meaning of "right":
$[\text{right}]^B =$

$$\left[\begin{array}{l} \text{w} = \begin{bmatrix} \mathbf{0.808} & \mathbf{0.200} \end{bmatrix} \\ \text{t} = 0.090 \\ \text{bg} = \left[\begin{array}{lcl} \text{sr}_{\text{pos}} & : & \text{RealVector} \\ \text{foo} & : & \text{Ind} \\ \text{spkr} & : & \text{Ind} \end{array}\right] \\ \text{f} = \lambda r : \text{bg}(\left[\text{c}^{\text{perc}}_{\text{right}} = \begin{bmatrix} \text{foo} = r.\text{foo} \\ \text{sr}_{\text{pos}} = r.\text{sr}_{\text{pos}} \end{bmatrix} : \begin{cases} \text{right}(r.\text{foo}) & \text{if } r.\text{sr}_{\text{pos}} \cdot \text{w} > \text{t} \\ \neg\, \text{right}(r.\text{foo}) & \text{otherwise} \end{cases} \right]) \end{array}\right]$$

A: "right"

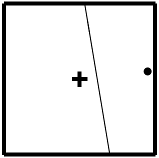

B: "okay"

A: "right"

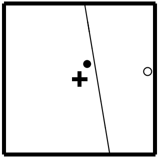

B: "aha"

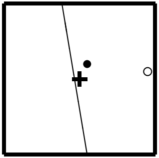

From learning to coordination

- In the left-or-right game, as described above, there is an asymmetry in that agent A is assumed to be fully competent at judging whether objects are to the right or not, whereas agent B is to learn this.
- By contrast, when humans interact they *mutually* adapt to each others' language use on multiple levels
  - alignment (Pickering and Garrod, 2004), entrainment (Brennan, 1996), negotiation (Mills and Healey, 2008) or coordination (Garrod and Anderson, 1987; Healey, 1997; Larsson, 2007)
- The LoR game could quite easily be altered to illustrate coordination directly
  - Let A and B switch roles after each round
  - In this symmetric LoR game, the agents would converge on a meaning of "right" that neither of them may subscribe to initially.

Outline

Compositionality

Nor can categorical representations yet be interpreted as "meaning" anything. It is true that they pick out the class of objects they "name," but the names do not have all the systematic properties of symbols and symbol systems (...). They are just an inert taxonomy. For systematicity it must be possible to combine and recombine them rulefully into propositions that can be semantically interpreted.

(Harnad, 1990)

Compositionality

- A crucial step in demonstrating the usefulness of the proposed approach is to show how the principle of compositionality can be applied also to subsymbolic aspects of meaning
- Exploring compositionality in something like the left-or-right game requires extending it.
  - add more words (e.g. "upper" and "lower") and some simple grammar ("upper left", "lower right" etc).
  - additional sensors and classifiers, e.g. for colour, shape and relative position, can be added, thus enabling meanings of colour and shape terms as well as complex phrases like "the green box is to the left of the upper red circle".

Compositionality: Basic Example

- Proof of concept of compositionality: show how to compute the meaning of "upper right" from the meanings of "upper" and "right".

$[\text{upper}]^B =$

$$\begin{bmatrix} \mathrm{w}_{\mathrm{upper}} = \ldots \\ \mathrm{t}_{\mathrm{upper}} = \ldots \\ \mathrm{bg} = \begin{bmatrix} \mathrm{sr}_{\mathrm{pos}} & : & \mathrm{RealVector} \\ \mathrm{foo} & : & \mathrm{Ind} \\ \mathrm{spkr} & : & \mathrm{Ind} \end{bmatrix} \\ \mathrm{f} = \lambda r : \mathrm{bg}(\begin{bmatrix} \mathrm{c}^{\mathrm{perc}}_{\mathrm{upper}} = \begin{bmatrix} \mathrm{sr}_{\mathrm{pos}} = r.\mathrm{sr}_{\mathrm{pos}} \\ \mathrm{foo} = r.\mathrm{foo} \end{bmatrix} : \pi_{\mathrm{upper}}(\mathrm{w}_{\mathrm{upper}}, \mathrm{t}_{\mathrm{upper}})(r) \end{bmatrix}) \end{bmatrix}$$

