Compositional and ontological semantics in learning from corrective feedback and explicit definition

Robin Cooper
University of Gothenburg
Sweden
cooper@ling.gu.se

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

Abstract

We present some examples of dialogues from the literature on first language acquisition where children appear to be learning word meanings from corrective feedback and argue that in order to be able to account for them all in a formal theory of semantic change and coordination, we need to make a distinction between compositional and ontological semantics. We suggest how TTR (Type Theory with Records) can be used in making this distinction and relating the two kinds of semantics.

1 Introduction

This paper concerns the semantics and pragmatics of semantic coordination in dialogues between adults and children. The overall goal of this research is to attempt a formal account of language coordination in dialogue, and semantic coordination in particular.

In Larsson and Cooper (2009), we provide a dialogue move analysis of some examples from the literature on corrective feedback. We also provide a fairly detailed discussion of one example using TTR (Cooper, 2005; Cooper, 2008) to formalize concepts. In this paper we argue that in order to be able to account for these examples in a formal theory of semantic change and coordination, we need to make a distinction between compositional and ontological semantics. Both these aspects of meaning need to be represented in the linguistic resources available to an agent. We suggest how TTR can be used in making this distinction and relating the two kinds of semantics.

We take the following view on first language acquisition: children learn the meanings of expressions by observing and interacting with others. We regard language acquisition as a special case of a more general phenomenon of language coordination, that is, the process of coordinating on a language sufficiently to enable information sharing and coordinated action. One thing which is special about language acquistion is that there can be a clear assymmetry between the agents involved with respect to expertise in the language being acquired when a child and an adult interact. However, we want to propose 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.

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; Garrod and Anderson, 1987; 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.

We want a formal semantics allowing for meanings that can change dynamically during the course of a dialogue as a result of meaning updates triggered by dialogue moves. In particular, innovative uses of linguistic expressions may trigger updates to lexical entries. To account for this we need to account for how agents detect expressions which are innovative with respect to the agent's current linguistic resources, either because the expression is entirely new to the agent or because it is a known expression which is used with a new meaning.

We also need an account of how agents assign meanings to innovative expressions relative to the context of use. It is important here to distinguish local coordination on situated meanings, which is part of conversational grounding (Clark and Brennan, 1990; Traum, 1994) from coordination on meanings which affects agent resources such as lexical entries. It is the latter that we are interested in here.

Finally, we need to account for how the lexicalised meaning of a non-innovative expression can be updated based on its previously assumed meaning and the meaning of an innovative use which

contrasts with it. For example, if we learn that an object is not an A but rather a B (where B is innovative for us) then we need not only to learn B but also to refine the meaning of A so that it does not apply to the object.

In the rest of this paper we will first present a view of how agents adjust their linguistic resources on the basis of dialogue interaction (section 2). We will then discuss how compositional semantics can be derived from corrective feedback (section 3) and then give a brief background to TTR (section 4). Finally, we will show how ontological semantics can be added to compositional semantics derived from corrective feedback and explicit definition.

2 Agents that coordinate linguistic resources

As in the information state update approach in general (Larsson and Traum, 2000), dialogue moves are associated with information state updates. For semantic coordination, the kind of update is rather different from the one associated with dialogue moves for coordinating on task-related information, and involves updating the linguistic resources available to the agent (grammar, lexicon, semantic interpretation rules etc.), rather than e.g. the conversational scoreboard as such. Our view is that agents do not just have monolithic linguistic resources as is standardly assumed. Rather they have generic resources which they modify to construct local resources for sublanguages for use in specific situations. Thus an agent A may associate a linguistic expression e with a particular concept (or collection of concepts if e is ambiguous) $[e]^A$ in its generic resource. In a particular domain α e may be associated with a modified version of $[e]^A$, $[e]_\alpha^A$ (Larsson, 2007). In some cases $[e]_\alpha^A$ may contain a smaller number of concepts than $[e]^A$, representing a decrease in ambiguity.

Particular concepts in $[e]_\alpha^A$ may be a *refinement* of one in $[e]^A$, that is, the domain related concepts have an extension which is a proper subset of the extension of the corresponding generic concept. This will, however, not be the case in general. For example, a black hole in the physics domain is not normally regarded as an object described by the generic or standard meaning of *black hole* provided by our linguistic resources outside the physical domain. Similarly a variable in the domain of logic is a syntactic expression whereas a variable in experimental psychology is not and quite possibly the word *variable* is not even a noun in generic linguistic resources.

