

Semantic negotiation in dialogue: the mechanisms of alignment

Gregory J. Mills

Interaction, Media and Communication
Research Group
Department of Computer Science
Queen Mary, University of London
London E1 4NS
gj@dcs.qmul.ac.uk

Patrick G.T. Healey

Interaction, Media and Communication
Research Group
Department of Computer Science
Queen Mary, University of London
London E1 4NS
ph@dcs.qmul.ac.uk

Abstract

A key problem for models of dialogue is to explain how semantic co-ordination in dialogue is achieved and sustained. This paper presents findings from a series of Maze Task experiments which are not readily explained by the primary co-ordination mechanisms of existing models. It demonstrates that alignment in dialogue is not simply an outcome of successful interaction, but a communicative resource exploited by interlocutors in converging on a semantic model. We argue this suggests mechanisms of co-ordination in dialogue which are of relevance for a general account of how semantic co-ordination is achieved.

1 Introduction

One of the first things apparent to European travellers on arriving at an American hotel is that the ground floor is also the first floor. Any confusion can be quickly corrected by an observant concierge, whether by explicitly stating the convention, or by implicitly bypassing the problem with a different description, such as “go up 5 flights of stairs”. Assuming this description is sufficient to guide the hapless traveller to the correct room, when the same traveller asks for assistance to find another part of the hotel, the concierge is faced with a choice of whether to give a description involving floor numbers or in terms of flights of stairs.

The immediate question that emerges is what motivates this choice between different semantic models of a domain, how they are deployed when interlocutors are faced with problematic understanding, and which semantic model is subsequently used once the problem is resolved. Although existing approaches to dialogue agree that answering this question necessarily involves focusing on the interactional devices available to interlocutors, their primary emphasis is on the information-exchange aspects of language use. Larsson (2007) provides a useful distinction between the co-ordination of information, i.e. establishing common ground (Clark, 1996) and the co-ordination of linguistic resources which are adapted to suit particular communicative situations in order to make such information-exchange possible. Part of this framework involves interlocutors negotiating which particular semantic model to use, and adapting their own interpretations on the basis of successful/unsuccessful use. However, although this framework sketches out a formal account of the mechanisms involved in this process, it is not concerned with predicting which particular semantic model will be adopted by interlocutors.

A model of dialogue which attempts to address this issue is the interactive alignment model of Pickering and Garrod (2004). In this model convergence on a semantic model is arrived at via tacit priming occurring at all levels of representation (phonetic, phonological, lexical, syntactic, semantic and situational): interlocutors are more likely to re-use the representations used by their partner, giving rise to a “winner-takes-all” dynamic (cf. Steels & Belpaeme, 2005) which leads to align-

ment of interlocutors' representations. This is further re-inforced by "percolation" occurring between levels, thus lexemes associated with particular semantic models will reinforce the use of these models.

The claims associated with the interactive alignment model (henceforth IM) are drawn from a series of maze task experiments (Garrod & Doherty 1994; Garrod and Anderson, 1987; Anderson and Garrod, 1987). This paper discusses some of the original findings of these experiments and a further set of maze task experiments conducted by Healey and Mills (2006), Mills and Healey (2006). These papers argued that the primary mechanisms provided by the IM are insufficient for explaining observed patterns in maze task dialogue; in particular how semantic co-ordination is achieved. The present paper argues that interlocutors in the Maze task exploit variation in usage in the service of semantic co-ordination. Furthermore we argue this suggests mechanisms which are relevant for a general account of how semantic co-ordination is achieved in dialogue. As the claims developed here are based on the maze task, we first explain the task in more detail. We then discuss a series of examples drawn from this task that raise basic issues for models of semantic co-ordination.




Figure 1: Example maze configuration. The solid black circle shows the player's current position, the cross represents the goal point that the player must reach, solid bars the gates and shaded areas the switch points.