Compositionality: Basic Example

Compositional meaning of "upper right" obtained by merging of meanings of "upper" and "right"

$$[\text{upper right}]_B = [\text{upper}]_B \dot{\wedge} [\text{right}]_B =$$

$$\left[\begin{array}{l}
\text{w}_{\text{upper}} = \ldots \\
\text{t}_{\text{upper}} = \ldots \\
\text{w}_{\text{right}} = \ldots \\
\text{t}_{\text{right}} = \ldots \\
\text{bg} = \left[\begin{array}{lcl} \text{sr}_{\text{pos}} & : & \text{RealVector} \\ \text{foo} & : & \text{Ind} \\ \text{spkr} & : & \text{Ind} \end{array}\right] \\
\text{f} = \lambda r{:}\text{bg}\left(\left[\begin{array}{l} \text{c}^{\text{perc}}_{\text{upper}} = \left[\begin{array}{l} \text{sr}_{\text{pos}} = r.\text{sr}_{\text{pos}} \\ \text{foo} = r.\text{foo} \end{array}\right] : \pi_{\text{upper}}(\text{w}_{\text{upper}}, \text{t}_{\text{upper}})(r) \\ \text{c}^{\text{perc}}_{\text{right}} = \left[\begin{array}{l} \text{sr}_{\text{pos}} = r.\text{sr}_{\text{pos}} \\ \text{foo} = r.\text{foo} \end{array}\right] : \pi_{\text{right}}(\text{w}_{\text{right}}, \text{t}_{\text{right}})(r) \end{array}\right] \right)
\end{array}\right]$$

Compositionality: Basic Example

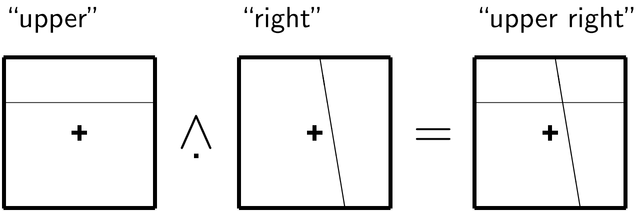

Compositionality: Degree modifiers

- What are the compositional semantics for degree modifiers, e.g. "far" in "far right"
- Proposal: "far" takes parameters of the "right" classifier and yields modified classifier for "far rightness" (increased threshold)

$$[\text{far}] = \begin{bmatrix} \alpha = 1.4 \\ \text{f} = \lambda m{:}\begin{bmatrix} \text{t} & : & \text{Real} \end{bmatrix} \\ \qquad (m \sqcap \begin{bmatrix} \text{t} & = & \alpha * m.\text{t} \end{bmatrix}) \end{bmatrix}$$

$$[\text{far right}] = [\text{far}].\text{f}([\text{right}]) =$$

$$\begin{bmatrix} \text{t} = 0.090 \\ \text{bg} = \ \ldots \\ \text{f} = \ldots \end{bmatrix} \sqcap \begin{bmatrix} \text{t} & = & 1.4{*}0.090 \end{bmatrix} = \begin{bmatrix} \text{t} = 0.126 \\ \text{bg} = \ \ldots \\ \text{f} = \ldots \end{bmatrix}$$

Compositionality: Degree modifiers

"right":

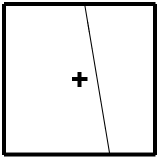

"far right":

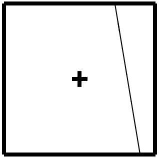

Vagueness

- A weakness of the perceptron classifier is that it does not allow modeling of vague concepts
- What is needed is a "noisy threshold" classifier
- In ongoing work, we are formulating a Bayesian noisy threshold classifier for vague concepts such as "tall"
- The classifier is trained on previous observations tall entities, and is sensitive to the kind of entity
  - skyscraper, human, basketball player, ...
- Instead of a binary judgement, the classifier returns an probabilistic Austinian proposition saying that a situation is of a certain type with a certain probability
- This account connects to the recently developed probabilistic version of TTR (Cooper *et al.*, 2014, Cooper *et al.*, 2015b)

Dialogue strategies for semantic coordination

- Currently, we only model the uncomplicated case where one agent defers to another
- Other dialogue strategies and their role in semantic coordination have been described, but they have not yet been connected to perceptual meanings and symbol grounding
  - Word meaning negotiation / litigation (Myrendal, 2013, Ludlow, 2014)
  - Corrective feedback (Larsson and Cooper, 2009)
  - Clarification requests (Cooper and Ginzburg, 2001; Cooper, 2010)
  - ...

Outline

Summary

- We take it that a central task of semantic theory is to model semantic plasticity and semantic coordination, as well as to connect language and the world
- By modelling how individuals (1) represent meanings, (2) use meanings to form judgements and (3) coordinate on meanings and judgements, we indirectly model the emergence, perpetuation and variation of meaning in a linguistic community.
- Although our representations concern individual agents, meaning itself is inherently social and dependent on learning and adaptation through interaction
- By incorporating classifiers into formal semantics as a way of representing perceptual meanings, and by training these classifiers in interaction, we show how these meanings are related to (perception of) the world and to interaction

Connection to workshop questions

- We provide a lexical semantics for terms referring to concrete and observable objects and properties (perceptual meanings), modeling their descriptive content in terms of classifiers
- This approach connects formal and cognitive semantics by modeling perceptual meanings as classifiers whose outputs are logical-inferential types

