Accommodating innovative meaning in dialogue

Staffan Larsson
University of Gothenburg
Sweden
sl@ling.gu.se

Abstract

Meaning accommodation occurs when a dialogue participant learns new meanings by observing utterances in context and adapting to the way linguistic expressions are used. This paper sketches a formal account of meaning accommodation, focusing on the detection of subjectively innovative uses of an expression. It also provides an example of adaptation of lexical meaning to accommodate innovation.

1 Introduction

Two agents do not need to share exactly the same linguistic resources (grammar, lexicon etc.) in order to be able to communicate, and an agent's linguistic resources can change during the course of a dialogue when she is confronted with a (for her) innovative use. For example, research on alignment shows that agents negotiate domain-specific microlanguages for the purposes of discussing the particular domain at hand (Clark and Wilkes-Gibbs, 1986; Pickering and Garrod, 2004; Brennan and Clark, 1996; Healey, 1997; Larsson, 2007). We use the term *semantic coordination* to refer to the process of interactively coordinating the meanings of linguistic expressions.

Several mechanisms are available for semantic coordination in dialogue. These include corrective feedback, where one DP implicitly corrects the way an expression is used by another DP, as well as explicit definitions and negotiations of meanings. However, it also possible to coordinate silently, by DPs observing the language use of others and adapting to it.

This paper aims towards a formal account of learning meanings from observation and accommodation in dialogue. The examples we present are from first language acquisition, where the child detects innovative (for her) uses and adapts her take on the meaning accordingly[1].

In the following, we will first introduce the basics of the TTR framework (section 2). We then review earlier work on learning meaning from corrective feedback, and discuss how various aspects of meaning can be represented using TTR (section 5). The following two sections show how we use TTR to represent multiple aspects of lexical meaning, including "perceptual type". Section 6 introduces the notion of meaning accommodation and briefly describes an experiment demonstrating that children can learn lexical meaning using meaning accommodation. Sections 7 and 8 explain how contexts are represented using TTR and provide an example of "normal" contextual interpretation. In section 9 we define a notion of innovation that we use to detect innovation in an example interaction where normal contextual interpretation breaks down, and specify the meaning update resulting from this interaction.

2 TTR

The received view in formal semantics (Kaplan, 1979) assumes that there are abstract and context-independent "literal" meanings (utterance-type meaning; Kaplan's "character") which can be regarded formally as functions from context to content; on each occasion of use, the context determines a specific content (utterance-token meaning). Abstract meanings are assumed to be static and are not affected by language use in specific contexts. Traditional formal semantics is thus ill-equipped to deal with semantic coordination, because of its static view of meaning.

[1] We regard semantic coordination in first language acquisition as a special case of semantic coordination in general, wherethere is a clear asymmetry between the agents involved with respect to expertise in the language being acquired when a child and an adult interact. However, we believe that the mechanisms for semantic coordination used in these situations are similar to those which are used when competent adult language users coordinate their language.

Aspects of Semantics and Pragmatics of Dialogue. SemDial 2010, 14th Workshop on the Semantics and Pragmatics of Dialogue

We shall make use of type theory with records (TTR) as characterized in Cooper (2005) and elsewhere. The advantage of TTR is that it integrates logical techniques such as binding and the lambda-calculus into feature-structure like objects called record types. Thus we get more structure than in a traditional formal semantics and more logic than is available in traditional unification-based systems. The feature structure like properties are important for developing similarity metrics on meanings and for the straightforward definition of meaning modifications involving refinement and generalization. The logical aspects are important for relating our semantics to the model and proof theoretic tradition associated with compositional semantics.

We will now briefly introduce the TTR formalism. 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:

$$\begin{bmatrix} l_1 & = & a_1 \\ l_2 & = & a_2 \\ \ldots & & \\ l_n & = & a_n \\ \ldots & & \end{bmatrix}$$

is of type:

$$\begin{bmatrix} l_1 & : & T_1 \\ l_2 & : & T_2(l_1) \\ \ldots & & \\ l_n & : & T_n(l_1, l_2, \ldots, l_{n-1}) \end{bmatrix}$$

Types constructed with predicates may also be *dependent*. This is represented by the fact that arguments to the predicate may be represented by labels used on the left of the ':' elsewhere in the record type.

