

Practices in Dialogue

David Schlangen

Department of Linguistics

University of Potsdam, Germany

david.schlangen@uni-potsdam.de

Abstract

We explore the idea that conversational episodes not only ground *facts*, but also establish *practices* (know how). We apply this idea to the well-studied phenomenon of lexical entrainment. We provide a formal model of situations giving rise to lexical entrainment, and use it to make precise implications of extant attempts to explain this phenomenon (the lexical pact model, and the interactive alignment model) and to make precise where our attempt differs: it is an automatic and non-strategic, but goal-driven account, which has a place for partner-specificity and group forming. We define a learning mechanism for practices, and test it in simulation. We close with a discussion of further implications of the model and possible extensions.

1 Introduction

Agnes and Bert are playing a game. Agnes is describing a figure printed on a card to Bert, and Bert is trying to find a copy of the card among several similar cards. They do this for a while, and begin to establish names for the figures, which they stick to during the course of the game.

This of course is the well-known reference game first described by (Krauss and Weinheimer, 1964), and made famous by (Clark and Wilkes-Gibbs, 1986; Brennan and Clark, 1996). The latter analyse this phenomenon as one of negotiating and forming a *pact*, where this metaphor is justified by observing that this ‘pact’ has parties for which it holds (in our example, Agnes and Bert, but not Agnes and Claire, who wasn’t party to the conversation), and that a party that assumes she has such a pact reacts to it being ‘broken’ by their partner (Metzing and Brennan, 2003).

A well-known alternative proposal, that of Pickering and Garrod (2004), would assume that

Agnes and Bert become, in some real sense, more like each other while playing the game: the representations that represent world and language to them become more alike. Unlike in Clark *et al.*’s proposal, the assumption here is that this is an automatic process that doesn’t require reasoning about the partner.

In this paper, we explore an idea that to some extent is a synthesis of elements of both these proposals, namely the idea that rounds in the reference game, and more generally, all kinds of interactional episodes, establish *practices* (a notion we take from anthropology, e.g. (Ortner, 1984; Bourdieu, 1990)), a form of know-how; they, automatically and unreflectedly, establish *a way of doing something with someone*, i.e., a way that is specific to a (type of) partner. We make precise what this can mean for the reference game, and speculate more generally on how this can complement models of interaction.

2 Lexical Entrainment in the Card Matching Game

By *lexical entrainment* (a term coined by (Garrod and Anderson, 1987)) we mean here the phenomenon that pairs of interlocutors tend to settle over the course of a conversation on a single description for a given object for which several descriptions would be possible (understandable, allowable by the language). This seems to be a very stable phenomenon, observed in symmetric conversations between players (e.g., (Krauss and Weinheimer, 1964; Clark and Wilkes-Gibbs, 1986; Garrod and Anderson, 1987; Brennan and Clark, 1996)), in conversations between novices and experts (where the experts’ terms come to dominate; (Isaacs and Clark, 1987)), in “conversations” between humans and computers (where the strength of the effect depends on the human’s beliefs about the capabilities of the computer; (Branigan *et al.*, forthcoming)), and in conversations between ‘nor-

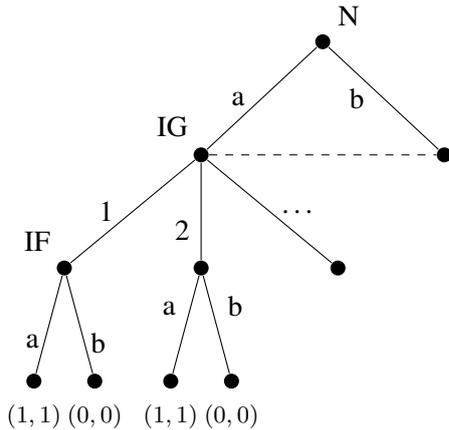


Figure 1: The game in extensive form

mal’ dialogue participants and participants suffering from severe amnesia (Duff et al., 2006).

In this section, we provide a formal model of a situation in which the phenomenon has often been studied, and note some characteristics of extant explanation attempts.

2.1 The Game

Nature picks one card from set \mathcal{O} . Player A, in role IG (instruction giver), observes this, but player B, in role IF (instruction follower), does not. IG plays from set \mathcal{V}_A , her vocabulary. (Note that this set is indexed by player, not role.) IF observes this, and as reaction picks one card out of set \mathcal{O} as well. If IF’s pick is equal to nature’s pick, both players win. Otherwise, both lose. Players are told whether they have won or lost, and the game is repeated.

