Studying the Co-Construction of Interaction Protocols in Collaborative Tasks with Humans

Anna-Lisa Vollmer*, Jonathan Grizou*, Manuel Lopes*, Katharina Rohlfing† and Pierre-Yves Oudeyer*
*Flowers Team, INRIA / ENSTA-Paristech, France
Email: {anna-lisa.vollmer, jonathan.grizou, manuel.lopes, pierre-yves.oudeyer}@inria.fr
†CITEC, Emergentist Semantics, Bielefeld University, Germany
Email: kjr@uni-bielefeld.de

***Abstract*—In interaction, humans align and effortlessly create common ground in communication, allowing efficient collaboration in widely diverse contexts. Robots are still far away from being able to adapt in such a flexible manner with non-expert humans to complete collaborative tasks. Challenges include the capability to understand unknown feedback or guidance signals, to make sense of what they refer to depending on their timing and context, and to agree on how to organize the interaction into roles and turns. As a first step in approaching this issue, we investigate here the processes used by humans to negotiate a protocol of interaction when they do not already share one. We introduce a new experimental setup, where two humans have to collaborate to solve a task. The channels of communication they can use are constrained and force them to invent and agree on a shared interaction protocol in order to solve the task. These constraints allow us to analyze how a communication protocol is progressively established through the interplay and history of individual actions. We report preliminary results obtained from a pilot study, and discuss how the understanding of strategies used by humans could be useful to achieve more flexible HRI.**

I. INTRODUCTION

In interaction, humans align and effortlessly, maybe even automatically, create common ground in communication [1], [2]. For this, they dispose of an immense amount of shared information. They make use of frames and interaction protocols established in the history of interaction. Frames create a common ground about the purpose of the interaction [3], [4] and include "predictable, recurrent interactive structures" ([5], p. 171). Frames thus provide interactants with guidelines about how to behave (a protocol for interaction) and also help interactants to understand the communicative intentions of their interaction partner. Interaction protocols comprise basic behavioral patterns like roles, turns, timing, and exchange mechanisms. We aim at investigating how these interaction protocols emerge, because it would shed light on the basic mechanisms underlying interaction and inform us about what are the main issues in building robots capable of a similar interactional flexibility to the one humans possess. We are for instance interested in what kind of strategies humans use to align and what kind of meanings of social signals they converge to. Whereas, in the long run, the obtained findings could be used as priors for a robotic system interacting with humans, in an initial step, we first need to conduct research into how interaction protocols are negotiated in human-human interaction.

We designed an experimental setup with which we aim at investigating the processes used by humans to negotiate a protocol of interaction, when they do not already share one. In the current paper, we present and justify the method used and mention the very first preliminary results obtained from a pilot study employing the setup.

Humans and robots view the world differently, so if we want to transfer our results to human-robot interaction (HRI), we should not assume that in the interactions we want to investigate, the partners see the world/interaction in the same way. To investigate the process of negotiating an interaction protocol, we thus consider a setup of a joint construction task in which participants assume asymmetric roles: the role of a builder and the role of an architect. With building blocks, the builder should assemble a target structure which is unknown to him/her but which the architect knows. This collaborative construction task with a joint goal renders the communication between participants indispensable and thus the game is not solvable by either one of the participants alone, e.g. with mere exploration. Thus, failing to complete the game successfully is equivalent to failing to communicate successfully. Communication is not face-to-face but channels are restricted, so that it is not possible for participants to communicate via familiar verbal or non-verbal communication channels, as for example speech or gestures. At the same time, the setup does not constrain all aspects of communication and thereby gives participants much freedom with respect to some features, including timing and rhythm or possible meanings (e.g. of button presses). The setup does not impose a predefined sequence of interaction upon participants, as it is often done in HRI scenarios [6], but still benefits from a laboratory setting in which we do not need to take the full complexity of natural social interaction into account. With the aim to simulate the sending of signals to an interaction partner who does not have the same perceptual capabilities – similar to an interaction with a robot – in our study the architect does not know how exactly his/her signals are perceived by the builder. For the successful completion of the thus highly challenging joint task of the game, both participants have to learn how to interact with each other.

The main contribution of the current paper is the presentation of the novel experimental method of our study. We would like to demonstrate that it allows studying important questions for the understanding of human negotiation of interaction protocols in joint construction tasks and that these questions are very important for HRI in the long-term. Also, we report preliminary results of a first pilot run of the study. We first briefly discuss related work, then present our method and the preliminary results of the pilot study and close with a conclusion in which we suggest first implications of these results for HRI.

II. RELATED WORK

In this section, we will briefly discuss the most relevant related work, which lie in the field of experimental semiotics, and highlight the novelty of our proposed experimental method. The field of experimental semiotics studies the emergence and evolution of communication systems [7]. Here, instead of computer simulations as conducted by others (see [8], [9]), controlled experiments in laboratory settings are designed to observe communication between participants who perform joint tasks. For instance, Galantucci et al. showed that pairs of participants performing a joint task could coordinate their behavior by agreeing on a symbol system [10].

