Towards a formal view of coordination and learning in dialogue

Robin Cooper, Staffan Larsson

University of Gothenburg
This research funded by The Swedish Bank Tercentenary Foundation

Coordination of Agents Workshop
November 4th-5th 2008, London

Learning from innovation

A frequent pattern in corrective feedback:

original utterance A says something

innovative utterance B says something parallel to A's utterance, containing a use which is innovative for A

learning step A learns from the innovative use

Gloves example

Naomi (2;7.16)

- 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.

Panda example (constructed)

- A: That's a nice bear.
- B: Yes, it's a nice panda.

Turn over example

Abe (2;10.20; Kuczaj 49: 206) [Need to make explicit what Kuczaj is]

- 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.

Breaking down the learning step

1. Align innovative utterance with original utterance
2. Use alignment to predict syntactic and semantic properties of innovative use
3. Integrate innovative element into *local grammar/lexicon* and *local ontology*.
4. Gradually refine syntactic and semantic properties of innovative use and incorporate into more general *linguistic resources* and more general *ontologies*.

Alignment in the gloves example

Naomi: mittens

|

Father: gloves

Alignment in the panda example

A: That's a nice bear

B: Yes, it's a nice panda

Alignment in the turn over example

Abe: Can you	tip	it	over		
	|	|	|		
Mother: Okay I'll	turn	it	over	for	you

Predicting syntactic properties in the glove example

Evidence from alignment.

Naomi: [$_N$ mittens]

↓

Father: [$_N$ gloves]

Predicting syntactic properties in the panda example

Evidence from

- alignment
- active chart edge
- (resulting passive edge + alignment)

A: That's $[_{NP}\ [_{Det}\ \text{a}]\ [_{A}\ \text{nice}]\ [_{N}\ \text{bear}]]$

B: Yes, it's $[_{NP}\ [_{Det}\ \text{a}]\ [_{A}\ \text{nice}]\ [_{N}\ \text{panda}]]$

Active edge: NP $\rightarrow [_{Det}\ \text{a}]\ [_{A}\ \text{nice}] \bullet$ N

Predicting syntactic properties in the turn over example

Evidence from

- alignment
- resulting passive edge (+ alignment)

Abe: Can you $[_{VP}$ $[_{V}$ tip$]$ it over $]$

Mother: Okay I'll $[_{VP}$ $[_{V}$ turn$]$ it over $]$ for you

Primacy of alignment

Hypothesis: Alignment evidence is primary in predicting syntactic properties of innovations when it available (as it is in corrective feedback). Other evidence supports or refutes.

Predicting compositional semantic properties

Following Montague, Blackburn + Bos, compositional semantics can be predicted from syntactic information such as category

- commonNounSemantics(N) $= \lambda x N'(x)$

or, using TTR,

- commonNounSemantics(N) $= \lambda r{:}\left[\text{x} : \mathit{Ind}\right] \left(\left[\text{c} : N'(r.\text{x})\right]\right)$

Deriving compositional semantics from record types

There's an obvious relationship between the function

- $\lambda r{:}\begin{bmatrix}\text{x} : \mathit{Ind}\end{bmatrix}\ (\begin{bmatrix}\text{c} : N'(r.\text{x})\end{bmatrix})$

and the record type

- $\begin{bmatrix}\text{x} : \mathit{Ind} \\ \text{c} : N'(\text{x})\end{bmatrix}$

The function can be derived from the record type.
This will be important for relating *compositional* semantics to *ontological* semantics.

Relating compositional semantics to domain ontologies

- We use local domain ontologies to *refine* compositional semantics.
- The same compositional semantics can get different refinements depending on context.

Ontological classes as record types

- Thing $= \left[\text{x} : \textit{Ind} \right]$
- $\{\text{Class } P\} = \begin{bmatrix} \text{x} : \textit{Ind} \\ \text{c}_P : P(\text{x}) \end{bmatrix}$
  "Create a class based on predicate P"
- $\{\text{SubClass } C_1\ C_2\} = C_1 \dot{\wedge} C_2$
  "Make C_1 be a subclass of C_2"

Example of subclasses

$$C_1 = \begin{bmatrix} \mathrm{x} : \mathit{Ind} \\ \mathrm{c}_{\text{clothing}} : \text{clothing(x)} \end{bmatrix}$$

$$C_2 = \begin{bmatrix} \mathrm{x} : \mathit{Ind} \\ \mathrm{c}_{\text{physobj}} : \text{physobj(x)} \end{bmatrix}$$

