

Topics in Cognitive Science 1 (2009) 393–410  
Copyright © 2009 Cognitive Science Society, Inc. All rights reserved.  
ISSN: 1756-8757 print / 1756-8765 online  
DOI: 10.1111/j.1756-8765.2009.01027.x

## Experimental Semiotics: A New Approach for Studying Communication as a Form of Joint Action

Bruno Galantucci<sup>a,b</sup>

<sup>a</sup>*Department of Psychology, Yeshiva University*

<sup>b</sup>*Haskins Laboratories, New Haven*

Received 8 July 2008; received in revised form 23 October 2008; accepted 8 January 2009

---

### Abstract

In the last few years, researchers have begun to investigate the emergence of novel forms of human communication in the laboratory. I survey this growing line of research, which may be called *experimental semiotics*, from three distinct angles. First, I situate the new approach in its theoretical and historical context. Second, I review a sample of studies that exemplify experimental semiotics. Third, I present an empirical study that illustrates how the new approach can help us understand the socio-cognitive underpinnings of human communication. The main conclusion of the paper will be that, by reproducing micro samples of historical processes in the laboratory, experimental semiotics offers new powerful tools for investigating human communication as a form of joint action.

*Keywords:* Experimental semiotics; Human communication; Joint action; Distributed cognition

---

### 1. Introduction

*We cannot hope to understand language use without viewing it as joint actions built on individual actions. The challenge is to explain how all these actions work.*

Clark, 1996

In the last 40 years, a number of students of human dialog have tackled empirically the challenge indicated by Herbert Clark of explaining how “all these actions work” (e.g., Clark & Wilkes-Gibbs, 1986; Garrod & Anderson, 1987; Goodwin, 2000; Krauss & Weinheimer, 1964). This line of research, which for convenience we may call *experimental*

---

Correspondence should be sent to Bruno Galantucci, Department of Psychology, Yeshiva University, 2495 Amsterdam Avenue, New York, NY 10033. E-mail: galantuc@yu.edu

*pragmatics*, has maintained its vitality in recent years (Barr & Keysar, 2007; Brennan, 2005; Pickering & Garrod, 2004) and continues to provide important insights into human communication, as illustrated by the contributions in this issue by Brennan & Hanna, Garrod & Pickering, and Shintel & Keysar.

Here I present a relatively new line of research that attempts to tackle Clark's challenge by focusing on human communication in general, rather than on the specifics of spoken conversation (Galantucci, 2005; Garrod, Fay, Lee, Oberlander, & MacLeod, 2007; Healey, Swoboda, Umata, & Katagiri, 2002b; Selten & Warglien, 2007). This new approach, which for convenience we may call *experimental semiotics*, adopts the core assumptions of experimental pragmatics. On the one hand, it adopts the assumption that communication is a real-time social process which must be understood at the level of dyadic interactions (Pickering & Garrod, 2004). On the other hand, it adopts the assumption that communicative interactions are embodied in the physical world (Goodwin, 2000) and embedded in fairly rich socio-cognitive contexts (Brennan & Clark, 1996; Hutchins, 1995; Krauss & Glucksberg, 1977; Suchman, 1987). At the same time, experimental semiotics differs from experimental pragmatics in two important ways.

The first difference is in theoretical focus. As mentioned before, experimental semiotics focuses on human communication in general rather than on spoken conversation, including for instance graphical communication. This shift in focus represents an important departure. In fact, as the current trend in the study of sign language suggests (Vermeerbergen, 2006), spoken language might be a highly specialized form of communication, and some general principles of human semiosis might not be readily transparent in speech. For example, spontaneously emerging novel sign-languages have deep and vast iconic roots (Fusellier-Souza, 2006) and, although iconicity in sign-languages tends to decrease over time (Frishberg, 1975), a fairly high degree of iconicity remains one of the defining features of historically established sign-languages (Taub, 2001). Granted, spoken languages do exhibit some iconicity (Haiman, 1985; Hinton, Nichols, & Ohala, 1994), but their degree of iconicity is minimal when compared to that of sign-languages. Research in experimental semiotics not only confirms that iconicity is an important feature of emerging human communication systems but it also explains why communication systems tend to retain part of their iconic roots (Fay, Garrod, & Roberts, 2008; see next section). Thus, the fact that iconicity plays a marginal role in spoken communication might be a by-product of adaptations specific to the acoustic medium, rather than a central design feature of human communication (Hockett, 1960). Other features of spoken language might be consequences of adaptations specific to speech as well. For example, a recent study by Selten and Warglien (2007) shows that the emergence of communication conventions is facilitated by inventories of signal units that are relatively large with respect to the number of meanings that people express by using the units. The fact that fully blown languages that express thousands of meanings can rely on inventories composed of as little as eleven phonemes (Firchow & Firchow, 1969) might again be the consequence of specific adaptations for audio-vocal communication.

The second difference between experimental pragmatics and experimental semiotics is in their objects of study. Experimental semiotics studies the emergence of new forms of communication; experimental pragmatics studies the spontaneous use of preexisting forms of

communication such as spoken English. The emergence of new forms of communication is not a novel object of study. It has been extensively studied both with sign-languages that emerge in relatively isolated populations (e.g., Sandler, Meir, Padden, & Aronoff, 2005; Senghas, Kita, & Ozyurek, 2004) and with home sign systems that emerge in families in which deaf children are raised by nonsigning parents (e.g., Goldin-Meadow & Feldman, 1977; Goldin-Meadow & Mylander, 1998). These lines of research have produced a wealth of knowledge about the origins of novel languages (Goldin-Meadow, 2003; Kegl, Senghas, & Coppola, 1999). However, because experimental semioticians observe the emergence of communication in the laboratory, they gain access to new opportunities for scientific inquiry. On the one hand, they have access to the complete history of the emergence of a communication system. As we shall see later, knowing the details of this history can greatly enhance our understanding of the emergence of communication. On the other hand, experimental semioticians can perform manipulations that would be very difficult to realize outside of the laboratory. For example, Healey, Swoboda, Umata, and King (2007) systematically manipulated the composition of the communities of people in their study, while Galantucci, Kroos, and Rhodes (2006) manipulated a physical property of the medium over which communication systems emerged. Outside of the laboratory, such manipulations would be problematic to realize.

