Coordinating on ad-hoc semantic systems in dialogue

Staffan Larsson

Dept. of linguistics

Göteborg University, Sweden

`sl@ling.gu.se`

Abstract

An exploratory study of a Map Task dialogue indicates that dialogue participants coordinate on an ad-hoc vocabulary and associated concepts (meanings) to enable information exchange, and that ad-hoc vocabularies can be cobbled together from a heterogeneous mix of "micro-vocabularies" borrowed from various other (a priori unrelated) domains. To account for these observations, we sketch a basic framework for formalising the process of coordination of semantic systems in dialogue, and relate this framework to some interactional processes of semantic coordination in dialogue, such as feedback, negotiation and accommodation.

1 Vocabulary in a Map Task dialogue

In the Map Task corpus[1], a GIVER explains a route, provided on the giver's map, to a FOLLOWER who has a similar (but slightly different) map but with no route marked. A map contains landmarks portrayed as labelled line drawings. In a route-giving task like that recorded in the Map Task corpus, expressions referring to landmarks, compass directions etc. can be a priori expected as a kind of "prototype" devices for talking about maps. A typical utterance may look as follows[2]:

GIVER: right **a camera shop**, right, head due **south** ... from that just ... **down** for about **twelve centimetres**, have you got **a parked van** at the bottom ?

Here, we may note two constructions expressing direction ("south", "down"), one expressing a distance ("twelve centimetres") and two referring to landmarks ("a camera shop", "a parked van"). A further example:

GIVER: go round the left hand side of the camera shop ... in between **the edge of the page** and the camera shop.

Whereas the previous expressions were completely expected given the general direction-giving task, the reference to an absolute position using "the edge of the page" is perhaps less expected. Clearly, this is a consequence of the dialogue participants (DPs) talking about a (paper) map rather than e.g. about some actual terrain.

GIVER: so you're ... you're going diagonally sort of north ... northeast ... it's not it's it's a sort of **two o'clock** almost **three o'clock** ... from the allotments ... over

Here, we have GIVER referring to map directions using the expressions "two o'clock" and "three o'clock". This is most likely an everyday variant of the practice of English-speaking pilots of using "o'clock" for directions[3]. Let's look at a final excerpt:

GIVER: right, you go ... down the side of the camera shop right for about twelve centimetres ... and do a sort of **a "u" shape** ... for and **the bottom of the "u" shape** should be about three centimetres long, right do you know what i'm meaning

---

[1] http://www.hcrc.ed.ac.uk/maptask/maptask-description.html

[2] The following excerpts are taken from Map Task dialogue q4nc4, available at the Map Task web site.

[3] Note the use of a hedging "sort of" before "two o'clock", which seems to indicate that the speaker is slightly unsure as to whether the following expression is quite appropriate. A similar observation is made by Brennan (To appear) (p. 11): "[h]edges seem to be one way of marking that a referring expression is provisional."

...

GIVER: you've worked it out already , eh we're doing **a "u" shape** round the parked van but it's a sort of three cent– see if you imagine a 'u' right ... **the stems of the "u" the ... vertical bits** are sort of three centimetres between

First, a trajectory is referred to using the expression "a 'u' shape". This trajectory is (or so we argue) then reified as an imagined 'u'-shape on the map, now acting more akin to a landmark with a concrete (if invisible) shape, size and even component parts ("the ... vertical bits"; "the stems of the 'u' ").

2 Micro-vocabularies used in Map Task dialogue

Based on the above excerpts (and others from the same dialogue), we are now able to provide a very tentative inventory of referring expressions used by GIVER and FOLLOWER in the Map Task dialogue. DPs refer to distances, absolute and relative locations, directions, and trajectories. Below, we list the sub-types of expressions used for each basic class.

