# Do Dialogues Have Content Yet?

Staffan Larsson
Centre for Language Technology
Dept of Philosophy, Linguistics and Theory of Science
University of Gothenburg
Sweden

October 10, 2019
DRiPS Workshop

# Do Dialogues Have Content?

Staffan Larsson
Centre for Language Technology
Dept of Philosophy, Linguistics and Theory of Science
University of Gothenburg
Sweden

June 30, 2011
LACL, Montpellier

# Outline

# Outline

# SDRT I

- Segmented Discourse Representation Theory (Asher and Lascarides, 2003)
- An approach to discourse interpretation
- An extension of Discourse Representation Theory (DRT)
- Combines the insights of dynamic semantics on anaphora with a richer theory of discourse structure.

# SDRT II

- SDRT offers a candidate account of what could be regarded as "the content of a dialogue"
  - Supplies a language for representing the logical form of discourse and of dialogue
  - Assigns to this language a dynamic semantic interpretation - the SDRS (Segmented Discourse Representation Structure)
- An SDRS...
  - ... results from parsing and integrating all the utterances in a dialogue
  - ... can be used for drawing inferences from dialogue, and for evaluating the truth of propositions
- A central aspect of SDRT is that it is based on a modal and dynamic variant of classical model-theoretic semantics
- The idea of using model-theoretic semantics for dialogue is not far-fetched, but something which is more or less standard practice.

# Classical model-theoretic semantics

A model for first-order logic: $\langle A, F \rangle$ consists of a domain $A$ consisting of a set of individuals, and a function $F$ that assigns values based on $A$ to constants and predicates:

- For a constant $c$, $F(c) \in A$; e.g. $F(\text{dog}) =$ the set of all dogs
- For a one-place predicate $P$, $F(P) \subseteq A$
- For a two-place predicate $R$, $F(R) \subseteq A \times A$
- ...

Truth is defined thus:

- $P(c)$ is true iff $F(c) \in F(P)$; e.g. dog(fido) is true iff $F(\text{fido}) \in F(\text{dog})$
- $R(a, b)$ is true iff $\langle F(a), F(b) \rangle \in F(R)$
- ...

# Variants of classical model-theoretic semantics I

- Modal logic: the truth of a proposition is defined relative to a possible world.
- Montagovian intensional model theory: adds a so-called intensional operator
- Property theory (Chierchia and Turner, 1988) allows different properties with identical extensions.

# Variants of classical model-theoretic semantics II

- Despite the refinements in dealing with intensional vocabulary in these extensions, they all share a common property with classical model-theoretic semantics, namely a fundamental *extensionality*.
- Extensions of predicates are simply postulated in terms of the interpretation function $F$ (in a more or less complicated way).
- Predicates are atoms, hence unanalysed and without structure.

# Outline

# Outline

# Classification

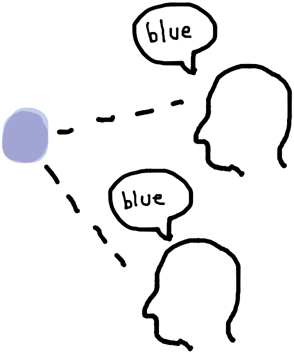

# Classification is subjective?

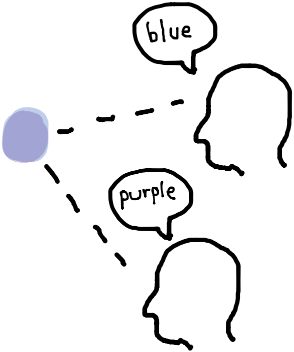

# Coordination process

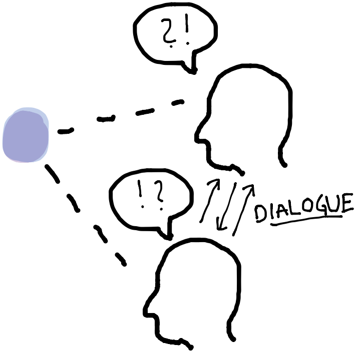

# Classification is coordinated

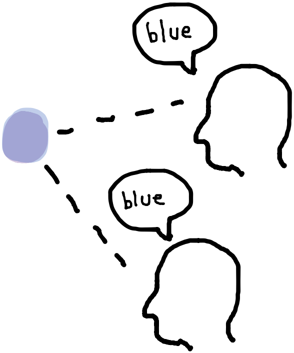

# Classification is coordinated

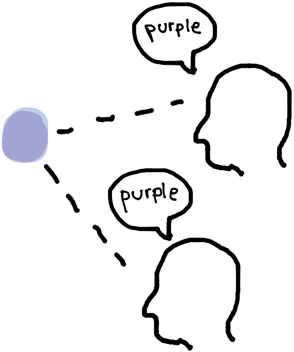

# Coordination can be creative

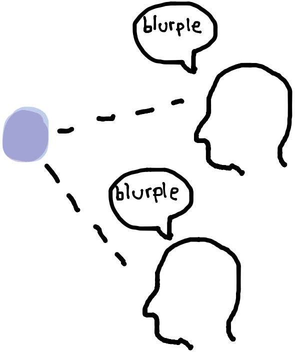

# Semantic coordination

- 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)
- *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/correction
  - Accommodation/deference: "silent" coordination where a DP observes the language use of another and adapts to it (Larsson, 2010)
- There are also dialogue strategies whose primary purpouse is to aid semantic coordination
  - In first language acquisition (Larsson and Cooper, 2009)
  - In online discussion forums (Myrendal, 2015; Noble *et al.*, 2019)

# Examples of *corrective feedback*

- C: That's a nice bear.
- D: Yes, it's a nice **panda**.
- C: Panda.

