

Reciprocal Learning via Dialogue Interaction: Challenges and Prospects

Raquel Fernández
University of Amsterdam
Netherlands

Staffan Larsson
Göteborg University
Sweden

Robin Cooper

Jonathan Ginzburg
Université Paris-Diderot
France

David Schlangen
Bielefeld University
Germany

Abstract

Humans learn to communicate with each other by engaging in *language coordination* during dialogue. In this position paper, we present the main ideas behind the challenge of language coordination in human-machine interaction. We review relevant empirical evidence and current approaches to learning conversational agents, and identify some of the problems that must be overcome for realising this challenge.

1 Introduction

Humans are remarkably flexible in their ability to cope with dynamic environments that require some form of learning. As social agents, humans typically learn by interacting and communicating with others by means of natural language. This is not surprising since language is essentially a device for achieving *informational coordination*, that is, a device for making sure that we share information sufficiently to be able to get things done together. Importantly, however, learning by talking also involves *language coordination*, that is, it involves not only learning about the subject matter we are talking about, but also learning about language itself – about which words we use to talk about a domain and what we mean by them. When we talk, we bring our own perspective to the conversation. Some of us may refer to a piece of music as “The Moonlight Sonata” but not recognize it as “Beethoven’s Op. 27, No. 2” and one person’s “comfort food” may be another’s “bland and uninteresting”. Despite these individual differences, research in psychology and cognitive science has shown that humans have a strong tendency to coordinate their language with their conversational partners: by participating in dialogues, speakers converge in the external formal features of their language (pronunciation, syntactic constructions, vocabulary) and crucially they also learn about each other’s underlying semantic distinctions (their ontologies).

Language coordination of this sort can be described as a case of *reciprocal learning* – a process whereby interacting agents learn to communicate with each other. Reciprocal learning should be distinguished from corpus-based learning, where samples are taken from a static body of examples and used to progressively adjust some internal representations in the learner. In reciprocal learning, agents incrementally learn

from each other. The coordinating process brought about by reciprocal learning is pervasive in human communication and a key aspect to achieve informational coordination. We thus argue that, if our goal is to create truly intelligent machines that can use language to collaborate with and learn from humans in a wide range of activities, we need to endow machines with the capability of language coordination. In this position paper, we present the main ideas behind the challenge of reciprocal learning for language coordination in human-machine interaction, identifying problems that must be overcome for realising this challenge. Our main claims are:

- humans learn to communicate with each other by learning *through* language and *about* language in dialogue interaction;
- state-of-the-art dialogue systems do not utilize incremental reciprocal learning: they are data-intensive while humans form hypotheses incrementally from single exposures to data (one-shot learning) and have techniques for correcting false hypotheses based on feedback;
- a bottleneck for designing conversational agents capable of reciprocal learning is the lack of a formal semantic theory of language coordination, which should be coupled with the right machine learning techniques.

In the next section, we give an overview of empirical findings related to language coordination. After that, in Section 3, we survey current approaches to conversational agents that attempt to integrate aspects of language coordination. Finally, in Section 4, we spell out the challenges of developing agents for natural dialogue interaction that can learn by talking.

2 Coordination and Learning in Dialogue

Psychologists and cognitive scientists working on linguistic interaction have long been aware that dialogue participants tend to adapt to each other in conversation. For instance, dialogue participants rapidly converge on the same vocabulary [Brennan, 1996], tend to use similar syntactic structures [Branigan *et al.*, 1995], adapt their pronunciation and speech rate to one another [Pardo, 2006], and even mimic their interlocutor’s gestures [Kimbara, 2006]. A number of researchers have also found experimental evidence that human users of computer dialogue systems adapt several features of their language to the productions of the system. [Branigan *et al.*,

2010] show that human users tend to align with the syntactic structures and the vocabulary used by a computer, while [Coulston *et al.*, 2002] found that children would adapt the amplitude of their speech with that of a spoken animated dialogue agent.