Thanks to its focus on novel forms of communication and to the new opportunities afforded by laboratory research, experimental semiotics permits us to address questions that are complementary to those typically addressed by experimental pragmatics. Experimental pragmatics typically focuses on how-questions such as *How does this aspect of communication work?* Experimental semiotics not only addresses the same how-questions from a different perspective, but it also offers historically based answers to why-questions such as *Why is this aspect of communication the way it is?* These answers might prove crucial for tackling Clark's challenge, as why-explanations are necessary parts of a scientific account of natural phenomena (Millikan, 1984; Tinbergen, 1963).

In the next section I introduce experimental semiotics in more detail and review a sample of studies that illustrate its relevance for students of human communication.

## 2. Experimental semiotics: A growing discipline

In the last few years, Healey and his colleagues (Healey, King, & Swoboda, 2004; Healey, McCabe, & Katagiri, 2000; Healey et al., 2002b) began using a graphical medium to study the development of novel communicative conventions among pairs of individuals<sup>1</sup>. Healey and colleagues adopted standard referential communication tasks such as those used in experimental pragmatics (e.g., Krauss & Weinheimer, 1964) but prevented the use of spoken or written language. In particular, they asked people to graphically describe a stimulus such as a piece of music or a concept to a partner, without allowing them to use letters or numbers (I will refer to this task as *graphical communication task*). The partner in the game was asked to recognize the stimulus among a set of stimuli in the case of concepts or, in the case of a piece of music, to say whether the description identified the same piece of music

s/he heard for that trial. Over a number of rounds of the game, Healey and colleagues observe people developing spontaneous communicative conventions to succeed at the graphical communication task.

Taken together, the studies performed by Healey and colleagues indicate that graphical communication and dialog rely on similar mechanisms for grounding communication (Healey, Garrod, Fay, Lee, & Oberlander, 2002a), organizing turn-taking (Umata, Shimojima, Katagiri, & Swoboda, 2003), and repairing communication failures (Healey et al., 2007). Moreover, these studies reveal that the real-time interactive processes that typically support successful joint action in spoken conversation play an important role in shaping novel communicative conventions as well. For example, when people are allowed to have freer communicative interactions, they develop conventions with a higher degree of abstraction than when the interactions are more constrained (Healey et al., 2002b). In particular, mutual-modifiability—the fact that people have opportunities to alter each other’s graphical productions—has been shown to be a key factor for the development of symbolic conventions (Healey et al., 2007). Another recent study, this time by Garrod and colleagues (Garrod et al., 2007), focuses on the processes through which an iconic form becomes symbolic over time, using a graphical communication task with a set of conceptual referents such as “drama” or “parliament.” The study suggests three main conclusions. First, symbols do not develop out of iconic forms merely because of repeated use. Some form of direct interaction between the producer of the sign and the receiver of the sign is necessary for symbolization to occur. Second, the process of symbolization is enhanced when people in a pair engage in richer interactions, exchanging roles as producers and receivers. Third, people that are not engaged in the interactions that give rise to a sign are less efficient in learning the sign than the people that developed it. Garrod and colleagues interpreted these results as indicating that symbols emerge from a process of grounding similar to the process described by students of experimental pragmatics (Clark & Brennan, 1991). Through this process, the informational content of a sign used by two people gradually shifts from the physical appearances of the sign to a representation level in which the sign grounds itself in the shared history of its use. That is, the first occurrences of a novel sign rely more heavily on the physical properties of the sign. At this stage the sign tends to identify one referent among *all of the possible referents* and iconicity is of much help. However, after a number of interactions, the sign begins to refer to previous communicative interactions rather than directly to its physical referent. At this stage, the sign identifies one element of *a small set of shared signs*, and iconicity is no longer crucial. Nonetheless, iconicity does not completely vanish. A follow-up study by Fay and colleagues (Fay et al., 2008) revealed that communication systems developed by a community of players tend to maintain a higher degree of iconicity than communication systems developed by isolated pairs. This occurs because of selective and adaptive processes that operate at the level of the community: On average, iconic signs are easier to learn and remember.

Thanks to the fact that they use referential communication tasks that are commonly used in experimental pragmatics, the studies here summarized have the important advantage that they can be readily compared to studies of experimental pragmatics. Indeed, all of the results presented in this section are not only in line with well-known results of experimental

pragmatics (Clark & Wilkes-Gibbs, 1986; Hupet & Chantraine, 1992; Krauss & Weinheimer, 1966) but also confirm one of the assumptions shared by experimental pragmatics and experimental semiotics: Human communication is a real-time social process that must be understood at the level of dyadic interactions. Moreover, as in the case of the study by Fay and colleagues, these studies offer new explanations for known phenomena such as the permanence of a relatively high degree of iconicity in signed languages.

However, for the purposes of experimental semiotics, the use of standard referential communication tasks also imposes some limitations. One of them is that the semiotic challenge for the participants is relatively moderate. On the one hand, participants in the studies described above know who must produce the forms for communicating and when. On the other hand, the meanings to be conveyed between participants in a pair are part of a set of meanings that the experimenter establishes for the participants ahead of time. For example, participants in the study by Healey and colleagues (Healey et al., 2004) had to identify a referent out of set of 12 possible referents<sup>2</sup> The use of moderate semiotic challenges facilitates the rapid emergence of communication but it also limits the possibility to observe the spontaneous emergence of communication out of fairly unstructured activities.

Another limitation of studies that use graphical communication tasks is that, although nontextual graphical communication is less constrained by prior conventions than spoken language, a number of preexisting communicative conventions remain. For example, participants in the studies by Healey and colleagues could use nonalphabetic graphic symbols (e.g., the \$ symbol for money) or pictorial representations (e.g., drawings of people or animals) and indeed used them frequently.