Most experimental semiotics studies developed to study joint action involve symmetric communication (cf. [11]). Two studies which do consider asymmetric communication are the studies conducted by de Ruiter et al. [12] and Griffiths et al. [13].

In their score- and round-based Tacit Communication Game, de Ruiter et al. investigated the cognitive processes responsible for the development and the recognition of new conventions by looking at reaction time. In a 3-by-3 grid world, two participants each manipulate a shape. For both of the shapes, the “sender” sees a target configuration. He/she first has to communicate the other player’s target configuration to the other player, the “receiver”, and second has to bring the own shape to his/her own respective target. De Ruiter et al. found that participants succeeded 83% of the time and that the timing of movements is used to indicate a position. When comparing success rates for when the sender saw versus did not see the receiver’s moves, the authors found that the game involves bidirectional communication and receiving information about the other player facilitates communication. The harder the communicative problem was, the more planning time was needed by both participants.

The setup of the study conducted by Griffiths et al. [13] is more directly related to our setup. It is based on the alien world game setup by Morlino et al., in which in a square world shown on a computer screen, positions (left or right) and movements (shake horizontally or shake vertically) of 16 objects have to be explored via a mouse to maximize a score [14]. It investigates the learning of categories, so the objects belonged to four categories which were defined by certain properties of the objects. Each category was associated with a target manipulation, i.e. shape and weight determined where an object should be positioned and how it should be moved. In the work by Griffiths et al. the learner could realize this task with the help of information given by a tutor who had prior knowledge about the categories the learner should explore. For this alteration, two players played the originally single player game simultaneously in separate rooms over a network connection. The computer screens in this setup additionally showed six buttons underneath the grid world. The tutor’s communication to the learner consisted of the pressing and releasing of these six buttons using a keyboard. This was the only action the tutor could perform on the world. The authors found that tutors most commonly used yes, no and concrete instructions, such as place and shake, as signals to the learners. Negative feedback was given least often and its amount correlated with task failure. Learners who ignored less signals performed the task better.

The main, very important difference between the two asymmetric setups described above and our setup concerns the very nature of the task. Whereas in Griffiths et al.’s study the task is solvable with mere exploration, in our setup the input of the architect is essential. The latter is also the case in the study by de Ruiter et al., but in our setup no score is displayed to either of the players who in our case are not separate learner and tutor, or receiver and sender, but they solve the task together assuming the roles of a builder and an architect. Correspondingly, in our setup, the game does not include multiple episodes or rounds but it is continuous with the builder deciding when the task is completed and the game ends. The game of the study by de Ruiter et al. is based on fixed turns which is not the case with our game, where participants can act simultaneously and react directly upon each others conduct.

By designing a continuous game without displaying a score, interaction remains natural (i.e., free) to a high degree.

Another important difference which makes our setup novel regards the restriction of communicative channels. In contrast to the other two works, in our setup, the architect is not aware of how his/her actions are presented on the builder side and how they will be perceived.

III. THE CO-CONSTRUCTION GAME

With the aim that improving our understanding of how humans negotiate protocols of interaction could provide hints on how robots could do it also, we designed a new experimental setup which allows constraining the communication channels between two partners in asymmetric roles who should collaborate in order to achieve a joint construction task. This section describes the details of the experimental setup, the participants we recruited for the current pilot study, and the protocol used for running the study.

A. Setup

Figure 1 gives an overview of the experimental setup which considers an architect and a builder that are each seated at a table in front of a computer screen in two separate rooms and can neither hear nor see each other.

The builder is equipped with a set of building blocks, in our case with 12 primary-colored Mega Bloks® toy blocks differing in shape and color (see Figure 2(b)). The goal of the