The similarity in the formulation above to signalling games (sequential games of incomplete information) studied in game theory (e.g., (Shoham and Leyton-Brown, 2009)), is of course not accidental. Figure 1 shows the game in extensive form, for $\mathcal{O} = \{a, b\}$, and $\mathcal{V}_A = \{1, 2, \dots\}$ (it doesn’t matter here what exactly is in this set). The setting is slightly different, though, and we are interested in slightly different questions. First, unlike in signalling games, where the question is how stable signalling systems evolve, we assume here that there are already conventions between A and B in place. Specifically, we assume that A and B, by virtue of belonging to the same language community (which they know they do), mutually believe that 1 and 2 are good labels for a , and not so good labels for b (and so on, for other combina-

tions of pairs of labels and entities). Second, the question that we are interested in is why it is, if they play this game repeatedly, possibly switching roles between episodes, that when nature picks an element of \mathcal{O} they’ve encountered before, IG becomes more likely to chose that element from her vocabulary that has been chosen in previous successful episodes—and not any other one which she has reason to believe could be successful as well. (Remember that they assume all normal language conventions hold.)

2.2 Lexical Pacts

The model of Brennan and Clark (1996) is characterised mostly by a list of features: it is *historical*, meaning that previous choices influence the current choice, mediated by *recency* and *frequency of use*, and also influenced by *provisionality*, that is the requirement of lexical choices to be sanctioned by the other party. An aspect that the authors stress is that the model also assumes *partner specificity*: speakers choose their wording “for the specific addressees they are now talking to” (Brennan and Clark, 1996, p.1484).

The authors do not say much about how exactly a model with these characteristics might be realised, apart from “[the results imply that] long-term memory representations are involved” (p.1486), and, from a follow-up paper: “entrainment is supported by an underlying episodic representation that associates a referent, a referring expression (and the perspective it encodes), and other relevant information about the context of use (such as who a partner is)” (Metzing and Brennan, 2003, p.203).

The following is our attempt at providing a formalisation of a model that at least is compatible with the general tone of the description in the papers cited above. It assumes that the episodic representations are explicitly encoded, perhaps along the lines of the reference diaries from (Clark and Marshall, 1978), and that the rule “reuse of labels maximises success” is explicitly represented and used in making the decision. (That explicit inference is involved in the model at least is a charge that opponents seem to make often, e.g. “on their account, alignment is therefore the result of a process of negotiation that is specialized to dialogue and involves inference.” (Garrod and Pickering, 2007, p.2).)

If there is reasoning involved, it must be a form

of practical reasoning, as the outcome is a way to act. The following formulates a schema in the form of an Aristotelic practical syllogism which could be assumed to underlie that reasoning.

- a. I want to refer to object x for you
- b. I believe that honouring our pact regarding the use of α as a label for x is the best way to do (a).

or

I believe that our last successful reference to x was with label α , and that re-using the last most recent successful label is the best way to do a reference.

- c. Therefore, I honour our pact, and use α .

This, then, is in this (interpretation of the) model the type of reasoning that IG must perform when it comes to choosing an action, and IF must follow the corresponding version of this scheme governing interpretation. In the course of playing several episodes of the game, the set of beliefs of the partners changes, to include beliefs about which terms were used successfully when.

2.3 Interactive Alignment

The *interactive alignment model* of Pickering and Garrod (2004), which is positioned as a counter-proposal to the one discussed above, is similarly only characterised indirectly and not provided with a formalisation. Its main features are nicely summarised in (Garrod and Pickering, 2007, p.2): “Our claim is that a major reason for this alignment is that the comprehension of *chef* (or alternatively *cook*) activates representations corresponding to this stimulus in B’s mind (roughly corresponding to a lexical entry). These representations remain active, so that when B comes to speak, it is more likely that he will utter *chef* (or *cook*). [...] We assume that this tendency to align is automatic.”

Applied to our formalisation of the game, it seems that this model assumes that IG bases her choice on whatever is most strongly activated in her mind, and, if understood at all, this representation will increase in activation in IF, making him more likely to use the label as well at the next opportunity. Hence, what changes here over the course of playing the game are the vocabularies of the players (or rather, the way they are represented in the minds of the players); beliefs with propositional content and reasoning do not enter in the description of the phenomenon. There does not seem to be room for partner-specificity in this model.