Recently, a number of researchers (De Ruiter, Noordzij, Newman-Norlund, Hagoort, & Toni, 2007; Galantucci, 2005; Galantucci, Fowler, & Richardson, 2003; Scott-Phillips, Kirby, & Ritchie, 2008) have begun to use experimental frameworks that (a) involve more challenging semiotic tasks than those typical of graphical communication studies, and (b) drastically reduce the possibility of using preexisting communicative conventions (I will refer to these tasks as *visual communication games*). Here I will focus in particular on the visual communication game that my colleagues and I recently developed (Galantucci, 2005; Galantucci et al., 2003). The task in this game is similar to a well-known task used for studies of experimental pragmatics (e.g., Garrod & Anderson, 1987). Pairs of participants must coordinate their moves in a virtual environment comprised of a set of rooms located on a grid and marked with icons. In such environment, players can succeed only when they communicate effectively. In particular, players have the task of moving to the same room with the minimum number of room changes, but they only have a local view of their room and they cannot see where the partner is located. As they need to know this in order to decide their next moves, they are encouraged to develop ways for describing their own positions, where they intend to move next, or what they suggest the partner should do. However, all of these things can be indicated in a variety of ways (cf. Garrod & Anderson, 1987). A room can be indicated as “the north-east room,” or as “the room that has a triangle on the floor,” or as “the room that has a door on the right side and a door on the lower side.” A sign by a player can indicate the current location of the player or the location the player suggests for the partner. Moreover, the task has no preestablished communication protocol (both players

can continuously exchange messages in real time) and establishing ‘‘who sends messages when’’ is part of the challenge of the task. In other words, players—who are in contact only through a graphical device—are faced with a number of severe semiotic challenges, and complete failures in developing even minimal communication systems sometimes occur. As we shall see in the next section, such failures can be informative about the minimal requirements for the establishment of a successful communication system.

A crucial element of the visual communication game that my colleagues and I developed is the unusual graphical device by which players can communicate. Players can send messages to each other by using a magnetic stylus on a small digitizing pad. The resultant tracings are relayed to the computer screens of both players. In particular, the horizontal location of the stylus on the pad controls the horizontal location of a tracing on the screen. The vertical location of the stylus on the pad is irrelevant. The tracing always appears at the top of the screen and then scrolls down the screen at a constant speed for 2.5 s, until it reaches the end of the screen and disappears. In such conditions, the use of standard graphic forms such as letters or numbers is practically impossible (Fig. 1). Nonetheless, most pairs of players manage to get a communication system started, giving us opportunities to observe the emergence of novel communication systems from their very foundations (Galantucci, 2005; Galantucci et al., 2003). These systems expose the bare bone structure of functional semiotic systems, revealing a number of features distinctive of human communication. For example, the signs of these novel communication systems are perceptually distinct and easy to produce (cf. Lindblom, 1990) and, in line with one of the core pragmatic requirements for human language (Clark, 1996), their use is tightly embedded within behavioral procedures that coordinate the actions of the partners. Because of their relative simplicity with respect to the behavioral procedures that support spoken language, these procedures can function as micro-models of the kind of phenomena one needs to study to tackle Clark’s challenge (Galantucci, 2005). Moreover, because successful pairs are then faced with new challenges

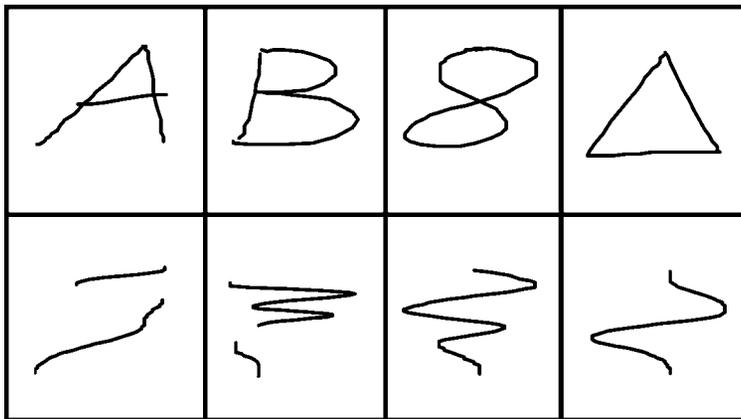


Fig. 1. How common graphic symbols appeared on the screen when traced on the digitizing pads used for the game.

in the game, the method allows us to study the further evolution of the communication systems that have emerged, showing that the systems are flexibly adapted by the players to accommodate for new needs (Galantucci, 2005) and that linguistic forms of communication such as combinatorial structures quickly emerge (Galantucci et al., 2006). In brief, the method allows us to collect data that are ideally suited for the goals of experimental semiotics. In the section that follows I will illustrate how the method also allows us to more directly appreciate the role that joint activities play during the emergence of communication, offering new tools to tackle Clark's challenge.

### **3. Joint action and the emergence of communication: A posthoc study**

When one observes the emergence of communication in real time, it is often difficult to understand what is happening and why. However, studies of experimental semiotics offer offline access to the complete record of the events that led to the emergence of a communication system. Post hoc analyses of these records provide opportunities for identifying the factors that are behind success, or failure, in the development of communication systems. In this section, I present one of such analyses. In particular, I focus on 16 pairs of participants that took part in an experiment conducted by my colleagues and me (Galantucci et al., 2006) using a visual communication game. Since the study was conducted with a method very similar to that illustrated in the previous section (Galantucci, 2005), I will not focus on its details here. Rather, I will focus on the properties of the game used for the experiment that are relevant for the current analyses.

#### *3.1. The game*

Two adults participated in a real-time videogame with interconnected computers at different locations. Each player controlled the movements of an agent in a shared virtual environment which, at the beginning of the session (Stage 1), was composed of four interconnected rooms (Fig. 2). Players engaged in a fully cooperative game. At the beginning of each round of the game, the agents were located in two different rooms at random, and the players' goal was to bring the agents into the same room without making more than a single room change per agent. Chance-level performance at Stage 1 of the game was 50% and could be improved only if information about location and intended movement of the agents was communicated via the graphic device illustrated above. Once communication occurred, however, the game reduced to a trivial task.