Some of our types will contain *manifest fields* (Coquand et al., 2004) like the ref-field in the following type:

$$\begin{bmatrix} \text{ref=obj123} & : & \text{Ind} \\ \text{c} & : & \text{glove(ref)} \end{bmatrix}$$

$\left[\text{ref=obj123:Ind}\right]$ is a convenient notation for $\left[\text{ref} : \text{Ind}_{\text{obj123}}\right]$ where $\text{Ind}_{\text{obj123}}$ is a *singleton type*. If $a : T$, then T_a is a singleton type and $b : T_a$ (i.e. b is of type T_a) iff $b = a$. Manifest fields allow us to progressively specify what values are required for the fields in a type.

An important notion in this kind of type theory is that of *subtype*. Formally, $T_1 \lhd T_2$ means that T_1 is a subtype of T_2. Two examples will suffice as explanation of this notion:

$$\begin{bmatrix} \text{ref} & : & \text{Ind} \\ \text{c} & : & \text{glove(ref)} \end{bmatrix} \lhd \begin{bmatrix} \text{ref} & : & \text{Ind} \end{bmatrix}$$

$$\begin{bmatrix} \text{ref=obj123} & : & \text{Ind} \end{bmatrix} \lhd \begin{bmatrix} \text{ref} & : & \text{Ind} \end{bmatrix}$$

Below, we will also have use for an operator for combining record types. The $\dot{\wedge}$ operator works as follows. Suppose that we have two record types C_1 and C_2:

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

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

$C_1 \wedge C_2$ is a type. In general if T_1 and T_2 are types then $T_1 \wedge T_2$ is a type and $a : T_1 \wedge T_2$ iff $a : T_1$ and $a : T_2$. A meet type $T_1 \wedge T_2$ of two record types can be simplified to a new record type by a process similar to unification in feature-based systems. We will represent the simplified type by putting a dot under the symbol $\wedge$. Thus if T_1 and T_2 are record types then there will be a type $T_1 \dot{\wedge} T_2$ equivalent to $T_1 \wedge T_2$ (in the sense that a will be of the first type if and only if it is of the second type).

$$C_1 \dot{\wedge} C_2 = \begin{bmatrix} \text{x} : \textit{Ind} \\ \text{c}_{\text{physobj}} : \text{physobj(x)} \\ \text{c}_{\text{clothing}} : \text{clothing(x)} \end{bmatrix}$$

The operation $\dot{\wedge}$ corresponds to unification in feature-based systems and its definition (which we omit here) is similar to the graph unification algorithm.

3 Learning compositional and ontological meaning from corrective feedback and explicit definition

This section gives a brief overview of the work presented in (Cooper and Larsson, 2009) , which will be useful as a backdrop to the description of meaning accommodation. It will also introduce some important concepts which we will be making use of later.

Recent research on first language acquisition (Clark, 2007; Saxton, 2000) argues that the learning process crucially relies on negative input, including corrective feedback. We see corrective feedback as part of the process of negotiation of a language between two agents. Here is one of the

examples of corrective feedback that we discuss in connection with our argument for this position:

"Gloves" example (Clark, 2007):

- 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.

We think that an important component in corrective feedback of this kind is syntactic alignment, that is, alignment of the correcting utterance with the utterance which is being corrected. This is a rather different sense of alignment than that associated with the negotiation of a common language, although the two senses are closely linked. By "syntactic alignment" here we mean something related to the kind of alignment that is used in parallel corpora. It provides a way of computing parallelism between the two utterances. Syntactic alignment may not be available in all cases but when it is, it seems to provide an efficient way of identifying what the target of the correction is.

Following Montague (1974) and Blackburn and Bos (2005) compositional semantics can be predicted from syntactic information such as category. For example, for common nouns we may use the formula

$$\text{commonNounSemantics}(N) = \lambda x N'(x)$$

or, using TTR,

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

There is an obvious relationship between this function and the record type

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

It is clear that the compositional function can be derived from the record type (but we will not provide a general definition of this derivation here due to space limitations). We will use the record type representation in the remainder of this paper.

