

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

An Experimental Investigation of the Role of Collaboration in the Evolution of Communication Systems

Permalink

<https://escholarship.org/uc/item/2046139d>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 31(31)

ISSN

1069-7977

Authors

Fay, Nicolas
Rogers, Shane
Swoboda, Nik
et al.

Publication Date

2009

Peer reviewed

An Experimental Investigation of the Role of Collaboration in the Evolution of Communication Systems

Bradley Walker (walkeb03@student.uwa.edu.au)

Nicolas Fay (nfay@psy.uwa.edu.au)

Shane Rogers (rogers03@student.uwa.edu.au)

School of Psychology, University of Western Australia, 35 Stirling Highway, CRAWLEY
WA 6009, Australia

Nik Swoboda (nswoboda@clip.dia.fi.upm.es)

Departamento de Inteligencia Artificial, Universidad Politecnica de Madrid, 28660-Boadilla del Monte,
Madrid, Spain

Abstract

Imitation alone cannot explain language evolution. Two additional ingredients have been proposed that may help explain the evolution of language systems: learning biases and social collaboration. An experimental method was developed that isolated the roles of collaboration and learning biases in the development of novel communication systems. Participants played a Pictionary-like task to develop ad hoc graphical communication systems in one of two conditions: one in which they interacted with a partner (Interaction condition), and one in which they received the same images from a “pseudo-partner” but did not interact (Pseudo-Interaction condition). Comparison of the resultant communication systems showed that the Interaction condition yielded higher identification accuracy, greater refinement of graphical signs and more alignment on a set of shared graphical signs (in fact, graphical alignment did not occur at all in the Pseudo-Interaction condition). Thus, collaboration plays a crucial role in the evolution of human communication systems.

Keywords: Language; collaboration; learning biases; cultural evolution; communication systems.

Background

It is widely believed that imitation alone can explain the evolution of human communication systems. Dawkins (1976) argued that memes, cultural units equivalent to genes, replicate and spread via imitation. As in biology, copying mistakes (mutations) promote variability and competition among memes, where more successful variants survive and drive memetic evolution. Blackmore (2000) argues that imitating communication techniques over a number of generations could have led, through the accumulation of such copying mistakes, to the emergence and development of language.

Until recently, imitation was considered a uniquely human ability, and therefore offered a plausible explanation as to why language is unique to humans. However, recent research shows that imitation also occurs in chimpanzees. Hopper et al. (2007) trained a high-ranking female chimpanzee to use a food-releasing mechanism by lifting a weight. The rest of her group was later allowed to observe her using the mechanism to obtain the food, and then were

able to try it out for themselves, most operating it successfully. Another group of chimpanzees in a “ghost” condition observed the mechanism being operated by invisible strings. None in this condition proceeded to use the mechanism successfully, indicating that observation of the moving mechanism was not enough; to adopt the techniques required to obtain the food, chimpanzees imitated each other, performing the same actions to achieve the same ends. Given this capability for imitation, why did language not arise in chimpanzees?

One answer is that imitation alone cannot explain the origin and evolution of language. The literature suggests two additional ingredients that may help explain the evolution of human communication systems: learning biases and social collaboration.

Learning Biases

Human learning biases may play a key role in the evolution of human communication systems. Like other cultural phenomena, communication systems spread through learning. Learning biases can therefore influence what is spread and how communication systems evolve.

The Iterated Learning Model (ILM) is a framework that has been used to investigate language evolution, highlighting the importance of constraints on learning. In a series of learner generations in the ILM, each learner receives input from the previous learner, develops hypotheses about the processes involved in its production, and then transmits their own output to the next learner based on these hypotheses. Learning plays a crucial role here, and any biases that restrict the learning of one thing over another will dictate what spreads through the population; learnable languages proliferate.

Computational studies using the ILM to simulate the evolution of language from earlier non-linguistic communication systems have isolated three aspects of human learning that are necessary for the development of linguistic structure: the capacity for generalization, bias against the acquisition of one-to-many meaning-signal mappings, and bias against the acquisition of many-to-one meaning-signal mappings (Brighton, Smith, & Kirby, 2005).

According to these simulations, without such biases language in its present state could not have evolved.