One of the most influential psycholinguistic theories put forward to explain such ubiquitous adaptation is the Interactive Alignment Model of [Pickering and Garrod, 2004]. This model assumes that, in dialogue, the linguistic representations employed by the interlocutors become aligned at many levels (phonological, lexical, syntactic) and that such alignment leads to coordination at the conceptual/semantic level. Pickering and Garrod assume alignment to be an automatic adaptation process, driven mostly by implicit priming mechanisms. In contrast, the seminal work of Clark and colleagues [Clark and Wilkes-Gibbs, 1986; Brennan and Clark, 1996] focuses on the explicit collaborative strategies used by speakers and hearers. This line of research sees dialogue as a form of joint action, where participants work together to make sure they understand each other by asking clarification questions, giving feedback, and establishing flexible “conceptual pacts” with specific dialogue partners.

Empirical evidence from research on first language acquisition by children is also very relevant to our concerns here. Recent research in this area [Clark, 2007; Saxton, 2000] emphasizes the role of dialogue interaction during learning, arguing that the language learning process does not depend solely on exposure to linguistic input, but it crucially relies on feedback given and received in interaction, including corrective feedback and clarification requests posed by the adults, which help the learners to identify and overcome particular problems and misalignments [Saxton *et al.*, 2005]. Here is an example of corrective feedback from Clark [2007]:

- (1) Naomi: mittens.
Father: gloves.
Naomi: gloves.
Father: when they have fingers in them they are called gloves and when the fingers are all put together they are called mittens.

Language acquisition may be regarded as a special case of the more general phenomenon of language coordination. One aspect that is special about language acquisition is that there is a clear asymmetry between the agents involved with respect to expertise in the language being acquired when a young child and an adult interact. However, the mechanisms for semantic coordination used in these situations seem similar to those which are used when competent adult language users coordinate their language. Competent agents do not need to share exactly the same linguistic resources (grammar, lexicon etc.) in order to be able to communicate. Our linguistic resources can change during the course of a single dialogue, for instance when we are confronted with a new word or an innovative use of a known word. The semantic learning that results from this process may be limited to a specific dialogue or a specific partner; it may become part of our long-term knowledge; or it may spread over a community and eventually become part of *the* language as it is represented in dictionaries.

Besides psycholinguistic research, corpus-based studies of

linguistic adaptation in dialogue also offer interesting insights. These studies use several measures to quantify the degree of alignment between dialogue participants and then use computational modelling techniques to reproduce it [Ritter *et al.*, 2006]. With regard to learning, for instance, [Ward and Litman, 2007] process a corpus of human-human tutoring dialogue between a teacher and a student and show that the alignment measures they develop are useful predictors of learning.

In summary, there is ample evidence that humans (adults and children) engage in language coordination in dialogue. They use both implicit mechanisms to align external features of their language, as proposed by Pickering and Garrod, and explicit collaborative strategies that lead to shared knowledge, as demonstrated by Clark and colleagues. We see these two perspectives on linguistic coordination as complementary. In our opinion, only conversational agents that are able to leverage implicit adaptation and intelligent use of dialogue strategies will be able to engage in effective reciprocal learning with humans.

3 Related Approaches

In this section, we review approaches to dialogue systems and other conversational agents that integrate aspects related to language coordination.

Several computational models of external linguistic adaptation to a human user have been created in recent years. For instance, text generation systems such as that of [Isard *et al.*, 2006] and [Walker *et al.*, 2007] adapt the surface linguistic form of the system’s productions to the individualities of a user. Both systems employ an ‘over-generation and rank’ approach consisting in generating a large number of alternative sentences that are then filtered according to individual preferences on the basis of training data collected prior to the system’s usage. Other existing systems for natural language generation focus on lexical alignment. For instance, the system described in [Janarthanam and Lemon, 2010] uses Reinforcement Learning to distinguish between expert and novice users of a broadband modem. Given a predefined set of synonym terms, it adapts its terminology for unknown users based on estimating their expertise as the dialogue progresses. [De Jong *et al.*, 2008] concentrate instead on style adaptation. They develop a virtual museum guide that is able to adapt to the level of formality and politeness of the user’s utterances by selecting a register for the system’s response from a predefined set of linguistic styles. Other researchers are starting to model not only alignment of speech, but gesture adaptation as well [Buschmeier *et al.*, 2010].