- $C_1 \wedge C_2$ is a type.
  $r : C_1 \wedge C_2$ iff $r : C_1$ and $r : C_2$
- $C_1 \dot{\wedge} C_2 = \begin{bmatrix} \mathrm{x} : \mathit{Ind} \\ \mathrm{c}_{\text{physobj}} : \text{physobj(x)} \\ \mathrm{c}_{\text{clothing}} : \text{clothing(x)} \end{bmatrix}$

$C_1 \dot{\wedge} C_2$ is a *refinement* of C_2 and also of C_1.
$C_1 \dot{\wedge} C_2 \sqsubseteq C_2$ and $C_1 \dot{\wedge} C_2 \sqsubseteq C_1$

Predicting ontological semantics properties in the glove example

Naomi's pre-gloves ontology
PhysObjClass = {Class physobj}
ClothingClass = {SubClass {Class clothing} PhysObjClass}
MittenClass = {SubClass {Class mitten} ClothingClass}

Naomi's post-gloves ontology
PhysObjClass = {Class physobj}
ClothingClass = {SubClass {Class clothing} PhysObjClass}
MittenClass = {SubClass {Class mitten} ClothingClass}
GloveClass = {SubClass {Class glove} ClothingClass} (from alignment of *mittens* and *gloves*)

Refining compositional semantics with the ontology

$$\text{GloveClass} = \begin{bmatrix} \text{x} & : & \mathit{Ind} \\ \text{c}_{\text{physobj}} & : & \text{physobj(x)} \\ \text{c}_{\text{clothing}} & : & \text{clothing(x)} \\ \text{c}_{\text{glove}} & : & \text{glove(x)} \end{bmatrix}$$

$$\text{GloveCompSem} = \begin{bmatrix} \text{x} & : & \mathit{Ind} \\ \text{c}_{\text{glove}} & : & \text{glove(x)} \end{bmatrix}$$

Refinement of GloveCompSem with GloveClass is

- GloveCompSem $\dot{\wedge}$ GloveClass
- i.e. GloveClass

Refined compositional semantics is derived from GloveClass

- $\lambda r{:}\begin{bmatrix}\text{x} : \mathit{Ind}\end{bmatrix} \left(\begin{bmatrix} \text{c}_{\text{physobj}} : \text{physobj}(r.\text{x}) \\ \text{c}_{\text{clothing}} : \text{clothing}(r.\text{x}) \\ \text{c}_{\text{glove}} : \text{glove}(r.\text{x}) \end{bmatrix} \right)$

Articulating Naomi's ontology

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

PhysObjClass = {Class physobj}
ClothingClass = {SubClass {Class clothing} PhysObjClass}

HandClothingClass = {SubClass {Class handclothing} ClothingClass}
WithFingersClass = {SubClass {Class withfingers} HandClothingClass}
WithoutFingersClass = {SubClass {Class withoutfingers} HandClothingClass}
MittenClass = WithoutFingersClass
GloveClass = WithFingersClass

Using *gloves* with different ontologies

Constructed examples.

- A: Do you have a pair of gloves I can borrow?
- B: *Hands A a pair of mittens.*
- A: Thanks.

- A: Do you have a pair of mittens I can borrow?
- B: *Hands A a pair of gloves*
- A: Don't you have any mittens?

Alternative ontologies and experience

GloveClass = HandClothingClass

but *not*

MittenClass = HandClothingClass

but when little children wear things on their hands it's normally mittens . . .

Some dialogue games for semantic coordination

- Corrective feedback: "Yes, it's a nice panda"
- Ostensive definition: "That's a bug"
- Explicit definition: "when they have fingers in them they are called gloves and when the fingers are all put together they are called mittens."
- *Meaning accommodation*: tacitly altering the meaning of an expression to make it fit with a novel use

A problem with studying meaning accommodation

[A methodological difficulty] is how to identify occurrences of accommodation, i.e. silent acceptance of a novel use of a word, in a corpus material with any degree of reliability; the researcher lacks access to the speaker's meanings.

Prof. Lars Ahrenberg, assessing the SemCoord project application

- This is indeed a problem! (How) can it be solved?