We now turn to our model, which, as we’ve mentioned, combines elements of both of these approaches.

3 A Learning Mechanism for Practices, Applied to the Card Matching Game

3.1 Modelling the Players

We assume that there is an additional structure that explicitly represents an agent’s associations between objects and labels. Formally, for each player there is a *map* \mathcal{M} , which is a matrix where the rows represent the objects from \mathcal{O} and the columns represent the labels from \mathcal{V} . Each cell $\mathcal{M}_{o,v}$ encodes the strength of association between object o and label v . Together with a strategy for using such associations, the maps determine the interpretation of an observed label (looking at the respective column), or the label to choose when wanting to refer to an object (looking at the row). Later, we will make the maps relative to *partner* and *situation types*, e.g. \mathcal{M}^E .



Figure 2: Two Maps (rows are objects, columns are labels, numbers and colours represent association strength)

Figure 2 shows two maps \mathcal{M}_A and \mathcal{M}_B (values are shown as number and are also coded in colour). It is read as follows: for object a (first row), both players rate label 1 (first column) highly; A prefers it, while B thinks label 5 is just as appropriate. All labels are somewhat acceptable to refer to a (and in fact, all labels are acceptable for any object). For objects c, d, e (rows 3 to 5), A and B have different preferred labels, and if their partner picks their preferred label to refer to any of these objects,

they would not pick out the same object, if they go for their strongest interpretation preference. E.g., if A says 3 with the intention to refer to object c , B will understand d , if both go for their strongest preference. (This is a fairly extreme example, with lots of ambiguity, and differences in preferences; we will use this below to show that our learning method can still lead to agreement even in such a situation.)

3.2 Amended Rules of the Game, and Learning

We define two versions of the game. In the simple version, the game is exactly as stated above in Section 2.1. At each decision point in an episode, the players choose *greedily*, that is, they go for the appropriate cell with the highest value. E.g., if IG observes nature picking object a , she checks what the highest value is in the corresponding row, and in this way finds the label she wants to use (as the one corresponding to the column of the highest value). IF now, on observing IG using the label, checks in the column corresponding to the label which row has the highest value, and takes the object corresponding to this row as his interpretation, which in this version of the game means that he solves and picks this object.

The episode then closes with *rewards* being distributed to the players, as follows. The choice made by IG is always reinforced (the weight of the cell used in deciding which label to use is increased by some factor), even if IF made the wrong choice, and for *both* players. The idea here is that the mapping IG has used is, once the game has been revealed, mutual knowledge, and so IG should be motivated to use it again, as should IF. Additionally, if IF made a wrong choice, this wrong choice is punished (the weight of the cell used in deciding which label to use is decreased by some factor).¹

There is an interesting aspect of the game that we have left implicit so far: when players switch roles, they will use the map which in the previous episode they used for generation (finding a label,

¹This game is very similar to the one analysed by (Argiento et al., 2009), who present a similar urn-based (i.e., numerical weight-based) learning scheme; our game, however, diverges in two important aspects: we assume that players start with non-uniform distributions (they already have preferences for labels/object combinations), and unlike (Argiento et al., 2009) we do not reward or punish both players equally. These authors are concerned with analysing whether a signal system can be learned this way, we are concerned with whether an existing signal system can be adapted.

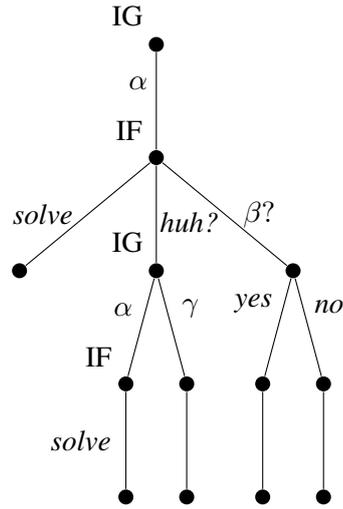


Figure 3: Extended action space for variant of game with Clarification Requests

given an object) or interpretation (finding an object, given a label) for interpretation or generation, respectively. This ensures, without further stipulations, that there is symmetry between interpretation and production: when they have used a label successfully with a certain interpretation, they will become more likely to use this interpretation when they hear this label.

