The Common Code: Diagnosis, a treatment and its prognosis

Patrick G. T. Healey
ATR International,
Media Integration and Communications Research Laboratories,
2-2, Hikaridai, Seika-cho, Soraku-gun,
Kyoto 619-02, Japan
ph@mic.atr.co.jp

Abstract

It is commonly, often tacitly, assumed that the mutual-intelligibility of natural language dialogue is underwritten by a common semantic code. Ultimately, this assumption is both emprically and conceptually unsustainable and adequate models of dialogue need to provide for ways in which semantic conflicts can be addressed. One possiblity for resolving such conflicts is by the use of some form of matching strategy to accomodate situations in which interlocuters do not share the same semantic model. The most clearly articulated proposal along these lines is the Input-Output coordination model. This paper examines this model and, on the basis of data obtained from a task-oriented dialogue, argues that it does not provide a sufficient mechanism for the resolution of semantic conflicts.

1 Diagnosis

Across a range of different approaches, the notion of 'speaking the same language' is commonly analysed in terms of some set of linguistic conventions or code that forms part of the common ground for interloctuors, and that the mutual-intelligibility of natural language depends on it. In many respects this idealisation is an elaboration of the common sense or pretheoretic view of communication, embodied in English by the conduit metaphor (Reddy, 1979). Speech communities are often defined by reference to a shared semantic code (see Taylor, 1992 for a discussion). Within Cognitive Science and AI, models of multi-agent systems assume, or perhaps presuppose, that the agents themselves are *semantically transparent* to one-another. That is, while admitting asymmetries between agents in, say, the type of speech acts assigned to an utterance or the plans derived from it, agents are considered to share, at some basic or literal level, a single semantic model. While for many purposes the idealisation to a common semantic code is sufficient, it is ultimately inadequate as a characterisation of dialogue.

Theoretical considerations suggest that, even in principle, this assumption is problematic; natural languages do not appear to be adequately characterised as codes in anything more that a metaphorical sense (e.g., Healey, in press; Nolan, 1994; Taylor, 1992). A number of studies also indicate that this assumption is empirically inadequate. For example, ethnographic analyses of technological interventions in the workplace frequently identify the existence of distinct semantic communities within companies, examining the consequences this has for (mis)communication (see e.g.; Robinson and Bannon, 1991; Schmidt and Bannon, 1992). Rather than assuming some semantic code as a fixed element of the initial common ground for parties to a dialogue, these examples suggest the need for some means by which a semantic model itself can be revised during the course of an interaction.

2 A Treatment

Currently there are two accounts of natural language dialogue which provide mechanisms that address this question; the collaborative model of dialogue (e.g., Clark, 1996; Clark and Wilkes-Gibbs, 1986) and the input-output coordination model (IOM) developed by Garrod and Anderson (1987) and Garrod and Doherty (1994). The collaborative model is examined in more detail in Healey (forthcoming), this paper focuses on the IOM.

The development of this model has been driven by the analysis of dialogues generated by the maze task (Garrod and Anderson, 1987, Garrod and Doherty, 1994), a variant of which is adopted in the experiment discussed below. This task is constructed so that individuals, working in pairs, are faced with a recurrent problem of describing to each other their positions within a maze-like grid, similar to that in fig 1). Garrod and Anderson (1987) demonstrated that for any given position in the grid, a range of semantically distinct description types can be produced, reflecting alternative possible conceptualisations of the area of the maze (the alternative description types are described in more detail below).

In Garrod's and Anderson's (1987) study, individuals were paired into dyads who repeatedly performed the task over a number of trials. It was observed that members of the same dyad were much more likely to use the

same description type than an arbitrary pair of individuals drawn from the sample population as a whole; indicating some dyad-based process of convergence.

Extending the analysis beyond isolated dyads, Garrod and Doherty (1994) investigated the pattern of convergence in a group of individuals. An important comparison in their study was between two conditions: a community group, in which dyads were always composed of individuals drawn from a single (sub)pool of subjects and a non-community group in which dyads were composed of individuals drawn from a wider group. This manipulation was designed to ensure that across a number of trials the community group could build up a common interaction history unavailable in the non-community group. Under these conditions Garrod and Doherty found that the community group converge rapidly, and strongly, on a specific description type whereas the non-community group do not. In neither condition was there any basis for particpants to hold explicit beliefs about group membership.