2 The maze task

The maze task developed by Garrod et al involves pairs of participants seated in separate rooms in front of a computer which displays a simple maze consisting of interconnected nodes (see Fig 1). Participants must move their respective position markers through the maze in order to reach a "goal" node. Some of the paths are blocked by gates, which are opened by participants guiding each other onto "switch" nodes (shaded areas). This provides participants with the recurrent co-ordination problem of collaboratively individuating and referring to maze locations in order to solve the maze. The descriptions used by participants to refer to maze locations are classified by Garrod et al. into four distinct types:

- Figural:** Picks out salient features of the maze:
 "The l-shape sticking out at the top"
 "The uppermost box"
- Path:** Traces a route along the connections between nodes:
 "Go 2 up, 1 down, 3 along, 5 up"
 "up, right, down, up"
- Line:** Treats the maze as consisting of horizontal or vertical vectors:
 "3rd row, 5th box"
 "4th column, second square"
- Matrix:** Cartesian co-ordinate system:
 "4,2"
 "A1"

It is assumed that these different description types correspond to different semantic models of the maze.

3 Conservatism

The first question, also raised by Healey and Mills (2006), concerns the tension between the interactive alignment model's inherently conservative primary co-ordination mechanism and the migration in description types commonly observed in the Maze task. To the extent that it relies on priming as its basic mechanism the IM cannot provide an account of how once a convention is established and used successfully, it might be supplanted by

another.. However, it is consistently observed that the description types used most frequently initially fall into disuse and are not converged on in later games. Across trials there is a general shift from more “concrete” (Figural and Path) descriptions towards more “abstract” (Line and Matrix) descriptions, which runs counter to precedence. A typical pattern of the shift is given in table 1, below:

0 mins:	The piece of the maze sticking out
2 mins:	The left hand corner of the maze
5 mins:	The northernmost box
10 mins:	Leftmost square of the row on top
15 mins:	3rd column middle square
20 mins:	3rd column 1st square
25 mins:	6th row longest column
30 mins:	6th row 1st column
40 mins:	6 r, 1 c
45 mins:	6,1

Table 1: Semantic shift from “Figural” and “Path” descriptions to “Line” and “Matrix” observed in maze task dialogues.

Garrod (1999) discusses this process as an “explosion” process. However, this, in itself, doesn’t explain the systematic patterns of change observed in the experiments.

4 Variation

The early explanations of co-ordination in the Maze Task also emphasized the importance of variation in the description types participants are exposed to. Garrod and Doherty (1994) assigned participants to one of three different groups: (1) isolated pairs who always interacted with the same partner in subsequent games, (2) a sub-community group whose members changed partners in each game, only interacting with members from the same sub-community, and (3) a non-community group whose members always interacted with a new partner who was not drawn from the same community. Although initially pairs in the sub-community group were less co-ordinated than the isolated pairs, using a wider variety of referring expressions, by the later trials, this pattern was reversed: participants in the sub-community group

had converged on a single Matrix scheme and consistently matched each other’s descriptions.

These findings present a problem for accounts of co-ordination which rely on priming, as they make the emphasis of the priority of alignment of representations at all levels problematic. The metaphor of two tightly-coupled production and comprehension systems is the paradigm case of successful co-ordination, as it allows rapid priming between interlocutors’ representations. However, these experiments show weaker semantic co-ordination in the isolated dyads than within the group. As Garrod and Doherty (1994) concur, this implies that variation, i.e. differences in interlocutors’ representations is important for establishing and sustaining semantic co-ordination.

5 Granularity of analysis

If variation of description types is intrinsic to the development of semantic co-ordination, this strongly suggests the importance of mechanisms involved in dealing with problematic understanding (Healey, forthcoming). All things being equal, variation increases the likelihood that interlocutors will encounter others whose use of language will differ more from their own. Further, any account of misunderstandings must also be able to address semantic differences between descriptions: participants in the maze task do not treat these four description types equally, and consequently are not appropriately modelled as co-ordination equilibria of the kind described by Lewis (1968) (Healey, 2004; forthcoming). Existing experimental data shows that participants systematically favour Figural and Path descriptions when encountering problematic dialogue (Mills and Healey, 2006; Healey, 1997) not the prior most frequently used semantic model as predicted by the IM.