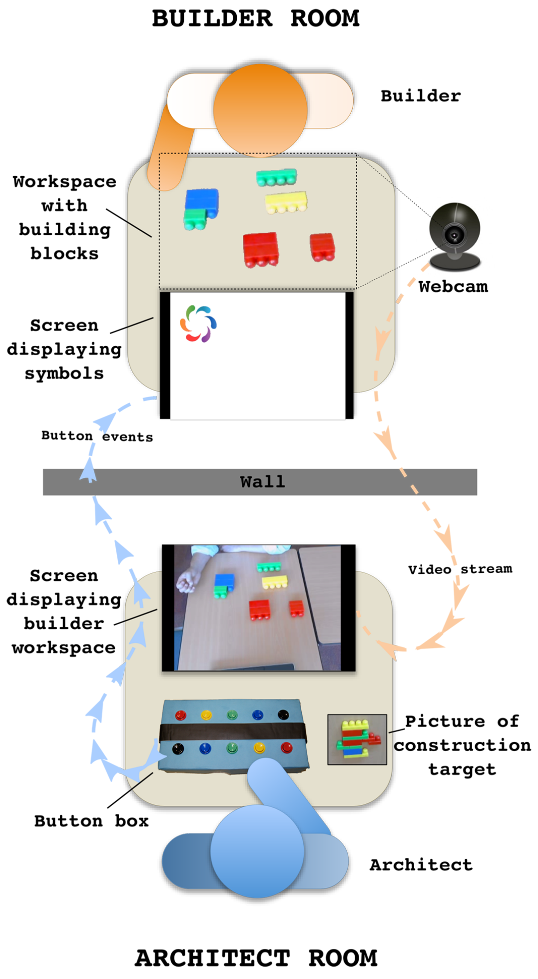

Fig. 1. Schematic view of our experimental setup. An architect (bottom) and a builder (top) should collaborate in order to build the construction target while located in different rooms. The architect has a picture of the targeted construction, while the builder has access to the construction blocks. The communication between them is restricted. The architect only sees a top view of the builder's workspace and can communicate with the builder only though the use of 10 buttons which, when pressed, display symbols on a screen on the builder side.

(a) The box and the buttons used as an interface for the architect to communicate with the builder.

(b) All toy blocks used in the collaborative construction task.

(c) Three examples of the 20 target structures presented to the architect.

Fig. 2.

Fig. 3. The ten signs displayed on the builder screen.

game is to assemble a specific construction yet unknown to the builder. As exemplified in Figure 2(c), a construction is a flat combination of several blocks at least linked to one another by one pad. It does not necessarily contain all available blocks.

The architect is given an image of the specific construction to be build and is told to guide the other player building it. A screen displays a live top view of the builder workspace. To communicate with the builder, the architect has access to a rudimentary interface made of 10 buttons, see Figure 2(a). Pressing a button displays a symbol on the screen located in the builder room. Each button is mapped to one of ten symbols and one of ten positions (two rows of five symbols) on the builder's screen, whereby the spacial organization of buttons differs from the spatial organization of displayed symbols. The mapping is randomized for each subject and fixed for the duration of one game. Figure 3 shows the different symbols.

B. *Participants*

We recruited 22 participants (19 m, 3 f) among students and staff at INRIA Bordeaux Sud-Ouest. Their age range was between 20 and 35 ($M = 25, SD = 3.91$) years. They played the collaborative game in pairs, where the two players in a pair were assigned randomly to the roles of a builder and an architect. Seven of the eleven pairs played the game together twice, such that each of the 14 participants involved assumed each role once. One second round of a dyad was excluded from the analyses, because the architect neglected the task instructions and altered the target structure during the game. This resulted in a total of 17 rounds.

C. *Procedure*

Participants were not given the chance to talk about the game before it began. Architect and builder were instructed about their respective roles separately in their respective rooms. We presented the architects with a set of 20 pictures of different constructions from which they chose one. The builder was informed about the constraint that applied on the construction, i.e. flat construction which does not necessarily contain all available blocks. Architect and builder were specifically told that the button positions did not directly map onto the symbols' positions displayed on the builder's screen, but that the mapping was fixed and arbitrary. Additionally, because the architect could see the hands of the builder during the game (see Figure 1), the builder is told to only use his/her hands to move blocks and not to use hand signs. In practice, this was well respected by participants.

The game was **not** preceded by any training sessions. We aimed at reducing the time between the instruction of the participants and the beginning of the game as much as possible, so that they did not have time to elaborate any concrete strategy before the game began.

Once the game started, we observed the behavior of the two players and asked them to speak aloud about the meaning associated to the symbols/buttons. The experimenters took notes on the participants' remarks. The experiment stopped only when the builder decided and told the experimenters that the structure he had build was correct.

IV. RESULTS

As stated before, the current pilot study serves as a proof of concept. We aimed at designing a setup allowing to study the processes involved in the formation of interaction protocols in asymmetric interaction with the particular constraint that the players could neither solve the task by themselves nor did they have access to any reward function.

Our first pilot study revealed a great potential in the use of our experimental method to study many aspects of communication relevant to HRI. With our setup, we will be able to study amongst others questions related to alignment, rhythm, contingency, and feedback, which have been in the focus of HRI research for some time [15]–[20].

Surprisingly, while the construction task in this setup seems really challenging on paper and participants thought they would never succeed, a majority of the architect-builder pairs succeeded on building the correct construction. We analyzed a total of 17 experiments, of which 13 were successful and 4 failed. The average duration of the runs was 18 minutes ($M = 18\ min, SD = 11\ min$) with a minimum of 7 minutes and a maximum of 45 minutes.

In what follows, we showcase very first results supporting our claim that our setup can be used to study the co-construction of meaning in restricted, asymmetric interaction. We will first show one run of the game in detail which should give the reader an idea about what happens in an interaction with our setup and the richness and aptness of the data to consider a variety of research questions. Then, we will continue with presenting preliminary results on the negotiation of signal meanings and with describing observations of the builder behavior. We will conclude with mentioning interesting, additional considerations which are beyond the scope of this paper but will be subject of future work.