Our idea is that the motor for generating new local resources in an agent lies in coordinating resources with another agent in a particular communicative situation s. The event s might be a turn in a dialogue, as in the examples we are discussing in this paper, or, might, for example, be a reading event. In a communicative situation s, an agent A may be confronted with an *innovative* utterance e, that is, an utterance which either uses linguistic expressions not already present in A's resources or linguistic expressions known by A but associated with an interpretation distinct from that provided by A's resources. At this point, A has to accommodate an interpretation for e which is specific to s, $[e]_s^A$, and which may be anchored to the specific objects under discussion in s.

Whereas in a view of semantics inherited from formal logic there is a pairing between a linguistic expression e and an interpretation e' (or a set of several interpretations if e is ambiguous), we want to see e as related to several interpretations: $[e]_s^A$ for communicative situations s, $[e]_\alpha^A$ for domains α (where we imagine that the domains are collected into a complex hierarchy or more and less general domains) and ultimately a general linguistic resource which is domain independent, $[e]^A$. We think of the acquisition of a pairing of an expression e with an interpretation e' as a progression from an instance where e' is $[e]_s^A$ for some particular communicative situation s, through potentially a series of increasingly general domains α where e' is regarded as being one of the interpretations in $[e]_\alpha^A$ and finally arriving at a state where e' is associated with e as part of a domain independent generic resource, that is, e' is in $[e]^A$.

There is no guarantee that any expression-interpretation pair will survive even beyond the particular communicative situation in which A first encountered it. For example, the kind of *ad hoc* coinages described in Garrod and Anderson (1987) using words like *leg* to describe part of an oddly shaped maze in the maze game probably do not survive beyond the particular dialogue in which they occur. The factors involved in determining how a particular expression-interpretation pair progresses we see as inherently stochastic with parameters including the degree to which A regards their interlocutor as an expert, how many times the pairing has been observed in other communicative situations and with different interlocutors, the utility of the interpretation in different communicative situations, and positive or negative feedback obtained when using the pairing in a communicative situation. For example, an agent may only allow a pairing to progress when it has been observed in at least n different communicative situations at least m of which were with an interlocutor considered to be an expert, and so on. We do not yet have a precise proposal for a theory of these stochastic aspects but rather are seeking to lay the groundwork of a semantic treatment on which such a theory could be built.

3 Learning compositional semantics from corrective feedback

Recent research on first language acquisition (Clark, 2007; Clark and Wong, 2002; Saxton, 1997; Saxton, 2000) argues that the learning process crucially relies on negative input, including corrective feedback. This research is often presented in the context of the discussion of negative evidence, which we believe plays an important role in language. However, we want to relate corrective feedback to the discussion of alignment. We see corrective feedback as part of the process of negotiation of a language between two agents. Here are the examples of corrective feedback that we discuss in connection with our argument for this position in Larsson and Cooper (2009):

"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.

Panda example (constructed)

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

"Turn over" example (Clark and Wong, 2002):

- 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.

A frequent pattern in corrective feedback is the following:

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

The learning step can be further broken down as follows:

1. Syntactically 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*.

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.

Syntactic alignment in the gloves example can be visually represented thus:

Naomi: mittens
|
Father: **gloves**

For the "panda" example, the corresponding representation is

A: That's   a nice bear
       | | |
B: Yes, it's a   nice **panda**

Finally, in the the "turn over" example:

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

We assume that in the "gloves" example, syntactic properties can be predicted from syntactic alignment:

Naomi: [$_N$ mittens]
        ↓
Father: [$_N$ **gloves**]

In the "panda" example, the syntactic category of the innovative expression *panda* can be predicted from alignment (*panda* is aligned with non-innovative *bear* which is known to be a noun). This conclusion could be confirmed by an active chart edge spanning the substring *a nice* analyzed as an NP needing a noun. More confirming information can be extracted from the parse chart by noting that the assumption that *panda* is

a noun allows us to complete an NP-structure parallel to the analysis of *a nice bear* with which it is aligned.

A: That's $[_{NP} [_{Det}$ a$] [_A$ nice$] [_N$ bear$]]$

B: Yes, it's $[_{NP} [_{Det}$ a$] [_A$ nice$] [_N$ **panda**$]]$

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

In the "turn over" example, evidence comes from alignment and the resulting passive edge (together with alignment) as in the panda-example. In this case, however, given normal assumptions about how the parsing works, there will not be an active edge available to confirm the hypothesis as there was in the panda-example.

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

Mother: Okay I'll $[_{VP} [_V$ **turn**$]$ it over $]$

for you

A possible hypothesis is that alignment evidence is primary in predicting syntactic properties of innovations when it is available (as it is in corrective feedback). Other evidence can be used to support or refute the analysis deriving from alignment.