- Naomi: mittens
- Father: **gloves**.
- Naomi: gloves.
- Father: when they have fingers in them they are called gloves and when the fingers are all put together they are called mittens.

- 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.

- Adam: Mommy, where my plate?
- Mother: You mean your saucer?

- The first one is made up, the others are quoted from various sources in Clark and Wong (2002) and Clark (2007).
- In general, corrective feedback can be regarded as offering an alternative form to the one that the speaker used.

# Word Meaning Negotiation (WMN)

- See Myrendal (2015).
- On occasion, DPs engage in word meaning negotiation (WMN):
  - they explicitly coordinate their respective takes on the meaning of a particular word (the *trigger word*)
- Occurs when a DP remarks on a particular word choice of another participant, thereby initiating a sequence focusing on the meaning of that word
- WMN can concern both
  - inferential meaning (e.g. e.g. by providing explicit definitions) and
  - referential meaning (e.g. by providing percetually salient examples)
- We want an account of WMN which covers negotiation of both types of meaning

# A preliminary taxonomy of WMN dialogue acts

From Myrendal (2015):

- Explicification: providing definitions
  - Generic vs. Specific
  - Positive vs. Negative
- Exemplification: providing examples
  - Positive vs. Negative
- Contrasting: providing alternative word
  - Related vs. unrelated word
- Opposing
- Requesting clarification
- Endorsing
  - Passive vs. active

For a more recent take on this, see Noble *et al.* (2019).

# Explicification example: “(child) abuse”

- This example is taken from a discussion about whether or not piercing the ears of young children is morally acceptable, or if it constitutes child abuse.
- First post in thread describes situation where an a friend of the thread starter has had their infant daughter’s ears pierced
  - *Ett klart ÖVERGREPP att ta hål i öronen på små barn! [...] - man förorsakar barnet smärta och en fysisk förändring som barnet självt inte har valt och som inte går att återst/"alla."*
  - *Clearly ABUSE to pierce the ears of young children! [...] - you inflict pain upon the child and a physical change which the child herself has not chosen and which cannot be made undone.*
- Type: explicification, specific

# Outline

# Semantic coordination and inference I

- How does semantic coordination affect inference and the notion of content of a dialogue?
- Example:
  - C: “I like all animals on TV”
  - C: “I saw a bear on TV yesterday”
  - ...
  - C: “Bear”
  - D: “Yes it’s a nice panda”
  - C: “Panda”
    - C updates meaning for “panda” and “bear”

# Semantic coordination and inference II

- What content can be inferred from this dialogue?
  - That C likes bears?
  - That C likes pandas?
- The problem is that since C's take on the meaning of "bear" and "panda" has changed, we don't really know whether C saw a bear or a panda on TV.

# Semantic coordination and inference III

- Another example:
  - C: "I have a pair of blue gloves"
  - ...
  - C: "Blue ball"
  - D: "Yes it's a nice green ball"
  - C: "Green ball"
- C updates perceptual type (classifiers) for "blue" and "green"
  - The particular ball in question is a positive training example for "green" and a negative training example for "blue"
- This can be seen as a case of C and D coordinating their takes on the meaning of "green" (and "blue").

# Semantic coordination and inference IV

- What content can be inferred from this dialogue?
  - That C has a pair of blue gloves?
  - That C has a pair of green gloves?
- Depends on how C classifies the colour of the gloves in question
  - Does C now classify as green everything (s)he previously classified as green?
  - Or has there been a more subtle change?
- There may be no clear-cut fact as to the correct name of a given colour.
  - Different speakers may have different takes, and these may change during interactions and over time.

# Ludlow on inference

Quoting Ludlow (2014), who notes the problem for inference caused by semantic coordination ("the dynamic lexicon"):

- To see the problem consider the most trivial possible logical argument:
  - $F(a)$, therefore $F(a)$
- If the meaning of $F$ shifts, the argument may not be sound even if $F$ is true — a kind of equivocation might have taken place.

# Outline

# Semantic coordination and classical model-theoretic semantics I

- Can we incorporate semantic coordination in our model-theoretical semantics for dialogue?
- Gloves example again
- We assume that the interpretation of the phrase "I have blue gloves" uttered by C involves using $F$ to yield the referent (some object in the model $A$, let's call it $a_{123}$).
- Roughly, we have $F(\text{C's blue gloves}) = a_{123}$.
- Now, the meaning of the predicate "blue" is given as a set of objects in $A$, including $a_{123}$: $F(\text{blue}) = \{\ldots, a_{123}, \ldots\}$.

# Semantic coordination and classical model-theoretic semantics II

- To update the meaning of "blue", what can we do?
- The interpretation of "blue" is a set, so what we can do is to add elements to, or remove elements from $F(\text{blue})$.
- As C's take on the meaning of "blue" changes as a result of D's utterance "Yes, it's a nice green ball", C should update here take on $F$ so that $F(\text{blue})$ after the update no longer contains any green objects
- $F(\text{blue})$ it should instead contain exactly all blue objects in $A$

# Semantic coordination and classical model-theoretic semantics III

- Seems to require that C is able to compute, for all objects in $A$, whether they should be included in the new meaning of "blue"
  - i.e., for all objects known to C, whether they are blue (in the new sense).
- This in turn requires some method of deciding, for each element in $A$, whether to include it in $F(blue)$ or not.
- What would this be?
  - Note that there is no "colour sample" to compare to, other than perhaps the objects in sight, e.g. the blue ball.
  - It seems we need some generalisation or description to compare to, but we only have an extension!
- Furthermore, we need to apply this method to all elements of $A$.
- Traditional model theoretical semantics seems to commit us to an unrealistic theory of conceptual learning

# Outline

# Outline

# An alternative: TTR

- What is needed to be able to account for semantic coordination and semantic coordination in a formal semantic theory?
- What seems to be missing in classical model-theoretical semantics is some *structured* representation of the *intensions* of linguistics expressions.
- In classical model theory, all we have is (minimally) structured representations of extensions, and as we have seen above this is not enough.