A. One experiment in detail

Figure 4 brings together information about button presses (logs), their intended and interpreted meanings (found by the experimenters from their notes and observations of logs), and the builder's actions (builder video of the construction workspace) and makes clear the bidirectionality of the interaction. On the bottom of the figure, we see that the builder proposes blocks to the architect (blocks not belonging to the target structure in black, blocks belonging to it in gray) (cf. Subsection IV-C) and on the top we see how the architect responds to the builder's actions in terms of button presses and meanings. Additionally, we see how the builder interprets these signals of button presses which he/she perceives as symbols on a screen (middle timeline of button presses and meanings) and how these interpretations and believes in turn again influence what the builder does next.

With respect to the meanings of the button presses, we observe the change of button meanings over the course of the interaction. The exact points in time when meaning changes occur have been matched to the button presses by hand and is therefore approximated. While this may be a problem for detailed analyses on a micro level, it is of little importance for the macro analysis presented here. During the first 4 minutes, the architect changes the intended meanings of signals many times and these meanings were not aligned with the builder's interpretation of signals. At 4 minutes, the architect presses all buttons at once, seemingly attempting to ask the builder to clear his/her mind and start over again. Right after this *Reset* signal, the architect changes to one simple *yes/no* strategy using button 1 and 6. On the builder's end, this Reset signal is followed by a pause of actions which hints at a direct confusion. It is only at 12 minutes into the game that the builder fully understands the intended meaning of the architect's button presses and can start joining two blocks correctly (green graph on the bottom). The experiment continues with the builder suggesting new blocks (bottom - black and gray events) and positions for new blocks (bottom - red and green events) one at a time which are validated or invalidated by the architect. After 19 minutes, the architect presses again all the buttons but this time with the aim of informing the builder that the construction is complete. The builder ended the experiment at that time. The *End* signal was well interpreted by the builder as the interaction was going smoothly until that time and the few remaining blocks were rejected (bottom - black event at 19 min). The final construction was indeed the target one intended by the architect, hence resulting in a successful experiment.

Our setup allows studying the evolution of meanings associated to each button and put it in relation with the current context in the interaction. We find that the constraints inherent to our setup allow analyzing communication, especially the interplay of individual actions and their interactional history, as well as their concrete timing, while lowering interactional complexity and thereby reducing communicative noise.

B. Meanings

Architects and builders start the game without having agreed on specific meanings the buttons should convey. We start by studying the associated meanings obtained from our notes on signal meanings reported by builder and architect. They seemed to initially consider a large set of possible meanings, but, in the end, were able to agree primarily on only a limited number.

