

# Modeling Intensions as Classifiers

Staffan Larsson

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

Logic and Algorithms in Computational Linguistics  
August 19, 2017

Introduction

Semantic coordination

Symbol grounding and perceptual meaning

Learning meanings from interaction

Compositionality

Vagueness

Other approaches, and desiderata on a solution

Summary

# Outline

Introduction

Semantic coordination

Symbol grounding and perceptual meaning

Learning meanings from interaction

Compositionality

Vagueness


Other approaches, and desiderata on a solution

Summary


# 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 (2013)


# 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

Introduction

Semantic coordination

Symbol grounding and perceptual meaning

Learning meanings from interaction

Compositionality

Vagueness

Other approaches, and desiderata on a solution

Summary

# 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

# Information coordination and language coordination

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

# Dialogue strategies for 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 purpose is to aid semantic coordination
  - ▶ In online discussion forums (Myrendal, 2015)
    - ▶ Explicification: giving definitions
    - ▶ Exemplification: providing examples
    - ▶ Contrast: rejecting one description and proposing another
    - ▶ ...
  - ▶ In first language acquisition



# Semantic coordination in first language acquisition

- ▶ “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)

# Semantic coordination in first language acquisition, 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

Introduction

Semantic coordination

Symbol grounding and perceptual meaning

Learning meanings from interaction

Compositionality

Vagueness

Other approaches, and desiderata on a solution

Summary

# 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 want to use a framework which also encompasses accounts of many problems traditionally studied in formal semantics<sup>1</sup>
- ▶ 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*

---

<sup>1</sup>Semantic phenomena which have been described using TTR include modelling of intensionality and mental attitudes (Cooper, 2005), dynamic generalised quantifiers (Cooper, 2004), co-predication and dot types in lexical innovation, frame semantics for temporal reasoning, reasoning in hypothetical contexts (Cooper, 2011), enthymematic reasoning (Breitholtz and Cooper, 2011), clarification requests (Cooper, 2010), negation (Cooper and Ginzburg, 2011), and information states in dialogue (Cooper, 1998; Ginzburg, 2012).

## 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.)

# Classifier example: the Perceptron

- ▶ The general account is intended to work for *any type of classifier* that takes low-level input and is trainable (using machine learning techniques)
- ▶ 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.


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

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

- ▶ 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





# The TTR perceptron cont'd

A TTR perceptron classifier can be represented as a record:

$$\kappa = \left[ \begin{array}{lcl} w & = & [0.800 \quad 0.010] \\ t & = & 0.090 \\ f & = & \lambda v : \text{RealVector} \left( \left\{ \begin{array}{ll} 1 & \text{if } v \cdot w > t \\ 0 & \text{otherwise} \end{array} \right\} \right) \end{array} \right]$$

Where  $\kappa.f$  will evaluate to

$$\lambda v : \text{RealVector} \left( \left\{ \begin{array}{ll} 1 & \text{if } v \cdot [0.800 \quad 0.010] > 0.090 \\ 0 & \text{otherwise} \end{array} \right\} \right)$$

- ▶ 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:

$$\kappa = \left[ \begin{array}{lcl} w & = & [0.800 \quad 0.010] \\ t & = & 0.090 \\ f & = & \lambda v : \text{RealVector} \left( \left\{ \begin{array}{ll} P & \text{if } v \cdot w > t \\ \neg P & \text{otherwise} \end{array} \right\} \right) \end{array} \right]$$

# The meaning of “(that is to the) right” in TTR

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


$\llbracket \text{right} \rrbracket^{Agt} =$

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

(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 ( $\text{obj}_{45}$ )



$$s = \left[ \begin{array}{ll} \text{sr}_{\text{pos}} = [0.900 & 0.100] & : \text{RealVector} \\ \text{foo} = \text{obj}_{45} & : \text{Ind} \\ \text{spkr} = A & : \text{Ind} \end{array} \right]$$

- ▶ Apply  $\llbracket \text{right} \rrbracket.f$  to  $s$ :