In the Gloves example, after the father's utterance of "gloves", Naomi could use syntactic alignment to understand this term as a noun with the corresponding kind of compositional semantics:

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

In this way, compositional semantics can be derived from corrective feedback in dialogue. However, compositional semantics of this kind does not reveal very much, if anything, about the details of word semantics unless we add ontological information.

The ontological semantics of a lexical expression describes how the concept associated with the expression relates to other relevant concepts in an ontology. Provided that Naomi learns from the interaction that gloves are also a kind of clothing, Naomi's glove class after the first utterance by the father is the following type

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

The father's second utterance contains a partial but explicit definition of the ontology of gloves and mittens:

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

When integrating this utterance, Naomi may modify her take on the ontological semantics so that after this update the meanings for "glove'" will be:

$$\begin{bmatrix} \text{x} : \mathit{Ind} \\ \text{c}_{\text{glove}} : \text{glove}'(\text{x}) \\ \text{c}_{\text{physobj}} : \text{physobj}(\text{x}) \\ \text{c}_{\text{clothing}} : \text{clothing}(\text{x}) \\ \text{c}_{\text{handclothing}} : \text{handclothing}(\text{x}) \\ \text{c}_{\text{withoutfingers}} : \text{withoutfingers}(\text{x}) \end{bmatrix}$$

4 Perceptual type

In this paper, we will add a further aspect of meaning, namely *perceptual type* (or perceptual meaning).

For our current purposes, we will represent perceptual meaning as a record type specifying and

individual and one or more propositions indicating that the individual is of a certain perceptual type, i.e., that it has certain physically observable characteristics.

The word "glove′", for example, may be associated with a certain shape:

GlovePercType =
$$\begin{bmatrix} \text{x} : \text{Ind} \\ \text{c}_{\text{glove-shape}} : \text{glove-shape(x)} \end{bmatrix}$$

Propositions, i.e., types which are constructed with predicates, are sometimes referred to as "types of proof". The idea is that something of this type would be a proof that a given individual has a certain property. One can have different ideas of what kind of objects count as proofs. For subsymbolic aspects of meaning, we are assuming that the proof-objects are readings from sensors. Thus, we assume that GlovePercType above will be inhabited if the agent detects a glove-shaped object in the context.

5 Aspects of meaning in TTR

Since all aspects of meaning can be modified as a result of language use in dialogue, we want our account of semantic innovation and semantic updates to include several aspects of lexical meaning. In this section, we will show how multiple aspects of lexical meaning can be represented using TTR.

Each lexical meaning representation in the lexicon can be structured according to the above aspects of meaning.

$$[\text{glove}] = \begin{bmatrix} \text{comp} & : & \text{GloveCompSem} \\ \text{class} & : & \text{GloveClass} \\ \text{perc} & : & \text{GlovePercType} \end{bmatrix}$$

A simplified representation can be obtained by collapsing the distinctions between the different aspects of meaning.[2].

$\bigwedge[\text{glove}] = [\text{glove}].\text{comp} \wedge [\text{glove}].\text{class} \wedge [\text{glove}].\text{perc} =$
$$\begin{bmatrix} \text{x} : \textit{Ind} \\ \text{c}_{\text{glove}} : \text{glove}'(\text{x}) \\ \text{c}_{\text{physobj}} : \text{physobj(x)} \\ \text{c}_{\text{clothing}} : \text{clothing(x)} \\ \text{c}_{\text{glove-shape}} : \text{glove-shape(x)} \end{bmatrix}$$

[2]We define the operator $\bigwedge$ as follows:

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

6 Meaning accommodation

To show how meaning can be learnt by observations of language use, (Carey and Bartlett, 1978) set up an experiment based on nonsense words to mimic the circumstances in which children naturally encounter new words. The subjects were 3- and 4-year olds. To enable testing for learning effects, they used the nonsense word: "chromium" to refer to the colour olive.They use it in the normal (nursery school) classroom activity of preparing snack time. In the experiment, one cup painted olive, and another was painted red. The adult test leader says to the child "Bring me 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".