#### *3.2. Experimental procedure*

Sixteen pairs of participants were recruited to play the game at Stage 1. Before playing the game, players were briefly instructed on the basic rules of the game and informed that their partners received the same instructions. During the game, players were encouraged to focus on the score—a numerical index that increased only when the pair won consistently

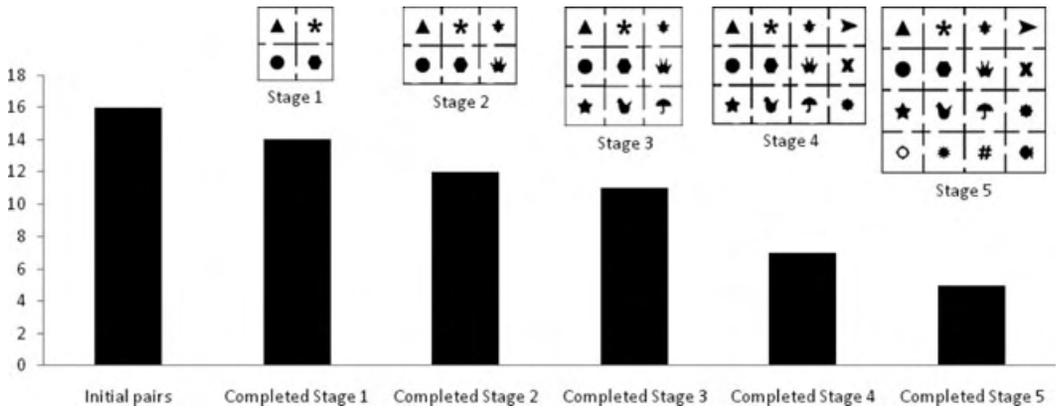


Fig. 2. The upper part of the figure illustrates the five game environments of the study. The lower part illustrates the number of pairs that completed each of the stages in the game.

in the game—as their primary goal. If the pair reached a threshold score that indicated successful communication, players were invited to play a new version of the game: The game environment was enlarged (6 rooms,  $2 \times 3$  grid, Stage 2; see Fig. 2) and an additional room change per round was allowed. For successful pairs, the size of the environment (and the number of room changes allowed) could grow three more times (Stages 3–5; Fig. 2) until the environment, at Stage 5, was composed of 16 rooms ( $4 \times 4$  grid). For the entire duration of the study, the movements of the agents and the activity on the digitizing pads were recorded at approximately 30 Hz. On termination of the experimental sessions, participants were interviewed for about 30 min during which they provided a detailed written description of the communication systems they developed for playing.

### 3.3. Results

Fig. 2 illustrates the final performance of the 16 pairs at the game. As illustrated in the figure, two pairs (12.5% of the sample) failed at the game, two pairs (12.5%) mastered only Stage 1, one (6.25%) mastered Stage 2, four (25%) mastered Stage 3, two (12.5%) mastered Stage 4, and five pairs (31.25%) mastered Stage 5, completing the game.

The communicative power of the communication systems developed by the pairs was determined by counting the number of locations in the game map that the systems allowed to discriminate. In particular, for each pair, my colleagues and I identified all the locations that were (a) consistently indicated as discriminable by both players in the descriptions provided at the end of the game; and (b) consistently discriminated by both players during the game (this was determined by inspecting the recordings of the game). For the two pairs that failed, communicative power was computed as zero.<sup>3</sup>

Not surprisingly, communicative power ( $M = 8$ ;  $SD = 5.66$ ) and performance in the game ( $M = 3.06$  stages,  $SD = 1.81$ ) were significantly correlated [ $r(16) = 0.91$ ,  $p < .001^4$ ].

This indicates that the method worked as intended: Success in the game implied more powerful communication systems. The remaining part of this section will focus on a general question about these systems: How did they come about?

As I mentioned before, when one observes the emergence of communication in real-time, it is often difficult to understand what is happening and why. This difficulty is an interesting phenomenon on its own. In a typical session of a visual communication game (Galantucci, 2005; Galantucci et al., 2006; Galantucci et al., 2003), each participant is directly observed by one experimenter, who takes note of the most salient behavioral and verbal reactions of the participant during the game. Moreover, the experimenters are in contact with each other in real time through a written messaging system and have access to a bird's-eye view of the game. Yet when participants begin to effectively communicate with each other, the experimenters have often little clue as to how the communication system works and how it was established. In other words, it seems that the well-known disadvantage of over-hearers in following a conversation with respect to the people involved in it (Schober & Clark, 1989) occurs also for experimenters observing the development of a novel communication system. However, analyses of the game records provide opportunities for identifying factors that might be behind success, or failure, in the development of communication systems. Before I present such analyses, I need to illustrate an important aspect of the game used by my colleagues and me. During a session of the game, there were two distinct modes of interactions in which players could exchange signals<sup>5</sup>. The first one occurred when a round of the game was ongoing. Players were always in different rooms and their views of the task environment had no overlap. In this context (which I will refer to as *online interaction*), hypotheses about the meaning of players' signals had to be tested by trial and error, keeping track of the successes and the failures at achieving the goal of finding each other. To illustrate this process let us imagine the case of a player (Player A, female) that, while her agent is alone in a room, sees a signal, say an S-shape, coming from the partner (Player B, male). Not knowing what Player B means with the signal and not having any other clue on how to make a successful move, Player A haphazardly crosses the closest door, ending up in the room where Player B is. As soon as player A finds the partner in the room, she has the opportunity to make an inference about what the S-shaped signal meant. The inference is by no means fail proof. The signal might refer to a number of things, including the icon on the floor of the current room, the location of the room on the game map, the movement that Player B had intended to do, etc. However, over a number of successful trials, Player A has the possibility to test the inferences she makes either by observing where Player B is when he uses the S-shaped again or by observing where Player B ends up when *she* produces an S-shaped signal. Considering that Player B can also adjust his signaling behavior to that of Player A, the pair can rely on a powerful mechanism for developing functional signs without ever explicitly negotiating their meaning (Galantucci, 2005).