Although adaptive systems like the ones described above implement some aspects of language coordination, these systems do not exhibit semantic learning (informational learning) as a result of language coordination since learning is restricted to user adaptation in the form of surface linguistic features. In recent years, some researchers have investigated semantic reciprocal learning for communication using artificial multi-agent systems. Relevant work here includes research on category formation and emergent vocabularies, often inspired by work on grounded language acquisition and lan-

guage evolution [Briscoe, 2002; Steels and Belpaeme, 2005; Macura and Ginzburg, 2006], as well as work on formal ontologies in the “semantic web” [van Diggelen *et al.*, 2007]. This line of research is important because it has shown that learning agents equipped with interaction protocols and strategies can coordinate on linguistic form and meaning. However, since this strand of work focuses on *synthetic* language coordination, with small vocabularies consisting of simple strings of characters, it hardly gets us closer to the goal of creating agents that can use *natural* language to coordinate with humans and learn from them. In the next section, we argue that a detailed theory of natural language dynamics is required to be able to apply similar techniques to natural language coordination.

4 Towards Reciprocal Learning

As mentioned in the introduction, we argue that an important bottleneck for creating artificial agents capable of reciprocal learning is the lack of a formal semantic theory of language coordination. Formal semantics and pragmatics – the linguistic disciplines that study language interpretation and that thus underpin the development of systems for natural language understanding [Portner and Partee, 2002] – offer precise formal analyses of meaning that can lend themselves to computational implementation [Larsson and Traum, 2001]. However, research within these disciplines has not yet paid much attention to the dynamics of language itself: language is typically assumed to be a static entity that does not change during the course of a dialogue. The empirical evidence reviewed in Section 2 has thus not yet been integrated into formal theories of meaning. We believe that a primary concern of contemporary linguistic theories of formal semantics should be to come to grips with the experimental findings regarding language coordination and to develop approaches that can account for the reciprocal learning processes that occur in natural language dialogue. Only by reorienting the focus of formal and computational semantics in this fashion will we be able to achieve a precise and deep understanding of natural language coordination processes that can underpin the development of learning conversational agents.

To address these issues, we have started to develop formal approaches to coordination of meaning in dialogue [Larsson and Cooper, 2009; Cooper and Larsson, 2009; Larsson, 2010]. This work takes an Information State Update approach to dialogue management [Larsson and Traum, 2001] where dialogue moves related to semantic coordination (such as corrective feedback) are associated with updates to linguistic resources, including ontologies and other aspects of lexical meaning. We make use of Type Theory with Records [Cooper, 2005] – a logical framework with fine-grained feature structures that allows for the definition of similarity metrics on meanings and meaning modifications involving refinement and generalization.

While foundational theoretical work of this sort is interesting in its own right, we believe that it is also important as a basis for implementations of conversational agents capable of language coordination with humans. For it to be readily useful, however, it needs to be coupled with suitable

machine learning techniques – a research avenue we are currently investigating in ongoing work. Language coordination, regarded as a case of reciprocal learning, imposes certain constraints on suitable learning algorithms. We can identify at least the following:

- learning algorithms for language coordination need to be able to operate on fine-grained linguistic representations, as they should afford semantic learning and not only adaptation to external linguistic features;
- they should be highly incremental, allowing for rapid learning from single (or very few) exposures to data;
- they should be reciprocal and interactive, being compatible with both explicit and implicit dialogue strategies.

As pointed out in Section 3, the techniques currently used in dialogue systems do not meet all the above requirements: corpus-based learning methods are not genuinely interactive; current systems based on Reinforcement Learning are interactive but not incremental enough since they require off-line training on large amounts of data (typically simulated users); and systems that adapt to the user productions given a pre-defined set of alternatives do not involve semantic coordination. The learning techniques that have been used in artificial multi-agent systems are clearly interactive and reciprocal but, as mentioned at the end of Section 3, have not yet been applied to sophisticated natural language coordination, due in part to the foundational problems pointed out above.