# Background

- 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.* (2015)
- Key idea:
  - modeling (perceptual) meanings as classifiers of real-valued (perceptual) data, and training these classifiers in interaction with the world and other agents

# Inferential and referential meaning

- Marconi (1997) distinguishes inferential and referential meaning
- *inferential* word meanings
  - enables inferences from uses of the word
  - "high level" or "symbolic"
  - can be modeled in formal semantics
- *referential* meaning
  - allows speakers to identify objects and situations referred to
  - "low-level" or "subsymbolic"
- Hypothesis: referential meaning can be modeled using classifiers outputting formal representations (Larsson, 2011; Larsson, 2013)

# The fundamental idea

- Agents associate linguistic expressions with structured information
- By interacting and updating the structured information associated with expressions, agents coordinate on meanings
  - Interaction in a shared perceptual environment essential to "symbol grounding"
- To the extent that they are sufficiently coordinated (wrt the task at hand) the expressions have meaning
- There is a multitude of procedures in natural dialogue to enable this coordination
- Meanings of complex expressions are derived from meanings of their constituents
  - including perceptual meanings modeled as classifiers

# Our task

- A formal semantics and pragmatics that accounts for all of this...
- ...while also keeping the insights gained from the previous 50 years oated f work in formal semantics...
- ...and connects to recent work on machine learning for perceptual (e.g. image) classification.

# Type Theory with Records

- We want to use a framework which also encompasses accounts of many problems traditionally studied in formal semantics[1]
- 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*
- Implemented (partially) in Python (Cooper, 2019)

---

[1]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).

# TTR: A brief introduction

- We will be formulating our account in a Type Theory with Records (TTR).
  - We can here only give a brief and partial introduction to TTR; see also Cooper (2005), Cooper (2012) and Cooper (in progress).
- $a : T$ is a judgment that $a$ is of type $T$.
- One *basic type* in TTR is *Ind*, the type of an individual
  - Another basic type is *Real*, the type of real numbers.
- Given that $T_1$ and $T_2$ are types, $T_1 \rightarrow T_2$ is a *functional type* whose domain is objects of type $T_1$ and whose range is objects of type $T_2$.

# Records and record types

- If
  - $a_1 : T_1$,
  - $a_2 : T_2(a_1)$,
  - $\ldots$,
  - $a_n : T_n(a_1, a_2, \ldots, a_{n-1})$,
- ...the record to the left is of the record type to the right.

$$\left[\begin{array}{lll} \ell_1 & = & a_1 \\ \ell_2 & = & a_2 \\ \ldots & & \\ \ell_n & = & a_n \\ \ldots & & \end{array}\right] : \left[\begin{array}{lll} \ell_1 & : & T_1 \\ \ell_2 & : & T_2(l_1) \\ \ldots & & \\ \ell_n & : & T_n(\ell_1, l_2, \ldots, l_{n-1}) \end{array}\right]$$

- $\ell_1, \ldots \ell_n$ are *labels* which can be used elsewhere to refer to the values associated with them.
- $T(a_1, \ldots, a_n)$ represents a type $T$ which depends on the objects $a_1, \ldots, a_n$

# Ptypes

- A sample record and record type:

$$\begin{bmatrix} \mathrm{ref} & = & \mathrm{obj}_{123} \\ \mathrm{c}_{\mathrm{man}} & = & \mathrm{s}_{456} \\ \mathrm{c}_{\mathrm{run}} & = & \mathrm{s}_{789} \end{bmatrix} : \begin{bmatrix} \mathrm{ref} & : & \mathit{Ind} \\ \mathrm{c}_{\mathrm{man}} & : & \mathrm{man}(\mathrm{ref}) \\ \mathrm{c}_{\mathrm{run}} & : & \mathrm{run}(\mathrm{ref}) \end{bmatrix}$$

- Types can be constructed from predicates, e.g., "run" or "man".
- Such types are called *ptypes* and correspond roughly to propositions in first order logic.
- A fundamental type-theoretical intuition is that something of a ptype $P(a_1, \ldots, a_n)$ is whatever it is that counts as a proof of $P(a_1, \ldots, a_n)$.
- One way of putting this is that "propositions are types of proofs"
- Above, we simply use $\mathrm{s}_{\ldots}$ as a placeholder for proofs; below, we will show how low-level perceptual input can be included in proofs.

# Meet types and merges of record types

- If $T_1$ and $T_2$ are types, $T_1 \wedge T_2$ is the meet type of $T_1$ and $T_2$
- For any types $T_1$, $T_2$ , $a : (T_1 \wedge T_2)$ iff $a : T_1$ and $a : T_2$
- If $T_1$ and $T_2$ are record types then there will always be a record type (not a meet) $T_3$ which is necessarily equivalent to $T_1 \wedge T_2$.
- $T_3$ is the *merge* of $T_1$ and $T_2$, written as $T_1 \dot{\wedge} T_2$
  - $\left[\mathrm{f} : T_1\right] \dot{\wedge} \left[\mathrm{g} : T_2\right] = \begin{bmatrix} \mathrm{f} : T_1 \\ \mathrm{g} : T_2 \end{bmatrix}$
  - $\left[\mathrm{f} : T_1\right] \dot{\wedge} \left[\mathrm{f} : T_2\right] = \left[\mathrm{f} : T_1 \wedge T_2\right]$

# Outline

# Classifiers, intensions and extensions

- Classifiers are functions whose domain is typically numerical (e.g. real-valued, integer or binary) vectors, and whose range is a set of categories (or probability distributions over categories)
- We will regard classifiers 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

# Classifying objects as being to the left or to the right

We will use this as a very basic example of perceptual classification (
Larsson, 2011;
Larsson, 2013).