The other kind of interaction occurred when a round was over. At that moment, agents could no longer leave their rooms until both players decided to terminate the round by moving the agents into one of four marked locations in the room. As soon as both agents entered one of such locations, a new round of the game resumed; agents were instantly relocated in two different rooms at random and players returned to an online interaction. In other words,

at the end of each round players gained control of the pace of the game and could decide to interact in absence of a direct pressure to win a round of the game (in what follows I will refer to these interactions as *offline interactions*). To illustrate an offline interaction, let us imagine a moment in which the players of the pair above have just won a round of the game. The agents are in the same room and the game is on hold until both players decide to resume it. One of the players, say Player B, decides not to enter into one of the locations that would cause the game to resume. Instead, he moves the agent close to the icon on the floor of the room and, while producing an S-shape signal, makes the agent bump a couple of times on the icon. At that point his partner's guesses about the meaning of the S-shape signal are considerably facilitated: The signal has likely something to do with the icon. Moreover, Player A might test her guesses by observing Player B's reaction when *she* bumps her agent on the icon while producing the S-shape signal. At this point Player B might express some form of consent (for example, by moving his agent up and down, as in a virtual nod) or initiate a new sequence of movements and signals. In this way, players can explicitly negotiate the meaning of the signals to be used in the game. One might think that such negotiations helped in the development of communication systems, because they offered opportunities to communicate in a context in which signaling had no direct consequence on the game. That is, during offline interactions, players knew that signals could not possibly be about moves to make at that time, since all doors were locked. In such conditions, players had the opportunity to use signals with a meta-communicative function, supporting the establishment of bits of a communication system that could be later used during the online phase of the game. Indeed, these interactions were allowed in the game precisely because they could facilitate the development of communication systems and pairs did take advantage of them. On average, pairs engaged in about 57 offline interactions<sup>6</sup> during the game time ( $M = 57.56$ ,  $SD = 64.28$ ). However, performance and the number of offline interactions were negatively correlated [ $r(16) = -.66$ ,  $p = .01$ ; Fig. 3B]. In other words, offline interactions had the opposite effect than one might have expected. A possible explanation for this result could be that the pairs that engaged in offline interactions slowed down their game pace and, hence, had fewer opportunities to improve their performance. However, such explanation is ruled out by the fact that, as mentioned in note 4, the correlation did not change [ $r(12) = -.67$ ,  $p = .01$ ] when the total time and the total number of rounds played by the pairs in the game were partialled out from the correlation reported above. Why, then, were offline interactions not beneficial for the pairs? The answer to the question requires some elaboration and will take most of the remainder of this section.

A first reason for the result is that players sometime engaged in an unexpected form of offline interaction, which was clearly detrimental for the development of a communication system. One might expect that players engaged in offline interactions only after a round was won. In fact, if the pair won the round, players completely shared their views of the task environment and could see each other's agents, which were in the same room. In this context (which I will refer to as *offline same view interaction*), hypotheses about the meaning of players' signals could be directly tested through the parallel communication channel provided by the movements of the agents in the room. In fact, players not only saw each other's agents' location in the room but they also saw their orientation in the game environment,

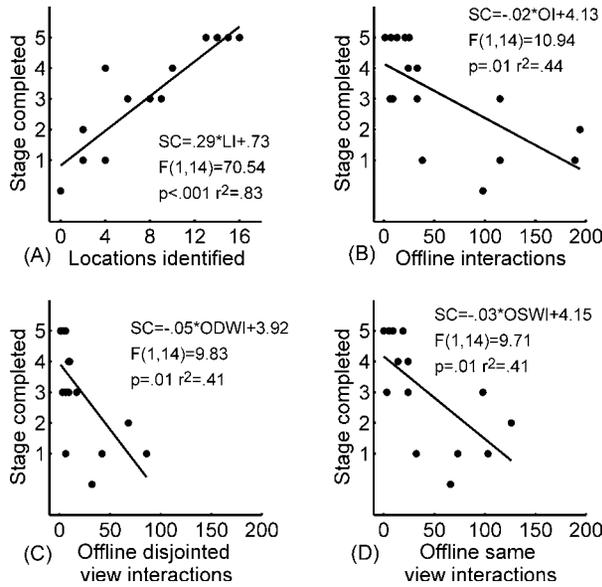


Fig. 3. Scatterplots of the main variables considered in the text. For all scatterplots, the vertical axis represents the maximum stage of the game completed by the pair while the horizontal axis represents: (A) the number of game locations identified by the pairs' communication systems at the end of the game; (B) the number of offline interactions; (C) the number of offline disjointed view interactions; and (D) the number of same view interactions. The lines in the graphs represent the linear regression through the data points.

given that the agents had human-like animated bodies. These movements were publicly visible and could be used to ground the meaning of the signals, as illustrated in the examples above.

In contrast, if the pair lost the round, there was no overlap between the players' views of the task environment and players could not see each other's agents, which were in different rooms. In this context (which I will refer to as *offline disjointed view interaction*), hypotheses about the meaning of players' signals could not be tested at all and any communication exchange between the players was doomed to remain ambiguous because it could never be grounded in a sharable experience. Clearly, such interactions could not be beneficial to the pair. Nevertheless, about 37% of the offline interactions across pairs were offline disjointed view interactions ( $M = .37$ ,  $SD = .22$ ). Not surprisingly, performance and number of offline disjointed view interactions ( $M = 18.94$ ,  $SD = 25.51$ ) were negatively correlated across pairs [ $r(16) = -.64$ ,  $p = .01$ ; Fig. 3C].

Before considering same view offline interactions, it is important to fully appreciate the significance of disjointed view offline interactions. This kind of interaction was not expected to occur at all. However, it occurred often, and did so for a simple reason. While playing the game, people often behaved in ways that were far from being optimal for communication, showing a surprisingly high degree of *communicational egocentrism* (cf. Keysar, 2007). For example, some players did not use the stylus for as long as two consecutive hours, often while the partner tried repeatedly to initiate some form of communication. For these players,

the idea of reciprocating acts of communication was not obvious. Also, some players developed signs that, while having the same exact appearance on the screen, were meant to indicate something different because they were drawn in different ways. For example, one player indicated the direction of the agent's vertical movements by drawing a vertical line from the pad's bottom to the top for upward movements, and from the pad's top to the bottom for downward movements. Given the properties of the communication device, the two lines appeared identical on the screen. Yet the player kept using them to differentially indicate direction and was frustrated by her partner's "lack of understanding." Given that the signs were drawn with intentions that were known to her, the player expected the partner to understand them accordingly. Incidentally, the partner understood what was possible to understand in such conditions. To her, the sign meant vertical motion, with no indication of direction.