Looking more closely at the dialogues, it is not clear that the co-ordination mechanisms actually operate directly at the level of the four basic semantic models. Consider the following excerpt in which a participant encounters difficulties with a Line description type and its associated counting conventions. The dialogue continues with more Figural descriptions, before resuming at turn (35) with a Line description:

- (1) A: go to the 1st row 2nd on the right
- (2) B: 2nd?
- (3) A: on the right
- (4) B: OK, I can only get to the left of the maze
- (5) A: go to the highest square on the left
- (6) B: yes. And then?
-
- (35) B: I'm on the top row 2nd square

Excerpt 1: Deletion of elements from problematic turn.

While superficially, A's turn at (3) appears simply as a repeat of (1), with "on the right" being omitted, the subsequent turns continue with Figural descriptions. On this basis, it is unclear whether (1) and (3) invoke the same Line model or whether (3) invokes a Figural description. There is a large class of similar clarification sub-dialogues which involve deletion of a problematic element and result in the continuation of the dialogue with more Figural descriptions.

This issue is of importance for any theory of semantic co-ordination as it raises the question of the granularity of the mechanisms involved in how interlocutors collaboratively change semantic model. Further, it strongly suggests that alignment is not simply an outcome of successful communication, but can provide the background against which other co-ordination mechanisms operate. Turns (1)-(6) demonstrate high levels of between-speaker alignment, while at the same time involving a shift in semantic model. Before returning to this below, we demonstrate further differences between the informational view of language and an account which focuses on semantic co-ordination.

6 Information vs. semantic co-ordination

From an informational perspective, if an utterance fails to secure reference, there is the general assumption that more information will be provided to allow resolution of the problem. However, in (3), no new information is provided by A. This is a counter-example to Clark and Marshall's (1981) model of definite reference repair, which states that to be effective "repair must add or alter descrip-

tors, but not delete them". Importantly, these CR responses that simply delete elements from the target turn are not treated by participants as repeats and queried again, but appear to promote resolution of the problematic understanding by engendering the use of more Figural descriptions. The words which are omitted do not appear, as with the level of description types, to be dictated by prior frequency of use (Mills, 2007). Instead, the data suggest that this pattern is motivated by a relaxation of the constraints of successful interpretation (Healey and Mills, 2006).

The example above raises a further question concerning the relationship between semantic co-ordination and the exchange of information. In existing "ladder models" of communication such as the collaborative model of Clark (1996) and Allwood (1995), there is the general expectation that on encountering and signalling problematic understanding, interlocutors enter a sub-dialogue to resolve the problem, which on completion proceeds at the same "level". From this perspective, B's turn-initial acknowledgment at (4) should demarcate the end of the sub-dialogue dealing with the problematic understanding. Focusing on the description types, however, shows that it is only at turn (35) that the interlocutors return to using the original problematic line description; the semantic effects persist beyond the immediate sub-dialogue. This highlights the inadequacy of a strict informational view of language as the response provides no additional information, yet still has the effect of resolving the misunderstanding.

7 Exploitation of alignment: patterns of deletion, modification and addition

In addition to deletion of elements contained in referring expressions, the maze task dialogues exhibit a multiplicity of ways in which interlocutors modify descriptions when dealing with problematic understanding, through the addition, substitution and (as described above) deletion of elements of semantic models. We argue that alignment is key to these patterns of modification, as it provides a backdrop against which changes can be made. The canonical example of this is embedded correction (Jefferson, 1983; Saxton, 2007) which exploits the structure provided by alignment to make a figure / ground distinction that allows the corrected element to be identified:

- (1) A: You need to go to the top of the 5th row
- (2) B: I can't get to the top of the 5th line

Excerpt 2: Substitution of problematic elements .

Embedded corrections in the maze task exhibit very high levels of between-speaker alignment, yet occur at points in the dialogue where there is problematic understanding. This indicates that alignment can not simply be reduced to an index of successful communication. While this particular conversational device which spans 2 turns (and possibly a third) has received much attention, closer inspection of the maze task dialogues reveal a far larger space of possible means of exploiting alignment. Excerpt 1 above showed deletions, Excerpt 2 substitutions, however a similar pattern also appears with the addition of Figural elements.

- (1) A: I'm in the 4th row 5th square
- (2) B: where's that ?
- (3) A: The end bit
- (4) B: cheers, I'm on the end bit right at the top
- (5) A: can you get to my switch?
-
- (23) B: am on the top row 3rd square

Excerpt 3: Addition of “Figural” elements.