The results indicated that a very low number of exposures (five) had influenced the child's naming of olive and 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. Based on this, it was concluded that acquisition proceeds in two phases. Firstly, *fast mapping* resulting from one or a few exposures to the new word; this includes only a fraction of the total information constituting full learning of the word, and typically includes hyponym relations. Secondly, *extended mapping* over several months, by which children arrive at full acquisition, including the ability to identify and name new instances. The experiment demonstrates learning from exposure without corrective feedback, and shows that both ontological information and perceptual type can be learnt in this way.

In accommodating innovative meanings, the innovation can concern either a new word for the DP, as in the experiment with "chromium", or a known word which is used in a new way. In the latter case, the needs to identify a way of using an expression that deviates from the way that expression is usually used (in the DP's subjective experience).

Below, we will propose a formalisation of the notion of innovative use of a known linguistic expression, in terms of a relation between the context of utterance and the meaning of an expression. We will also provide an illustrative example which includes the semantic updates resulting from the accommodation.

7 Representing contexts in TTR

To represent individual dialogue participants' takes on contexts[3], we will use record types with manifest fields. This allows our context to be underspecified, reflecting the fact that an agent may not have a complete representation of the environment. For example, an agent may entertain some propositions without necessarily having a proof for them, or know of the existence of certain objects without having identified them.

In first language acquisition, learning of perceptual type typically takes place in contexts where the referent is in the shared focus of attention and thus perceivable to the dialogue participants, and for the time being we limit our analysis to such cases. For our current purposes, we assume that our DPs are able to establish a shared focus of attention, and we will designate the label "focobj" for the object or objects taken by a DP to be in shared focus.

We assume that a DP's take on the context may be amended during interpretation by checking the environment (Knoeferle and Crocker, 2006). That is, the take on the context against which an utterance is interpreted and evaluated may not be identical to the DP's take on the context prior to the interpretation process. We will not be going into details of incremental interpretation here, however.

8 Contextual interpretation for non-innovative uses

As a backdrop for our discussion of how to detect innovation, we will first show how "normal" contextual interpretation, in the absence of innovations, is assumed to work.

For our current purposes, we will not go into details of compositional analysis, but instead provide the resulting semantic representation of the meaning of a whole utterance. We will follow Kaplan in assuming this result to be a function from context to content. Parts of the meaning is *foregrounded*, and which parts are *backgrounded*. Background meaning (BG) represents constraints on the context, whereas foreground material (FG) is the information to be added to the context by the utterance in question. We can represent this either as a record or as a function:

$$\begin{bmatrix} \text{BG} = \ldots \\ \text{FG} = \ldots \end{bmatrix}$$

$\lambda t{:}\text{BG(FG)}$

The functional version takes as argument a record type t which is a subtype of the background meaning of the uttered expression (typically a context containing manifest fields representing objects in the environment and propositions about these objects), and returns a record type corresponding to the foreground meaning. Contextual interpretation amounts to applying this function to the context. After contextual interpretation, we assume that the content of the utterance in question is to be added to the information state (Traum and Larsson, 2003) of the interpreting agent, which we here will take to be identical the agent's take on the context[4].

To illustrate contextual interpretation, we will use a modified version of the "gloves" example, where Naomi simply observes an utterance by Father:

Modified "Gloves" example:

- (Naomi is putting on her new gloves)
- Father: Those are nice gloves!

In this section, we will assume that Naomi is familiar with the term "gloves" and assigns it the standard (in the community of adult speakers of English) lexical meaning.

$$\bigwedge_{\cdot}[\text{glove}]^{Naomi} = \begin{bmatrix} \text{x} : \textit{Ind} \\ \text{c}_{\text{glove}} : \text{glove}'(\text{x}) \\ \text{c}_{\text{physobj}} : \text{physobj(x)} \\ \text{c}_{\text{clothing}} : \text{clothing(x)} \\ \text{c}_{\text{glove-shape}} : \text{glove-shape(x)} \end{bmatrix}$$

We assume that Naomi's take on the meaning of the sentence uttered by Father is the following (using record type notation):

$[\text{Those are nice gloves}]^{Naomi} =$

[3] Occasionally and somewhat sloppily referred to as "contexts" below.