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

# Outline

Introduction

Semantic coordination

Symbol grounding and perceptual meaning

Learning meanings from interaction

Compositionality

Vagueness

Other approaches, and desiderata on a solution

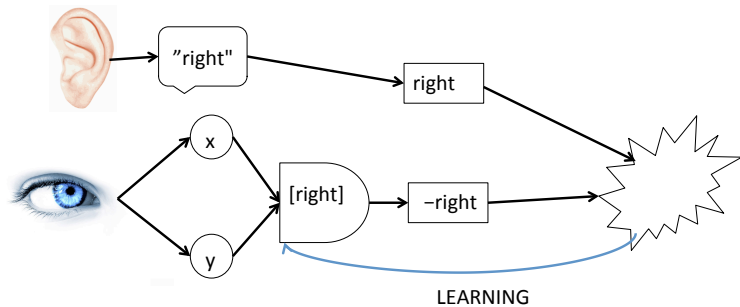
Summary

# Information coordination 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 purpose 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 A's utterance is consistent with B's take on the situation.
  5. If an inconsistency results from interpretation, B assumes A is right (B defers to A), says "aha", and learns from this exchange; otherwise, B says "okay"





# 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, 2013)
- ▶ 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":


$\llbracket \text{right} \rrbracket^B =$

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



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":

$\llbracket \text{right} \rrbracket^B =$


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

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 (semantic coordination, as above)
- ▶ 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

Introduction

Semantic coordination

Symbol grounding and perceptual meaning

Learning meanings from interaction

**Compositionality**

Vagueness

Other approaches, and desiderata on a solution

Summary

# 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”.

$\llbracket \text{upper} \rrbracket^B =$

$$\left[ \begin{array}{l} w_{\text{upper}} = \dots \\ t_{\text{upper}} = \dots \\ \text{bg} = \left[ \begin{array}{ll} \text{sr}_{\text{pos}} & : \text{RealVector} \\ \text{foo} & : \text{Ind} \\ \text{spkr} & : \text{Ind} \end{array} \right] \\ \text{f} = \lambda r : \text{bg} \left( \left[ \begin{array}{l} c_{\text{upper}}^{\text{perc}} = \left[ \begin{array}{l} \text{sr}_{\text{pos}} = r.\text{sr}_{\text{pos}} \\ \text{foo} = r.\text{foo} \end{array} \right] : \pi_{\text{upper}}(w_{\text{upper}}, t_{\text{upper}})(r) \end{array} \right] \right) \end{array} \right]$$


# Compositionality: Basic Example

Compositional meaning of “upper right” obtained by merging of meanings of “upper” and “right”

$$\llbracket \text{upper right} \rrbracket_B = \llbracket \text{upper} \rrbracket_B \wedge \llbracket \text{right} \rrbracket_B =$$

$$\left[ \begin{array}{l} w_{\text{upper}} = \dots \\ t_{\text{upper}} = \dots \\ w_{\text{right}} = \dots \\ t_{\text{right}} = \dots \\ \text{bg} = \left[ \begin{array}{ll} \text{sr}_{\text{pos}} & : \text{RealVector} \\ \text{foo} & : \text{Ind} \\ \text{spkr} & : \text{Ind} \end{array} \right] \\ \text{f} = \lambda r : \text{bg} \left( \left[ \begin{array}{l} c_{\text{upper}}^{\text{perc}} = \left[ \begin{array}{ll} \text{sr}_{\text{pos}} = r.\text{sr}_{\text{pos}} \\ \text{foo} = r.\text{foo} \end{array} \right] : \pi_{\text{upper}}(w_{\text{upper}}, t_{\text{upper}})(r) \\ c_{\text{right}}^{\text{perc}} = \left[ \begin{array}{ll} \text{sr}_{\text{pos}} = r.\text{sr}_{\text{pos}} \\ \text{foo} = r.\text{foo} \end{array} \right] : \pi_{\text{right}}(w_{\text{right}}, t_{\text{right}})(r) \end{array} \right] \right) \end{array} \right]$$

# 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)


$$\llbracket \text{far} \rrbracket = \left[ \begin{array}{l} \alpha = 1.4 \\ f = \lambda m: \left[ \begin{array}{l} t : \text{Real} \\ (m \sqcap \left[ \begin{array}{l} t = \alpha * m.t \end{array} \right]) \end{array} \right] \end{array} \right]$$

$$\llbracket \text{far right} \rrbracket = \llbracket \text{far} \rrbracket.f(\llbracket \text{right} \rrbracket) =$$


$$\left[ \begin{array}{l} t = 0.090 \\ \text{bg} = \dots \\ f = \dots \end{array} \right] \sqcap \left[ \begin{array}{l} t = 1.4 * 0.090 \end{array} \right] = \left[ \begin{array}{l} t = 0.126 \\ \text{bg} = \dots \\ f = \dots \end{array} \right]$$