Behavioral studies employing the ILM framework support the notion that learning biases impose constraints that steer the course of cultural evolution. Kalish, Griffiths and Lewandowsky (2007) created a computer task in which participants learned one of four functions (positive linear, negative linear, non-linear and random) from corrective feedback presented after guessing the y -value associated with a given x -value. A test phase, with no feedback, followed this training phase. Participants' responses in the test phase served as input in the training of the next generation of learners in the chain. Each generation thus learned the function purely from the output of the previous generation, for a total of nine generations. Regardless of the original function, most chains of learners eventually settled on a positive linear function of response, suggesting that the information presented at the beginning of a chain has less impact than learning biases on the eventual form of this information at the end of the chain.

Collaboration

An alternative view is that social processes contribute to cultural and linguistic evolution. Tomasello et al. (2005) argue that the capacity to understand others as having goals and intentions is shared by humans and apes, but the motivation to share emotions and experiences with others is unique to humans. When individuals engage in activities that involve shared goals and require coordination of their actions to achieve those goals, they develop shared practices and beliefs. Like other cultural phenomena, language exists as shared knowledge and changes through interaction, causing cultural evolution (Fay et al., under review; Garrod & Doherty, 1994).

One mechanism for this change is alignment, where interlocutors start to think about and express concepts similarly during communication (i.e., they align). Studies of dialogue have found that grounding, the process of ensuring mutual understanding, leads to alignment (Clark & Wilkes-Gibbs, 1986). Alignment expands the information shared by interlocutors and thus lessens the effort required in production and comprehension (Pickering & Garrod, 2004). More shared information, or common ground, reduces the effort needed to achieve mutual understanding.

To study the origin and evolution of communication systems, several behavioral studies have used the Pictionary task (e.g., Fay, Garrod, & Roberts, 2008; Fay et al., under review; Garrod et al., 2007). Like the board game *Pictionary*, participants communicate concepts to one another by drawing, without using letters or numbers, and without speaking. When participants repeatedly communicate the same concepts, novel graphical communication systems emerge and evolve.

Participants in the Pictionary task take on director and matcher roles, where the director attempts to graphically communicate concepts to the matcher, and the matcher attempts to identify drawings produced by the director (from

a list of alternatives). Once each of the concepts on the director's list have been drawn, partners swap roles and repeat the game, using the same set of concepts (presented in a different random order in each game). Typically six games are played. By allowing the matcher to annotate the director's drawings, participants can engage in a graphical "discussion".

In this task, participants engaged in a collaborative activity with the shared goal of successfully communicating each concept. Through feedback and interaction, participants are able to coordinate their actions to achieve this goal. Garrod et al. (2007) found that over the course of the task the interacting participants' ability to identify each other's drawings improved, their drawings became less complex and more schematic, and more similar. In other words, communication became more effective, efficient, and aligned.

Aim and Hypotheses

Learning biases and collaboration have the capacity to explain language evolution. All humans have learning biases and all humans engage in collaborative activity; the question is whether one or the other, or both, is a prerequisite for language evolution. The aim of the present study is to determine the relative contribution of collaboration to the evolution of communication systems.

To do so, two conditions were contrasted. In condition one, the Interaction condition, pairs of participants were able to interact freely while engaged in the Pictionary task. In condition two individual participants were led to believe that they were interacting with a partner while playing the Pictionary task, but they were not. Instead, drawings produced by a member of an interacting dyad were played to them, and they attempted to identify the director's intended referent. These individual participants then produced drawings they believed their "pseudo-partner" would try to identify. This will be referred to as the Pseudo-Interaction condition, as participants thought they were interacting with their partner but only received and responded to drawings produced by their pseudo-partner. Thus the Interaction condition permits feedback and interaction within a given trial, whereas participants in the Pseudo-Interaction condition are restricted to simple production and reception, with no interaction.

If collaboration is not important to the evolution of communication systems, its potential benefits (more effective, efficient and aligned communication) may stem from underlying learning biases. In this case, communication systems developed in the Interaction and Pseudo-Interaction conditions will show similar levels of identification accuracy, graphical refinement and graphical alignment. Alternatively, if collaboration is important to the evolution of communication systems, then communication systems developed in the Interaction condition will be superior to those emerging from the Pseudo-Interaction condition.

Method

Participants

Forty-five students from the University of Western Australia participated in exchange for course credit. Ages ranged from 17-61 years, with a median of 20. Ten participants were male. Participants were randomly assigned to conditions, 15 pairs to the Interaction condition and 15 individuals to the Pseudo-Interaction condition.