- distances on page, in centimetres ("about twelve centimetres")
- absolute locations
  - landmarks ("the camera shop")
  - page edges ("the edge of the page"; "at the bottom"; "the far right-hand side")
  - typography on page ("the words 'yacht club'")
  - (imagined) letter shapes ("the bottom of the 'u' shape"; "the stems of the 'u' the ... vertical bits")
- relative locations
  - relative to landmark ("left hand side of (landmark)")
  - relative to sheet of paper ("the other side of the page")
- directions
  - compass directions ("head due south")
  - left, right, up, down, diagonally, etc.
  - clock directions ("sort of two o'clock")
- trajectories
  - imagined/drawn lines ("a straight line up the ...")
  - letter shapes as trajectories ("do sort of a 'u' shape")

3 Interleaving resource registers

How can we account for this diversity in the range of linguistic expressions used in a simple direction-giving dialogue? In this section, we will propose a basic terminology intended to form a basis for a formal account of what we see happening in dialogues such as the one quoted above.

3.1 Perspectives

In the Map Task dialogue, the DPs need to coordinate on a way of talking about the map. What the above excerpts show is that there are several ways of talking about a map; this is also shown in the Maze Game experiments (Garrod and Anderson, 1987; Healey, 1997) where DPs alternative between an abstract "coordinate system" perspective on a maze ("Go to the fourth row down and the second from the right"; "Six three"), and more concrete perspectives involving e.g. corridors ("Go forward, then turn left at the junction") or shapes ("the bit sticking out on the right"). In our view, a way of talking about X involves *taking a perspective*[4] on X and selecting a vocabulary associated with that perspective. Taking a perspective P on subject matter X in dialogue involves an analogue - "talking about X as P" - e.g. talking about directions on a map as clock arms. Different perspectives have different advantages and disadvantages; for example, an abstract perspective is compact but error-prone; a clock perspective on directions may e.g. enable shorter utterances. One plausible reason for interleaving and switching several perspectives and associated vocabularies thus seems to be that it increases the efficiency of communication.

3.2 Resource and ad-hoc registers

On a fundamental level, we believe that a language can be regarded as consisting of a multitude of activity-specific "language games" involving activity-specific *registers*. A register is a an *activity-specific semantic system* (a "microlanguage"), consisting minimally of a set of linguistic signs, i.e., linguistic expressions and associated concepts (meanings)[5]. In dialogue, registers may be used as *resources* which can be borrowed or appropriated into a new activity and

---

[4]Garrod and Anderson (1987) and Healey (1997) instead talk about adopting "description types".

[5]A *compositional* register will more generally contain *mappings* between expressions and meanings.

adapted to the domain at hand. Putting it differently, an *ad-hoc register* is assembled to be able to talk about some subject matter from one or more perspectives. In the map-task dialogue, several different resource registers are introduced and accepted[6]; often, both introduction and acceptance are implicit, but sometimes verbal signals (including feedback) are used to manage semantic coordination. For example, one could imagine "sort of" being used to signal introduction of new register.

As mentioned, in the Map Task dialogue we find some resource registers that can be regarded as "standard" or "default" ways of talking about maps, whereas others are more unexpected. First, the standard map registers subsumes (1) a *landmarks* register provided to DPs as pictures and text on map, (2) a *compass directions* register, and (3) a *(metric) distance* register. The non-standard parts of the ad-hoc register are:

- *clock* register: map directions as clock hands "two o'clock" etc.

- *sheet-of-paper* register perspective: map as a sheet of paper edges of page distances on page relations between pages (e.g. "opposing page")

- *letter shape* register perspective: Viewing map as a piece of paper where letter shapes can be drawn letter shapes ("a 'u' shape") parts of letter shapes ("stems")

3.3 Appropriating and interleaving registers

To describe the dynamics of registers in the above dialogue, we can say that the clock, sheet-of-paper and letter-shape registers are *appropriated* into the map task activity, where it is *interleaved* with landmark, compass direction, and metric distance registers to form an ad-hoc register[7]. This involves adapting the meanings associated with resource register vocabularies to the current situation.

4 Meaning potentials

To describe how linguistic expressions can be interactively (in dialogue) appropriated into a new activity, we need an account of semantics which (1) allows several activity-specific meanings for a single expression, and (2) allows open and dynamic meanings which can be modified as a consequence of language use. 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 believe that the idea of "meaning potentials" may offer a more dynamic view of meaning. The term originates from "dialogical" approaches to meaning (Recanati, 2003). On the "dialogical" view, language is essentially dynamic; meaning is negotiated, extended, modified both in concrete situations and historically. Interaction and context are essential for describing language, and there is a general focus on the context-dependent nature of meaning. Linguistic expressions have meaning potentials, which are not a fixed and static set of semantic features, but a dynamic potential which can give rise to different situated interpretations. Different contexts exploit different parts of the meaning potential of a word.