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

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

or, using TTR,

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

Thus, we see how 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. Before we proceed to ontological semantics we shall give a brief background on some aspects of TTR.

4 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.

We shall make use of type theory with records (TTR) as characterized in Cooper (2005; 2008) 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 meanings 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. Below is an example of a record type:

$$\left[\begin{array}{lll} \text{ref} & : & Ind \\ \text{size} & : & \text{size(ref, MuchBiggerThanMe)} \\ \text{shape} & : & \text{shape(ref, BearShape)} \end{array}\right]$$

A record of this type has to have fields with the same labels as those in the type. (It may also include additional fields not required by the type.) In place of the types which occur to the right of ':' in the record type, the record must contain an object of that type. Here is an example of a record of the above type:

$$\left[\begin{array}{lll} \text{ref} & = & \text{obj123} \\ \text{size} & = & \text{sizesensorreading85} \\ \text{shape} & = & \text{shapesensorreading62} \\ \text{colour} & = & \text{coloursensorreading78} \end{array}\right]$$

Thus, for example, what occurs to the right of the '=' in the ref field of the record is an object of type *Ind*, that is, an individual. Types which are constructed with predicates like *size* and *shape* are sometimes referred to as "types of proof". The idea is that something of this type would be a proof that a given individual (the first argument) has a certain size or shape (the second argument). One can have different ideas of what kind of objects count as proofs. Here we are assuming that the proof-objects are readings from sensors. This is a second way (in addition to the progression of local resources towards general resources) that our theory interfaces with an analogue non-categorical world. We imagine that the mapping from sensor readings to types involves sampling of analogue data in a way that is not unsimilar to the digitization pro-

cess involved, for example, in speech recognition. Again we have nothing detailed to say about this at the moment, although we regard it as an important part of our theory that it is able to make a connection between the realm of feature vectors and the realm of model-theoretic semantics.

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. This means, for example, that in considering whether a record is of the record type, you will need to find a proof that the object which is in the ref-field of the record has the size represented by *MuchBiggerThanMe*. That is, this type depends on the value for the ref-field.

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

$$\begin{bmatrix} \text{ref=obj123} & : & \text{Ind} \\ \text{size} & : & \text{size(ref, MuchBiggerThanMe)} \\ \text{shape} & : & \text{shape(ref, BearShape)} \end{bmatrix}$$

$\begin{bmatrix} \text{ref=obj123:Ind} \end{bmatrix}$ is a convenient notation for $\begin{bmatrix} \text{ref} : \text{Ind}_{\text{obj123}} \end{bmatrix}$ 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*. For example,

$$\begin{bmatrix} \text{ref} & : & \text{Ind} \\ \text{size} & : & \text{size(ref, MuchBiggerThanMe)} \end{bmatrix}$$

is a subtype of

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

as is also

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

5 Learning ontological semantics from corrective feedback and explicit definition

As a (modest) "proof of concept" of our approach, we will in this section provide a TTR analysis of updates to compositional and ontological semantics for the "mittens" example above. As pointed out by one of the reviewers, our approach to coordination of ontological semantics bears resemblances to work on ontology mapping and ontology negotiation on the semantic web (van Diggelen et al., 2007).

Using TTR, we can formalise ontological classes as record types:

$$\text{Thing} = \begin{bmatrix} \text{x} : \textit{Ind} \end{bmatrix}$$

$$\{\text{Class } P\} = \begin{bmatrix} \text{x} : \textit{Ind} \\ \text{c}_P : P(\text{x}) \end{bmatrix}$$

We will use a function SubClass which creates a class based on a predicate P:

$\{\text{SubClass } C_1 \ C_2\} = C_1 \wedge C_2$ ("Make a subclass of C_2 based on C_1")

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

Given this formal apparatus, we can show how ontological semantics properties can be predicted in the glove example. Naomi's pre-gloves ontology contains (we assume) the following:

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

This ontology is shown in Figure 1, where the arrow represents the subclass relation. Provided that Naomi learns from the interaction, Naomi's post-gloves ontology may include the following (see also Figure 2):

PhysObjClass = {Class physobj}
ClothingClass = {SubClass {Class clothing} PhysObjClass}
MittenClass = {SubClass {Class mitten} ClothingClass}
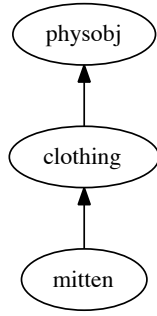
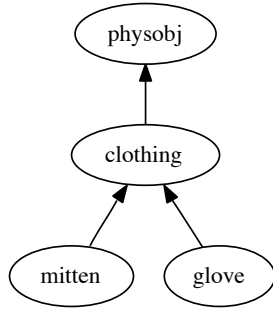Figure 1: Naomi's "pre-gloves" ontology

