

Integrating the Horizontal and Vertical Cultural Transmission of Novel Communication Systems

Carrie Theisen-White (carrietheisenwhite@gmail.com)

Department of Psychology, Yeshiva University
2495 Amsterdam Avenue, New York, NY 10033, USA

Simon Kirby (simon@ling.ed.ac.uk)

Language Evolution and Computation Research Unit, University of Edinburgh
Dugald Stewart Building, 3 Charles Street, Edinburgh EH8 9AD, UK

Jon Oberlander (jon@inf.ed.ac.uk)

School of Informatics, University of Edinburgh
Informatics Forum, 10 Crichton Street, Edinburgh EH8 9AB, UK

Abstract

In recent experimental work on the emergence of novel communication systems, the cultural transmission has tended to be either horizontal or vertical. We present an experiment that integrates the two. Results show that both methods of transmission do some work.

Keywords: horizontal transmission; vertical transmission; iterated learning; systematic compositionality; systematicity; compositionality.

Introduction

A new field which may be called ‘experimental semiotics’ (Galantucci, 2009) involves the experimental investigation of novel forms of human communication. These studies have examined how the negotiation, repeated use, and transmission of a communication system between individuals through a process of ‘iterated learning’ (Kirby & Hurford, 2002) leads to adaptation of that system. This work has shed light on how language might have evolved to have its unique features such as systematic compositionality and arbitrariness.

For example, Kirby, Cornish, and Smith (2008) demonstrated how systematic compositionality could emerge from the vertical transmission of communication systems through generations of participants. Each participant learned an ‘alien language’ and then was tested on it. The language consisted of words for a highly structured set of items (27 items, where each was one of three shapes, was one of three colors, and had one of three motions). The words that the first participant learned were randomly generated strings of syllables. Each subsequent participant learned the words that the previous participant produced during his or her test. There was a catch: each participant was trained on only half of the language but tested on the whole language. (Participants were unaware of this.) Kirby et al. (2008) found that the languages became increasingly systematically compositional (i.e. increasingly,

words for similar items shared syllables) through the generations.

Garrod, Fay, Lee, Oberlander, and MacLeod (2007) showed how the horizontal transmission of signs back and forth between participants in a closely interacting dyad can lead to the emergence of arbitrariness. Each round, one of the participants (the ‘Drawer’) has an ordered list of items and the other (the ‘Matcher’) has an unordered list of the same items. The Drawer takes each of the items in turn and produces a signal for it on a virtual whiteboard the two share, so that the Matcher can identify the item. Participants can draw whatever they like but may not write anything, such as letters or numbers. Pairs play for a number of blocks, so that each item is drawn and identified several times over the course of the game. Among other measures of how the signals for the items changed with interaction, Garrod et al. (2007) found that graphical complexity (roughly, the amount of virtual ink used) decreased. They argued that this complexity is a measure of iconicity – all else equal, the less information there is in the signal, the less the signal can resemble its meaning – thus showing that signs become more arbitrary with interaction.

These two studies are quite representative of the field. Most experimental semiotics studies thus far have involved either horizontal transmission of signs (i.e. a closed set of individuals, usually a dyad, producing signs for each other) or vertical transmission of signs (i.e. chains of individuals or groups of individuals, in which the first produces signs for the second, the second produces signs for the third, and so on). Recent reviews (Galantucci & Garrod, 2010, 2011; Scott-Phillips & Kirby, 2010) have noted this dichotomy and identified the understanding of the relative contribution of the different types of cultural transmission as a key goal for future work.

Fay, Garrod, Roberts, and Swoboda (2010) and Garrod, Fay, Rogers, Walker, and Swoboda (2010) began to bridge the gap between horizontal and vertical transmission by pitting them as competing theories of the evolution of language. For example, Garrod et al. (2010) present a

graphical communication task with two conditions: one implementation of vertical transmission and one implementation of horizontal transmission. They found that, of these implementations, only the horizontal transmission led to simplification of signs (i.e. to arbitrariness).

We propose that *integrating* horizontal transmission and vertical transmission in one experiment would be informative as well. In fact, it may even be necessary. Even if it could be shown that horizontal transmission resulted in a communication system with all the crucial features of language, the story cannot end there – the communication system must be transmitted beyond the closed set of individuals who developed it, and work has shown that this is far from straightforward (Galantucci, Theisen, Gutierrez, Kroos, & Rhodes, in press). Similarly, even if it could be shown that vertical cultural transmission resulted in a communication system with all the crucial features of language, the story cannot begin there – unless we want to assume that one individual invented a communication system on his own and transmitted it to the next generation, we must allow that horizontal transmission shaped at least the beginning of language. In addition, it is clear that in the real world, generations are not neatly divided and populations are not static, so both processes are undoubtedly in operation.