# Compositionality: Degree modifiers

“right”:



“far right”:



# Outline

Introduction

Semantic coordination

Symbol grounding and perceptual meaning

Learning meanings from interaction

Compositionality

Vagueness

Other approaches, and desiderata on a solution

Summary

# 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 Fernández and Larsson (2014), we formulate 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 a 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)



# Vagueness for scalar predicates

- ▶ Case study: scalar predicates
  - ▶ e.g. '*tall*', '*long*' and '*expensive*'
  - ▶ Interpreted with respect to a scale, i.e., a dimension such as height, length, or cost along which entities for which the relevant dimension is applicable can be ordered.
  - ▶ Have a relatively simple semantics (they are often uni-dimensional) and thus constitute a perfect case-study for investigating the properties and effects of vagueness on language use.
- ▶ (However, our account should also work for n-dimensional concepts, e.g. colours, shapes)

# Modeling vagueness using a noisy threshold

- ▶ There are several ways in which one can account for vagueness
- ▶ Here, in line with Lassiter (2011), we opt for substituting the precise threshold with a noisy, probabilistic threshold.
- ▶ We consider the threshold to be a normal random variable, which can be represented by the parameters of its Gaussian distribution, the mean  $\mu$  and the standard deviation  $\sigma$  (the noise width).
  - ▶ Which noise function may be the most appropriate is an empirical question we do not tackle here.
  - ▶ Our choice of Gaussian noise follows Schmidt *et al.* (2009).

# The meaning of ‘*Tall*’

$$\mathbf{tall} = \left[ \begin{array}{l} T_{ctxt} = \left[ \begin{array}{l} c : \text{Type} \\ x : c \\ h : \mathbb{R}^+ \end{array} \right] \\ \mu = \mu_{tall} \\ \sigma = \sigma_{tall} \\ f = \lambda r : T_{ctxt}. \left[ \begin{array}{l} \text{sit} = r \\ \text{sit-type} = [c_{tall} : \text{tall}(r.x)] \\ \text{prob} = \kappa_{tall}(\sigma, \mu, r) \end{array} \right] \end{array} \right]$$

- ▶  $T_{ctxt}.c$  is the comparison class (allowing us to model context sensitivity)
- ▶  $T_{ctxt}.x$  is an individual of type  $T_{ctxt}.c$
- ▶ The output of the function  $\mathbf{tall}.f$  is now a *probabilistic* Austinian proposition (Cooper *et al.*, 2014).

# A classifier for tallness

- ▶ We define a tallness classifier  $\kappa_{tall}$  that takes as parameters  $\mu_{tall}$  and  $\sigma_{tall}$ , both of them dependent on a comparison class and hence of type  $Type \rightarrow \mathbb{R}^+$ .
  - ▶ The comparison class here specifies a type, e.g. *Human*, *Child* or *BasketballPlayer*
- ▶ The output of the classifier is a probability


$$\kappa_{tall}(\mu, \sigma, r) = \frac{1}{2} \left[ 1 + \operatorname{erf} \left( \frac{r.h - \mu(r.c)}{\sigma(r.c)\sqrt{2}} \right) \right]$$

$$\kappa_{tall} : (Type \rightarrow \mathbb{R}^+, Type \rightarrow \mathbb{R}^+, T_{ctxt}) \rightarrow [0, 1]$$

- ▶ Here erf is the error function, defined as

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_{t=0}^x e^{-t^2} dt$$

- ▶ The error function defines a sigmoid shape, in line with the upward monotonicity of 'tall'.
- ▶ The output of  $\kappa_{\text{tall}}(\mu, \sigma, r)$  corresponds to the probability that  $h$  will exceed the normal random threshold with mean  $\mu$  and deviation  $\sigma$ .