Figure 2: Naomi's "post-gloves" ontology

GloveClass = {SubClass {Class glove} ClothingClass} (from alignment of *mittens* and *gloves*)

This means that the glove class is the following type

$$
\begin{bmatrix}
\text{x} & : & Ind \\
c_{\text{physobj}} & : & \text{physobj(x)} \\
c_{\text{clothing}} & : & \text{clothing(x)} \\
c_{\text{glove}} & : & \text{glove(x)}
\end{bmatrix}
$$

which can be used as a refinement of the type corresponding to the compositional semantics:

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

Thus we can obtain the new function below as a refined compositional semantics:

$$
\lambda r{:}\begin{bmatrix} \text{x} : Ind \end{bmatrix} \ (\begin{bmatrix}
c_{\text{physobj}} & : & \text{physobj}(r.\text{x}) \\
c_{\text{clothing}} & : & \text{clothing}(r.\text{x}) \\
c_{\text{glove}} & : & \text{glove}(r.\text{x})
\end{bmatrix})
$$

In the "glove" example, the father's second ut-

Figure 3: Naomi's ontology after explicit definition

terance 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 (see also Figure 3):

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

In TTR, after this update the meanings for "glove" and "mitten" will be respectively:

$$
\begin{bmatrix}
\text{x} : Ind \\
c_{\text{physobj}} : \text{physobj(x)} \\
c_{\text{clothing}} : \text{clothing(x)} \\
c_{\text{handclothing}} : \text{handclothing(x)} \\
c_{\text{withoutfingers}} : \text{withoutfingers(x)} \\
c_{\text{glove}} : \text{glove(cntxt.x)}
\end{bmatrix}
$$

and

$$\begin{bmatrix} \text{x} : Ind \\ c_{physobj} : \text{physobj(x)} \\ c_{clothing} : \text{clothing(x)} \\ c_{handclothing} : \text{handclothing(x)} \\ c_{withfingers} : \text{withfingers(x)} \\ c_{mitten} : \text{mitten(cntxt.x)} \end{bmatrix}$$

6 Conclusion and future work

By providing a basic compositional semantic resource and providing the ability to refine this with local ontologies, which may be associated with given domains or even specific dialogues, we allow for an extremely flexible view of word meaning that provides mechanisms for associating a central core of meaning with situation specific meanings that can be generated on the fly.

Future work includes exploring the relation to work on ontology negotiation on the semantic web, as well as extending our account to cover further aspects of meaning, including perceptually grounded meaning and connotations. We also wish to relate detailed accounts of semantic updates to other kinds of dialogue strategies, such as ostensive definitions and meaning accommodation (Larsson, 2008).

Acknowledgments

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.

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

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

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

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

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

Robin Cooper. 2008. Type theory with records and unification-based grammar. In Fritz Hamm and Stephan Kepser, editors, *Logics for Linguistic Structures*, pages 9 – 34. Mouton de Gruyter.

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

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

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.

Staffan Larsson and Robin Cooper. 2009. Towards a formal view of corrective feedback. In Afra Alishahi, Thierry Poibeau, and Aline Villavicencio, editors, *Proceedings of the Workshop on Cognitive Aspects of Computational Language Acquisition, EACL*, pages 1–9.

Staffan Larsson and David Traum. 2000. Information state and dialogue management in the trindi dialogue move engine toolkit. *NLE Special Issue on Best Practice in Spoken Language Dialogue Systems Engineering*, pages 323–340.

Staffan Larsson. 2007. Coordinating on ad-hoc semantic systems in dialogue. In R. Artstein and L. Vieu, editors, *Proceedings of the 10th workshop on the semantics and pragmatics of dialogue*, pages 109–116.

Staffan Larsson. 2008. Formalizing the dynamics of semantic systems in dialogue. In Robin Cooper and Ruth Kempson, editors, *Language in Flux - Dialogue Coordination, Language Variation, Change and Evolution*, pages 121–142. College Publications, London.

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. 1997. The contrast theory of negative input. *Journal of Child Language*, 24:139–161.

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

David R. Traum. 1994. *A Computational Theory of Grounding in Natural Language Conversation*. Ph.D. thesis, Department of Computer Science, University of Rochester. Also available as TR 545, Department of Computer Science, University of Rochester.

Jurriaan van Diggelen, Robbert-Jan Beun, Frank Dignum, Rogier M. van Eijk, and John-Jules Meyer. 2007. Ontology negotiation in heterogeneous multi-agent systems: The anemone system. *Applied Ontology*, 2:267–303.