How to study meaning accommodation

Sue shared my view of young children as sucking up words like little vacuum cleaners. Her idea was to drop a nonsense syllable into their intake and follow what happened to it. She did not want to teach them the word, in any ordinary sense of the word "teach", but simply to insert it into conversation in as natural a way as possible... The advantage of using a nonsense syllable was that we could be sure it was unfamiliar and that the children would not encounter it outside the playroom - we would have a complete record of all their experience with it.

G.A. Miller: "Spontaneous Apprentices" (1976)

- Can the existence of meaning accommodation be proven experimentally?
- Carey & Bartlett 1978: Acquiring a single new word. Papers and Reports on Child Language Development.

Outline of the experiment

- Idea: mimic the circumstances in which children naturally encounter new words (subjects: 3- and 4-year olds)
- To enable test for learning effects, use nonsense word: "chromium" to refer to the colour olive
- Avoid pointing to a sample of the colour olive and say "this is called chromium" (this would be an ostensive definition)

Outline of the experiment

- Avoid pointing to a sample of the colour olive and say "this is called chromium" (this would be an ostensive definition)
- Instead, use it in the normal (nursery school) classroom activity of preparing snack time:
  - One cup painted olive, another painted red
  - (adult to child) "Bring the the chromium cup; not the red one, the chromium one."
  - All children picked the right cup; however, this could be done by focusing on the contrast "not the red one" without attending to the word "chromium"
  - In pilot test, most children asked for confirmation: "Do you mean this one?"; 4 of 14 repeated "chromium"

Pre- and post-exposure tests

- Pre-exposure test: Before exposure to "chromium", the children called olive "green" or sometimes "brown" in production test
- Post-exposure comprehension tests
  - Pointing test: children were asked to point to various colour samples as they were named
  - Hyponym test, since earlier work indicated that one of the earlier things children learn about colour words is that they are hyponyms of the word "colour"
- Cycles:
  - Cycle 1: one exposure
  - Cycle 2, ten weeks later: two further exposures
  - After cycle 2, total of 5 exposures (including exposures in tests)

Results

- Comprehension test: after two exposures, significant effect (35% to 63% increase in correctness)
- Naming task: by cycle 2, significantly more children (than in control group) showed evidence of learning that there is a word for olive, but could not remember it; one child had learned "chromium"
- Hyponym task: after first exposure, 50% (of children who understood the task) judged "chromium" a colour.
- 5 exposures had influenced the child's naming of olive; had effected a lexical entry for "chromium" which in many cases included that it was a colour term, and in some cases knowledge of its referent
- Some learning seems to occur after a single exposure, at least sometimes; one child named olive "cram" after a single exposure

Results

- Acquisition in two phases:
  - *fast mapping* resulting from one or a few exposures to the new word; includes only a fraction of the total information constituting full learning of the word; typically includes hyponym relations
  - extended mapping, over several months, by which children arrive at full acquisition, including the ability to identify and name new instances

- Lessons from the Carey & Bartlett experiment
  - Shows learning from exposure without corrective feedback
  - Learning of ontological class (fast mapping)
  - Learning of *perceptual type* (extended mapping)
  - Transfer of learning from snack-preparation activity to naming and hyponym tests
- Raises two (sets of) issues
  - Which aspects of meaning are there? How do they differ in how they are learnt, and how they can change over time?
  - How do changes in meaning transfer between activities and activity types? Which factors influence transfer? Are there any general rules of transfer?
- Also shows how learning from interaction can be experimentally tested; in this case, for meaning accommodation

Aspects of meaning

Word meanings can be defined using one or more of the following:

- CompSem - compositional semantics
  - Represented e.g. as lambda expressions
- Class - ontological class: relations to other concepts
  - Represented e.g. as relations in OWL-style ontology
- PerceptualType - *perceptual type*: a classifier of sensory input as indicating presence of referent or not
  - Can be realised in many different ways (neural network, statistical model, etc.)
- Connotation - connotational type: non-truth-conditional aspects of concept, e.g. positive or negative value
- ...