- 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 a transformation 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
  -

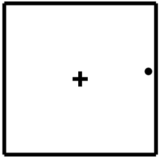

# Classifying objects as being to the left or to the right

We will use this as a very basic example of perceptual classification.

- 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 transformation 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 returns 1 or 0 depending on whether the object is to the right or not

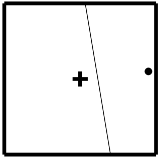

$\Rightarrow 1$

# Classifier example: the Perceptron I

- 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)

# Classifier example: the Perceptron II

- 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.

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

  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.

# The left-or-right game

This simple "game" is intended to capture some of the properties of first language acquisition.

- $A$ 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 and forms a take on the perceived situation
  3. $A$ says either "left" or "right"
  4. $B$ interprets $A$'s utterance based on $B$'s take on the situation, yielding a ptype $p$
  5. $B$ judges whether ($B$'s take on) the situation is of type $p$
  6. If not, $B$ assumes $A$ is right ($B$ defers to $A$), says "aha", and learns from this exchange; otherwise, $B$ says "okay"

For an earlier version of this account, see Larsson (2011), Larsson (2013).

# Representing (takes on) situations

- In first language acquisition, training of perceptual meanings typically takes place in situations where the referent is in the shared focus of attention and thus perceivable to the dialogue participants
- For the time being we limit our analysis to such cases.
- We assume that our DPs (dialogue participants) are able to establish a shared focus of attention
- A (simple) sensor collects some information (sensor input) from the environment and emits a real-valued vector.
- The sensor is assumed to be oriented towards the object in shared focus of attention.

# Example

- Example of B's take on a situation in the left-or-right game:

$$s_1 = \begin{bmatrix} \mathrm{sr}_{pos} & = & \begin{bmatrix} 0.900 & 0.100 \end{bmatrix} \\ \mathrm{foo} & = & \mathrm{obj}_{45} \end{bmatrix}$$

- In the left-or-right game, we will assume that B's take on the situation includes
  - a reading from a position sensor (labelled "$\mathrm{sr}_{\mathrm{pos}}$")
  - a field foo for an object in shared focus of attention.
- The position sensor returns a two-dimensional real-valued vector representing the horizontal vertical coordinates of the focused object:
  - $\begin{bmatrix} x & y \end{bmatrix}$ where $-1.0 \leq x, y \leq 1.0$ and $\begin{bmatrix} 0.0 & 0.0 \end{bmatrix}$ represents the center of the frame.
- In $\mathrm{s}_1$, $B$'s sensor is oriented towards $\mathrm{obj}_{45}$ and $\mathrm{sr}_{\mathrm{pos}}$ returns a vector corresponding to the position of $\mathrm{obj}_{45}$.

# TTR perceptron in LoR game

- We formulate a classifier as a function $\pi$ with a well defined TTR type
- Instead of a Boolean output, we want a ptype (or the negation thereof):

$$\pi_{right} : \begin{bmatrix} \text{w} & : & \mathit{RealVector} \\ \text{t} & : & \mathit{Real} \end{bmatrix} \rightarrow \begin{bmatrix} \text{foo} & : & \mathit{Ind} \\ \text{sr} & : & \mathit{RealVector} \end{bmatrix} \rightarrow \mathit{Type}$$

such that if $par : \begin{bmatrix} \text{w} & : & \mathit{RealVector} \\ \text{t} & : & \mathit{Real} \end{bmatrix}$ and $r : \begin{bmatrix} \text{foo} & : & \mathit{Ind} \\ \text{sr} & : & \mathit{RealVector} \end{bmatrix}$,

then $\pi_{right}(par,\ r) = \begin{cases} \text{right}(r.\text{foo}) & \text{if } r.sr \cdot par.\text{w} > par.\text{t} \\ \neg\ \text{right}(r.\text{foo}) & \text{otherwise} \end{cases}$

- Note that the function itself is defined outside TTR
- This allows any classifier to used with TTR, no matter how it is implemented.

# Meaning entries for predicates

- Several types of expressions in natural language (nouns, verbs, adjectives) can be modelled semantically using predicates
- We will represent the (perceptual) meaning of predicates as records containing
  - Classifier parameters (param): a (possibly empty) record
  - Background meaning (bg): a record type representing assumptions about the context of utterance (presuppositions)
  - Interpretation function (interp)
  - Classification function (clfr)

$$Mng = \begin{bmatrix} \text{params} & : & Rec \\ \text{bg} & : & RecType \\ \text{intrp} & : & \text{bg} \rightarrow Type \\ \text{clfr} & : & \text{bg} \rightarrow Type \end{bmatrix}$$

# Help function *Pred*

- Predicate meanings are defined for a predicate with a certain arity.
- Sometimes we want to know the meaning of the predicate used in a ptype.

$Pred(P(a_1, \ldots, a_n)) = P_{\langle T_1, \ldots, T_n \rangle}$ for $P(a_1, \ldots, a_n) \in$ **PType** s.t.