At a first glance, this excerpt looks like a straightforward clarification request followed by the provision of more details, specifying that the “5th square” is also “the end bit”. B's use of “cheers” in (4) and subsequent provision of her own maze location would appear to demarcate the end of the clarification sequence, as they provide an acknowledgment and a “next relevant contribution” (Clark, 1996). However, focusing on the ensuing turns yields a pattern that parallels the first example. The semantic effects stretch beyond the immediate clarification sub-dialogue: both interlocutors continue with more Figural descriptions

until turn (23) where the original, problematic Line description is attempted again.

A further issue emerges when interlocutors finally re-use the original description, as in turn (23) of Excerpt 1, and (35) above: although the surface form of the descriptions are similar, this does not necessarily entail that they individuate the same locations. For example, the counting conventions associated with squares may change, such as counting from the left instead of the right or counting from 0 as opposed to 1, similar to the concierge example above. The axes may also change, with “rows” referring to vertical vectors (i.e. columns).

This raises important questions of the relationship between the problematic utterance, the signalling of the problem, the response, the ensuing figural sub-dialogue and the subsequent return to the superficially similar but potentially altered description type. It appears that alignment is not simply an outcome but an interactional resource that is exploited to facilitate the continuation with more Figural descriptions (cf. Saxton, 2007).

In the first excerpt, turns (1) and (3) only differ minimally from each other, while in the second example, turn (3) can be seen to be operating elliptically on turn (1). However, both engender similar semantic shifts towards Figural descriptions and result in a return to the originally problematic Line description.

This leads to the immediate question of what motivates interlocutors' patterns of alignment and modification, and how they reflect differences of understanding and diagnosis of the problem. The tacit and fine-grained nature of these modifications exacerbates the problem of arriving at a preliminary taxonomy, as these dialogue sequences are not readily categorizable as either “elaborations” or “reformulations” (cf. Purver et al., 2004, Schlangen 2004).

8 Boundary of (mis)communication

During the course of maze task dialogues, participants shift seamlessly and tacitly from one description type to another. This occurs both within problematic and unproblematic dialogue. From an informational perspective, miscommunication is readily describable as a form of mismatch, yet from a semantic perspective, participants match each other more when encountering difficulties.

Thus alignment cannot be taken as a straightforward index of successful interaction.

This also raises a methodological point. Measures of matching of representations, whether at the level of description type or its constituent elements are only an approximate index of semantic co-ordination. The excerpts above demonstrate the importance of the interplay between what is retained and what is modified. What is required is a measure that is sensitive to the kind of model being used and the kind of repair being performed.

In addition, more frequent repair does not necessarily entail that a dialogue is unsuccessful. It is not the case that interlocutors introduce their utterances carefully, and once they are sufficiently co-ordinated, move on. The general pattern is that when participants introduce abstract (Line and Matrix) descriptions, they do so opportunistically. At the start of the games they frequently attempt both Line and Matrix descriptions, which are associated with higher co-ordination. However, there is evidence that it is only where they can go through the process of detecting and responding to differences in usage, i.e. repair, that co-ordination develops (Healey and Mills, 2006).

If the boundary between description types and also the boundary between successful and unsuccessful use can be as porous as demonstrated in the excerpts above, this also suggests a more complex picture of referential contraction (Krauss and Weinheimer, 1966) than provided by current models of dialogue. In current models this is primarily associated with successful use: in the collaborative model, interlocutors follow the principle of “least collaborative effort” (Clark and Wilkes-Gibbs, 1986), whereby successful use sets a precedent for an expression; co-ordination on precedence allows interlocutors to delete elements of the description on successive mention. It is assumed that the information associated with these deleted elements that are no longer on the conversational surface can be re-accessed in the common ground and mentioned explicitly, e.g. to assist disambiguation.

By contrast, the phenomena from the maze task show how similar processes are operative during problematic dialogue, raising further questions concerning the difference between elements that are removed in successful, as opposed to problematic dialogue and where this boundary lies.