We refer to the dynamic aspect meaning potentials as *semantic plasticity*. Semantic plasticity will be central to our account of how activity-specific abstract[8] meanings are updated and gradually change as a consequence of use.

5 Towards a formalisation of semantic plasticity and meaning potentials

To describe in more detail how DPs coordinate on registers (e.g. when adapting a resource register to a new domain), we need a dynamic account of meanings and registers allowing incremental modifications (updates) to semantic systems. We also need a description of possible dialogue strategies for register coordination. Describing this process *formally* requires formalising the dynamics of registers and meaning potentials, and the dia-

---

[6]Often, several resource registers are used in a single phrase, as e.g. in "in between the edge of the page and the camera shop".

[7]This "interleaving strategy" can be compared with the "switching strategies" evident in maze game experiments (Healey, Garrod), where speakers switch between perspectives (description types). Presumably, both interleaving and switching are possible.

[8]We use "abstract meaning" to refer to utterance-type meanings, either activity-specific or activity-independent.

logue protocols involved in negotiating semantic systems. In this section, we will take some initial steps towards this goal by sketching a formal account of semantic plasticity.

We propose to regard the meaning of a linguistic construction or word[9] to depend on previous uses of that word. This makes it possible to model how meanings change as a result of using language in dialogue. The basic idea is that speakers have internalised (potentially complex) dispositions, or *usage patterns*, governing the use of specific words. These dispositions depend, among other things, on observations of previous situations where the word in question has been used, and on specific generalisations over these situations.

Semantic plasticity is described in terms of updates to individual usage patterns associated with words (in general, linguistic constructions) triggered by observations of their use in dialogue. When a usage pattern $[c]$ is sufficiently coordinated[10] (shared) within a community, we will talk about $[c]$ as the meaning potential of a word c. By modelling plasticity of usage patterns of individuals, we thus indirectly model semantic plasticity in a linguistic community.

5.1 Usage sets and usage patterns

To get a handle on semantic plasticity, we will start by positing for each language user A and word c a *usage-set*[11] S_c^A containing all situations where A has observed a use (token) of c. Formally, $S_c^A = \{s \mid A$ has observed a use of c in situation $s\}$. This should be regarded merely as an abstract theoretical entity. .

We assume that A generalises over S_c^A; this generalisation we call the usage pattern (or usage disposition) $[c]^A$. In cognitive terms one can think of the usage pattern as the "memory trace" of observed uses of c.

That c has been used in a situation simply means

that someone has uttered a token of c in that situation[12].

5.2 Situated meanings and interpretations

On each occasion of use of c in situation s, c has a specific situated utterance-token meaning which derives partly from the shared abstract utterance-type meaning (meaning potential) $[c]$ and partly from s. We write this meaning formally as $[c]_s$. The subjective counterpart of a situated meaning is a *situated interpretation*, written as $[c]_s^A$ for an agent A; this is the interpretation that A makes of c in s based on A's usage pattern $[c]^A$. A situated meaning $[c]_s$ arises in a situation when the DPs in s make sufficiently similar situated interpretations of c in s.

5.3 Appropriate and non-appropriate uses

We will assume that new uses of a word c can be classified as appropriate or inappropriate given an existing usage pattern[13] for c[14]. The formal notation we will use to express that a use of c in situation s is appropriate with regard to A's usage pattern for c is $[c]^A \vdash s$. Correspondingly, $[c]^A \nvdash s$ means that s is not an appropriate situation in which to use c given $[c]^{A}$[15].

On the whole, if a token of c uttered in a situa-

---

[12]It is important to point out that the notion of "situation" we are using here is an abstract one; the reason is that we want to keep the framework general. In more concrete instantiations of this abstract framework, the notion of a situation will be specified based on the activity in which an agent acts and the requirements on the agent in this activity, as well as the representations and sensory-motor machinery of the agent. As a simple example, in the work of Steels and Belpaeme (2005) the situation is limited to a colour sample, perceived by a robot through a camera and processed into a representation of colours in the form of three real-valued numbers.