- $\langle T_1, \ldots, T_n \rangle \in Arity(P))$
- $a_1 : T_1, \ldots, a_n : T_n$

- $Pred(\text{right}(obj_{45})) = \text{right}_{\langle Ind \rangle}$

# Lookup function for predicate meanings

$$Mng = \left[\begin{array}{lll} \text{params} & : & \mathit{Rec} \\ \text{bg} & : & \mathit{RecType} \\ \text{intrp} & : & \text{bg} \rightarrow \mathit{Type} \\ \text{clfr} & : & \text{bg} \rightarrow \mathit{Type} \end{array}\right]$$

*PredMng* is a function

- whose domain is $\{P_A \mid P \in \textbf{Pred},\ A \in \mathit{Arity}(P)\}$ and
- whose range is in $\{r \mid r : \mathit{Mng}\}$

$$\mathit{PredMng}(\text{right}_{\langle \mathit{Ind} \rangle}) = \left[\begin{array}{lll} \text{params} & = & \left[\begin{array}{l} \text{w} = \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} \\ \text{t} = 0.090 \end{array}\right] \\ \text{bg} & = & \left[\begin{array}{lll} \text{sr}_{\text{pos}} & : & \mathit{RealVector} \\ \text{foo} & : & \mathit{Ind} \end{array}\right] \\ \text{intrp} & = & \lambda r : \text{bg} \cdot \text{right}(r.\text{foo}) \\ \text{clfr} & = & \lambda r : \text{bg} \cdot \pi_{\text{right}}(\text{params}, r) \end{array}\right]$$

We also define $[\![\ \text{right}\ ]\!] = \mathit{PredMng}(\text{right}_{\langle \mathit{Ind} \rangle}).\text{intrp}$

# Classifiers and witness conditions

- According to Cooper (in progress):
  - for $T \in \mathbf{PType}$, $s : T$ iff $s \in F(T)$
  - ... where **PType** is the set of ptypes
- We extend this definition to include witness conditions (Cooper, in progress) that call classifiers
- for $T \in \mathbf{PType}$, $s : T$ iff
  - $(\text{PredMng}(\text{Pred}(T)).\text{clfr})(s) = T$ or
  - $s \in F(T)$
- New judgements are made using the clfr function
- Previous judgements are stored in $F(T)$, the *witness cache* for $T$
- Classifiers are at the core of TTR

# Interpretation I

Assume that an agent $A$ places an object on the surface and says "That one is to the right", or just "Right".

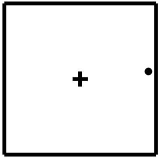

Agent $B$ watches and gets a position sensor reading $\begin{bmatrix}0.900 & 0.100\end{bmatrix}$ which is part of $B$'s take on the current situation ($s_1$):

$$s_1 = \begin{bmatrix} \mathrm{sr}_{\mathrm{pos}} & = & \begin{bmatrix}0.900 & 0.100\end{bmatrix} \\ \mathrm{foo} & = & obj_{45} \end{bmatrix}$$

$B$ now interprets $A$'s utterance in context by computing

$[\![ \text{ right } ]\!](s_1)=\text{right}(obj_{45})$

# Interpretation II

Recall

$$PredMng(\text{right}_{\langle Ind \rangle}) = \begin{bmatrix} \text{params} & = & \begin{bmatrix} \text{w} = \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} \\ \text{t} = 0.090 \end{bmatrix} \\ \text{bg} & = & \begin{bmatrix} \text{sr}_{\text{pos}} & : & RealVector \\ \text{foo} & : & Ind \end{bmatrix} \\ \text{intrp} & = & \lambda r : \text{bg} \cdot \text{right}(r.\text{foo}) \\ \text{clfr} & = & \lambda r : \text{bg} \cdot \pi_{\text{right}}(\text{params}, r) \end{bmatrix}$$

$[\![ \text{ right } ]\!] = PredMng(\text{right}_{\langle Ind \rangle}).\text{intrp}$

Derivation : $[\![ \text{ right } ]\!](s_1) = (PredMng(\text{right}_{\langle Ind \rangle}).\text{intrp})(s_1) =$

$$(\lambda r : \begin{bmatrix} \text{sr}_{\text{pos}} : RealVector \\ \text{foo} \ : Ind \end{bmatrix} \cdot \text{right}(r.\text{foo}))( \begin{bmatrix} \text{sr}_{\text{pos}} = \begin{bmatrix} 0.900 & 0.100 \end{bmatrix} \\ \text{foo} \ = obj_{45} \end{bmatrix}) =$$

$\text{right}(obj_{45})$

# Classification I

Next, $B$ decides if $A$'s utterance correctly describes (her take on) the situation, i.e. if

$s_1 : [\![ \text{ right } ]\!](s_1)$,

i.e., if $s_1 : \text{right}(obj_{45})$

Recall that

- for $T \in \textbf{PType}$, $s : T$ iff
  - $(\text{PredMng}(\text{Pred}(T)).\text{clfr})(s){=}T$ or
  - $s \in F(T)$

For $T{=}\text{right}(obj_{45})$, we get

- $s{:}\text{right}(obj_{45})$ iff
  - $(\mathit{PredMng}(\text{right}_{\langle \text{Ind} \rangle}).\text{clfr})(s){=}\ \text{right}(obj_{45})$ or
  - $s \in F(\text{right}(obj_{45}))$

So let's check this for $s_1$!

# Classification II

$$(\mathit{PredMng}(\text{right}_{\langle \text{Ind} \rangle}).\text{clfr})(s_1)$$

$$= (\lambda r : \begin{bmatrix} \text{sr}_{\text{pos}} : \mathit{RealVector} \\ \text{foo} : \mathit{Ind} \end{bmatrix} \cdot \pi_{\text{right}}(\begin{bmatrix} \text{w} = \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} \\ \text{t} = 0.090 \end{bmatrix}, r))(\begin{bmatrix} \text{sr}_{\text{pos}} = \begin{bmatrix} 0.900 & 0.100 \end{bmatrix} \\ \text{foo} = \mathit{obj}_{45} \end{bmatrix})$$

$$= \pi_{\text{right}}(\begin{bmatrix} \text{w} = \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} \\ \text{t} = 0.090 \end{bmatrix}, \begin{bmatrix} \text{sr}_{\text{pos}} = \begin{bmatrix} 0.900 & 0.100 \end{bmatrix} \\ \text{foo} = \mathit{obj}_{45} \end{bmatrix})$$

$$= \begin{cases} \text{right}(\mathit{obj}_{45}) & \text{if } \begin{bmatrix} 0.900 & 0.100 \end{bmatrix} \cdot \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} > 0.090 \\ \neg \text{ right}(\mathit{obj}_{45}) & \text{otherwise} \end{cases}$$

$$= \text{right}(\mathit{obj}_{45})$$

Hence, $s_1 : \text{right}(\mathit{obj}_{45})$

Hence, $s_1 : [\![ \text{ right } ]\!](s_1)$

# Round 1

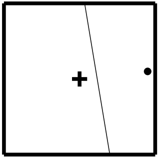

A: “right”
B: “okay”