5 Conclusions

This position paper has presented the main ideas behind the challenge of reciprocal learning for language coordination in human-machine interaction. A key way to move forward, we claim, is to make progress on the development of formal theories of language dynamics and coordination, and to combine the insights of these theories with suitable machine learning techniques for reciprocal, incremental, and interactive learning. From the dialogue systems research community, we are making first steps in this direction. A lot of interesting relevant research on learning is currently being explored within the robotics and the human-robot interaction communities. Learning techniques such as one-shot learning, bootstrap learning, and active learning [Fei-Fei *et al.*, 2006; Modayil and Kuipers, 2007; Chao *et al.*, 2010] might be able to open the door to developing conversational agents whose learning capabilities meet the requirements of language coordination. An important challenge is to adapt learning methods that have not been designed to deal with the intricacies of natural language to the problem of language learning via dialogue interaction. We believe that much is to be achieved from exchange between these different research communities.

Acknowledgements

Raquel Fernández acknowledges support by NWO (VENI grant) and David Schlangen by DFG (Emmy Noether grant).

References

- [Branigan *et al.*, 1995] H.P. Branigan, M.J. Pickering, S.P. Liv-
ersedge, A.J. Stewart, and T.P. Urbach. Syntactic priming: Investigating the mental representation of language. *Journal of Psycholinguistic Research*, 24(6):489–506, 1995.
- [Branigan *et al.*, 2010] H.P. Branigan, M.J. Pickering, J. Pearson, and J.F. McLean. Linguistic alignment between people and computers. *Journal of Pragmatics*, 42(9):2355–2368, 2010.
- [Brennan and Clark, 1996] S. Brennan and H. Clark. Conceptual Pacts and Lexical Choice in Conversation. *Journal of Experimental Psychology*, 22(6):1482–1493, 1996.
- [Brennan, 1996] S.E. Brennan. Lexical entrainment in spontaneous dialog. In *Proc. of the International Symposium on Spoken Dialogue*, pages 41–44, 1996.
- [Briscoe, 2002] E. J. Briscoe, editor. *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge University Press, 2002.
- [Buschmeier *et al.*, 2010] H. Buschmeier, K. Bergmann, and S. Kopp. Adaptive expressiveness – Virtual conversational agents that can align to their interaction partner. In *Proc. of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 91–98, Toronto, Canada, 2010.
- [Chao *et al.*, 2010] C. Chao, M. Cakmak, and A.L. Thomaz. Transparent active learning for robots. In *Proc. of the 5th International Conference on Human-Robot Interaction*, pages 317–324, 2010.
- [Clark and Wilkes-Gibbs, 1986] H. Clark and D. Wilkes-Gibbs. Referring as a collaborative process. *Cognition*, 22:1–39, 1986.
- [Clark, 2007] E.V. Clark. Young children’s uptake of new words in conversation. *Language in Society*, 36(02):157–182, 2007.
- [Cooper and Larsson, 2009] Robin Cooper and Staffan Larsson. Compositional and ontological semantics in learning from corrective feedback and explicit definition. In *Proc. of SemDial 2009: 13th Workshop on the Semantics and Pragmatics of Dialogue*, pages 59–66, 2009.
- [Cooper, 2005] Robin Cooper. Austinian truth, attitudes and type theory. *Research on Language and Computation*, 3:333–362, 2005.
- [Coulston *et al.*, 2002] R. Coulston, S. Oviatt, and C. Darves. Amplitude convergence in childrens conversational speech with animated personas. In *Proc. of the 7th International Conference on Spoken Language Processing*, volume 4, pages 2689–2692, 2002.