[13]It may be thought that appropriateness should be defined in terms of collective meaning potentials rather than individual usage pattens, to make sense of talk of "incorrect use of words." However, we believe that such talk is better regarded as one of many strategies for explicit negotiation of meanings, which always occurs in concrete situations and between individual DPs with their respective usage patterns. A theoretical notion of correct or incorrect use of words (independent of individual usage patterns) runs into several problems, such as defining how many DPs must share a usage pattern in order for it to be deemed "correct." This does not mean we cannot make sense of talk of incorrect and correct use of words; it only means that regard such notions primarily as devices in negotiations of shared meanings.

[14]In general, appropriateness is not necessarily a Boolean property, but rather a matter of degree. This is a simplification in the current theory.

[15]The exact method of deciding whether a new token is appropriate or not will depend on the specific kinds of representations, learning algorithms, and measures of similarity that are assumed (or, in an artificial agent, implemented).

---

[9]Although we intend this account to cover not only words but also other constructs phrases, syntactic categories, and other linguistic elements, we will henceforth (for simplicity) use "word" instead of "linguistic construction".

[10]Roughly, a usage pattern connected to an expression is sufficiently coordinated in a community when speakers and hearers are able to use that expression to exchange information sufficiently to enable them to achieve their shared and private goals. For example, in the Map Task dialogues an expression is sufficiently coordinated when DPs are able to make use of it in carrying out the route-giving tasks assigned to them.

[11]An alternative term is *situation-collocation*.

tion s is consistent with $[c]^A$, A is likely to understand c and to judge s to be an appropriate situation of use of c. However, it is important to leave open the possibility that a DP may not understand, or understand but reject, a token of c even if this token of c in the current situation is appropriate with respect to A's usage pattern for c. Similarly, a DP may choose to use a word in a situation where she judges it inappropriate given previous uses; we call this a *creative use* (in contrast to conservative uses which are appropriate given previous uses).

5.4 Usage-pattern updates

It follows from the definition of $[c]^A$ that whenever A observes or performs a use of c, S_c^A will be extended, and so the usage pattern $[c]^A$ may change. This is a *usage pattern update*. *Prima facie*, there are many different possible kinds of ways that a usage pattern may be modified, depending on assumptions regarding semantic representation.

Usage-pattern updates can be distinguished according to several dimensions; we will start by making a rough distinction between *reinforcements* and *revisions*.

If a use of c in situation s is consistent with A's usage pattern for c, i.e., c is appropriate in s ($[c]^A \vdash s$), there is no drastic change; the previous disposition is reinforced by extending $[c]^A$ with A's situated interpretation of c in s, $[c]_s^A$. We will write this formally as $[c]^A \circ_= [c]_s^A$). However, if the current use of c is not consistent with usage disposition ($[c]^A \nvdash s$), there will be a relatively drastic revision of the disposition (formally, $[c]^A \circ_* [c]_s^A$).

5.5 Situation-types and structured meaning potentials

To account for how registers can be appropriated (borrowed) from one activity (e.g. telling the time) to another (e.g. direction-giving) we need a formalisation which allows new meanings of existing words to be created as a result of observed novel (at least subjectively) language use. Meaning potentials, which in addition to being dynamic can also be *structured*, and thus allow for different contexts to exploit different meaning potential *components*, seem useful.

We will use *situation-type* as a general term for contexts, activities, institutions etc. where words take on specific meanings. A register, or "microlanguage", is the lexicon used in a situation-type,

pairing the words used (vocabulary) with meanings (what can be talked about; ontologies; coordinated usage patterns) in the situation-type[16]

In general, a situation-type may be associated with several registers (corresponding to different perspectives on the situation-type), each providing a mapping from a vocabulary to (abstract) meanings specific to the situation-type. Conversely, the meaning potential for a word is often structured into several situation-type-specific components.

We have established that $[c]^A$ is agent A's usage pattern for word c, and that $[c]_s^A$ is the interpretation that agent A makes of c in s; this interpretation is a function of s and $[c]^A$. We will now extend our notation with $[c]_\alpha^A$ - an agent A's situation-type-specific usage pattern for c in situation-type α. In general, any aspect of the utterance situation-type may activate usage pattern components. A structured meaning potential exists in a linguistic community with coordinated structured usage patterns. A component of structured meaning potential for c in situation-type α is written as $[c]_\alpha$[17].