# Outline

# TTR and Semantic Coordination

- TTR allows structured meanings which can be updated as a result of interaction in dialogue
- All represented aspects of meaning may be modified
  - "propositional" (symbolic) information may be added and subtracted
  - perceptual (subsymbolic) classifiers may be retrained, e.g. by modifying weight vectors

# Another round of the LoR game

Recall agent $B$'s initial take on the meaning of "right":

$$PredMng(\text{right}_{\langle Ind \rangle}) = \begin{bmatrix} \text{params} & = & \begin{bmatrix} \text{w} = \begin{bmatrix} 0.800 & 0.010 \end{bmatrix} \\ \text{t} = 0.090 \end{bmatrix} \\ \text{bg} & = & \begin{bmatrix} \text{sr}_{\text{pos}} & : & RealVector \\ \text{foo} & : & Ind \end{bmatrix} \\ \text{intrp} & = & \lambda r : \text{bg} \cdot \text{right}(r.\text{foo}) \\ \text{clfr} & = & \lambda r : \text{bg} \cdot \pi_{\text{right}}(\text{params}, r) \end{bmatrix}$$

# Round 2

In a second round, things get a bit more complicated...

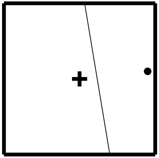

A: "right"
B: "okay"

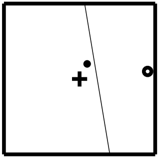

A: "right"

# Learning perceptual meaning from interaction

- In this round, *A* places another object in a different position in the frame and again says "right".
- Now, B's take on the situation is as follows:

$$s_2 = \begin{bmatrix} \mathrm{sr}_{pos} = \begin{bmatrix} 0.100 & 0.200 \end{bmatrix} \\ \mathrm{foo} = \mathrm{obj}_{46} \end{bmatrix}$$

- Note that foo has been updated and that there is a new sensor reading.

# Interpretation I

- As before, $B$ interprets $A$'s utterance by computing $[\![$ right $]\!](s_2)$=right$(obj_{45})$
- Next, $B$ decides if $A$'s utterance correctly describes (her take on) the situation, i.e. if

  $s_2 : [\![$ right $]\!](s_1)$,

  i.e., if $s_1$ : right$(obj_{45})$
- This time, however, $(PredMng(\text{right}_{\langle \text{Ind} \rangle}).\text{clfr})(s_2) = \neg\text{right}(obj_{45})$
- So $s_2$ is not of type $[\![$ right $]\!](s_2)$, and $A$'s utterance does not correctly describe $B$'s take on the situation

# Updating perceptual meaning I

- According to the rules of the game, *B* resolves this conflict by trusting A's judgement over B's own classification
- *B* learns from this exchange by updating the weights used by the classifier perceptron associated with "right"

# Updating perceptual meaning II

Perceptrons are updated using the *perceptron training rule*:

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

where

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

where in turn

- $o_t$ is the target output,
- $o$ is the actual output, and
- $x_i$ is the input

- The perceptron training rule can be formulated as a function operating on TTR predicate meanings (Larsson, 2013)
- In the LoR-game, training results in moving the line dividing the surface

Agent B's revised on the meaning of "right":

$$PredMng(\text{right}_{\langle Ind \rangle})' = \begin{bmatrix} \text{params} & = & \begin{bmatrix} \text{w} = \color{red}{\begin{bmatrix} 0.808 & 0.2002 \end{bmatrix}} \\ \text{t} = 0.090 \end{bmatrix} \\ \text{bg} & = & \begin{bmatrix} \text{sr}_{\text{pos}} & : & \mathit{RealVector} \\ \text{foo} & : & \mathit{Ind} \end{bmatrix} \\ \text{intrp} & = & \lambda r : \text{bg} \cdot \text{right}(r.\text{foo}) \\ \text{clfr} & = & \lambda r : \text{bg} \cdot \pi_{\text{right}}(\text{params}, r) \end{bmatrix}$$

A: "right"

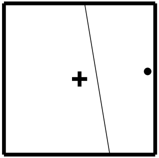

B: "okay"

A: "right"

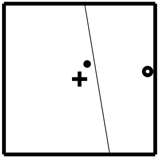

B: "aha"

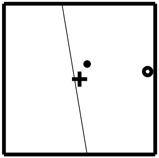

# From learning to coordination

- 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 may converge on a meaning of "right" that neither of them may subscribe to initially.