Intensional meaning

- Each meaning representation in the lexicon can be structured according to the above aspects of meaning
- $[\text{“glove”}] = \begin{bmatrix} \text{compsem} & : & \text{GloveCompSem} \\ \text{class} & : & \text{GloveClass} \\ \text{perc} & : & \text{PercType} \end{bmatrix}$
- The intensional meaning of “glove” can be defined as

$$\dot{\bigwedge}[\text{“glove”}] = [\text{“glove”}].\text{compsem} \mathrel{\dot{\wedge}} [\text{“glove”}].\text{class} \mathrel{\dot{\wedge}} \ldots$$

  - if $T = \begin{bmatrix} \ell_1 & : & T_1 \\ \vdots & & \\ \ell_n & : & T_n \end{bmatrix}$, then $\dot{\bigwedge} T = T_1 \mathrel{\dot{\wedge}} \ldots \mathrel{\dot{\wedge}} T_n$.

Intensional meaning

$$\bigwedge\!\!\!\!\cdot\,[\text{“glove”}] =$$

$$[\text{“glove”}].\text{compsem} \wedge\!\!\!\!\cdot\; [\text{“glove”}].\text{class} \wedge\!\!\!\!\cdot\; [\text{“glove”}].\text{perctype} =$$

$$\begin{bmatrix} \text{x} & : & \mathit{Ind} \\ \text{c}_{\text{physobj}} & : & \text{physobj(x)} \\ \text{c}_{\text{clothing}} & : & \text{clothing(x)} \\ \text{c}_{\text{glove}} & : & \text{glove(x)} \\ \text{c}_{\text{glove-perc}} & : & \text{glove-perc(x)} \end{bmatrix}$$

- This is just one way of combining meaning aspects, which makes the assumption that they all have the same status.

Intension and appropriateness

- Intuition: the intensional meaning of an expression say something about the kind of context in which it can be appropriately used.
- "glove" can be appropriately used in c iff $c : \dot{\bigwedge}[\text{“glove”}]$
- In general:
  - An expression e can be appropriately used in c iff $c : \dot{\bigwedge}[e]$
  - The intension $\dot{\bigwedge}[e]$ of an expression e is a *context type*.

Intension and extension

- Theoretically, the extension of an expression represented by a type is the set of (potential) contexts of that type; the extension $EXT(e)$ of an expression e is the set of contexts which are of type $\dot{\bigwedge}[e]$:
  - $EXT(e) = \{c \mid c : \dot{\bigwedge}[e]\}$
- We are still making the assumption that they all have the same status
- However, meaning is complex, involving several types corresponding to different aspects of intensional meaning

Intension and extension

- Are all aspects of equal importance?
  - Probably not; some are obligatory, others optional (knowing what a glove looks like is part of knowing the meaning of "glove", but one can use "glove" even if there is not glove in the context)
- Does a meaning correspond to a single intension and a single extension (as assumed above), or does each meaning aspect determine its own intension and extension?
  - An argument for the latter: the connotational extension is typically a proper subset of the truth-conditional extension ("tibia" and "shinbone" have same TC extensions but different connotational extensions, both proper subsets of the TC extension)
- Perhaps we need intensions, extensions and judgements of appropriateness on each level individually

Semantic change in terms of extension

- In general, the extension of an expression can expand, contract, or shift (both expand and contract; in the extreme case, the intersection of the old and the new extension is empty)
- These changes correspond to changes in the intensional type
  - Expansion of extension = generalisation of type
  - Contraction of extension = refinement of type
  - Shift of extension = both generalisation and refinement of type
- All aspects of lexical meaning are subject to these operations (updates)

Semantic change in terms of intension

- How do these aspects of meaning differ in how they are learnt, and how they can change over time?
- Idea:
  - Symbolic learning (cf. adding facts to database) is fast
    - Compositional semantics and ontological class are symbolic
  - Subsymbolic learning (cf. training neural nets) is slower
    - Perceptual type is (partly) subsymbolic
  - So compositional semantics and ontological class can be learnt quickly (fast mapping), whereas perceptual type is learnt slowly (extended mapping)
  - Seems to fit with results from Carey&Bartlett example

Extending the experiment: learning and context-types

- How do changes in meaning transfer between activities and activity types? Which factors influence transfer? Are there any general rules of transfer?
- In Carey & Bartlett, the test were done in a completely different context from the one in which the word had been introduced
- Would the effects have been stronger if tests had been done in the same type of context?

Extending the experiment: learning and context-types, cont'd

- Idea for experiment: learning within and across context types
  - Are concepts immediately equally available in all context types?
  - Are concepts initially context-type-specific and then get gradually generalised as they are encountered in other context types?
- If the latter, does this generalisation follow any specific patterns, e.g.
  - If c used in X and Y, and Z is the most specific supertype of both X and Y, does c become more easily available in all other context types which are subtypes of X?