As a simple example inspired by the Map Task dialogue above, the meaning potential ["two o'clock"] can be described as structured into

- ["two o'clock"]$_{clock}$, where *clock* stands for an activity type involving telling the time; this meaning potential component can be paraphrased "02:00 AM or PM"

- ["two o'clock"]$_{direction-giving}$, where α has been assigned a situation type index corresponding to direction-giving activities; this meaning potential component is paraphraseable as "east-northeast direction"

5.6 Interpretation and update involving structured usage patterns

A token c_s of a word c in situation s is interpreted by B as $[c]_s^B$. If $[c]^B$ is a complex usage pattern, some component of $[c]^B$ must be selected as the abstract meaning to be used for contextual interpretation. Now, assume that situation s is classified by B as being of situation-type α. This triggers a component of $[c]^B$ - the *activated usage pattern component* $[c]_\alpha^B$.

---

[16]This terminology builds on (and modifies slightly) that of Halliday (1978).

[17]An obvious extension to this formalism, which we will not develop further here, would be to index meaning potentials (and their components) by the linguistic community in which they exist.

In this case, $[c]_\alpha^B$ is a likely candidate for which part of $[c]$ gets updated. (If B is not able to find a relevant usage pattern component, B may create a new ad-hoc component, which can be updated during the dialogue. This pattern may or may not be retained afterwards; it may be assimilated into some existing component of $[c]$, or the start of a new usage pattern component.)

Let's take an example. Assume ["two o'clock"] is structured into ["two o'clock"]$_{clock}$ and ["two o'clock"]$_{direction-giving}$, as above. Now assume we get the following utterance:

GIVER: "sort of two o'clock"

Because the activity is direction-giving, FOLLOWER activates ["two o'clock"]$_{direction-giving}^{follower}$. FOLLOWER then instantiates ["two o'clock"]$_{direction-giving}^{follower}$ to arrive at a contextual interpretation ["two o'clock"]$_s^{follower}$ (roughly, a 60 degree angle on FOLLOWER's map). Insofar as ["two o'clock"]$_{direction-giving}^{follower} \vdash s$, we get a reinforcing update ["two o'clock"]$_{direction-giving}^{follower}$ $\circ_= $ ["two o'clock"]$_s^{follower}$.

6 Semantic coordination

This section sketches a framework for modelling *negotiation of meaning in dialogue*, i.e. the social processes (dialogue games) involved in the explicit and implicit negotiation of meaning in dialogue, and their relation to the cognitive processes (semantic updates).

After discussing the basic devices available to speakers for conducting semantic negotiation, we will give examples of how the theory sketched above can be used to analyse short dialogue excerpts in terms of semantic updates. As yet, the theory does not include a taxonomy of dialogue moves involved in semantic negotiation, and therefore the analysis does not include dialogue moves; instead, utterances are analysed directly in terms of their associated semantic updates. Coming up with a general taxonomy of such moves and their associated updates is a major future research goal.

6.1 Basic devices for coordination in dialogue

We assume (provisionally) three basic devices available to dialogue participants for negotiating (and, typically, achieving coordination of) linguistic resources: feedback, explicit negotiation, and accommodation. "Negotiation" is used here in a weak sense of "interactive achievement of coordination".

Feedback (Allwood, 1995; Clark, 1996) involves signals indicating perception, understanding, and acceptance of utterances in dialogue, as well as failure to perceive or understand; clarification requests; and rejections. It is well known that feedback governs that coordination of the dialogue gameboard ("informational coordination"); however, it also guides coordination of language use ("language coordination").

For example, *corrective* feedback is common in adult-child interaction. Below is an example; A is the child, B the adult, and as part of the common ground there is a topical object in the situation s visible to both A and B. We also assume that A is not familiar with the word "panda".[18]

A: Nice bear

B: Yes, it's a nice panda

Here, B rejects this use of "bear" by providing negative feedback in the form of a correction (and in addition, B gives positive feedback accepting the assertion that the focused object (animal) "is nice"). For an account of this example in terms of semantic plasticity and coordination, see Larsson (2007).