Connection to distributional semantics

- Classifiers can perhaps be regarded as models of distributions in terms of co-occuring low-level (sensory) data derived from an observable situation
- We do not currently model co-occurring language distributionally
- We model compositionality not on the level of low-level data (as in standard distributional semantics) but on the level of classifiers
  - object classified as being "upper right" if it is "upper" and "right"
  - object classified as being "far right" using the "right" classifier modified by the perceptual meaning of "far"

Connections to Barbara's talk: Semantic primitives

- A problem with Leibniz' characteristica universalis (or the idea of a fixed set of static semantic primitives in general):
  - It may be the case that the "primitive" (most basic) features underlying perceptual meanings are themselves dynamic and trained (indirectly or directly) in social interaction as part of semantic coordination
- If we take the "primitive features" to be low-level features used (and possibly discovered) by classifiers (cf. deep learning), they may not necessarily make sense to us
- If we take "primitive features" to be the lowest level of logical-inferentual types (e.g. the features detected by the classifiers that represent perceptual meanings of concrete words), they will make sense to is but they will still be dynamic

Connections to Barbara's talk: Beech vs. Elm

- These may have the same perceptual semantics (for me) in the sense that I use a single classifier for both; modeling perceptual meanings as classifiers in fact offers a way of modeling this difference between speakers
- ...but I can still be aware that they are different things (perceptual meanings are not the only kinds of meanings in TTR)...
- ...and other speakers may have more elaborate (perceptual and logical-inferentual) takes on the meanings of "beech" and "elm"
- In concrete interactions, the requirements of the situation and activity at hand will affect whether we can still use these words to communicate; if not, semantic coordination may ensue.
- In this coordination process, it is more likely that the expert will inform the novice than the other way around; power matters in semantic coordination (as in any negotiation)

Barsalou, L. W.; Simmons, W. K.; Barbey, A. K.; and Wilson, C. D. 2003.
Grounding conceptual knowledge in modality-specific systems.
Trends in Cognitive Sciences 7:84–91.

Brennan, S. E. and Clark, H. H. 1996.
Conceptual pacts and lexical choice in conversation.
Journal of Experimental Psychology: Learning, Memory and Cognition 22:482–493.

Brennan, Susan E. 1996.
Lexical Entrainment in Spontaneous Dialog.
In *International Symposium on Spoken Dialog*.
41–44.

Clark, H. H. and Brennan, S. E. 1990.
Grounding in communication.
In Resnick, L. B.; Levine, J.; and Behrend, S. D., editors 1990, *Perspectives on Socially Shared Cognition*. APA.
127 – 149.

Clark, H. H. and Gerrig, R. J. 1983.
Understanding old words with new meanings.
Journal of Verbal Learning and Verbal Behavior 22:591–608.

Clark, Herbert H. and Schaefer, Edward F. 1989.
Contributing to discourse.
Cognitive Science 13:259–294.
Also appears as Chapter 5 in Clark (1992).

Clark, H. H. and Wilkes-Gibbs, D. 1986.
Refering as a collaborative process.
Cognition 22:1–39.

Clark, Herbert H. 1992.
Arenas of Language Use.
University of Chicago Press.

Clark, H. H. 1996.
Using Language.
Cambridge University Press, Cambridge.

Cooper, Robin and Ginzburg, Jonathan 2001.
Resolving ellipsis in clarification.
In *Proceedings of the 39th meeting of the Assocation for Computational Linguistics, Toulouse*.
236–243.

Cooper, Robin and Larsson, Staffan 2009.
Compositional and ontological semantics in learning from corrective feedback and explicit definition.
In Edlund, Jens; Gustafson, Joakim; Hjalmarsson, Anna; and Skantze, Gabriel, editors 2009, *Proceedings of DiaHolmia, 2009 Workshop on the Semantics and Pragmatics of Dialogue*.

Cooper, Robin; Dobnik, Simon; Lappin, Shalom; and Larsson, Staffan 2014.
A probabilistic rich type theory for semantic interpretation.
In *Proceedings of the EACL Workshop on Type Theory and Natural Language Semantics (TTNLS)*.

Cooper, Robin; Dobnik, Simon; Lappin, Shalom; and Staffan, Larsson 2015a.
Probabilistic type theory and natural language semantics.
Under review.

Cooper, Robin; Dobnik, Simon; Larsson, Staffan; and Lappin, Shalom 2015b.
Probabilistic type theory and natural language semantics.
LiLT (Linguistic Issues in Language Technology) 10.

Cooper, Robin 2005.
Austinian truth, attitudes and type theory.
Research on Language and Computation 3:333–362.

Cooper, Robin 2010.
Generalized quantifiers and clarification content.
In Łupkowski, Paweł and Purver, Matthew, editors 2010, *Aspects of Semantics and Pragmatics of Dialogue. SemDial 2010, 14th Workshop*

on the Semantics and Pragmatics of Dialogue, Poznań. Polish Society for Cognitive Science.