The second version of the game is somewhat more complicated, as it gives the players more options for acting. (The extended action space is shown graphically in Figure 3, for one branch after IG’s first action.) After IG has played (has provided a label α for object o chosen by Nature), IF can decide whether to solve (pick an object), or make a clarification request (CR). It uses the following rules to decide what to do: if there are “distractor objects” for the label IG used, that is, if the highest weight in the appropriate column is not sufficiently far away from other weights in the column, then IF rejects, and ask a CR that could be paraphrased as “huh?”. If this is not the case, but there is another label that IF deems more appropriate for the object that is the best interpretation given this label (in the row in which the highest weight of the column corresponding to IF’s label is found, there is another cell with a higher weight), then IF asks a reformulation CR, proposing this other label.

IG now reacts as follows to clarification questions: In reply to a reject-CR (“huh?”), IG will use the second best label for the object that is to be

named, if there is one; otherwise, it will just repeat the previous label. In reply to a reformulation-CR, it will say “yes” if the proposed alternative label is within a certain range of the original label (for object o), “no” otherwise. Back to IF: on hearing “yes”, IF will solve with the mapping it used for the reformulation. If IG said “no”, IF picks the best interpretation for the original label (even though for that object IF thinks there is a better name). If IG offers a reformulation, IF tries to find the object that is best for both labels, but if there is none, he picks the first one.

Rewarding the players after an episode is somewhat more involved in this variant of the game. We now potentially have several different cells involved in one episode (for making the initial proposal, for formulating a CR, for answering it). Rewards are now computed according to formula (1), where c_n denotes the cell used at step n , counting backwards and starting with 0, so step 0 is the last utterance of the player being rewarded. We use hyperbolic discounting on the reward, with some discount factor δ ; the rationale here is that the exchange towards the end of the episode is more important and impresses itself more on the players.

$$(1) \quad c_n \leftarrow c_n + \frac{1}{(1+\delta*n)}$$

We again follow the principle that IG is always rewarded, and IG’s choices—they are now transparent for IF, since IF knows what IG’s goal was—are also reinforced for IF. IF’s own choices are rewarded or punished depending on the success of the episode.

Before we turn to the experiments we performed with simulations of these games, two more remarks on the second variant of the game. First, what is the idea behind these rules (which seem positively baroque compared to the austere rules of Game Theory games)? They are meant to be a relatively plausible model of how clarification actually works (Purver et al., 2001; Purver, 2004; Schlangen, 2004).² They have the effect of letting the participants more quickly explore their semantic maps, as we will see below. (This perhaps incidentally offers corroborating evidence for the thesis that clarification behaviour supports language construction (Ginzburg and Macura, 2006).)

²Of course, ideally, these rules would be learned as well; though not in the same game. We are assuming adult-level language competence here. Learning clarification strategies, perhaps modelling (Matthews et al., 2007), is an orthogonal problem.

As a second remark, the reader is advised to note that the players at no point keep a model of their partner. They always only keep their own map; a feature of this model that brings it closer to the interactive alignment model than to the lexical pacts model.

We now turn to simulation experiments with this game and their results.³

3.3 Experiments, and Results

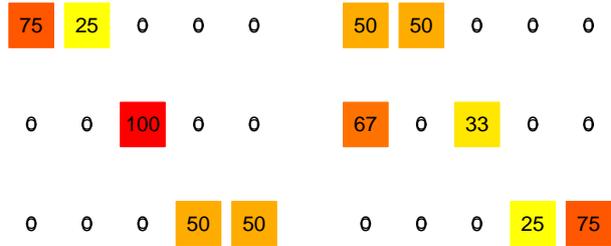


Figure 4: Two maps with less ambiguity. Three objects, five labels.

We implemented the variants of the game in a computer program and ran experiments with two different map sets, the one from Figure 2, and a smaller map set with less ambiguity (Figure 4). As evaluation measures we use *task success*, and *average map distance*, which we define as the mean root squared distance between cell values: $ad(\mathcal{M}^\alpha, \mathcal{M}^\beta) = \sum_i \sum_k \sqrt{(\mathcal{M}_{i,k}^\alpha - \mathcal{M}_{i,k}^\beta)^2} * \frac{1}{i*k}$

Figure 5 shows a plot of the average map distance in Experiment 1 (with the maps from Figure 2), for the variant with clarification requests (solid lines) and that without (dashed lines). Superimposed are the successful episodes (green circles) and the unsuccessful ones (blue boxes). We see that in both variants the distance between maps decreases steadily (faster for the variant without CRs), and after some episodes, successes become much more likely than failures (much later for the variant without CRs). What may the reason be for the differences? In the variant with CRs it takes the players longer to align their maps, because there is more room for misunderstanding, as there is more exploration. The flipside of this is that they are faster to explore the whole map, and so have heard all relevant labels used earlier. In the variant without CRs, each episode only teaches them

³We would like to, but cannot at the moment offer a formal analysis of properties of this learning method, and so can only look at how it fares on the task.