[4] In a complete model including both information sharing, dialogue management and semantic coordination the information state would probably need to be more complex than this.

$$\left[\begin{array}{ll} & \left[\begin{array}{l}\text{focobj} : \textit{Ind} \\ \text{c}_{\text{glove}} : \text{glove}'(\text{focobj}) \\ \text{c}_{\text{physobj}} : \text{physobj}(\text{focobj}) \\ \text{c}_{\text{clothing}} : \text{clothing}(\text{focobj}) \\ \text{c}_{\text{glove-shape}} : \text{glove-shape}(\text{focobj})\end{array}\right] \\ \text{FG} = & \left[\text{c}_{\text{nice}} : \text{nice}'(\text{BG.focobj})\right]\end{array}\right]$$

Of course, we are skipping over a lot here, including a compositional analysis and TTR analyses of "those", "are" and "nice". We also ignore the complication that "gloves" refers to a pair of gloves. Furthermore, we assume that the deictic "those" picks out the object in shared focus of attention (focobj).

Now for the context. We assume that when interpreting Father's utterance, Naomi perceives her gloves to be in shared focus. Her take on the situation includes the fact that they are gloves, that they have a certain appearance, and a certain ontological status.

$$c^{Naomi} = \left[\begin{array}{lcl}\text{focobj=a} & : & \text{Ind} \\ \text{c}_{\text{glove}} & : & \text{glove}'(\text{focobj}) \\ \text{c}_{\text{physobj}} & : & \text{physobj}(\text{focobj}) \\ \text{c}_{\text{clothing}} & : & \text{clothing}(\text{focobj}) \\ \text{c}_{\text{glove-shape}} & : & \text{glove-shape}(\text{focobj})\end{array}\right]$$

This allows us to do contextual interpretation by applying (the function version of) $\dot{\bigwedge}$[Those are nice gloves]^{Naomi} to c^{Naomi}:

$$[\text{Those are nice gloves}]^{Naomi}_{c} = [\text{Those are nice gloves}]^{Naomi}(c^{Naomi}) = \left[\text{c}_{\text{nice}} : \text{nice}'(\text{focobj})\right]$$

This record type can then be used to extend the information state (context). Note that nice$'$ will now be predicated of the individual a labeled by focobj in the context.

9 Formalising innovation

This section provides a TTR analysis of the first part of meaning accommodation, i.e. detection of an innovative use. Detecting innovation when a word is completely novel is relatively trivial; one only needs to recognise that it is not a word one has heard before[5]. We will instead focus on the case where a known expression is used with a (subjectively) innovative meaning.

[5]This is not to say that *assigning a meaning* to the new word is in any sense trivial.

We want a formal notion of innovativeness which we can use to detect innovative uses of linguistic expressions. The underlying intuition is that the meaning of an expression should say something about the kind of context in which it can be (non-innovatively) used. But how, exactly?

Intuitively, contextual interpretation can go wrong in at least the following ways:

1. Some information presupposed by the expression contradicts some information in the context; we refer to this as *background inconsistency*.

2. Some content conveyed by the utterance of the expression contradicts something in the context; we refer to this as *foreground inconsistency*.

Formally, these intuitions can be captured as follows. An expression e is *innovative* in context c if there is a mismatch between e and c in either of the following ways[6]:

1. c is inconsistent with the backgrounded meaning BG of e, formally $[e].\text{BG}\dot{\bigwedge}c \approx \bot$ (background inconsistency)

2. the content $[e]_c$ of e in c is inconsistent with c, formally $c\dot{\bigwedge}[e]_c \approx \bot$ (foreground inconsistency)

This definition follows naturally from how contextual interpretation works. Recall that meaning can be seen as a function from context to content, where background meaning serves as a constraint in the context. The definition of innovation checks that it will be possible to apply the meaning-function to the context, by checking that the context is consistent with the constraints imposed by the backgrounded meaning, and that the resulting contextual interpretation will be consistent with the context.