Cooper, Robin 2012.
Type theory and semantics in flux.
In Kempson, Ruth; Asher, Nicholas; and Fernando, Tim, editors 2012, *Handbook of the Philosophy of Science*, volume 14: Philosophy of Linguistics. Elsevier BV.
General editors: Dov M. Gabbay, Paul Thagard and John Woods.

Dobnik, Simon and Cooper, Robin 2013.
Spatial descriptions in type theory with records.
In *Proceedings of IWCS 2013 Workshop on Computational Models of Spatial Language Interpretation and Generation (CoSLI-3)*, Potsdam, Germany. Association for Computational Linguistics.
1–6.

Dobnik, Simon; Larsson, Staffan; and Cooper, Robin 2011.
Toward perceptually grounded formal semantics.

In *Proceedings of the Workshop on Integrating Language and Vision at NIPS 2011*, Sierra Nevada, Spain. Neural Information Processing Systems Foundation (NIPS).

Feldman, Jerome 2010.
Embodied language, best-fit analysis, and formal compositionality.
Physics of Life Reviews 7(4):385 – 410.

Garrod, Simon C. and Anderson, Anthony 1987.
Saying what you mean in dialogue: a study in conceptual and semantic co-ordination.
Cognition 27:181–218.

Harnad, Stevan 1990.
The symbol grounding problem.
Physica D: Nonlinear Phenomena 42(1990):335–346.

Healey, P.G.T. 1997.
Expertise or expertese?: The emergence of task-oriented sub-languages.

In Shafto, M.G. and Langley, P., editors 1997, *Proceedings of the 19th Annual Conference of the Cognitive Science Society*.
301–306.

Kelleher, J; Costello, F; and Vangenabith, J 2005.
Dynamically structuring, updating and interrelating representations of visual and linguistic discourse context.
Artificial Intelligence 167(1-2):62–102.

Kennington, Casey and Schlangen, David 2015.
Simple learning and compositional application of perceptually grounded word meanings for incremental reference resolution.
In *Proceedings of the Conference for the Association for Computational Linguistics (ACL)*.
292–301.

Larsson, Staffan and Cooper, Robin 2009.
Towards a formal view of corrective feedback.

In Alishahi, A; Poibeau, T; and Villavicencio, A, editors 2009, *Proceedings of the Workshop on Cognitive Aspects of Computational Language Acquisition, EACL*.
1–9.

Larsson, Staffan 2007.
Coordinating on ad-hoc semantic systems in dialogue.
In *Proceedings of the 10th workshop on the semantics and pragmatics of dialogue*.

Larsson, Staffan 2011.
The ttr perceptron: Dynamic perceptual meanings and semantic coordination.
In *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA).

Larsson, Staffan 2015.
Formal semantics for perceptual classification.
Journal of Logic and Computation 25(2):335–369.
Published online 2013-12-18.

Ludlow, Peter 2014.
Living Words: Meaning Underdetermination and the Dynamic Lexicon.

Oxford University Press.

Mills, Gregory J. and Healey, Patrick G. T. 2008.
Semantic negotiation in dialogue: The mechanisms of alignment.
In *Proceedings of the 9th SIGdial Workshop on Discourse and Dialogue*, SIGdial '08, Stroudsburg, PA, USA. Association for Computational Linguistics.
46–53.

Myrendal, Jenny 2013.
From meaning potential to situated meaning - word meaning negotiation in asynchronous cmc.
University of Gothenburg.

Myrendal, Jenny 2015.

Word Meaning Negotiation in Online Discussion Forum Communication.
Ph.D. Dissertation, University of Gothenburg.

Pickering, Martin J. and Garrod, Simon 2004.
Toward a mechanistic psychology of dialogue.
Behavioral and Brain Sciences 27(02):169–226.

Rosenblatt, F 1958.
The perceptron: a probabilistic model for information storage and organization in the brain.
Psychological review 65(6):386–408.

Roy, Deb 2005.
Grounding words in perception and action: computational insights.
Trends in Cognitive Sciences 9(8):389–396.

Skočaj, Danijel; Janiček, Miroslav; Kristan, Matej; Kruijff, Geert-Jan M.; Leonardis, Aleš; Lison, Pierre; Vrečko, Alen; and Zillich, Michael 2010.

A basic cognitive system for interactive continuous learning of visual concepts.
In *ICRA 2010 workshop ICAIR - Interactive Communication for Autonomous Intelligent Robots*, Anchorage, AK, USA.
30–36.

Steels, Luc and Belpaeme, Tony 2005.
Coordinating perceptually grounded categories through language: A case study for colour.
Behavioral and Brain Sciences 28(4):469–89.
Target Paper, discussion 489-529.