Explicit negotiation is the overt meta-linguistic negotiation of the proper usage of words, including e.g. cases where explicit verbal or ostensive definitions are proposed (and possibly discussed). Although semantic negotiation typically has the goal of coordinating language use, it may in general be both antagonistic and cooperative. In Steels and Belpaeme (2005), robot agents play a language game of referring to and pointing to colour samples. The colour-language system of an individual agent is modelled as a set of categories in the form of neural nets that respond to sensory data from colour samples, and a lexicon connecting words to categories. This is clearly a case of explicit semantic plasticity and semantic negotiation, as categories are updated as a result of language use. Semantic negotiation here takes the form of explicit and cooperative negotiation. For an account of a dialogue taken from the this exper-

---

[18]This example from Herb Clark, p.c.; similar examples can be found in Clark (2003)

iment in terms of semantic plasticity, see Larsson (2007).

By *accommodation* we refer to adaptations to the behaviour of other DPs. For example, one may adapt to the presuppositions of an utterance of "The King of France is bald" by modifying the dialogue gameboard to include the existence of a king of France. We want to extend the notion of accommodation beyond the dialogue gameboard, to include changes in the language system.

For each word used in an utterance u, the addressee (here, B) in a dialogue is (usually) expected to react if he thinks a word in u was inappropriately used. If B is able to construct a situated interpretation $[c]_s^B$ (which may involve more or less effort) but finds this use inappropriate ($[c]^B \nvdash s$), this may be due to a mismatch between s (as perceived by B) and $[c]^B$. B may now reject this use of c explicitly using negative feedback, or quietly alter $[c]^B$ ($[c]_B \circ_* [c]_s^B$) so that this use of c can be counted as appropriate after all.

6.2 Coordination through accommodation

We will now give an example of semantic coordination in dialogue, where meaning accommodation leads to updates to complex usage patterns.

Assume we get the following utterance in a Map Task dialogue in a situation s:

GIVER: "sort of two o'clock"

Assume[19] also that FOLLOWER is not familiar with the "direction-giving" use of "two o'clock". More precisely, ["two o'clock"]^{fol} only contains ["two o'clock"]$_{clock}$, so ["two o'clock"]$^{fol} \nvdash s$.

By analogical reasoning using contextual features, FOLLOWER is nevertheless able to correctly understand A's utterance and arrives at a contextual interpretation ["two o'clock"]$_s^{fol}$. Now, since ["two o'clock"]$^{fol} \nvdash s$, FOLLOWER needs to revise ["two o'clock"]^{fol} by creating a new activity-specific component ["two o'clock"]$_{d-g}^{fol}$. We get an overall update ["two o'clock"]^{fol} $\circ_=$ ["two o'clock"]$_s^{fol}$ which can be decomposed as two updates, (1) creation of ["two o'clock"]$_{d-gg}^{fol}$, followed by ["two o'clock"]$_{d-g}^{fol}$ $\circ_=$ ["two o'clock"]$_s^{fol}$. After this update, ["two o'clock"]$^{fol} \vdash s$, i.e., the novel (for FOLLOWER)

use of "two o'clock" by GIVER has been accommodated.

7 Kinds of coordination in dialogue

On our view, two kinds of coordination happen in everyday human-human dialogue. *Informational coordination* has successfully been studied using the concepts of dialogue games and updates to a shared dialogue gameboard. One of the goals of the research presented here is to extend this approach to describing *language coordination* (and more specifically, semantic coordination) in terms of the dynamics of updates to language systems.

The framework sketched here aims at describing all kinds of semantic coordination[20]. In the "two o'clock" example given above, coordination is essentially a matter of mapping an expression ("two o'clock") to a pre-existing meaning (denoted in the compass directions register as "east-northeast"). For this kind of coordination, some version of traditional formal semantics may suffice, provided it is extended with a dynamic mapping between linguistic expressions and their meanings[21].

However, in other cases the dynamics go beyond word-meaning mappings. Specifically, to account for cases where an expression is used to denote a *new* concept, such as "the u-shape" above, we need to describe the dynamics of *concept creation*. Similarly, existing concepts may be affected by their use in dialogue, e.g., by subtly modifying values of usage-governing conceptual features by small increments. For example, in Steels and Belpaeme (2005), concepts are represented as neural nets which are updated by small adjustments to network weights, according to a standard back-propagation algorithm.