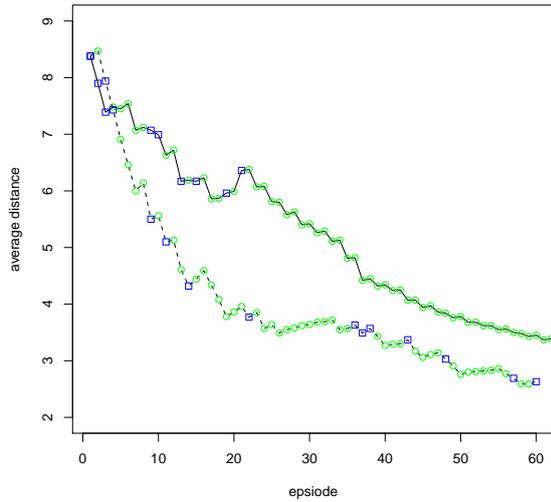


Figure 5: Results of Experiment 1

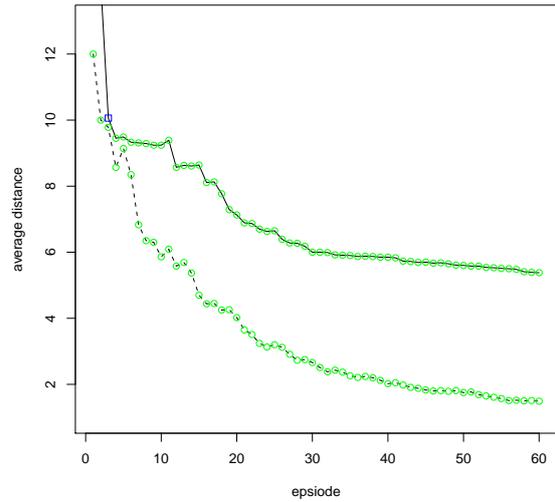


Figure 7: Results of Experiment 2

about one mapping, and so it can happen even after many episodes that they hit a pairing that they haven't sufficiently learned.



Figure 6: The maps after 20 episodes

Figure 6 shows the two maps from Figure 2 after 20 episodes. The figure shows nicely how similar the maps have become. Note that the players have aligned on player A's map; this seems to be due simply to the fact that A was IG first, and so had a slight, but decisive, advantage in spreading her preferences.

Figure 7 shows the results for Experiment 2 with the simpler maps. We see the same tendency for the variant with CR to decrease less quickly; as the maps are so simple (but arguably much

more representative of normal Matching Game settings), the partners very quickly stop making any mistakes.

We take these results as evidence that the learning mechanism does indeed lead to “lexical entrainment” between the partners: they come to be much more likely to re-use successful labels, and this indeed increases their success. So, from the looks of it they do indeed form “conceptual pacts”, but without any model of their partner, just by accidentally becoming more like each other.

“Is this not just priming, then?”, one may be tempted to ask. If we understand the idea behind priming correctly, then that is a mechanism driven by occurrence, not by success. In our model, reward is given (to IF) only when success is reached. In that sense, our model is goal-driven, whereas priming is not. Alternatively, our model may be seen as a way to spell out what priming is meant to achieve—we are not aware of much work on computational models of priming with respect to this phenomenon. (There is work by Reitter (2008), but that is modelling corpora and not providing simulations. More closely related is (Buschmeier et al., 2009), where alignment and priming is modelled by activation functions. However, these do not take into account the role of clarification requests and of success in reaching a communicative goal.)

In any case, however, what we have done so far only provides half of the story. What will happen when a player encounter a new, naive agent,

or when a previous partner starts to ‘misbehave’? This part we have not explored with experiments yet. As the setup is at the moment, players will react in the same way to a new partner using different terms and to their old partner suddenly using different terms. This contradicts the findings of Metzinger and Brennan (2003) discussed above. We discuss in the next section how the model could be made to account for this phenomenon as well.