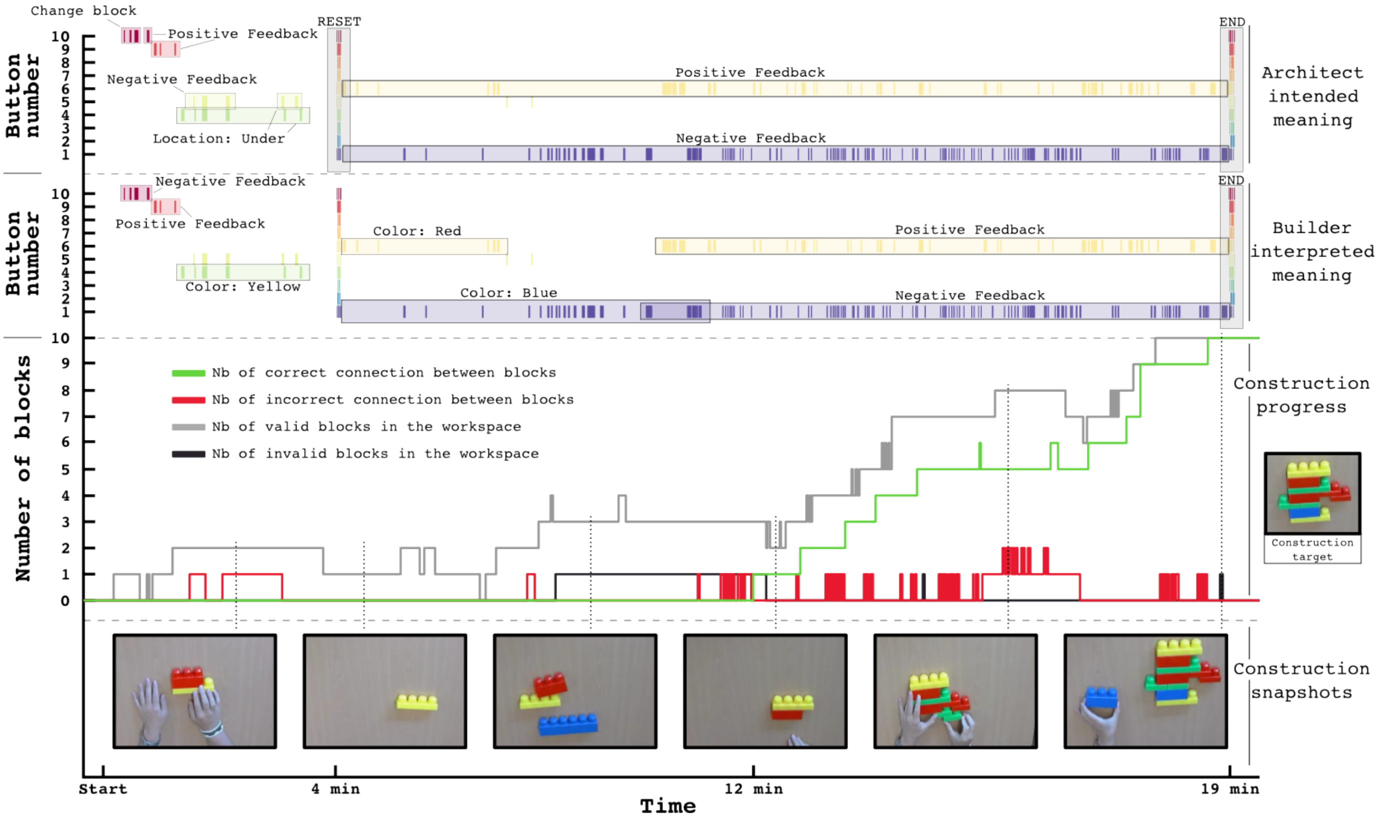

Fig. 4. Timeline for one experiment of an architect and a builder collaborating towards building the construction target (right hand side). The top and middle part show the timeline of button presses associated with the intended meaning from the architect (top) and the understood meaning from the builder (middle). There were 10 buttons, for which we logged all button presses for each experiment and here display all occurrences as colored dashes. Using the signal meanings participants reported during the game, the events are annotated with the meaning the architect intended or the builder understood. Events that are not annotated were not mentioned by the participants. At the bottom, the figure additionally visualizes the progress made by the builder in assembling the target structure and also shows incorrect block propositions, joining of incorrect blocks and mistakes. These events were annotated by hand using the video annotation tool ELAN developed by the Max Planck Institute for Psycholinguistics, The Language Archive, Nijmegen, The Netherlands [21]. A block proposition here started, when the transportation of the block towards the workspace ended and the block lay still on the table. It ended when the block was again picked up and subsequently removed from the workspace. These presentation events were classified into correct and incorrect propositions by determining whether the proposed block was part of the target structure. Equivalently, a joining event started, when two blocks were successfully joined at either a correct or incorrect position (again depending on whether the resulting configuration was part of the target structure). It ended before the first frame in which the two previously joined blocks were again pulled apart.

Types of Meanings When analyzing the notes on the participants' explanation of signal meanings (see Subsection III-C), we identified nine different categories of meanings:

1) **Positive Feedback**
2) **Negative Feedback**
3) **End**: The construction is finished.
4) **Reset**: Start over.
5) **Guidance**: Instruction on what to do. It includes *change, invert, revert, new block, continue, stack*.
6) **Color**: Reference to the color of a block. It includes *yellow, blue, red, green*.
7) **Size**: Reference to the size of a block. It includes *small, medium, big*.
8) **Location**: Reference to the location of a block. It includes *under, above, left, right*.
9) **Group**: Reference to a group of blocks. It includes *in, out, group_X*.

Importantly, those categories where **not** suggested to the participants beforehand, but only identified by us in a posteriori analysis.

For each experiment, we determined if the architect or the builder considered each type of meaning (see Figure 5). In every single experiment, positive and negative feedback were considered on both architect and builder side. The *End* meaning has been considered on both sides in 14 experiments. More concrete instructions such as *Guidance, Color, Size, or Location* were less often considered, especially by the builder.

This is in line with the findings in [13], where *"yes"* and *"no"* were also identified to be among the most common types of signal meanings.

Matching of meanings between architect and builder Knowing which meaning categories were considered by each of the participants does not tell us if a particular pair of players understood each other. We therefore compared the associated meanings reported by architect and builder for all signals. Similarly to [13], we then determined the number of signals that were understood, misinterpreted, or ignored. We

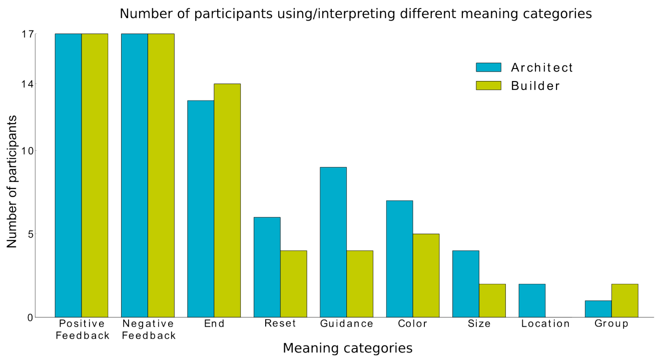

Fig. 5. Number of participants that used (architect) or interpreted (builder) signals as conveying different types of meaning. All participants considered positive and negative feedback types of meaning.