These dynamics, which we refer to as *concept-level* dynamics, are an important motivation for the introduction of meaning potentials. They are also our main reason for believing that traditional formal semantics will not suffice to account for semantic plasticity coordination.

To deal with concept-level dynamics in a general way, one will probably need to keep track of of semantic features connected to expressions in

---

[19]In this example, we will use the following abbreviations: fol = follower, d-g = direction-giving.

[20]A typology of variants of semantic coordination is a future research goal.

[21]Note that "dynamic semantics" (Groenendijk and Stokhof, 1988) is not dynamic in this sense, as it follows traditional formal semantics in assuming a static mapping between words and meanings.

the lexicon (Pustejovsky, 1991) and allow these feature matrices to be updated as a result of semantic negotiation and coordination subdialogues. Work in this direction may benefit from ideas put forward by Gärdenfors (2000), as well as in work on machine learning (Mitchell, 1997) and Latent Semantic Analysis (Landauer and Dumais, 1997). One version of formal semantics which seems promising for the illumination of concept-level dynamics is the record-type theoretic approach which Cooper has been developing (Cooper, 2005a; Cooper, 2005b). This formal approach allows for both underspecification or uncertainty of meaning by the use of types of meaning and also a structured approach to meaning analysis which allows for modification of meaning in a way which is not possible, for example, in the classical formal semantics analysis of meaning as functions from contexts to intensions.

8 Conclusion

To account for the observed dynamics of semantic systems in dialogue, we have sketched a formalisation of the notion of meaning potential, in the form of dynamic structured usage patterns which are shared within a linguistic community through a process of semantic coordination in dialogue. This process can be described as updates to structured usage patterns resulting from language use. We have also outlined some basic mechanisms of coordination: feedback, explicit negotiation, and accommodation.

This paper presents preliminary work aiming towards a unified theoretical account of semantic coordination. Apart from developing the theory and the formal framework further, we want to extend the coverage of this theory by further empirical studies, and to start implementing strategies for semantic coordination in practical dialogue systems.

References

Jens Allwood. 1995. An activity based approach to pragmatics. Technical Report (GPTL) 75, Gothenburg Papers in Theoretical Linguistics, University of Göteborg.

S. E. Brennan. To appear. The vocabulary problem in spoken language systems. In S. Luperfoy, editor, *Automated spoken dialog systems*. Cambridge, MA: MIT Press.

H. H. Clark. 1996. *Using Language*. Cambridge University Press, Cambridge.

E. V. Clark. 2003. *First language acquisition*. Cambridge: Cambridge University Press.

Robin Cooper. 2005a. Austinian truth, attitudes and type theory. *Research on Language and Computation*, 3(4):333–362, December.

Robin Cooper. 2005b. Records and record types in semantic theory. *J. Log. and Comput.*, 15(2):99–112.

Peter Gärdenfors. 2000. *Conceptual Spaces: The Geometry of Thought*. MIT Press, Cambridge, MA, USA.

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.

J. A. G. Groenendijk and M. J. B. Stokhof. 1988. Context and information in dynamic semantics. In *Working models of human perception*. Academic Press.

M.A.K Halliday. 1978. *Language as Social Semiotic: The Social Interpretation of Language and Meaning*. Baltimore: University Park Press.

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.

Thomas K Landauer and Susan T. Dumais. 1997. A solution to plato's problem: The latent semantic analysis theory of the acquisition, induction and representation of knowledge. *Psychological Review*, 104:211–240.

Staffan Larsson. 2007. A general framework for semantic plasticity and negotiation. In H. C. Bunt, editor, *Proceedings of the Seventh International Workshop on Computational Semantics (IWCS-7)*.

Tom M. Mitchell. 1997. *Machine Learning*. McGraw-Hill, New York.

J. Pustejovsky. 1991. The generative lexicon. *Computational Linguistics*, 17(4):409–441.

Francois Recanati. 2003. *Literal Meaning - The Very Idea*. Cambridge University Press.

Luc Steels and Tony Belpaeme. 2005. Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, 28(4):469–89, August. Target Paper, discussion 489-529.