4 Necessary Extensions

We list in this section extensions to the model that are necessary to bring it towards more fully capturing the phenomenon (and, more ambitiously, to making sense of the label *practice*).

The first of those extensions will need to capture the partner-specificity observed by Metzinger and Brennan (2003) (reactions to previous partner switching terms) and Branigan et al. (forthcoming) (different strength of alignment depending on beliefs about capabilities of (artificial) partner). At the moment, we can only outline the likely shape of this part of the model, and mention some requirements that we would like to see met.

First, we widen the scope of the discussion. A different view of the maps from the previous section: We have introduced them as linking objects and labels, but, viewed in a more general way, what they link is a way of performing an act of referring (or an act of understanding) to the goal of making the partner pick out an object (or the goal of picking out the object the partner wants me to pick out). The maps were a special case of what more generally we will call a *policy*, which is, then, a structure that links goals and actions.

The idea now is that what is being modified during the course of an activity is a “local” policy, particularised to the current situation. The policy from which one started remains stable, and a new copy is branched off for the current conversation. Some process of abstraction over situations then needs to categorise and generalise: will the local policy be of use again, do the changes need to merged with the starting policy? In such inconsequential situations as the matching game, the policy (the map) might be associated in memory with a particular speaker, and then slowly be forgotten (or rather, regress to the base line map). But in more weighty situations, a policy will become associated with a group of partners, or a type of situation; and conversely, a group is formed through the association with a policy. Membership in the

group (being in a situation / doing a certain activity) then entails following this policy, this practice; and deviating from the practice rules out membership, or leads to irritation when done by a member of the group. (I.e., practices are a form of situated normativity, (Rietveld, 2008).) This is what a practice is, then: a policy that holds in certain situations, for certain groups of agents.

Returning to the more concrete subject matter of this paper: how does this account for the phenomenon of partner-specificity of lexical entrainment? It can explain the effect that “breaking the pact” has, the irritation that the partner exhibits at this, if we assume that breaking the pact simply is not behaving in the way that the chosen policy for this situation (which includes this partner) predicts. If we assume as a further element that the mechanism for copying and adapting maps is sensitive to aspects of the situation, then we might have a handle on the fact that beliefs about the partner can influence alignment with it. The results of (Branigan et al., forthcoming) might then be explained by differences in the assumptions of group membership: the computer believed to be more advanced is assumed to conform stronger to normal practices, whereas the simpler one does not trigger strong presuppositions. Working out these rather general ideas will have to remain for future work, however.

We close this section by noting that the general ideas sketched here are far from being new. The godfather of this line of thinking of course is Wittgenstein (1953) with the notion of *language game*. The way we have restricted policies to situation types (or rather, have stipulated that this should be done) could be seen as one way of spelling out this notion. The particular line on Wittgensteinian thought I’m following here is further represented by (Levinson, 1992) and (Bourdieu, 1990) (from whom the label *practice* is taken). More recently, the idea of conversations establishing micro-languages explored by Larsson (2008) is very closely related; our understanding is that the ideas sketched here should be complementary to this approach. The precise connections to this and other related work, however, need to be worked out more clearly in future work.

5 Conclusions

We have provided a formal model of a game which has often been used to study the phenomenon of lexical entrainment (Garrod and Anderson, 1987).

We have used this model to make precise some implications of the lexical pact model (Brennan and Clark, 1996), and the interactive alignment model (Pickering and Garrod, 2004), and our own practice-based proposal, which we have also tested in simulation, showing that it can model the process of aligning on language use in this very simple game.

What we have provided here is quite clearly only a first sketch of such a practice-based account of lexical entrainment (and, much more so, a first sketch of a more general theory of the role of practices in conversation). It is our hope, however, that our formalisation of the problem as outcome of a kind of a signalling game already is a useful contribution, and helps to better understand what the debate actually is about—and that our sketch of a practice-based account has at least succeeded in positioning it as an alternative that would be promising to work out further.

Acknowledgments This work was funded by DFG Grant SCHL845/3-1, Emmy Noether Programme.