In cases where an innovative use is detected, "normal" contextual interpretation breaks down and needs to be fixed, either by objecting to or questioning this usage[7], or by adapting one's

[6]Due to space limitations we do not formally define inconsistency here. For current purposes, it suffices to say that a record type is inconsistent (equivalent to $\bot$) if it contains a proposition and its negation.

[7]Background and foreground inconsistencies differ in how they can be rejected. While foreground inconsistencies can be rejected by saying simply "that's not true", background inconsistencies cannot be rejected in this way. A more

take on the meaning[8]. We refer to the latter as meaning-accommodation. Modification by *addition* or *deletion* of fields may be necessary to make the meaning non-innovative.

As an example of meaning accommodation that shows the relation between this kind of learning and the case of corrective feedback described above, we will again use the modified "gloves" example. Here, we wish to illustrate what happens when a previously known word is encountered with a different meaning. We therefore assume, for the sake of argument, that Naomi initially has a concept of gloves. We will assume that Naomi takes "gloves" as having a perceptual type distinct for that of "mittens". However, again for the sake of argument, we assume that she is mistaken as to the nature of this difference; for example, she may disregard the difference in shape and instead think that mittens and gloves have different textures (e.g. that gloves are shiny whereas mittens are woolly).

$\dot{\bigwedge}[\text{glove}]_{Naomi}$ =

$$\begin{bmatrix} \text{x} : \textit{Ind} \\ \text{c}_{\text{glove}} : \text{glove}'(\text{x}) \\ \text{c}_{\text{physobj}} : \text{physobj}(\text{x}) \\ \text{c}_{\text{clothing}} : \text{clothing}(\text{x}) \\ \text{c}_{\text{shiny-texture}} : \text{shiny-texture}(\text{x}) \\ \text{c}_{\text{handclothing-shape}} : \text{handclothing-shape}(\text{x}) \end{bmatrix}$$

That is, Naomi thinks that mittens and gloves both have a common shape, but that they differ in texture. This means that the meaning of Father's utterance will be

$[\text{Those are nice gloves}]^{Naomi}$ =

$$\begin{bmatrix} \text{BG} = & \begin{bmatrix} \text{focobj} : \textit{Ind} \\ \text{c}_{\text{glove}} : \text{glove}'(\text{focobj}) \\ \text{c}_{\text{physobj}} : \text{physobj}(\text{focobj}) \\ \text{c}_{\text{clothing}} : \text{clothing}(\text{focobj}) \\ \text{c}_{\text{shiny-texture}} : \text{shiny-texture}(\text{focobj}) \\ \text{c}_{\text{handclothing-shape}} : \text{handclothing-shape}(\text{focobj}) \end{bmatrix} \\ \text{FG} = & \begin{bmatrix} \text{c}_{\text{nice}} : \text{nice}'(\text{FG.focobj}) \end{bmatrix} \end{bmatrix}$$

When encountering Father's utterance, we take it that the relevant take on the context for evaluating and understanding the utterance is something like

c^{Naomi} =

$$\begin{bmatrix} \text{focobj=a} : \text{Ind} \\ \text{c}_{\text{physobj}} : \text{physobj}(\text{focobj}) \\ \text{c}_{\text{clothing}} : \text{clothing}(\text{focobj}) \\ \text{c}_{\text{woolly-texture}} : \text{woolly-texture}(\text{focobj}) \\ \text{c}_{\text{handclothing-shape}} : \text{handclothing-shape}(\text{focobj}) \\ \text{c}_{\text{not-shiny-texture}} : \text{not}(\text{shiny-texture}(\text{focobj})) \end{bmatrix}$$

The $\text{c}_{\text{not-shiny-texture}}$ field can either result from consulting the environment by checking whether a shiny texture cannot be detected on focobj, or by inference from the proposition in $\text{c}_{\text{woolly-texture}}$.

Now, according to our definition of innovation, Naomi will detect a background inconsistency in that [Those are nice gloves].BG $\dot{\bigwedge} c_{Naomi} = \bot$. The inconsistency of course stems from the presence of a proposition (shiny-texture(focobj)) and its negation in the combined record. Contextual interpretation will thus fail, since the meaning-function cannot be applied to the context[9].