Offline disjointed view interactions were but another example of communicational egocentrism. When players initiated one of such interactions, what mattered to them was to express meanings that they had in mind, regardless of the fact that the chance that those meanings could be understood by the partner was almost null. For this reason it is no surprise that, the more often pairs engaged in offline disjointed view interactions, the worse they performed in the game. The presence of such interactions was a clear indication of communicational egocentrism.

Surprisingly, however, also offline same view interactions were not beneficial for the development of communication systems. Performance and number of offline same view interactions ( $M = 38.63$ ,  $SD = 40.92$ ) were negatively correlated [ $r(16) = -.64$ ,  $p = .01$ ; Fig. 3D]. Moreover, although the number of offline same view interactions and the number of offline disjointed view interactions were positively correlated [ $r(16) = .87$ ,  $p < .001$ ], the correlation between performance and number of offline same view interactions remained negative [ $r(13) = -.22$ ,  $p = .43$ ] when the number of offline disjointed view interactions was partialled out. In other words also offline same view interactions did not seem to be beneficial for the development of a communication system. Why?

If one watches the game recordings, one finds that the answer is, in the end, rather simple. Offline same view interactions were helpful only when they were *well-grounded*, that is, scaffolded by a rich mesh of joint actions (Sebanz, Bekkering, & Knoblich, 2006). A detailed analysis of what constitute a well-grounded interaction goes beyond the scope of this paper (see Galantucci & Steels, 2008, for such analysis). In essence, well-grounded interactions were interactions in which the partners either entered into *frames of mutual attention*, explicitly attending and responding to their respective behaviors (cf. the "nodding" example illustrated above), or entered into *frames of joint attention*, explicitly focusing their attention toward a same third entity (cf. the "pointing to the icon" example illustrated above). When well-grounded interactions occurred, the number of offline interactions the pair engaged into dramatically decreased. That is, when well-grounded, little bits of metacommunication suffice for the set-up of a novel communication system. Indeed, some of the most successful pairs in the study had almost no offline interactions at all, demonstrating that a communication system can be set-up with little, if any, metacommunication (see Garrod & Anderson, 1987, for a parallel observation in experimental pragmatics). In

contrast, during non-well-grounded interactions, neither of the two frames mentioned above was established. In such conditions, any exchange of signals was doomed to be irrelevant, if not confusing. In fact, the functioning communication system of a few pairs lost its functionality because of such confusion. For these pairs, metacommunication was a problem rather than a resource. Across pairs, this was the main reason for the negative correlations between performance and the number of offline interactions.

### *3.3. Conclusions*

In sum, the increase in semiotic challenge that is characteristic of visual communication games allows us to observe vast differences in pairs' success at establishing communication systems. The analysis of these differences suggests three overarching conclusions. First, metacommunication might help for setting up a communication system, but it is not strictly necessary. Second, when well grounded, a little metacommunication can go a long way. Third, non-well-grounded metacommunication leads to excessive metacommunication, which, in turn, hinders the development of a communication system. These conclusions suggest a simple historical constraint that is of relevance for tackling Clark's challenge. Whichever way all the actions mentioned by Clark work, it is unlikely that they work the way they do because of prior explicit negotiations.

## **4. General conclusions**

This paper comprised three main sections. The first section was aimed at introducing experimental semiotics, a recently developed approach to study human communication. In that section, I situated experimental semiotics in its theoretical and historical context, illustrating how it adopts important assumptions from previous approaches to human communication, which were collectively labeled experimental pragmatics. In particular, experimental semiotics adopts the assumptions that (a) language is a real-time social process that must be understood at the level of dyadic interactions (Pickering & Garrod, 2004) and (b) that communicative interactions are embodied in the physical world (Goodwin, 2000) and embedded in fairly rich socio-cognitive contexts (Brennan & Clark, 1996; Hutchins, 1995; Krauss & Glucksberg, 1977; Suchman, 1987). At the same time, I illustrated how experimental semiotics differs from experimental pragmatics. It differs because it addresses human communication in a general sense (hence the term "semiotics"), rather than focusing specifically on interactions driven primarily by spoken language. It also differs from experimental pragmatics because it focuses on forms of communication that emerge anew in the laboratory, allowing us to manipulate and record the historical processes that ground the birth of a new communication system.

The second section of the paper offered a brief review of a sample of studies in experimental semiotics. In that section I showed how experimental semiotics (a) provides results that corroborate and complement results from experimental pragmatics and (b) allows us to develop a deeper understanding of well-known phenomena of human communication

such as the presence and permanence of relatively high degrees of iconicity in sign-languages. Moreover, I argued that the recent emergence of methods in experimental semiotics that depart from the traditional methods of experimental pragmatics has created new opportunities to understand the fundamental socio-cognitive mechanisms that are behind the emergence of communication. In that context, I introduced a method that my colleagues and I recently developed and argued that the method, due to its substantial semiotic challenge, is well suited for studying essential preconditions for human communication.

In the last section of the paper I presented a posthoc study conducted with the method mentioned above. In particular, the study focused on the role played by metacommunication for the emergence of novel communication systems. The results of the study suggested that metacommunicatio plays a small (if any) role. This conclusion suggested in turn a simple historical constraint: Explicit negotiation is not a likely way for natural human conventions to develop.

Here I conclude by noting that experimental semiotics opens a new perspective for understanding communication as a form of joint action. This perspective hinges onto the theoretical tenet that, in order to understand how human communication works, we must have direct access to the historical processes that support the development of communication systems as socio-cultural forms of joint action *and* be able to manipulate them. In linguistics, there is a traditional distinction between what is called *synchronic linguistics*—the study of the present day structure of a language—and *diachronic linguistics*—the study of how languages came to have the structure that they have because of historical and geographic processes (de Saussure, 1916/1983). The same distinction can be applied to the study of communication as a form of joint action. Experimental pragmatics studies the present day structure of language as a form of joint action. Experimental semiotics studies how communication systems arise as elaborated forms of joint action via historical processes that can be observed in the laboratory. Until now, such processes have been largely inaccessible to scientific inquiry. It is my hope that, by studying them in the laboratory, experimental semiotics will soon become a fully mature discipline.