Larsson's model of semantic co-ordination places a strong emphasis on the role of feedback in

negotiating this boundary in terms of appropriateness gleaned from feedback (e.g. repair, acknowledgements etc.), and provides a schema which analyzes the effects of novel uses of a word and the subsequent update of interlocutors' representations.

Findings from the maze task experiments augment this approach as they suggest that evidence of appropriateness is also derived in the absence of overt repair from semantic change alone. The excerpts indicate that interlocutors are sensitive to which particular tacit shift in model leads to a relaxation of the constraints on successful communication, and consequently can be exploited to indicate problematic understanding (Mills, 2007). For example, consider the following two excerpts:

- (1) A: It's on the 5th row 4th square
- (2) B: Huh?
- (3) A: The last square

- (1) A: It's on the 5th row 4th square
- (2) A: The last square

Excerpts 4, 5: Provision of feedback

If the dialogue continues successfully in both these instances, it is unclear how to adequately capture the differences between them, in particular, how both patterns affect subsequent use of the description types,

One of the main challenges facing an account of semantic co-ordination is teasing apart how interlocutors' models are affected by both semantic change exploited as a resource using the mechanisms of alignment outlined above, and feedback concerning that change, as both aspects inhabit the boundary between successful and unsuccessful use.

Evidence from the maze task suggests this boundary is one of the important loci in the development of semantic co-ordination.

9 Semantic plasticity

To describe how interlocutors dynamically adapt the meanings of the words they use to the communicative situation and how they are shaped throughout the course of the dialogue, Larsson (2006) introduces the notion of “semantic plasticity”

ty". This model is sensitive to the fact that descriptions can involve a plethora of different "ad-hoc registers", which resonates strongly with the empirical phenomena described here. However, the data from all the maze task experiments presents a further problem for attempts to model these phenomena, as successful co-ordination on the more specific abstract levels appears to be predicated upon prior successful use of less specific Figural descriptions: the Figural descriptions are highly specific to individual mazes and allow participants to co-ordinate on their salient features, whereas the Line and Matrix descriptions abstract away from each individual instance to form dyad-specific conceptualizations of vectors and their associated counting conventions.

While Larsson's account highlights the sheer flexibility of ways in which linguistic resources are mobilized and adapted to particular interaction settings, the data from the maze task suggest an additional level of complexity. Namely that the semantic resources can not be treated as separate, essentially equal encyclopaedias that interlocutors draw on. One way in which the cumulative shift toward Matrix descriptions is achieved is by the combination of different "registers" (Larsson 2007) to form a super-ordinate one. Here the question concerns which specific features of each semantic model are included in the final one, in particular when there are problems of commensurability. For example, as table 1 shows, a common pattern in maze task dialogues is that approximately half-way through the dialogues participants use "Line" descriptions. It can occur that they alternate between describing the maze as consisting of vertical and horizontal vectors, say with one participant favouring horizontal and the other favouring vertical vectors (space considerations preclude a thorough examination of this process, described in Mills, 2007). It frequently occurs that Matrix descriptions emerge when these two different Line models are combined to form a Matrix description. This process, however, is not as a rule simply a matter of combining the two. Frequently, the two types of Line description employ different counting conventions, as in the example of the concierge above, giving rise to the problem of whether to retain different counting conventions for the different axes, or employ the same one. The question then emerges as to how this super-ordinate, more abstract semantic model affects the original models.

Results from the maze task suggest this is achieved tacitly by interlocutors, employing similar patterns of modification to those described in the excerpts above (Mills, 2007).

10 Conclusion

The phenomena described here demonstrate the need for an account of semantic co-ordination that explains how interlocutors converge on a semantic representation. Dialogues from the maze task provide compelling evidence that such an account must necessarily be able to account for how variation, and hence differences in semantic models are resolved. This approach necessarily involves shifting the focus from an informational view of language towards a focus on how interlocutors actually address these differences.

In a sense, this presents a reversal of the priorities of existing models. For the interactive alignment model, as well as the collaborative model, misunderstanding is seen as a secondary problem that emerges as a complication of communication which is ordinarily successful (Healey, 2004; forthcoming). The collaborative model explicitly states that in order for communication to be successful, positive evidence of understanding must be demonstrated.