Figuring out exactly how to modify "glove′" in the above example is non-trivial and potentially depends on a complex mix of different contextual factors and other clues. It is fairly clear to English speakers what Naomi *should* do: the perceptual type assigned to "glove′" needs to be altered by taking the difference in shape between gloves and mittens into account, and by disregarding the mistaken importance of difference in texture. Naomi would then arrive at the concept of "glove′" presented above in section 5, which would yield a meaning of Father's utterance which would be consistent with Naomi's take on the context (which would now, as a consequence, disregard texture but include that focobj is glove-shaped). In fact, both meaning and context would be as described in section 8.

Note that as a result of contextual interpretation, Naomi has not only interpreted Father's utterance but as an effect of fixing an context-meaning mismatch she has also arrived at a new potential lexical meaning of "glove′".

The different aspects of meaning of an expres-

appropriate reaction to a background inconsistency is to issue a clarification request regarding the faulty expression, in our example something like "glove?".

[8] Arguably, a further option should be included: to adapt one's take on the context.

[9] In this description, we are leaving out a complicating aspect of contextual interpretation, namely how background fields may be coerced to the foreground if they do not correspond to anything in the context (Cooper, 2005). If one tried to apply a function with background inconsistency to a context this would trigger coercion of the inconsistent material to the foreground (since it does not match anything in the context) which would result in a foreground inconsistency.

sion are not independent. For example, additions to ontological semantics may be echoed by additions to perceptual type. For example, if the interaction in our modified "gloves" example continued with Father (anticipating Naomi's problems with figuring out the difference between gloves and mittens) saying (as in the original example) "when they have fingers in them they are called gloves and when the fingers are all put together they are called mittens", Naomi could use this as a basis for constructing a "glove-detector" based on existing knowledge of the appearance of fingers. Nevertheless, learning to properly identify gloves may take some time and require many exposures to instances of gloves.

10 Conclusion and future work

The work presented here is part of a research agenda aiming towards a formal account of semantic coordination in dialogue. Here, we have focused on formalising a notion of (semantic) innovation which is central to meaning accommodation. We have provided an example involving the detection of a subjectively innovative use of an expression, and the adaptation of a lexical meaning to accommodate this innovation.

Apart from filling in the details missing from the present account (such as a full compositional analysis), future work on meaning accommodation includes providing general formal accounts of (1) how to determine when to do meaning accommodation in reaction to a detected innovative use (as opposed to rejecting it), and (2) deciding exactly how to modify a concept to accommodate an innovative use.

Acknowledgments

This research was supported by The Swedish Bank Tercentenary Foundation Project P2007/0717, Semantic Coordination in Dialogue.

References

Patrick Blackburn and Johan Bos. 2005. *Representation and Inference for Natural Language: A First Course in Computational Semantics*. CSLI Studies in Computational Linguistics. CSLI Publications, Stanford.

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

Samuel Carey and E. Bartlett. 1978. Acquiring a single new word. *Papers and Reports on Child Language Development*, 15:17–29.

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

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

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

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

Thierry Coquand, Randy Pollack, and Makoto Takeyama. 2004. A logical framework with dependently typed records. *Fundamenta Informaticae*, XX:1–22.

P.G.T. Healey. 1997. Expertise or expertese?: The emergence of task-oriented sub-languages. In M.G. Shafto and P. Langley, editors, *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, pages 301–306.

D. Kaplan. 1979. Dthat. In P. Cole, editor, *Syntax and Semantics v. 9, Pragmatics*, pages 221–243. Academic Press, New York.

Pia Knoeferle and Matthew W. Crocker. 2006. The coordinated interplay of scene, utterance, and world knowledge: Evidence from eye tracking. *Cognitive Science*, 30(3):481–529.

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

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

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

Matthew Saxton. 2000. Negative evidence and negative feedback: immediate effects on the grammaticality of child speech. *First Language*, 20(3):221–252.

David Traum and Staffan Larsson. 2003. The information state approach to dialogue management. In Ronnie Smith and Jan Kuppevelt, editors, *Current and New Directions in Discourse & Dialogue*. Kluwer Academic Publishers.