## Example

- ▶ Assume that we have  $\mu_{tall}(\text{Human})=1.87$  and  $\sigma_{tall}(\text{Human})=0.05$ .
- ▶ Let's also assume  $ctxt = \begin{bmatrix} c = \text{Human} \\ x = \text{john\_smith} \\ h = 1.88 \end{bmatrix}$
- ▶ In this case, **tall.f(ctxt)** will compute as follows:

$$\lambda r : T_{ctxt}. \begin{bmatrix} \text{sit} = r \\ \text{sit-type} = [c_{tall} : \text{tall}(r.x)] \\ \text{prob} = \kappa_{tall}(\mu_{tall}, \sigma_{tall}, r) \end{bmatrix} \left( \begin{bmatrix} c = \text{Human} \\ x = \text{john\_smith} \\ h = 1.88 \end{bmatrix} \right) =$$

$$\begin{bmatrix} \text{sit} = \begin{bmatrix} c = \text{Human} \\ x = \text{john\_smith} \\ h = 1.88 \end{bmatrix} \\ \text{sit-type} = [c_{tall} : \text{tall}(\text{john\_smith})] \\ \text{prob} = 0.579 \end{bmatrix}$$

$$\text{since } \kappa_{tall}(\mu_{tall}, \sigma_{tall}, \begin{bmatrix} c=\text{Human} \\ x=\text{john\_smith} \\ h=1.88 \end{bmatrix}) = \frac{1}{2} \left[ 1 + \text{erf} \left( \frac{1.88-1.87}{0.05\sqrt{2}} \right) \right] = 0.579$$

- ▶ This probability can now be used in further probabilistic reasoning, to decide whether to refer to an individual  $x$  as tall, or to evaluate someone else's utterance describing  $x$  is tall.
- ▶ For example, an agent may map different probabilities to different adjective qualifiers of tallness to yield compositional phrases such as '*sort of tall*', '*quite tall*', '*very tall*', '*extremely tall*', etc.
- ▶ The meanings of these composed adjectival phrases could specify probability ranges trained independently.
- ▶ Compositionality for vague perceptual meanings, and the interaction between compositionality and learning, is an important area for future research.

# Computing the Noisy Threshold

- ▶ for a vague scalar predicate like '*tall*', we assume that an agent will have at its disposal a set of *observations*  $\Omega_{tall}^T$  consisting of entities of a particular type  $T$  (a comparison class such as *Human*) that have been judged to be tall, together with their observed heights.
- ▶ Different functions can be used to compute  $\mu_{tall}$  and  $\sigma_{tall}$  from  $\Omega_{tall}^T$ .
- ▶ What constitutes an appropriate function for a certain predicate is an empirical matter; Schmidt *et al.* (2009) collect judgements of people asked to indicate which items are tall given distributions of items of different heights.



## Computing the Noisy Threshold, cont'd

- ▶ The best performing threshold model in their study is the *relative height by range* model, where (in our notation):

$$\mu_{tall}(T) = \max(\Omega_{tall}^T) - k \cdot (\max(\Omega_{tall}^T) - \min(\Omega_{tall}^T))$$

- ▶  $\max(\Omega_{tall}^T)$  and  $\min(\Omega_{tall}^T)$  stand for the maximum and the minimum height, respectively
- ▶ The model includes two parameters,  $k$  and a noise-width parameter that in our approach corresponds to  $\sigma_{tall}$ .
  - ▶ Any item within the top  $k\%$  of the range of heights that have been judged to be tall counts as tall.
  - ▶ Schmidt *et. al.* report that the best fit of their data was obtained with  $k = 29\%$  and  $\sigma_{tall} = 0.05$ .

# Updating Vague Meanings

- ▶ How is the vague meaning of 'tall' updated as an agent is exposed to new judgements via language use?
- ▶ If a new entity  $x : T$  with height  $h$  is referred to as tall, the agent adds  $h$  to its set of observations  $\Omega_{tall}^T$  and recomputes  $\mu_{tall}(\text{Human})$ , for instance using RH-R
- ▶ This in turn will trigger an update to the probability outputted by  $\kappa_{tall}$ .