Perhaps the most striking finding of both Garrod's and Anderson's (1987) and Garrod's and Doherty's (1994) studies was that the that the observed convergence in use of description types was not achieved through explicit negotiation. If anything, it was found that negotiation, where it did occur actually interfered with the inter-speaker coordination of description types.

2.1 The Input-Output Model

The basic mechanism by which the IOM aims to account for these findings is by appeal to input-output coordination. In essence, the proposal is that the parties to a dialogue attempt to coordinate their production and interpretation of utterances by each maintaining, as far as possible, a single underlying representation for both. This representation governs each individual's lexical, syntactic and semantic choices in their processing of maze-related utterances. Adhereing to this strategy in conversation will have the result that each individual's output will tend to match the last relevant input; i.e., their partners last description. The main advantage of this account is that following such a matching strategy provides an efficient means of narrowing down the common ground between interlocutors without recourse to higher-order beliefs or explicit negotiation.

Where some conflict in description types occurs, for example, where the interlocutor's last turn is simply uninterpretable on the current understanding, additional mechanisms are required to resolve the problem. Garrod and Doherty (1994) supplement this basic account of coordination by deriving general, cognitive, constraints on the stability of the underlying conceptual representation. They suggest that where a conflict arises the most stable representation will predominate. In terms of the maze game, the prediction is that where conflict between individuals' description types arises, they will switch to the description scheme most commonly used by both of them in their immediately preceding maze games.

Elaborating the conditions which can promote representational stability, they argue by analogy with exemplar based concept learning that the greater the range of description types to which an individual is exposed, the more stable the derived representation will be. In the case of the community group in Garrod's and Doherty's study, individuals are exposed to a greater range of exemplars in virtue of performing the task with a range of different individuals and, combined with the development of a common interaction history within the group, this underpins both the strong convergence of the community group and the contrast with the non-community group. The IOM thus provides an proposal for how semantic variation could be overcome in dialogue through a low-level matching process together with mechanisms for conflict resolution.

3 Experiment

The current study was initiated with the aim of examining in more detail the resolution of semantic conflicts by creating conditions under which the emergent coordination of description types in different 'community' groups could be brought into conflict. A paper-based version of the maze task was employed which preserved the requirement to generate spatial descriptions in cooperation with another individual.

Design

This study was designed with the goal of promoting the emergence of a number of sub-communities comparable to the single community group reported by Garrod and Doherty (1994). Thus in the first phase (trials 1-5), dyads were composed of individuals drawn from the same pool of participants. Pairing different individuals from the same pool on each trial ensures that a common interaction history develops within that pool of individuals. The experimental manipulation of dyad composition is made in the second phase (trial 6) where half the dyads are composed, as before, of individuals drawn from the same pool or subgroup of participants; the *Homogenous Condition* and half are composed of individuals drawn from different subgroups; the *The Mixed condition.* The rationale being that crossing between subgroup should interfere with whatever degree of group-based linguistic coordination has built up in the preceeding trials.

The resulting design was a simple factorial with dyad composition (mixed vs. homogenous) as a between-subjects independent variable and trial number as a within-subjects independent variable. In order to ensure that there were enough participants to provide new pairings of individuals, in each sub-group, on each trial a total of 24 participants were used. For trials 1-5 they were divided into three subgroups of eight On trial six half the subjects were paired with individuals from the same subgroup and half from different subgroups, with all the combinations equally represented resulting in 6 Homogenous dyads and 6 Mixed Dyads.

Subjects

The study used 24 participants randomly assigned into the three sub-groups of eight. They were recruited from

Figure 2: Distribution of Description Types According to Experience

The overall pattern, similar to that found in Healey (forthcoming), shows Figural description types falling across trials, matched by a rise in the more abstract Line and Matrix description types. Pearson's product-moment correlation calculated between the proportions of Figural and Matrix-type descriptions in each trial indicates a strong negative relationship, r= -0.74, and there is a reliable difference in the pattern of raw frequencies of Figural and Matrix-type descriptions across trials: $\chi^2_{(5)}$=35.068, p=0.000.

Figure 3: Distribution of Description Types in Homogenous (H) and Mixed (M) Dyads

An additional comparison was made between the Mixed dyads (trial six) and the naïve dyads performing the task for the fist time on trial 1. Although comparison of the frequencies of all description types in the mixed dyads and all dyads on trial 1 were reliably different: $\chi^2_{(3)}$=26.28, p=0.000, both groups display a similar preference for Path and Figural description types while differing in the relative proprtions of Line and Matrix-type descriptions (see table 1). Suggesting a shift, in the mixed dyads, toward the description types used by naïve pairs.