Task and Procedure

Interaction Condition Sitting at separate computers, pairs of participants completed a computerized version of the Pictionary task. The task consisted of six games. In each game the director tried to communicate 16 concepts to the matcher, who selected a response from 20 items (16 targets plus 4 distracters). The same concepts were used in every game; for a complete list, see Table 1. Participants alternated roles between each game, one directing in games 1, 3 and 5 (matching in 2, 4 and 6), and the other directing in games 2, 4 and 6 (matching in 1, 3 and 5).

In this condition, matchers were able to freely annotate the directors' drawings (e.g., request clarification by circling part of the director's drawing). Additional feedback was given when a participant imitated or partly imitated their partner's prior drawing, suggesting partner comprehension. The trial ended when the matcher selected the concept they believed the drawing represented. The game ended when each of the 16 target items had been communicated.

A virtual whiteboard tool was used to manage the task and record all drawing activity (Healey, Swoboda, & King, 2002). A screenshot of the director's screen is shown in Figure 1. The list of concepts was displayed at the top of both participants' screens, with the drawing area directly below. For directors, the concepts were organized sequentially (though in a different, randomized order in each game), with the current to-be-communicated concept written in white. The concepts were displayed in a different random order on the matcher's screen. As participants drew, by pressing the left mouse button, their drawing activity was displayed simultaneously on both screens. To avoid confusion, directors drew in black ink and matchers in green ink. Both participants were able to erase parts of the drawing after clicking the "erase" button. Item selection

was made with a mouse; once done, both screens cleared and the game progressed to the next trial. Participants did not receive any feedback regarding their communicative success.



Figure 1: Screenshot of the director's computer screen from the Interaction condition of the Pictionary task. Gray parts of the drawing, which depicts a microwave, were drawn by the matcher and were green in the original.

Pseudo-Interaction Condition Here individual participants played six games of the Pictionary task, alternating between director and matcher roles from game to game. In games 1, 3 and 5, participants were matchers and attempted to identify the images drawn by their pseudo-partner. In games 2, 4 and 6 they directed, drawing images they believed would be matched by their partner. Each participant who directed in games 1, 3 and 5 in the Interaction condition served as the pseudo-partner of a participant in the Pseudo-Interaction condition.

Participants in this condition were unable to annotate the drawings they received, and, as no partner actually received their images, no feedback was given in later games by reuse of their drawing (unless fortuitously). As no partner attempted to identify the participants' drawings in this condition, the game progressed to the next concept when the director decided they had finished drawing (and clicked a "send" button), rather than when the matcher made their selection. To simulate the pseudo-partner selecting an answer, there was a random delay of 2 to 8 seconds before progressing to the next trial.

To ensure that matchers in the Pseudo-Interaction condition received the same information as those in the Interaction condition, they were not able to select an answer

Table 1: List of concepts.

	Places	People	TV Shows	Objects	Abstract
Targets	Theater	Russell Crowe	Drama	Microwave	Loud
	Art Gallery	Arnold Schwarzenegger	Soap Opera	Computer Monitor	Homesick
	Museum	Brad Pitt	Cartoon	Television	Poverty
	Parliament				
Distractors		Hugh Grant	Sci Fi	Refrigerator	Sadness

from the list of concepts until the image had been completely played back. Thus, matchers in the Pseudo-Interaction condition received exactly the same information as matchers in the Interaction condition. That Pseudo-Interaction matchers had no control over when the drawing was finished added consistency to the task, in that directors appeared to have full control over when they had finished drawing, irrespective of whether the participant was directing or matching.

Results

An example of the drawings produced in each condition is shown in Figure 2, where arrows indicate the sequence of images in the two conditions (solid arrows indicate communication between interacting partners and dashed arrows between pseudo-partners). Interaction images from games 1, 3 and 5 are part of both sequences. Parts of the game 3 and 4 drawings (Interaction condition) were produced by the matcher; these were green in the original but are shown here in gray.

As can be seen in Figure 2, the drawings produced in the Interaction condition became increasingly similar and simplified as the task progressed. Although the Pseudo-Interaction participant's drawings are similar to each other, they are dissimilar to their pseudo-partner's drawings. Drawings produced in the Pseudo-Interaction condition were simplified across the task, although less than those produced in the Interaction condition. Though not shown in Figure 2, identification accuracy increased in both conditions, but to a greater degree in the Interaction condition. These observations were confirmed by ANOVA.