## Connection to probabilistic TTR

- ▶ Generally, we want classifiers for vague perceptual terms which take real-valued input (derived from sensor input) and give probabilistic judgements as output
- ▶ These judgements can be used as input to probabilistic reasoning
- ▶ For example, we can imagine an agent having vague and context-sensitive classifiers for shape and colour, taking real-valued vector input derived from digitized pictures
- ▶ The output of these classifiers can be used as input to a classifier of objects, e.g. fruits, in a Bayes net
- ▶ The fruit classifier would be used to specify the perceptual meanings of words denoting fruits ('*apple*', '*pear*', '*orange*' etc.)
- ▶ All classifiers are continually updated as interaction proceeds (semantic coordination again)

## Learning in probabilistic TTR

The fruit classifier would be trained from interaction using the learning theory of probabilistic TTR (Cooper *et al.*, 2014)

$\kappa: \text{Sit} \rightarrow \text{Set}\left(\begin{array}{ccc} \text{sit} & : & \text{Sit} \\ \text{sit-type} & : & \text{Type} \\ \text{prob} & : & [0,1] \end{array}\right)$  such that if  $s:\text{Sit}$  then

$$\kappa(s) = \left\{ \begin{array}{lcl} \text{sit} & = & s \\ \text{sit-type} & = & T \\ \text{prob} & = & \frac{p_{A,\mathfrak{J}}(s:T \mid s:T_{e_1}, \dots, s:T_{e_n})}{p_{A,\mathfrak{J}}(s:T_{e_1}) \dots p_{A,\mathfrak{J}}(s:T_{e_n})} \end{array} \mid T \in \langle T_{c_1}, \dots, T_{c_m} \rangle \right\}$$

where

- ▶ An agent,  $A$ , makes judgements based on a finite string of probabilistic Austinian propositions,  $\mathfrak{J}$
- ▶ For a type,  $T$ ,  $\mathfrak{J}_T = \{j \mid j \in \mathfrak{J} \text{ and } j.\text{sit-type} = T\}$
- ▶  $p_{A,\mathfrak{J}}(r:T_c \mid r:T_{e_1}, \dots, r:T_{e_n}) = \text{prior}_{\mathfrak{J}}(T_c) \frac{p_{A,\mathfrak{J}}(s:T_{e_1} \mid s:T_c) \dots p_{A,\mathfrak{J}}(s:T_{e_n} \mid s:T_c)}{\text{prior}_{\mathfrak{J}}(T_{e_1}) + \dots + \text{prior}_{\mathfrak{J}}(T_{e_n})}$
- ▶  $p_{A,\mathfrak{J}}(s:T_1 \mid s:T_2) = \frac{\|T_1 \wedge T_2\|_{\mathfrak{J}}}{\|T_2\|_{\mathfrak{J}}}$ , if  $\|T_2\|_{\mathfrak{J}} \neq 0$ , and 0 otherwise.
- ▶  $\text{prior}_{\mathfrak{J}}(T) = \frac{\|T\|_{\mathfrak{J}}}{\mathcal{P}(\mathfrak{J})} = \frac{\sum_{j \in \mathfrak{J}_T} j.\text{prob}}{\sum_{j \in \mathfrak{J}} j.\text{prob}}$  if  $\mathcal{P}(\mathfrak{J}) > 0$ , and 0 otherwise.

# Outline

Introduction

Semantic coordination

Symbol grounding and perceptual meaning

Learning meanings from interaction

Compositionality

Vagueness

Other approaches, and desiderata on a solution

Summary

# Compositionality for perceptual meaning

- ▶ Problem of compositionality for perceptual classification:  
*Given an NL expression (phrase or sentence)  $e$  that includes  $n \geq 2$  perceptual words (or subexpressions)  $e_1, \dots, e_n$ , how can an agent decide whether  $e$  correctly describes a visual scene (perceptually given situation)  $s$ ?*
- ▶ This minimally requires
  - ▶ computing the meaning  $c$  of  $e$
  - ▶ using  $c$ , classify a situation  $s$  as being described (or not) by  $e$
- ▶ Different approaches to compositionality for classifiers give different solutions to how this is to be done.

# Desiderata on solution

- ▶ Possible desiderata on solutions to the problem of compositionality for perceptual meanings.
- ▶ For each desideratum/feature, each approach will be marked with "+", "-" or "?".
  - ▶ "+": a proof of concept solution addressing the desideratum in question exists
  - ▶ "-": no proof of concept exists
  - ▶ "?": unclear (to me) if a proof of concept exists

A solution to the problem of compositionality for perceptual classification should...