Type:	Figural	Path	Line	Matrix
Mixed:	16%	40%	39%	5%
Trial 1:	18%	48%	15%	19%

Table 1: Description types in Trial 1 and Mixed (Trial 6)

Turning to the IOM, entrainment scores, following the method described in Garrod and Doherty (1994), were also calculated for each member of a dyad in each trial. This is an index, varying between one and zero where 1=perfect entrainment, of the tendency for individuals to generate descriptions of the same type as those their partners have just generated. It is calculated as the number of description types produced by an individual that match the preceding description type produced by their partner, divided by the total number of exchanges of description types in that trial.

The average entrainment score was 0.57, and although above chance (0.32), lower than the average of 0.9 reported by Garrod and Doherty (1994). Examination of the entrainment scores between members of each dyad revealed a strong positive correlation; Pearson's product-moment: r=0.90, but there was apparently no reliable increase in entrainment across trials; 1:0.58, 2:0.58, 3: 0.54, 4:0.58, 5:0.56 and 6:0.54. An analysis of variance on the average entrainment score for each pair confirmed this; omnibus $F_{(5,66)}$=0.054, p=0.99, linear trend: $t_{(66)}$=0.311, p (one-tailed)=0.378.

A comparison of entraintment scores was also made for the experimental manipulation of group composition. The average scores on trial 6 were: Mixed: 0.536 and Homogenous: 0.542. The scores were entered into an analysis of variance with group composition, mixed versus homogenous, as a single, between-subjects, factor. This confirmed that there was no reliable difference; $F_{(1,22)}$=0.003, p=0.957.

4 Discussion

The results of this study reveal an interesting tension between the type of semantic coordination observed, as indexed by description types, and the degree of matching, as indexed by entrainment scores. The experimental manipulation of group composition clearly intereferred with the degree of semantic coordination achieved within the subgroup prior to trial 6; relative to the Homogenous dyads, the Mixed dyads relied far more on Figural description types and far less on Matrix. There was, however, no difference between these two conditions in the degree of entrainment displayed. Furthermore, the degree of coordination, as measured by entrainment scores, was effectively constant across trials. Entrainment therefore appears to be independent of choice of description type. In one sense this is unsurprising since entrainment scores calculate the frequency of matching regardless of the type of description produced. Thus, perfect entrainment on Figural descriptions is, on this index, equivalent to perfect entrainment on Matrix descriptions.

However, it can be argued the different description types entail quite different degrees of cordination, in particular; that matching of Figural descriptions represents the weakest degree of coordination and Matrix the strongest. The reasoning is that the generalisations possible from one Figural description to another are very weak. An expression like "the arm sticking out on the left" is unlikely to apply sucessfully to more than a few instances of the maze configuration. By contrast, Matrix descriptions can potentially invoke the same order of axes, the same origin and the same counting scheme on each occassion of use. This is supported by the fact that where Matrix description types tend to be highly elliptical, utterances often amounting to just two numbers, e.g., "three four", Figural description types are almost always extended, produced in installments over a number of turns and involve several stages of checking for comprehension. In this task, Figural descriptions are, in a sense, the lowest common denominator, calling only on the pre-established linguistic coordination that each individual brings to the task in the first place whereas Matrix descriptions call on local, more specific, con-

ventions established during the course of the task (cf. Garrod and Anderson,1987). As a result, entrainment scores provide a relatively coarse index of coordination which is not sensitive to the reliable pattern, observed in the Mixed dyads, of switching toward more primitive description types, since both individuals tend to make the same switch. It seems that, as indexed by entrainment, input-output coordination cannot easily account for this switch.

As noted above, where a conflict in description types occurs, the IOM predicts that the subsequent shift in description types will be goverened by the relative strength of the underlying conceptual representation of the Maze. The strength of the representations is understood here as a function of the range of examplars to which an individual is exposed. As a result, where individuals are using different description types they should shift to the one most commonly used by both players in all previous games. This is supported by Garrod's and Doherty's (1994) data, in their community group this pattern held in 8 out of 9 cases of possible conflict, the apparent exception being an artifact of the coding scheme[1]. However, this assumption does not hold true for the mixed dyads, where conflicts are most likely to occur. Combining the data from this study and that in Healey (forthcoming) out of 18 mixed dyads; 7 go against this pattern and 11 conform. Ignoring (like Garrod and Dohery) those Mixed dyads that had used the same description type in the previous game, gives 7 pairs deviating from this pattern and 8 conforming.[2] For this data, it appears that mixed dyads were most likely to shift not to the most commonly used previous scheme, but to the most basic, or least 'coordinated' scheme, with the two frequently coinciding. The IOM does not explicitly address this question as it does not discriminate between the degree of semantic coordination implied by the different description types.