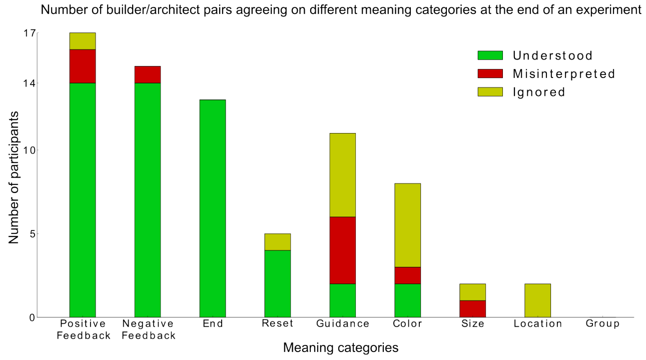

Fig. 7. Number of builder/architect pairs agreeing or disagreeing on different meaning categories at the end of an experiment.

define signals which were understood as signals where both architect and builder agree on a common meaning. For signals which were misinterpreted, the builder reported a different associated meaning than the one the architect intended. The signals which were mentioned by the architect, but not by the builder, we counted as ignored signals. We then averaged the results for successful and failed experiments, see figure 6. For successful experiments, the average number of signals understood is $M = 3.6, SD = 0.7$ which mostly corresponds to *Positive feedback, Negative feedback, End*, and occasionally *Reset* when needed (see Figure 7). Interestingly for failed experiments, this number drops to $M = 1.3, SD = 1.1$, with a larger amount of signals misinterpreted and ignored.

Fig. 6. Distribution of meaning categories that were understood, misinterpreted, and ignored by the builders. Data average across all builders for successful (blue) and failed (yellow) experiments.

It seems that even though many different signal meanings are initially considered by the architect, the players agree only on very few specific ones (positive feedback, negative feedback and End). The question of what are the main factors determining which meanings are considered by participants arises. This leads over to the next subsection in which we will consider the builder behavior to explore its role in which signal meanings are considered and in the ultimate outcome of the game.

C. Builder Strategies

For the builder, we aimed at identifying common actions across participants in an attempt to quantify the builders' strategies from the video data showing a top-down view of the workspace. What follows is a description of observations on the builders' behaviors.

We identified two main strategies the builders embarked on (for an overview see Table I). For these two strategies, the builders began by presenting only one block at a time. When they presented several blocks at once throughout the game, they did not seem to embark on a successful strategy.

The most common strategy for builders was to determine one correct brick at a time and to subsequently join it with the already assembled structure (see Figure 4). Figure 4 is a case example of one game/run of the study in which this strategy is used successfully. The builders in 12 (five first rounds and their five respective second rounds, one independent single first round, and one second round) of the 17 runs pursued the same strategy. Only one game (a first round with a successful corresponding second round) of these 12 failed.

The other strategy was to find all blocks belonging to the target structure. Blocks identified as correct were not joined right away, but in a first step all blocks belonging to the target structure were determined and were then subsequently joined one at a time in a second step. This strategy also involved the presentation of only one block at a time and was eventually pursued by two builders who both started out with a different strategy involving the presentation of multiple blocks. One builder initially tried to find which forms belonged to the target structure. Ultimately, he then identified all blocks belonging to the target structure by one at a time dividing all blocks into two groups. This builder played in a second round, for which in its corresponding first round multiple blocks were presented at a time by the builder and the game failed. Another builder at the beginning tried to elicit a label for either color or form from the architect. In this case, all blocks of one specific color or of one specific shape were presented at a time. This strategy was only pursued by one builder at the beginning of the game, but was not successful and then therefore discontinued in favor of the strategy of finding which blocks belong to the target structure. This builder played in a first round. In the corresponding second round, the builder embarked on the first strategy.

The remaining three builders (in three first rounds) also presented multiple blocks at once but the set of blocks presented did not have any common properties and seemed random. These builders did not have any apparent systematic strategy and their games did not come to a successful end.

Presentation of blocks	Strategy	Number of games	Successful	Failed
Present one block at a time	Find one block and join right away, repeat	12	11	1
	Find all blocks belonging to the structure, then start joining	2	2	0
Present multiple blocks at a time	No strategy	3	0	3

TABLE I

Taking a closer look at the four failed experiments, we find that in one of them, where the builder presented one block at a time, in the end the target construction was almost finished. Architect and builder understood each other, but what happened is that an early mistake in the position of one block was not signaled and thus corrected by the architect right away. He waited until the rest of the structure was completed and then tried to address the mistake by means of the introduction of a new signal. This new signal was interpreted by the builder as an End signal, leading to the end of the game with one block in a position next to the target one. With respect to the other failed experiments, the structure at the time the game ended was far from the target construction and there was no noticeable progress in all three cases.

Whereas, with the current data and analysis, we cannot yet draw any conclusions, still this observation suggests that the way the builders propose next steps and ask for information from the architect is important for the success of the game. Builders seem to build frames and create slots for the architect's input. These frames form the context which shapes the interpretation of the signals. This is similar to how in other cases of asymmetric or restricted communication, as for example in interactions with preverbal infants or in interactions with impaired persons, people provide frames to understand what their interaction partners with their different or limited conversational abilities want to communicate [22], [23].