## ...handle intersective compositionality [COM].

- ▶ "light green" interpreted as "light and green"
- ▶ "upper right" interpreted as "upper and right"
- ▶ Definition, discrete case:  $c$  is obtained by intersective composition from  $c_1$  and  $c_2$  provided that  $s : c$  iff  $s : c_1$  and  $s : c_2$ 
  - ▶  $s : \llbracket \text{light green} \rrbracket$  iff  $s : \llbracket \text{light} \rrbracket$  and  $s : \llbracket \text{green} \rrbracket$
- ▶ Definition, generalisation:  $c$  is obtained by intersective composition from  $c_1$  and  $c_2$  provided that

$$\delta(s : c) = f(\delta(s : c_1), \delta(s : c_2))$$

for some  $f$ , where  $\delta(x : t)$  is a measure of the degree to which  $x$  is judged to be  $t$  (e.g. the probability that  $x$  is of type  $t$ )

- ▶  $\delta(s : \llbracket \text{light green} \rrbracket) = f(\delta(s : \llbracket \text{light} \rrbracket), \delta(s : \llbracket \text{green} \rrbracket))$
- ▶ Most work on compositionality for classifiers focuses on intersective compositionality



## ...handle non-intersective compositionality [NON].

- ▶ There seem to be cases where intersective compositionality does not work.
- ▶ For example, “sort of green” probably does not mean “sort of and green”
- ▶ Definition:  $c$  is obtained by non-intersective composition from  $c_1$  and  $c_2$  provided that
  - ▶ there is no  $f$  such that  $\delta(s : c) = f(\delta(s : c_1), \delta(s : c_2))$
  - ▶ but there is a function  $f$  such that  $\delta(s : c) = f(s, c_1, c_2)$
  - ▶ For example, e.g.  $\delta(s : c) = \delta(s : (c_1(c_2)))$
  - ▶ For “sort of green”,  $\delta(s : \llbracket \text{sort of green} \rrbracket) = \delta(s : \llbracket \text{sort of} \rrbracket(\llbracket \text{green} \rrbracket))$

...account for learning of perceptual meanings [LEA].

- ▶ This amounts to accounting for how classifiers are updated based on sensory observations of visual scenes and associated linguistic descriptions.
- ▶ This can be done in different ways, e.g. from corpora or from interaction with humans.

...account for vagueness [VAG].

- ▶ This is especially relevant for vague judgements (which may possibly include all perceptual judgements).

...work with state of the art classifiers [SOA].

- ▶ There are many approaches to visual classification, and recently deep learning approaches have made great strides.
- ▶ It is of course an advantage if an account of compositionality for visual classifiers can benefit from these advances
- ▶ Hence, it is desirable that the account is neutral to the type of classifier, as far as possible (or, if a particular type of classifier is deemed to be the best, that it is compatible with that type).

...connect perceptual meanings to other semantic phenomena [SEM].


- ▶ Of course, perceptual meaning is but one of a multitude of semantic phenomena
- ▶ Inference, quantification, modality, intensionality, etc.
- ▶ An account of perceptual semantics is more useful if it is formulated in a framework where many other semantic phenomena are also accounted for

# Approaches to compositionality for visual classifiers


- ▶ Meanings as sets
- ▶ Meanings as transparent functions
- ▶ Meanings as opaque functions

# Meanings as sets

- ▶ Formal semantics in the context of perception in robots
- ▶ E.g. Matuszek *et al.* (2012), Krishnamurthy and Kollar (2013)
- ▶ Using classifiers in conjunction with NL semantics based on first order logic (FOL) in the Possible Worlds Semantics (PWS) tradition (Montague, 1974).
- ▶ Basic method: Apply all classifiers to all objects in the scene, producing a first order model where meanings of predicates are sets of referents (or  $n$ -tuples of referents in the case of  $n$ -place relations).




**Environment:**  
(image on left)



**Grounding:**  $\{(2, 1), (3, 1)\}$   
**Denotation:**  $\{2, 3\}$

**Query:**  
“things to the right  
of the blue mug”



Language	Denotation
The mugs	$\{1, 3\}$
The objects on the table	$\{1, 2, 3\}$
There is an LCD monitor	$\{2\}$
Is the blue mug right of the monitor?	$\{\}$
The monitor is behind the blue cup.	$\{2\}$

# Transparent vs. opaque functions

- ▶ In general, a parameterised function takes an input (domain) and a set of parameters, and yields an output (range).
- ▶ We distinguish *transparent* and *opaque* functions:
- ▶ A *transparent* function is a parameterised function where
  - ▶ the inputs and parameters have clear interpretations understandable to humans
  - ▶ the effects on the output of manipulating the parameters are predictable.