- [De Jong *et al.*, 2008] M. De Jong, M. Theune, and D. Hofs. Politeness and alignment in dialogues with a virtual guide. In *Proc. of the 7th international joint conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 207–214, 2008.
- [Fei-Fei *et al.*, 2006] L. Fei-Fei, R. Fergus, and P. Perona. One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 594–611, 2006.
- [Isard *et al.*, 2006] A. Isard, C. Brockmann, and J. Oberlander. Individuality and alignment in generated dialogues. In *Proc. of the 4th International Natural Language Generation Conference*, pages 25–32, 2006.
- [Janarthanam and Lemon, 2010] S. Janarthanam and O. Lemon. Learning to adapt to unknown users: Referring expression generation in spoken dialogue systems. In *Proc. of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 69–78, 2010.
- [Kimbara, 2006] I. Kimbara. On gestural mimicry. *Gesture*, 6(1):39–61, 2006.
- [Larsson and Cooper, 2009] S. Larsson and R. Cooper. Towards a formal view of corrective feedback. In *Proc. of the Workshop on Cognitive Aspects of Computational Language Acquisition*, pages 1–9, 2009.
- [Larsson and Traum, 2001] S. Larsson and D. Traum. Information state and dialogue management in the TRINDI dialogue move engine toolkit. *Natural Language Engineering*, 6(3&4):323–340, 2001.
- [Larsson, 2010] Staffan Larsson. Accommodating innovative meaning in dialogue. In *Proc. of SemDial 2010, 14th Workshop on the Semantics and Pragmatics of Dialogue*, pages 83–90, 2010.
- [Macura and Ginzburg, 2006] Z. Macura and J. Ginzburg. Lexicon convergence in a population with and without metacommunication. In *Symbol Grounding and Beyond: Proc. of the 3rd International Workshop on the Emergence and Evolution of Linguistic Communication*, pages 100–112, 2006.
- [Modayil and Kuipers, 2007] J. Modayil and B. Kuipers. Autonomous development of a grounded object ontology by a learning robot. In *Proc. of the 22nd Conference on Artificial Intelligence (AAAI-07)*, volume 22, pages 1095–1101, 2007.
- [Pardo, 2006] J.S. Pardo. On phonetic convergence during conversational interaction. *The Journal of the Acoustical Society of America*, 119:2382–2393, 2006.
- [Pickering and Garrod, 2004] M.J. Pickering and S. Garrod. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27(02):169–190, 2004.
- [Portner and Partee, 2002] P Portner and B. H. Partee, editors. *Formal Semantics: The Essential Readings*. Malden, Mass., 2002.
- [Reitter *et al.*, 2006] D. Reitter, F. Keller, and J.D. Moore. Computational modelling of structural priming in dialogue. In *Proc. of the Human Language Technology Conference of the NAACL, Companion Volume*, pages 121–124, 2006.
- [Saxton *et al.*, 2005] M. Saxton, C. Houston-Price, and N. Dawson. The prompt hypothesis: clarification requests as corrective input for grammatical errors. *Applied Psycholinguistics*, 26(03):393–414, 2005.
- [Saxton, 2000] M. Saxton. Negative evidence and negative feedback: Immediate effects on the grammaticality of child speech. *First Language*, 20(60):221, 2000.
- [Steels and Belpaeme, 2005] Luc Steels and Tony Belpaeme. Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, 28(4):469–89, 2005.
- [van Diggelen *et al.*, 2007] Jurriaan van Diggelen, Robbert-Jan Beun, Frank Dignum, Rogier M. van Eijk, and John-Jules Meyer. Ontology negotiation in heterogeneous multi-agent systems: The anemone system. *Applied Ontology*, 2:267–303, 2007.
- [Walker *et al.*, 2007] M. Walker, A. Stent, F. Mairesse, and R. Prasad. Individual and domain adaptation in sentence planning for dialogue. *Journal of Artificial Intelligence Research*, 30(1):413–456, 2007.
- [Ward and Litman, 2007] A. Ward and D. Litman. Dialog convergence and learning. In *Proc. of the 2007 conference on Artificial Intelligence in Education: Building Technology Rich Learning Contexts That Work*, pages 262–269. IOS Press, 2007.