D. *Additional Observations*

This subsection briefly indicates interesting, additional observations we made with our pilot study, as well as interesting considerations for future work.

First of all, we would like to state that the history of the interaction is crucial for understanding meanings. A person who has not witnessed the course of the interaction, is not able to fill in and complete the task without special instructions. We observe a phase of confusion and negotiation at the beginning of the interactions and after that a completion phase in which signal meanings have been constituted. The latter seems to be characterized by smooth, consistent patterns. In the initial phase of negotiation, we observed instances where the players adapted to their partners by changing the meaning of a button when they noticed the other player understands it differently (cf. Figure 4 in Subsection IV-A). There were for example cases in which the meaning of signals meaning yes and no were reversed. In contrast, we also observed that some players, both architects and builders, insisted on their strategies, even though the interaction with their respective partner did not work, i.e. they did not agree on any meaning and the task did not progress. Thus, there seem to be leaders and followers in terms of strategies, which could be personality-dependent, but could also manifest their ability to employ a theory of mind.

We also note that when builder and architect switched roles after a first round, their behaviors and performances were influenced (e.g., builder strategies were adopted across rounds). If a second round was systematically part of the experimental procedure, it would be interesting to see whether participants succeed faster in the second game they play with reversed roles and if they adopt similar strategies.

Another interesting aspect concerns timing, not only at which points in time the architect gives feedback and instructions, but also the interplay between the builder's and the architect's actions. The rhythm of the interaction partners' actions might be an important low-level feature in determining whether a certain signal means positive or negative feedback.

While the above points are highly relevant and worth investigating, their examination is beyond the scope of this paper and will be subject of future work.

V. CONCLUSION

We presented a new experimental method which allows studying important aspects of human communication with high relevance to HRI. In a first pilot study, we show that two players that never had a chance to interact by the means of a restricted interface before were able to communicate and act upon communicative acts whose meanings were never explicitly negotiated between interaction partners. From our preliminary results, we can already suggest first implications for HRI. These implications are twofold.

(a) Both builder and architect have preconceptions of what interaction frames the other player is likely to understand, trying to use or interpret signals with respect to those frames. The "feedback frame" seems the most commonly thought about and the easiest to understand in the context of our experiment.

Humans are capable of solving the kind of communication problem robots can have with humans. We have learned an interesting lesson: humans can solve such restricted asymmetric interaction problems by projecting the interaction into different common frames of interaction and selecting the one that is more coherent with the history of interaction. In this setup such an interaction frame provides information about many properties of the interaction including for example (1) a context (e.g. building something with a limited number of blocks and specific constraints (on a table, flat, ...)) and (2) a set of possible meanings (e.g. evaluations, guidance signals, references to colors or shapes), which we acquired from our experience interacting with others.

Based on such observations, by replacing the builder by an artificial agent (e.g. a robot), we can aim at constructing robots capable of learning a task from human instructions without programming it in advance to understand the human communicative acts and without preprogramming a specific rigid interaction protocol. To do so, we should equip our robot with a set of common interaction frames on which it can rely to find one that is coherent with the interaction history. We have begun to explore this direction using a pick and place experiment where a robot is instructed to reach a specific configuration of objects using raw vocal utterances whose mapping to their associated meaning is unknown at the beginning [24]. We further extended this work to brain-computer interaction (BCI) scenarios enabling calibration-free BCI control of an artificial agent in a reaching task [25].

(b) The builder's actions play an important role in the understanding of signals. Meaning is co-constructed by the interaction partners. With his/her propositions of blocks and positions, the builder provides frames in which he/she creates slots for the architect to provide information. And thus the builder's frames constitute the meaning of the architect's input to a large extent. This has also been found in asymmetric and restricted interactions with interaction partners with limited communicational abilities, as for example preverbal infants or impaired persons [22], [23]. A similar mechanism of proposition in a learning robot could be a means to elicit appropriate signals from a human tutor in HRI [18], [26], [27].

ACKNOWLEDGMENT

We would like to thanks Chloé Rozenbaum for her help in designing the final version of the setup and in conducting the pilot study. This work has been partially supported by INRIA and the ERC grant EXPLORERS 24007.

REFERENCES

[1] H. H. Clark and S. E. Brennan, "Grounding in communication," *Perspectives on socially shared cognition*, vol. 13, no. 1991, pp. 127–149, 1991.

[2] M. J. Pickering and S. Garrod, "Toward a mechanistic psychology of dialogue," *Behavioral and brain sciences*, vol. 27, no. 2, pp. 169–189, 2004.

[3] M. Tomasello, *The cultural origins of human cognition*. Harvard University Press, 2009.

[4] K. J. Rohlfing, J. S. Poblete, and F. Joublin, "Learning new words in unfamiliar frames from direct and indirect teaching," in *17th Workshop on the Semantics and Pragmatics of Dialogue (SemDial)*, 2013, p. 121130.

[5] A. Ninio and C. E. Snow, *Pragmatic development.* Westview Press, 1996.