Communicative Effectiveness

The effectiveness of the communication systems developed during the task was measured by the extent to which matchers were able to correctly identify the directors' intended referents (percentage of drawings correctly identified by matchers). As Pseudo-Interaction partners only matched in games 1, 3 and 5, the conditions were compared across those games.

Identification accuracy is plotted in Figure 3, where gaps in the Pseudo-Interaction condition at games 2, 4 and 6 correspond to games in which the participants were directing. As can be seen, matcher accuracy improved across games in both conditions, but improved more dramatically for matchers in the Interaction condition compared to those in the Pseudo-Interaction condition.

Accuracy scores were entered into a 2x3 mixed design ANOVA that treated condition (Interaction, Pseudo-Interaction) as a between-subjects factor and game (1, 3, 5) as within. This returned a main effect of game, $F(2, 56) = 33.89, p < .01$, partial $\eta^2 = .55$, and condition, $F(1, 28) = 7.01, p < .05$, partial $\eta^2 = .20$, but no game by condition interaction, $F(2, 56) = 3.01, p > .05$. Communication systems developed in both conditions became more effective across games, but matcher accuracy was consistently higher in the Interaction condition.




Figure 3: Accuracy of matcher responses over the task for both conditions.

Graphical Efficiency

Communication system efficiency was measured by calculating the complexity of each drawing. Drawing complexity was quantified using Pelli et al.'s (2006) perimetric complexity measure, calculated by dividing the




Figure 2: Arnold Schwarzenegger, as depicted by an Interaction pair and the corresponding Pseudo-Interaction individual.

squared perimeter of an image by the inked area. This has been shown to reliably correlate with subjective ratings of drawing complexity (Garrod et al., 2007). As Pseudo-Interaction participants only directed in games 2, 4 and 6, the analysis was restricted to those games.

Complexity data are plotted in Figure 4, gaps in the Pseudo-Interaction condition corresponding to games in which the participants were matching. Graphical complexity is consistently lower in the Interaction condition, although the complexity of graphical signs is reduced across games in both conditions.

Graphical complexity scores were entered into a 2x3 mixed design ANOVA treating condition (Interaction, Pseudo-Interaction) as a between-subjects factor and game (2, 4, 6) as within. This yielded a main effect of game, $F(2, 56) = 26.09, p < .01$, partial $\eta^2 = .48$, and condition, $F(1, 28) = 48.51, p < .01$, partial $\eta^2 = .63$. The game by condition interaction was not significant, $F < 1$. Drawings in both conditions became simpler, and therefore more efficient, but those produced in the Interaction condition were significantly more efficient throughout the task.




Figure 4: Perimetric complexity of images drawn in each condition over the task.

Graphical Alignment

As participants interact they adopt each other's expressions. In graphical terms, they adopt aspects of each other's drawings, which become increasingly similar. Alignment was measured as the rated similarity between the drawings produced by each participant and their partner. The similarity of drawings produced in adjacent games (e.g., Brad Pitt at games 1&2, 2&3, 3&4, 4&5, and 5&6) was rated by a judge (BW) on a scale of 0 (totally dissimilar) to 8 (practically identical). A second judge (NF) rated 20% of these images to test reliability. Their scores were highly correlated, $r = .83, p < .01$.

The similarity of drawings in adjacent games is shown in Figure 5, where scores above the neutral line at 4 are similar and scores below are dissimilar. Drawings produced in the

Interaction condition became increasingly similar across games, whereas drawings produced in the Pseudo-Interaction condition did not. A 2x5 mixed design ANOVA was conducted, treating condition (Interaction, Pseudo-Interaction) as a between-subjects factor and game (1&2, 2&3, ..., 5&6) as within. This returned a main effect of game, $F(4, 112) = 27.20, p < .01$, partial $\eta^2 = .49$, condition, $F(1, 28) = 39.75, p < .01$, partial $\eta^2 = .59$, and a game by condition interaction, $F(4, 112) = 11.24, p < .01$, partial $\eta^2 = .29$. The interaction effect is explained by the increasing similarity of drawings across games in the Interaction condition and the lack of change in the Pseudo-Interaction condition, where drawings remained equally dissimilar across games.




Figure 5: Similarity ratings of drawings across adjacent games in the Interaction and Pseudo-Interaction conditions. Scores below 4 indicate drawing dissimilarity whereas scores above 4 indicate drawing similarity.