By contrast, the view presented here brings problematic understanding into the foreground, as it is in such instances, when conventions don't work as expected, that interlocutors gain a sense of their applicability. The phenomena presented here suggest that the processes operating in instances of misunderstanding are as much progenitors of semantic co-ordination, as their traditional counterpart of displays of positive understanding. Interlocutors' separate interaction histories inescapably give rise to problems concerning the development and sustenance of mutual-intelligibility, intrinsically requiring interlocutors to resolve differences of semantic model in interaction. The data from the maze task experiments demonstrate how this can be achieved through tacitly modifying the constituents of semantic models. This modification involves the exploitation of alignment, and has the effect of relaxing the constraints on successful understanding.

Any theory of dialogue must, in the first instance be concerned with what interlocutors actually do. The phenomena presented here demonstrate

mechanisms of semantic co-ordination that have previously fallen under the category of information-exchange, and the questions raised present rich opportunities for further experimental investigation.

References

- Allwood, J. (1995). An activity based approach to pragmatics. *Technical Report (GPTL) 75*, Gothenburg Papers in Theoretical Linguistics, University of Göteborg.
- Clark, H. H. *Using Language*. Cambridge University Press, Cambridge.
- Garrod, S. and Doherty, G. 1994. Conversation, co-ordination and convention: an empirical investigation of how groups establish linguistic conventions. *Cognition*, 53:181-215.
- Healey, P.G.T. 1997. "Expertise or expert-ese: The emergence of task-oriented sub-languages." *Proceedings of the Nineteenth Annual Conference of The Cognitive Science Society*. M.G. Shafto and Langley, P. (Ed.s) August 7th-10th, Stanford University, CA. pp. 301-306.
- Healey, P.G.T. 2004. "Dialogue in the degenerate case?" Peer Commentary on Pickering and Garrod: 'The Interactive Alignment Model'. *Behavioral and Brain Sciences* 27(2) p. 201.
- Healey P.G.T. (forthcoming) "Interactive Misalignment: The Role of Repair in the Development of Group Sub-languages" in Cooper R. and Kempson R. (eds) *Language Change*.
- Healey, P.G.T. and Mills, G.J. 2006. "Participation, Precedence and Co-ordination in Dialogue" in Sun R. and Miyake, N. (eds.) *Proceedings of Cogsci06: The 28th Annual Conference of the Cognitive Science Society*. Vancouver, BC, Canada. 26-29th July. pp. 1470-1475.
- Jefferson, G. (1983). On exposed and embedded correction in conversation. *Studium Linguistik*, 14, 58-68.
- Krauss, R. M. and Weinheimer, S. 1966. Concurrent feedback, confirmation and the encoding of referents in verbal communication. *Journal of Personality and Social Psychology*, 4:343-346.
- Larsson, S. 2007. A general framework for semantic plasticity and negotiation. In Bunt, H. C., and Thijsse, E. C. G. (eds): *Proceedings of the 7th International Workshop on Computational Semantics (IWCS-7)*.
- Larsson, S. 2006. Semantic plasticity. Paper presented at LCM 2006. (Language, Culture and Mind), July 2006, Paris, France.
- Lewis, D. 1969. *Convention: A philosophical study*. Harvard University Press.
- Mills, G.J. and Healey, P.G.T. (2006) Clarifying Spatial Descriptions: Local and Global Effects on Semantic Co-ordination. In Schlangen, D. and Fernandez, R. (eds.) *Proceedings of Brandial06 The 10th Workshop on the Semantics and Pragmatics of Dialogue*. University of Potsdam, Germany; September 11th-13th. pp.122-129.
- Mills, G. J. 2007. The development of semantic co-ordination in dialogue: the role of direct interaction. Unpublished PhD. Thesis.
- Purver, M., Healey, P. G. T., King, J., and Mills, G. 2003. Answering clarification questions. In *Proceedings of the 5th Workshop of the ACL SIG on Discourse and Sialogue (SIGdial03)*, Sapporo, Japan.
- Saxton, M. 1997. The contrast theory of negative input. *Journal of Child Language*, 24, 139-161.
- Schlangen, D. 2004. Causes and strategies for requesting clarification in dialogue. In *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue*.
- Steels, L. and Belpaeme T. 2005. Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, 28(4):469-89