## Notes

1. A few years before the studies cited in this paragraph occurred, Goldin-Meadow, McNeill, and Singleton (1996) studied the generation of novel forms of human communication in the laboratory. In particular, they asked people to describe visually presented scenes twice, first using speech and then using exclusively gesturing. The results of the study provided valuable information about the differences between gesturing as an accompaniment of speech and gesturing as an exclusive communicative device. However, I will not consider this study as a full instance of experimental semiotics because it focused on unidirectional communication.
2. Selten and Warglien (2007) limited the challenge further, asking people to communicate by selecting forms from a small set of preestablished tokens.

3. One of these two pairs actually developed a partial communication system that, after much struggle, allowed them to reach the threshold score to complete Stage 1 of the game. However, the system collapsed as soon as the pair moved to Stage 2 and players did not reestablish functional communication by the end of the fourth hour of playing. At that point, their participation was interrupted because of the players' frustration. Moreover, the communication power of the system that players temporarily used to reach Stage 2 could not be determined because the descriptions provided by the players at the end of the game did not meet the required consistency criteria. In order to compare the results of this pair with those of the other pairs, the communicative power for the pair was computed as zero and, in consequence, the performance was considered a complete failure. The results presented here and their relative levels of significance do not change regardless of whether the data for this pair are included in the analyses.
4. For reasons related to the goals of the original study (Galantucci et al., 2006), pairs did not play for the same amount of time nor they played the same number of rounds in the game. In consequence, correlations involving performance were computed partialling out the total time and the total number of rounds played by the pairs. However, time and rounds played accounted for almost none of the variance in performance [ $F(2,13) = .03$ ,  $p = .97$ ,  $r^2 = .01$ ] and the analyses reported here produced equivalent pattern of results (and levels of significance) when the effects of the two variables were not partialled out from the performance scores. For ease of exposition of the regressions analyses reported in Fig. 3, I report only the results of the analyses conducted with the original performance scores. Additionally, the original study included an experimental manipulation of the communication medium. The manipulation had no detectable effect on performance [ $F(1,14) < 1$ ,  $\eta^2 = .03$ ] and all correlations presented in the paper followed the same pattern in the two experimental conditions; see Fig. 3A].
5. Throughout the remaining part of the paper, I will distinguish between signals, that is, the perceivable products of the physical activity on the digital scratchpad, and signs, that is, the meaningful units of functional communication systems.
6. In order for an interaction to count as an offline interaction, it had to occur when a round of the game was over, be longer than 5 s, and contain some signaling activity from one of the players.

## Acknowledgments

Had Luc Steels not insisted that pairs' failures in visual communication games could become useful data for research, many of the analyses presented here would have never been performed. Thanks, Luc. Also, many of the phenomena illustrated in the empirical section of this paper were identified over a number of conversations I had with Christian Kroos and Theo Rhodes, most often while watching players struggling in the game. Thanks Christian, thanks Theo. Finally, thanks to Simon Garrod, Andrea Scarantino, and two anonymous reviewers for their helpful comments.

The research presented in Section 3 was supported by NICHD grant HD-01994 and NIDCD grant DC-03782 to Haskins Laboratories.

## References

- Barr, D. J., & Keysar, B. (2007). Perspective taking and the coordination of meaning in language use. In M. J. Traxler & M. A. Gernsbacher (Eds.), *Handbook of psycholinguistics* (2nd ed., pp. 901–938). Amsterdam: Academic Press.
- Brennan, S. E. (2005). How conversation is shaped by visual and spoken evidence. In J. Trueswell & M. K. Tanenhaus (Eds.), *Approaches to studying world-situated language use: Bridging the language-as-product and language-action traditions* (pp. 95–129). Cambridge, MA: MIT Press.
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology-Learning Memory and Cognition*, 22(6), 1482–1493.
- Clark, H. H. (1996). *Using language*. Cambridge, England: Cambridge University Press.
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). Washington, DC: American Psychological Association.
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22(1), 1–39.
- De Ruiter, J. P., Noordzjij, M., Newman-Norlund, S., Hagoort, P., & Toni, I. (2007). On the origin of intentions. In P. Haggard, Y. Rossetti & M. Kawato (Eds.), *Attention and performance XXII: Sensorimotor foundation of higher cognition* (pp. 593–610). Oxford, England: Oxford University Press.
- Fay, N., Garrod, S., & Roberts, L. (2008). The fitness and functionality of culturally evolved communication systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509), 3553–3561.
- Firchow, I., & Firchow, J. (1969). An abbreviated phoneme inventory. *Anthropological Linguistics*, 11(9), 271–276.
- Frishberg, N. (1975). Arbitrariness and iconicity: Historical change in American Sign Language. *Language*, 51, 696–719.
- Fusellier-Souza, I. (2006). Emergence and development of signed languages: From a semiogenetic point of view. *Sign language studies*, 7(1), 30–56.
- Galantucci, B. (2005). An experimental study of the emergence of human communication systems. *Cognitive Science*, 29(5), 737–767.
- Galantucci, B., Fowler, C. A., & Richardson, M. J. (2003). Experimental investigations of the emergence of communication procedures. In R. Sheena & J. Effken (Eds.), *Studies in Perception and Action VII* (pp. 120–124). Mahwah, NJ: Erlbaum.
- Galantucci, B., Kroos, C., & Rhodes, T. (2006). Rapidity of fading and the emergence of duality of patterning. In A. Cangelosi, A. D. M. Smith & K. Smith (Eds.), *The evolution of language—Proceedings of the 6th International Conference on the Evolution of Language* (pp. 413–415). London: World Scientific.
- Galantucci, B., & Steels, L. (2008). The emergence of embodied communication in artificial agents and humans. In I. Wachsmuth, M. Lenzen & G. Knoblich (Eds.), *Embodied communication in humans and machines* (pp. 229–256). Oxford, England: Oxford University Press.
- Garrod, S., & Anderson, A. (1987). Saying What you mean in dialog—a study in conceptual and semantic coordination. *Cognition*, 27(2), 181–218.
- Garrod, S., Fay, N., Lee, J., Oberlander, J., & MacLeod, T. (2007). Foundations of representation: Where might graphical symbol systems come from? *Cognitive Science*, 31(6), 961–987.
- Goldin-Meadow, S. (2003). *The resilience of language: What gesture creation in deaf children can tell us about how all children learn language*. New York: Psychology Press.
- Goldin-Meadow, S., & Feldman, H. (1977). Development of language-like communication without a language model. *Science*, 197(4301), 401–403.