- In contrast to classical model theory, we are not updating sets (extensions), but structured records (intensions)
- In terms of learning, modifying structured intensions seems more sensible than re-classifying all known objects according to an some unknown algorithm

# Outline

# Indeterminate extensions I

- This view of intensional meaning as involving sensors interfacing with the real world has important consequences for how we think about extensions
- A re-trained sensor cannot immediately tell us which objects and situations will trigger it – the only way to find out is to apply it in various situations and see what happens
- There is no longer a definite extension of an expression apart from specific situations of language use
- Instead, in each instance of use of an expression, the situation of use is classified as falling under the intension of the expression or not

# Indeterminate extensions II

- In general, especially in humans, a classifier does not operate in a vacuum and may be sensitive to a multitude of aspects of the perceived situation, including
  - shared and individual goals
  - various social aspects of the situation
  - perceptual factors (light conditions etc)
  - priming effects
  - ...
- Classification is thus a situated, complex and stochastic process
- This is also true of classifiers involving "propositions", i.e. classifiers whose extensions are (sets of) situations
- Whether a certain classifier ("proposition") classifies a situation ("is true") or not cannot be determined in the abstract, but only by applying the classifier to a situation

# Outline

# Ludlow's (2014) attempt I

Quoting Ludlow (2014) again:

- To see the problem consider the most trivial possible logical argument:
  - $F(a)$, therefore $F(a)$
- If the meaning of $F$ shifts, the argument may not be sound even if $F$ is true — a kind of equivocation might have taken place.
- Does this mean that logic goes out the window? Not at all.

# Ludlow's (2014) attempt II

- For expository purposes lets sequentially number occurrences of terms in an argument, so that, for example, in the argument we just gave the form is the following.
  - $F1(a)$, therefore $F2(a)$
- Again, we are saying that the term is $F$, and that $F1$ and $F2$ are occurrences of the term $F$ within the argument.
- It appears that the argument above is sound if the meaning is stable between $F1$ and $F2$ but also if $F2$ is a broadening of $F1$ (a narrowing or a lateral shift in meaning will not preserve truth).

# Ludlow's (2014) attempt III

- Lets take a concrete example.
  - Jones is an athlete$_1$. Therefore Jones is an athlete$_2$.
- If a shift has taken place between premise and conclusion (between the meaning of "athlete$_1$" and "athlete$_2$") it cannot be a shift that rules out individuals that were recognized semantic values of "athlete$_1$".
- If "athlete$_1$" admits racecar drivers and "athlete$_2$" does not, then the argument is not sound.
- If the second occurrence broadens the meaning of the term "athlete", the argument is sound.

End of quote!

# The problem with Ludlow's attempt

- Ludlow defines broadening and narrowing extensionally:
  - "If $A < B$ is a general way of indicating that all $A$s are $B$s, or that every $A$ is (a) $B$, or that all $A$ is $B$ (in effect moving from a narrower to a broader range)..."
- But as we have seen, re-training a classifier does not, in general, allow us to draw any conclusions about changes in extensions
- Again, the only way to compute the extension before and after the change in meaning would be to perceive and classify each item
- This means we cannot, in general, tell whether a meaning update is a broadening, a narrowing, or a lateral shift (neither broadening or narrowing)
- Hence, we cannot in general tell whether a meaning shift preserves and inference or not.

# Outline

# Conclusion

The phenomenon of semantic coordination poses some general requirements on any formal theory purporting to account for it:

- Needs to treat intensions of linguistic expressions as first-class objects, i.e., it needs something equivalent to types
- Types need to be structured, since accounting for semantic coordination in a way that makes sense requires the possibility of modifying intensions, and only structured objects can be modified
- Needs to allow for a fundamental indeterminacy of extensions of linguistic expressions, as
    - some meanings involve classifying situations in the world based on perceptions thereof
    - classification of real-world situations is (in general) a complex stochastic process

# Structured intensions in SDRT?

- While SDRT currently relies on classical model-theoretic semantics, it also has some features which may be useful if one would extend it to account for semantic coordination
  - It insists on the utility of a level of representation between language and model, namely the language of SDRSs
  - Also, SDRSs are structured meaning representations.
- It would perhaps be possible to recast SDRT in a type-theoretic framework, thereby making it better equipped to deal with semantic coordination
- Work by Asher (Asher, 2010) on a type-theoretic account of word meaning is encouraging in this respect

# So, do dialogues have content yet?

- Classifical model-theoretic semantics (still) has great difficulty accounting for semantic coordination.
- TTR can account for semantic coordination, but
  - Indeterminacy complicates the idea that the truth value of every proposition is at every point in time determinate (true or false)
  - Inferences involving "propositions" containing predicates whose meaning has changed during the dialogue are not always reliable
- Semantic coordination seems to undermine both classical model-theoretic interpretation and the ability to draw inferences from whole dialogues

# Where do we go from here?

- When meanings have *not* changed, inferences work as before
- (How) can we tell whether meanings have changed?
- When meanings do change, (how) can we tell which inferences are preserved?
- Well, how do people do this?
- Need to study the interaction between semantic coordination and inference in dialogues between humans!

# Bibliography I

Asher, Nicholas and Lascarides, Alex 2003.
*Logics of Conversation*.
Cambridge University Press, http://www.cambridge.org/.

Asher, Nicholas 2010.
*Lexical Meaning in Context*.
Cambridge University Press, http://www.cambridge.org/.

Breitholtz, Ellen and Cooper, Robin 2011.
Enthymemes as rhetorical resources.
In Artstein, Ron; Core, Mark; DeVault, David; Georgila, Kallirroi; Kaiser, Elsi; and Stent, Amanda, editors 2011, *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA). Institute for Creative Technologies. 149–157.

# Bibliography II

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.

Chierchia, Gennaro and Turner, Raymond 1988.
Semantics and property theory.
*Linguistics and Philosophy* 11(3).

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

# Bibliography III


Clark, Eve V. and Wong, Andrew D. W. 2002.
Pragmatic directions about language use: Offers of words and relations.
*Language in Society* 31:181–212.

Clark, E. V. 2007.
Young children's uptake of new words in conversation.
*Language in Society* 36:157–82.