The deeper issue is that a principly cognitive, individualistic mechanism, such as that offered by the IOM, does not adequately attentend to the distribution of semantic resources in a community as a whole. THe argument of this paper has been that this is a critical factor in determining successful coordination. Where the semantic resources developed within, and specific to, a subcommunity are withdrawn, as in the mixed dyads, interlocutors switch, not to the strongest representation -which may still conflict with their partner's representation- but to the description scheme supported by the widest semantic community of which they are both members.

[1] For this pair, although Matrix was the most common previous scheme type, they had adopted versions with different, conflicting, labelling systems for the axes.

[2] The chance level of independently switching to the same scheme is difficult to calculate precisely here but, intuitively, will be high since there are only four description types. Therefore if both players shift there are only two possible alternatives. If only one shifts there are 3 possible alternatives.

5 Prognosis

The question of how semantic coordination is achieved has important consequences for, amongst other things, theories of language development and the understanding of communication more generally. In most accounts, the common code plays the role of what amounts to a semantic 'meta-language' which can be used to resolve problems of interpretation. Although this will at least approximate what occurs in many situations, there is a clear threat of infinite regress. Once the assumption of a shared semantic code is abandoned the problem is to identify alternative possible explanations that can adequately ground the mutual-intelligibility of dialogue.

The model examined in this paper includes an explicit proposal that might provide a means of underpinning semantic coordination. However, the appeal to a matching principle in conjunction judgements about the likely relative strengths of the conceptual representations used in production and interpretation does not account for the data reported here. Although semantic changes in a community surely must be associated with conceptual changes for the individual members of that community, appeal to the relative strength, or persistence, of those changes does not seem to characterise how semantic conflicts are resolved. What appears to be needed is a more interactive account of how semantic conflicts or variations could be resolved, one that is both sensitive to the semantic resources in the community and the way that disparities in those resources are detected and repaired.

Acknowledgments

The author wishes to thank the EPSRC for as research studentship which provided the financial support that made this research possible. Thanks are also due to Rosemarie McCabe, Carl Vogel and Kritstiina Yokinnen for helpful encouragement. Something rather more than thanks is due to Jimmy Cuthbert for vital help in carrying out this study.

References

[Clark, 1996] H.H. Clark. *Using Language* Cambridge: Cambridge University Press, 1996.

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

[Garrod and Anderson, 1987] S.C. Garrod and A. Anderson. Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. *Cognition*, 27:181–281, 1987.

[Garrod and Doherty, 1994] S.C. Garrod and G. Doherty. Conversation, coordination and convention: an empirical study of how groups establish linguistic conventions. *Cognition*, 53:181–218, 1994.

[Healey, forthcoming] P. Healey. Expertise or Expertese? The emergence of task-oriented sublanguages.

Paper to be presented at *The Nineteenth Annual Conference of the Cognitive Science Society* Stanford University, August 7–10,1997.

[Healey, in press] P. Healey. The Public and The Private: Two domains of analysis for semantic theory. In J. Ginzburg, Z. Khaihavilli, C. Vogel and E. Vallduví (eds.) *The Tbilisi Symposium on Logic, Language and Computation: Selected Papers* Stanford CA: SiLLI/CSLI Publications.

[Reddy, 1979] M.J. Reddy. The conduit Metaphor: A Case of Frame Conflict in our Langauge about Language. In *Metaphor and Thought*, pages 284-234, Cambridge: Cambridge University Press, 1979.

[Robinson and Bannon, 1991] M. Robinson and L. Bannon. Questioning representations. In Bannon, L., M. Robinson, & K. Schmidt, (eds). *Proceedings of the Second European Conference on CSCW.* Dordrecht: Kluwer, pages 219–233.

[Schmidt and Bannon, 1992] K. Schmidt and L. Bannon. Taking CSCW seriously: Supporting Articulation Work. *Computer-Supported Cooperative Work*, 1(1-2), pages 7–40.

[Taylor, 1992] T.J. Taylor. *Mutual Misunderstanding: Scepticism and the Theorising of Langauge and Interpretation.* London: Routledge, 1992.