References

- Raffaele Argiento, Robin Pemantle, Brian Skyrms, and Stanislav Volkov. 2009. Learning to signal: analysis of a micro-level reinforcement model. *Stochastic Processes and their Applications*, 119(2):373–390.
- Pierre Bourdieu. 1990. *The Logic of Practice*. Stanford University Press, Stanford, CA, USA. transl. R. Nice; orig. *Le sens pratique*, Les Éditions de Minuit, 1980.
- Holly P. Branigan, Martin J. Pickering, J. Pearson, and J.F. McLean. *forthcoming*. Linguistic alignment between humans and computers. *Journal of Pragmatics*.
- Susan E. Brennan and Herbert H. Clark. 1996. Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6):1482–1493.
- Hendrik Buschmeier, Kerstin Bergmann, and Stefan Kopp. 2009. An alignment-capable microplanner for natural language generation. In *Proceedings of the 12th European Workshop on Natural Language Generation*, pages 82–89, Athens, Greece.
- Herbert H. Clark and Catherine R. Marshall. 1978. Reference diaries. In D. L. Waltz, editor, *Theoretical Issues in Natural Language Processing (Vol. 2)*, pages 57–63. ACM, New York, USA.
- Herbert H. Clark and Deanna Wilkes-Gibbs. 1986. Referring as a collaborative process. *Cognition*, 22:1–39.
- Melissa C. Duff, Julie Hengst, Daniel Tranel, and Neal J. Cohen. 2006. Development of shared information in communication despite hippocampal amnesia. *Nature Neuroscience*, 9(1):140–146.
- Simon Garrod and Anthony Anderson. 1987. Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. *Cognition*, 27:181–218.
- Simon Garrod and Martin J. Pickering. 2007. Automaticity in language production in monologue and dialogue. In A.S. Meyer, L.R. Wheeldon, and A. Krott, editors, *Automaticity and control in language processing*, pages 1–21. Psychology Press, Hove, UK.
- Jonathan Ginzburg and Zoran Macura. 2006. Lexical acquisition with and without metacommunication. In *The Emergence of Communication and Language*, pages 287–301. Springer, Heidelberg, Germany.
- Ellen A. Isaacs and Herbert H. Clark. 1987. References in conversation between experts and novices. *Journal of Experimental Psychology: General*, 116(1):26–37.
- Robert M. Krauss and Sidney Weinheimer. 1964. Changes in reference phrases as a function of frequency of usage in social interaction: A preliminary study. *Psychonomic Science*, 1:266–278.
- 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*. College Publications, London, UK.
- Stephen C. Levinson. 1992. Activity types and language. In P. Drew and J. Heritage, editors, *Talk at work: Interaction in Institutional Settings*, pages 66–100. Cambridge University Press, Cambridge, UK.
- Danielle Matthews, Elena Lieven, and Michael Tomasello. 2007. How toddlers and preschoolers learn to uniquely identify referents for others: A training study. *Child Development*, 78:1744–1759.
- Charles Metzger and Susan E. Brennan. 2003. When conceptual pacts are broken: Partner-specific effects on the comprehension of referring expressions. *Journal of Memory and Language*, 49:201–213.
- Sherry Ortner. 1984. Theory in anthropology since the sixties. *Comparative Studies in Society and History*, 26(1):126–166.
- Martin Pickering and Simon Garrod. 2004. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27:169–226.
- Matthew Purver, Jonathan Ginzburg, and Patrick Healey. 2001. On the means for clarification in dialogue. In *Proceedings of the 2nd SIGdial Workshop on Discourse and Dialogue*, Aalborg, Denmark.
- Matthew Purver. 2004. *The Theory and Use of Clarification Requests in Dialogue*. Ph.D. thesis, King’s College, University of London, London, UK, August.
- David Reitter. 2008. *Context Effects in Language Production: Models of Syntactic Priming in Dialogue Corpora*. Ph.D. thesis, School of Informatics, Edinburgh, UK.
- Erik Rietveld. 2008. Situated normativity: The normative aspect of embodied cognition in unreflective action. *Mind*, 117(468):973–1001.
- David Schlangen. 2004. Causes and strategies for requesting clarification in dialogue. In *Proceedings of the 5th Workshop of the ACL SIG on Discourse and Dialogue*, Boston, USA, April.
- Yoav Shoham and Kevin Leyton-Brown. 2009. *Multia-gent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, Cambridge, UK.
- Ludwig Wittgenstein. 1953. *Tractatus Logicus Philosophicus und Philosophische Untersuchungen*, volume 1 of *Werkausgabe*. Suhrkamp, Frankfurt am Main. this edition 1984.