# Meanings as transparent functions

- ▶ simple threshold classifier (parameter: threshold)
- ▶ noisy threshold classifier (parameters: threshold, standard deviation)
- ▶ simple perceptron (parameters: weights and threshold)
- ▶ cubic spline function (multiple parameters) (Gapp, 1994)



We count the above TTR account as transparent functions.

# Opaque functions

- ▶ An *opaque* function is a parameterised function where
  - ▶ the inputs and parameters do *not* have clear interpretations understandable to humans.
  - ▶ the effects on the output of manipulating the parameters are *not* predictable.
  - ▶ Examples include most neural networks, including deep neural nets, and probability distributions collected from observations
  - ▶ (Note that if a probability distribution derived from observations is analysed using some kind of curve-fitting, it may become transparent)
  - ▶ (Also, in some cases it may be possible to retroactively analyse and understand the roles of the parameters (weights) of a neural network.)
- ▶ Note that the distinction is slightly vague and that there are borderline cases.
- ▶ For example, a  $n$ -input neuron with a threshold can implement a transparent linear classifier function in  $n$ -dimensional space.

## Meanings as opaque functions

- ▶ Opaque functions are functions whose parameters (if any) are *not* understandable to human interpreters
- ▶ A couple of related approaches to meanings as classifiers can be seen as examples of this overall approach.
  - ▶ functions which are defined extensionally (e.g. as a table)
  - ▶ (most) neural networks with hidden layers
- ▶ Early example (from psychology): Logan and Sadler (1996), modeling "degrees of goodness" of spatial descriptions based on informants' judgements



# Opaque Bayesian and RNN models of colour word meanings

- ▶ McMahan et. al. (2015) model meanings of colour words as probability distributions over colour spaces, derived from a corpus of colour descriptions
- ▶ Monroe (2016) build on this but instead encode such distributions implicitly in a recurrent neural network (RNN) sequence decoder.
  - ▶ Input: word sequence and a colour sample
  - ▶ Output: probability of the sequence as describing the colour sample, equal to the product of probabilities of each successive word in the sentence conditioned on the colour sample input and the preceding words.

# How the different approaches satisfy the desiderata

Meanings as...	INT	NON	LEA	VAG	SOA	SEM
... sets	+	-	?	?	+	+
... transparent functions	+	+	+	+	-	+
... opaque functions	+	-	+	+	+	?

- ▶ INT=intersective compositionality
- ▶ NON=non-intersective compositionality
- ▶ LEA=learning perceptual meanings
- ▶ VAG=accounting for vagueness
- ▶ SOA=work with state of the art classifiers
- ▶ SEM=connect to other semantic phenomena

For details, see 2017 IWCS paper (Montpellier).

# Current work on non-intersective compositionality for classifiers

- ▶ The lack of an account of non-intersective compositionality is, on our view, a serious shortcoming of both the “meanings as sets” and the “meanings are opaque functions” approaches
- ▶ The main drawback of the “transparent functions” approach is that it seems to exclude state of the art classifiers, such as deep neural nets.
- ▶ This means that none of the current approaches fulfill all our desiderata.
- ▶ The question is then if any of the approaches can be improved to satisfy all the desiderata, or if some kind of hybrid approach is needed.
- ▶ We are working on a hybrid approach

# Perceptrons and other neural classifiers

- ▶ A possible additional desideratum is *biological plausibility*
- ▶ The perceptron is a very simple classifier, yet biologically plausible
- ▶ Over the last 5 years or so, (much) more complex neural classifiers have been used with great success for image classification, captioning, and visual question answering
- ▶ Classifiers based on neural networks also have the benefit of being (more or less) biologically plausible
- ▶ Neural TTR (cf. Robin's talk) provides a mapping from regular TTR types to biologically plausible representations
- ▶ Neural classifiers for perceptual input can straightforwardly be connected to neural TTR
- ▶ This combination offers the possibility of a biologically plausible model of how human language is grounded in perception and interaction