Discussion

Communication systems evolve differently when partners are able or unable to interact. While identification accuracy in the Pseudo-Interaction condition did increase across games, it plateaued at around 80%. This contrasts with the Interaction condition, where identification accuracy rose to above 95% (i.e., ceiling). Similarly, although graphical refinement did occur in the Pseudo-Interaction condition, the signs remained significantly more complex than those in the Interaction condition, making them less efficient to produce and interpret. The effect of interaction on graphical alignment was striking; only in the Interaction condition did participants' drawings become increasingly similar across games. Graphical alignment was absent in the Pseudo-Interaction condition. In light of the communication benefits that arise from using the same signs (Pickering & Garrod, 2004), this finding represents a problem for accounts of language evolution that do not permit interaction between agents.

The benefits of collaboration to the evolution of communication systems suggest that language evolution cannot be explained by learning biases alone (see also Fay et al., under review). Social collaboration promoted more accurate communication systems that were less effortful to use. Furthermore, collaboration underpinned agents' ability to negotiate a 'shared' sign system (i.e., a single system used by both interlocutors). Presumably collaboration is responsible for much of the efficiency and effectiveness seen in today's languages (e.g., Clark & Wilkes-Gibbs, 1986).

Learning biases may still play a role in the evolution of communication systems, but this may be a smaller role than first anticipated. This has implications for the ILM, which relies heavily on the evolutionary influence of learning biases but has no collaborative component. Thus, it may not provide a complete picture of language evolution. A collaborative account may offer a better explanation of lexical evolution, whereas the ILM may better describe the evolution of syntax and grammar, which may be more strongly influenced by learnability.

We now return to the role of imitation. Drawing alignment is only observed in the Interaction condition, indicating that collaboration is crucial to the imitation of graphical signs. As partners interacted they established shared sign systems, such that over time it became increasingly unclear who was imitating whom. That is, it is unclear if participants were imitating the drawings produced by their partner, or the drawings produced by themselves, or a combination of the two.

The finding that collaboration plays a key role in the evolution of communication systems lends experimental support to theory concerning the role of collaboration in cultural evolution (e.g., Tomasello et al., 2005). Furthermore, there is no reason to believe that the effect of collaboration seen here is limited to the evolution of communication systems. Although communication may be the most significant of collaborative human activities, it is far from being the only one. It is quite likely that collaboration plays a similarly important role in cultural evolution more generally (e.g., the cultural evolution of concrete artifacts; see Caldwell & Millen, 2008).

Acknowledgments

This research was supported by an ARC Discovery Grant (grant DP0556991) awarded to N.F.

References

- Blackmore, S. (2000). The power of memes. *Scientific American*, 283, 64-73.
- Brighton, H., Smith, K., & Kirby, S. (2005). Language as an evolutionary system. *Physics of Life Reviews*, 2, 177-226.
- Caldwell, C. A., & Millen, A. E. (2008). Experimental models for testing hypotheses about cumulative cultural evolution. *Evolution and Human Behavior*, 29, 165-171.
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22, 1-39.

- Dawkins, R. (1976). *The selfish gene*. Oxford, UK: 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*, 363, 3553-3561.
- Fay, N., Garrod, S., Roberts, L., & Swoboda, N. (under review). The interactive evolution of human communication systems.
- Garrod, S., & Doherty, G. (1994). Conversation, coordination and convention: An empirical investigation of how groups establish linguistic conventions. *Cognition*, 53(3), 181-215.
- Garrod, S., Fay, N., Lee, J., Oberlander, J., & MacLeod, T. (2007). Foundations of representation: Where might graphical symbol systems come from? *Cognitive Science*, 31, 961-987.
- Healey, P. G. T., Swoboda, N., & King, J. (2002). A tool for performing and analyzing experiments on graphical communication. *Proceedings of the HCI02 Conference on People and Computers XVI* (pp. 65-68). London, UK: Springer.
- Hopper, L. M., Spiteri, A., Lambeth, S. P., Schapiro, S. J., Horner, V., & Whiten, A. (2007). Experimental studies of traditions and underlying transmission processes in chimpanzees. *Animal Behaviour*, 73, 1021-1032.
- Kalish, M. L., Griffiths, T. L., & Lewandowsky, S. (2007). Iterated learning: Intergenerational knowledge transmission reveals inductive biases. *Psychonomic Bulletin & Review*, 14, 288-294.
- Pelli, D. G., Burns, C. W., Farell, B., & Moore-Page, D. C. (2006). Feature detection and letter identification. *Vision Research*, 46, 4646-4674.
- Pickering, M., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27, 169-226.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28, 675-735.