- Goldin-Meadow, S., McNeill, D., & Singleton, J. (1996). Silence is liberating: Removing the handcuffs on grammatical expression in the manual modality. *Psychological Review*, 103(1), 34–55.
- Goldin-Meadow, S., & Mylander, C. (1998). Spontaneous sign systems created by deaf children in two cultures. *Nature*, 391(6664), 279–281.
- Goodwin, C. (2000). Action and embodiment within situated human interaction. *Journal of Pragmatics*, 32(10), 1489–1522.
- Haiman, J. (Ed.). (1985). *Iconicity in syntax*. Amsterdam: John Benjamins.
- Healey, P. G. T., Garrod, S., Fay, N., Lee, J., & Oberlander, J. (2002a). Interactional context in graphical communication. In B. Bel & I. Marlien (Eds.), *Proceedings of the 24th Annual Conference of the Cognitive Science Society* (pp. 441–446). Mahwah, NJ: Erlbaum.
- Healey, P. G. T., King, J., & Swoboda, N. (2004). Co-ordinating conventions in graphical dialogue: Effects of repetition and interaction. In A. Blackwell, K. Marriott & A. Shimojima (Eds.), *Diagrams 2004: Third International Conference on Theory and Application of Diagrams*. March 22–24, 2004. Cambridge, England.
- Healey, P. G. T., McCabe, R., & Katagiri, Y. (2000). A comparison of graphics and speech in a task-oriented interaction. In M. Anderson, P. Cheng & V. Haarslev (Eds.), *Theory and Application of Diagrams, Proceedings, 1889*, pp. 245–256.
- Healey, P. G. T., Swoboda, N., Umata, I., & Katagiri, Y. (2002b). Graphical representation in graphical dialogue. *International Journal of Human-Computer Studies*, 57(4), 375–395.
- Healey, P. G. T., Swoboda, N., Umata, I., & King, J. (2007). Graphical language games: Interactional constraints on representational form. *Cognitive Science*, 31(2), 285–309.
- Hinton, L., Nichols, J., & Ohala, J. J. (Eds.) (1994). *Sound symbolism*. Cambridge, England: Cambridge University Press.
- Hockett, C. F. (1960). Logical considerations in the study of animal communication. In W. Lanyon & W. Tavolga (Eds.), *Animal sounds and communications* (pp. 392–430). Washington, DC: American Institute of Biological Sciences.
- Hupet, M., & Chantraine, Y. (1992). Changes in repeated references—Collaboration or repetition effects. *Journal of Psycholinguistic Research*, 21(6), 485–496.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA, US: The MIT Press.
- Kegl, J., Senghas, A., & Coppola, M. (1999). Creation through contact: Sign language emergence and sign language change in Nicaragua. In M. DeGraff (Ed.), *Language creation and language change: Creolization, diachrony, and development* (pp. 179–237). Cambridge, MA: MIT Press.
- Keysar, B. (2007). Communication and miscommunication: The role of egocentric processes. *Intercultural Pragmatics*, 4(1), 71–84.
- Krauss, R. M., & Glucksberg, S. (1977). Social and non-social speech. *Scientific American*, 236(2), 100–105.
- Krauss, R. M., & Weinheimer, S. (1966). Concurrent feedback confirmation and encoding of referents in verbal communication. *Journal of Personality and Social Psychology*, 4(3), 343–346.
- Krauss, R. M., & Weinheimer, S. (1964). Changes in reference phrases as a function of frequency of usage in social interaction—a preliminary study. *Psychonomic Science*, 1(5), 113–114.
- Lindblom, B. (1990). On the notion of ‘possible speech sound’. *Journal of Phonetics*, 18(2), 135–152.
- Millikan, R. G. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.
- Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27(2), 169–226.
- Sandler, W., Meir, I., Padden, C., & Aronoff, M. (2005). The emergence of grammar: Systematic structure in a new language. *Proceedings of the National Academy of Sciences of the United States of America*, 102(7), 2661–2665.
- de Saussure, F. (1916/1983). *Course in general linguistics* (R. Harris, Trans.) London: Duckworth & Co.
- Schober, M. F., & Clark, H. H. (1989). Understanding by addressees and overhearers. *Cognitive Psychology*, 21, 211–232.
- Scott-Phillips, T. C., Kirby, S., & Ritchie, G. R. S. (2008). Signalling signalhood and the emergence of communication. In A. D. M. Smith, K. Smith & R. Ferrer i Cancho (Eds.), *The evolution of language – Proceedings*

- of the 7th International Conference on the Evolution of Language (pp. 497–498). Singapore: World Scientific.
- Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: Bodies and minds moving together. *Trends in Cognitive Sciences*, 10(2), 70–76.
- Selten, R., & Warglien, M. (2007). The emergence of simple languages in an experimental coordination game. *Proceedings of the National Academy of Sciences of the United States of America*, 104(18), 7361–7366.
- Senghas, A., Kita, S., & Ozyurek, A. (2004). Children creating core properties of language: Evidence from an emerging sign language in Nicaragua. *Science*, 305(5691), 1779–1782.
- Suchman, L. A. (1987). *Plans and situated actions : The problem of human-machine communication*. Cambridge [Cambridgeshire]; New York: Cambridge University Press.
- Taub, S. F. (2001). *Language from the body*. Cambridge, England: Cambridge University Press.
- Tinbergen, N. (1963). On aims and methods in ethology. *Zeitschrift für Tierpsychologie*, 20, 410–433.
- Umata, I., Shimojima, A., Katagiri, Y., & Swoboda, N. (2003). *Interaction organisation in graphical communication*. Paper presented at the 25th Annual Conference of the Cognitive Science Society.
- Vermeerbergen, M. (2006). Past and current trends in sign language research. *Language & Communication*, 26(2), 168–192.