# Outline

Introduction

Semantic coordination

Symbol grounding and perceptual meaning

Learning meanings from interaction

Compositionality

Vagueness

Other approaches, and desiderata on a solution

Summary



# Summary

- ▶ Central tasks of semantic theory:
  - ▶ model semantic plasticity and semantic coordination
  - ▶ connect language and the world
- ▶ We model how individuals
  - ▶ represent meanings
  - ▶ use meanings to form judgements
  - ▶ coordinate on meanings and judgements
- ▶ By incorporating classifiers into formal semantics as a way of representing perceptual (intensional) meanings, and by training these classifiers in interaction, we show how these meanings are related to (perception of) the world and to interaction
- ▶ This model is useful for understanding 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

## Future work areas

- ▶ Classifiers in probabilistic TTR: Bayesian vs. neural
- ▶ Compositionality (intersective and non-intersective) for perceptual meanings
- ▶ Dialogue strategies for semantic coordination, and how they update (takes on) meanings
- ▶ Exploit potential of TTR for combining perception and inference, including hybrid inference rules
- ▶ Implement and connect to dialogue system
- ▶ Apply to Visual Question Answering task
- ▶ Connect to neural TTR
- ▶ Philosophical consequences of intensions as classifiers (what happens to extensions?)



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.



Breitholtz, Ellen and Cooper, Robin 2011.

Enthymemes as rhetorical resources.

In *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA).  
149–157.



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.

Referring 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 2011.  
Negation in dialogue.  
*In Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA).



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 1998.

Information states, attitudes and dependent record types.

In *ITALLC98*.

85–106.



Cooper, Robin 2004.

Dynamic generalised quantifiers and hypothetical contexts.

In *Ursus Philosophicus, a festschrift for Björn Haglund*. Department of Philosophy, University of Gothenburg.



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 2011.

Copredication, quantification and frames.

In Pogodalla, Sylvain and Prost, Jean-Philippe, editors 2011, *LACL*, volume 6736 of *Lecture Notes in Computer Science*. Springer. 64–79.



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.



Fernández, Raquel and Larsson, Staffan 2014.

Vagueness and learning: A type-theoretic approach.  
In *Proceedings of the 3rd Joint Conference on Lexical and Computational Semantics (\*SEM 2014)*.



Gapp, Klaus-Peter 1994.

Basic meanings of spatial relations: computation and evaluation in 3d space.

In *Proceedings of the Twelfth AAAI National Conference on Artificial Intelligence*. AAAI Press.  
1393–1398.



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.



Ginzburg, Jonathan 2012.

*The Interactive Stance.*

Oxford University Press, New York.



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.



Krishnamurthy, Jayant and Kollar, Thomas 2013.

Jointly Learning to Parse and Perceive : Connecting Natural Language to the Physical World.

*Transactions of the Association for Computational Linguistics* 1:193–206.



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 2013.

Formal semantics for perceptual classification.

*Journal of Logic and Computation* 25(2):335–369.

Published online 2013-12-18.



Lassiter, Dan 2011.

Vagueness as probabilistic linguistic knowledge.

In Nowen, R.; Rooij, R.van; Sauerland, U.; and Schmitz, H. C., editors 2011, *Vagueness in Communication*. Springer.



Logan, Gordon D and Sadler, Daniel D 1996.

A computational analysis of the apprehension of spatial relations.



Ludlow, Peter 2014.

*Living Words: Meaning Underdetermination and the Dynamic Lexicon.*

Oxford University Press.



Matuszek, Cynthia; FitzGerald, N; and Zettlemoyer, Luke 2012.

A joint model of language and perception for grounded attribute learning.

*In Proceedings of the 29th International Conference on Machine Learning.*



McMahan, Brian and Stone, Matthew 2015.

A bayesian model of grounded color semantics.

*Transactions of the Association for Computational Linguistics*  
3:103–115.



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.



Monroe, Will; Goodman, Noah D; and Potts, Christopher 2016.  
Learning to generate compositional color descriptions.  
*arXiv preprint arXiv:1606.03821*.



Montague, Richard 1974.  
*Formal Philosophy: Selected Papers of Richard Montague*.  
Yale University Press, New Haven.  
ed. and with an introduction by Richmond H. Thomason.



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.



Schmidt, L.A.; Goodman, N.D.; Barner, D.; and Tenenbaum, J.B.  
2009.

How tall is tall? compositionality, statistics, and gradable adjectives.  
In *Proceedings of the 31st annual conference of the cognitive science society*.



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.