[6] B. Akgun, M. Cakmak, J. W. Yoo, and A. Thomaz, "Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective," in *Proceedings of the International Conference on Human-Robot Interaction (HRI)*, 2012.

[7] B. Galantucci, "Experimental semiotics: A new approach for studying communication as a form of joint action," *Topics in Cognitive Science*, vol. 1, no. 2, pp. 393–410, 2009.

[8] A. Cangelosi and D. Parisi, *Simulating the evolution of language*. Springer London, 2002.

[9] L. Steels, *Experiments in cultural language evolution*. John Benjamins Publishing, 2012, vol. 3.

[10] B. Galantucci, "An experimental study of the emergence of human communication systems," *Cognitive science*, vol. 29, no. 5, pp. 737–767, 2005.

[11] B. Galantucci and S. Garrod, "Experimental semiotics: a review," *Frontiers in human neuroscience*, vol. 5, 2011.

[12] J. P. De Ruiter, M. L. Noordzij, S. Newman-Norlund, R. Newman-Norlund, P. Hagoort, S. C. Levinson, and I. Toni, "Exploring the cognitive infrastructure of communication," *Interaction Studies*, vol. 11, no. 1, pp. 51–77, 2010.

[13] S. Griffiths, S. Nolfi, G. Morlino, L. Schillingmann, S. Kuehnel, K. Rohlfing, and B. Wrede, "Bottom-up learning of feedback in a categorization task," in *Development and Learning and Epigenetic Robotics (ICDL), 2012 IEEE International Conference on*. IEEE, 2012, pp. 1–6.

[14] G. Morlino, C. Gianelli, A. M. Borghi, and S. Nolfi, "Developing the ability to manipulate objects: A comparative study with human and artificial agents," in *Processings of the Tenth International Conference on Epigenetic Robotics*, 2010, pp. 169–170.

[15] S. Kopp, "Social resonance and embodied coordination in face-to-face conversation with artificial interlocutors," *Speech Communication*, vol. 52, no. 6, pp. 587–597, 2010.

[16] M. P. Michalowski, S. Sabanovic, and H. Kozima, "A dancing robot for rhythmic social interaction," in *Human-Robot Interaction (HRI), 2007 2nd ACM/IEEE International Conference on*. IEEE, 2007, pp. 89–96.

[17] K. Fischer, K. Lohan, J. Saunders, C. Nehaniv, B. Wrede, and K. Rohlfing, "The impact of the contingency of robot feedback on hri," in *Collaboration Technologies and Systems (CTS), 2013 International Conference on*. IEEE, 2013, pp. 210–217.

[18] A.-L. Vollmer, M. Mühlig, J. J. Steil, K. Pitsch, J. Fritsch, K. J. Rohlfing, and B. Wrede, "Robots show us how to teach them: Feedback from robots shapes tutoring behavior during action learning," *PloS one*, vol. 9, no. 3, p. e91349, 2014.

[19] K. Pitsch, A.-L. Vollmer, and M. Muhlig, "Robot feedback shapes the tutor's presentation how a robot's online gaze strategies lead to micro-adaptation of the human's conduct." *Interaction Studies*, vol. 14, no. 2, 2013.

[20] B. Wrede, S. Kopp, K. Rohlfing, M. Lohse, and C. Muhl, "Appropriate feedback in asymmetric interactions," *Journal of Pragmatics*, vol. 42, no. 9, pp. 2369–2384, 2010.

[21] P. Wittenburg, H. Brugman, A. Russel, A. Klassmann, and H. Sloetjes, "Elan: a professional framework for multimodality research," in *Proceedings of LREC*, vol. 2006, 2006.

[22] E. Ochs, B. B. Schieffelin, and M. Platt, "Propositions across utterances and speakers," *Developmental pragmatics*, pp. 251–268, 1979.

[23] C. Goodwin, "Co-constructing meaning in conversations with an aphasic man," *Research on language and social interaction*, vol. 28, no. 3, pp. 233–260, 1995.

[24] J. Grizou, M. Lopes, and P.-Y. Oudeyer, "Robot Learning Simultaneously a Task and How to Interpret Human Instructions," in *Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics (ICDL-EpiRob)*, Osaka, Japan, 2013.

[25] J. Grizou, I. Iturrate, L. Montesano, P.-Y. Oudeyer, and M. Lopes, "Calibration-Free BCI Based Control," in *AAAI Conference on Artificial Intelligence*, Quebec, Canada, 2014.

[26] M. Cakmak and A. L. Thomaz, "Designing robot learners that ask good questions," in *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 2012, pp. 17–24.

[27] A. Cangelosi, G. Metta, G. Sagerer, S. Nolfi, C. Nehaniv, K. Fischer, J. Tani, T. Belpaeme, G. Sandini, F. Nori *et al.*, "Integration of action and language knowledge: A roadmap for developmental robotics," *Autonomous Mental Development, IEEE Transactions on*, vol. 2, no. 3, pp. 167–195, 2010.