Context-type trees

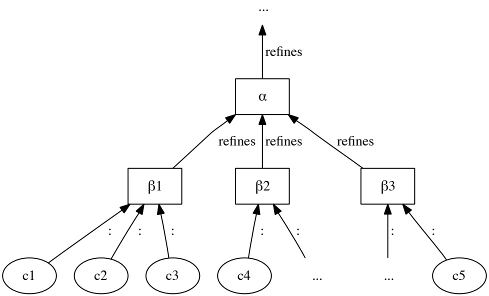

- $\alpha, \beta_1, \ldots, \beta_n$ are context types; $c_1, \ldots, c_m$ are contexts
- We can order context types by the refinement relation into (multiple inheritance) trees

Meanings and context-type trees

Assumption:
If $c : \alpha$ (c is a context of type α), $[\mathrm{e}]_s$ should be more refined (more specific) than $[\mathrm{e}]_\alpha$.

Instantiation of context type: $\mathrm{s} : \alpha \rightarrow [\mathrm{e}]_c \sqsubseteq [\mathrm{e}]_\alpha$

Also, if situation type α is a subtype of situation type β, the α-specific meaning should be a refinement of the β-specific one:

Context subtypes: $\beta \sqsubseteq \alpha \rightarrow [\mathrm{c}]^A_\beta \sqsubseteq [\mathrm{c}]^A_\alpha$

This relation should probably be transitive.

Direct transfer between context types

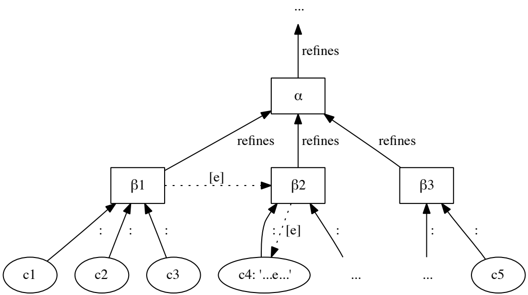

- No generalisation stored
- Meaning in c_5 requires (1) generalisation over $[\mathrm{e}]_{\beta_1}$ and $[\mathrm{e}]_{\beta_2}$, and (2) refinement to $[\mathrm{e}]_{c_5}$, possibly via $[\mathrm{e}]_{\beta_3}$

Indirect transfer between context types

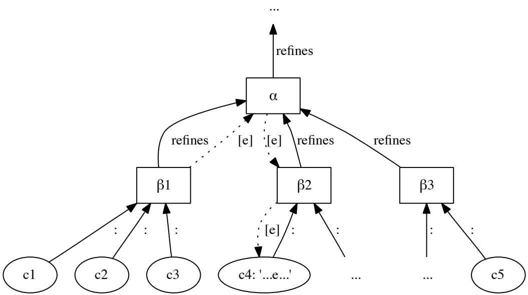

- Transfer through generalisation and refinement
- Generalisation stored and available for future use
- Meaning in c_5 requires only refinement of $[\mathrm{e}]_\alpha$ to $[\mathrm{e}]_{c_5}$, possibly via $[\mathrm{e}]_{\beta_3}$

Generalisation and context types

- In principle, any aspect of the context may be used as basis for generalisation from context to meaning (context type):
  - Jointly perceived environment (e.g. shared focus of attention)
  - Linguistic incl. grammatical/syntactic context
  - Activity (as encoded e.g. in scripts or frames)
  - Current dialogue partner or language community
  - Physical and temporal setting
  - Institutional setting
- When generalising, there are many alternative context types to generalise to
  - Which principles guide this selection?
  - How are different aspects of context related to aspects of meaning?

Parameters for experiments

We need to know more about how all these parameters interact:

- Kind of dialogue game
  - corrective feedback, ostensive definition, accommodation, ...
- Kinds of evidence available for learner
  - syntactic alignment, ...
- Kind of meaning learnt
  - compositional, ontological, perceptual, connotational, ...
- Kind of learning mechanisms and representations
  - symbolic, subsymbolic
- Kind of transfer in learning across contexts and context types
  - direct, indirect
- Aspect of context used as basis for generalisation

Example: Does corrective feedback involving syntactic alignment affect learning of (symbolic) ontological meaning directly or indirectly across activity types?