Cooper, Robin and Ginzburg, Jonathan 2011.
Negation in dialogue.
In Artstein, Ron; Core, Mark; DeVault, David; Georgila, Kallirroi; Kaiser, Elsi; and Stent, Amanda, editors 2011, *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA). Institute for Creative Technologies.

# Bibliography IV


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 Staffan, Larsson 2015.
Probabilistic type theory and natural language semantics.
Under review.

Cooper, Robin 1998.
Information states, attitudes and dependent record types.
In *ITALLC98*.
85–106.

# Bibliography V


Cooper, Robin 2004.
Dynamic generalised quantifiers and hypothetical contexts.
In Svennerlind, Christer, editor 2004, *Ursus Philosophicus, a festschrift for Björn Haglund*. Department of Philosophy, University of Gothenburg, Gothenburg (Sweden).

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.

# Bibliography VI


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.

# Bibliography VII

Cooper, Robin 2019.
Types as learnable cognitive resources in pyttr.
In Condoravdi, Cleo and King, Tracy Holloway, editors 2019, *Tokens of Meaning: Papers in Honor of Lauri Karttunen*. CSLI Publications. 569–586.

Cooper, Robin gress.
*Type theory and language - From perception to linguistic communication*.

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.

# Bibliography VIII

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).

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.

# Bibliography IX


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.

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*, Athens, Greece. Association for Computational Linguistics.
1–9.

# Bibliography X

Larsson, Staffan 2007.
Coordinating on ad-hoc semantic systems in dialogue.
In Artstein, Ron and Vieu, Laure, editors 2007, *Proceedings of the 11th workshop on the semantics and pragmatics of dialogue*, Rovereto, Italy. SemDial.
109–116.

Larsson, Staffan 2010.
Accommodating innovative meaning in dialogue.
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.
83–90.

# Bibliography XI


Larsson, Staffan 2011.
The ttr perceptron: Dynamic perceptual meanings and semantic coordination.
In Artstein, Ron; Core, Mark; DeVault, David; Georgila, Kallirroi; Kaiser, Elsi; and Stent, Amanda, editors 2011, *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA). Institute for Creative Technologies.

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

# Bibliography XII


Ludlow, Peter 2014.
*Living Words: Meaning Underdetermination and the Dynamic Lexicon.*
Oxford University Press.

Marconi, Diego 1997.
*Lexical competence.*
MIT press.

Myrendal, Jenny 2015.
*Word Meaning Negotiation in Online Discussion Forum Communication.*
Ph.D. Dissertation, University of Gothenburg.

# Bibliography XIII


Noble, Bill; Sayeed, Asad; and Larsson, Staffan 2019.
Towards a formal model of word meaning negotiation.
In *Proceedings of the 23rd Workshop on the Semantics and Pragmatics of Dialogue - Poster Abstracts*, London, United Kingdom. SEMDIAL.

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.

# Extra slides

# Aside: perception and truth

- This does not mean that perception and truth are conflated
- If $N$ and her friend $M$ were at the zoo and they both referred to the panda as a bear, it would still be a panda
- However, instead of saying that they were wrong, one could instead say that they (inadvertedly) established a local semantic convention of not making any distinction between pandas and bears
- Given this convention, they were perfectly right to call the panda a bear

# First language acquisition vs. coordination in general

- First language acquisition examples: child detects innovative (for her) uses and adapts her take on the meaning accordingly.
- Semantic coordination in first language acquisition is a special case of semantic coordination in general
  - Asymmetry between the agents involved with respect to expertise in the language being acquired.
- Working hypothesis: the mechanisms for semantic coordination used in these situations are similar to those which are used when competent adult language users coordinate their language.

# What about meaning postulates?

- Extending and reducing sets can be done using meaning postulates
- For example, we can learn that bears are brown by adding the meaning postulate $\forall x.\text{bear}(x) \rightarrow \text{brown}(x)$
- In the model, this amounts to an update [bear]'=[bear]$\setminus$[$\neg$brown]

# Problem 1 for meaning postulates: Subsymbolic updates

- This may work for symbolic updates *but* does not handle perceptual/subsymbolic meanings insofar as these are defined in terms of classifiers of sensor data which can be updated in small increments which do not correspond to meanings of any other expression; there is no $P$ such that [e]'=[e]$\backslash P$
- Example: bear-shaped, blue
- If we believe that a large portion of NL expressions "bottom out" in perceptual meanings, then we face the problem of how we are able to know the extensions of these expressions

# Problem 2 for meaning postulates: The universal base vocabulary

- for the meaning postulate account to hold up as a general account of sem coord, *all* expressions must
  - ... have determinate extensions
  - ... be learnt in a categorial fashion; essentially by being given definitions stated in terms of already known expression
- Since all concepts are defined in terms of other concepts, to avoid infinite regress some meanings must be given before learning starts, and all subsequent meanings must be recombinations of these initial meanings
- We can call this *the universal base vocabulary* - universal (at least) in the sense of underlying all other meanings
- Can there be such a vocabulary?
- How does it get instilled in the mind?

# Meanings are social (II)

- How can speakers can be coordinated with respect to something which potentially has rather different physio-biological realisations in each speaker?
  - speakers are coordinated on a perceptual type insofar as they agree on their classifications of instances as belonging to that type or not,
  - regardless of how how the classifier is implemented in each speaker.
- On this view, meanings are neither purely abstract entities nor purely psychological entities
- Instead, they are social entities which have have both an abstract and a psychological side.

# Aside: Predicate modifiers

- "Blue" can also be analysed as a predicate modifier, i.e. a function from predicates to extensions: blue(glove)(g).
- All this does is say that "blue" can mean different things depending on the object it is ascribed to.
- However, we still end up with extensions, i.e. sets, which can only be modified by adding or deleting elements