Current Work

Our previous work (Theisen, Oberlander, & Kirby, 2010) is a good start to integrating horizontal and vertical transmission. In it, we showed that systematic compositionality spontaneously arises from horizontal transmission. Pairs of participants played a Pictionary-style game much like that of Garrod et al. (2007). There were two main differences. First, the set of items to be communicated was structured: each item could be thought of as one of five entity types (such as person or building) that relates to one of ten themes (such as education or agriculture). This was to allow systematic compositionality to emerge. Second, neither participant in a pair learned in advance what items they would communicate about; the items simply appeared in the game. This was to ensure that the game introduced no direct pressure for systematic compositionality. Figure 1 shows a subset of one pair’s final signs, i.e. what was drawn on the whiteboard for each item the last time that item appeared in the game. Notice how systematically compositional they are – signals for similar items share an element. For example, the signals for activities all include rows of squiggly lines. In fact, the systematic compositionality of the pairs’ final sets of signs was significantly higher than the systematic compositionality of ‘mixed final’ sets we constructed – sets of final drawings taken from *different* pairs of participants – showing that pairs’ final drawings of items systematically re-use drawing elements more than can be attributed to a tendency across pairs to draw these items a certain way. Interestingly, pairs’ initial sets of signs (what was drawn on the whiteboard for

each item the first time that item appeared in the game for them) were also significantly more systematically compositional than mixed initial sets. Further, the systematic compositionality of pairs’ signs did not increase significantly over the course of their games. Rather, it seemed that systematic compositionality spontaneously arose in pairs’ sets of signs and was maintained through their games.




Figure 1. A subset of the signs developed by one pair in the experiment reported in Theisen et al. (2010). Italics distinguish which participant was drawing. The sets of signs were systematically compositional – signals for similar items (e.g., activities) shared an element (e.g., rows of squiggly lines).

It can be hard to imagine how this could happen. As discussed in Theisen et al. (2010), the shared interaction history between partners in a pair seems to play a large role. Consider the drawings of *school bus* in Figure 2. For pairs A – D, *school bus* was the first primary education item to be drawn, so the drawings reflect what people draw for *school bus* when uninfluenced by previous drawings of primary education items – roughly, what motivated signals (i.e. signals that have an inherent connection to their meanings) for *school bus* look like. In contrast, Pair E drew *school bus* after they had already drawn another primary education item: *teacher*. Their drawing of *school bus* appears to re-use elements from their drawing of *teacher*, i.e. to be strongly influenced by signs in the pair’s shared interaction history. In particular, note that a chalkboard in a drawing for *school bus* is not strongly motivated – none of the uninfluenced drawings of *school bus* include a chalkboard. While one might expect that the first time someone draws *school bus* with their partner, they would draw it no differently than if they were drawing with a new partner, instead it appears

that the history of interaction shared by two people influences even new signs they create for use with each other.


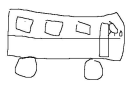




	teacher	school bus
Pair A		
Pair B		
Pair C		
Pair D		
Pair E		

Figure 2. The first drawings of school bus from five different pairs. For pairs A – D, each was the first primary education item to be drawn in their game. In contrast, Pair E drew another primary education item (teacher) before drawing school bus. Their drawing for school bus appears to re-use elements from their drawing for teacher, viz. the chalkboard.

The aim of the current work was to integrate horizontal and vertical transmission in one experiment. Since we knew that systematic compositionality can emerge from vertical transmission (Kirby, et al., 2008) or horizontal transmission (Theisen, et al., 2010), we added vertical transmission to the design of our previous experiment and measured systematic compositionality. In particular, a pair played a Pictionary-style game for a period of time. The next pair played the same game but before starting was taught a subset of the signs the first pair used at the end of their game, the third pair was taught a subset of the signs the second pair used at the end of their game, and so on. Four such chains of four generations (pairs) each were formed.

This design allowed us to isolate the effect of horizontal transmission from that of vertical transmission. We could identify any effect of horizontal transmission by simply looking at the first generation's signs and by comparing each generation's initial signs with their final signs. We could identify any effect of vertical transmission by seeing whether anything changes over the four generations. We

hypothesized that, in addition to the effect of horizontal transmission on systematic compositionality seen in Theisen et al. (2010), we would see an effect of vertical transmission on systematic compositionality. The alternative hypothesis is that, given the effectiveness of horizontal transmission on systematic compositionality, there is no role for vertical transmission to play.

Method

Participants 24 University of Edinburgh students, both male and female, participated in exchange for £12. All were native British English speakers. Participants who played together did not know each other.

Apparatus Partners were seated in separate soundproof booths with computers. The game was run using the Pigeon software (Healey, Swoboda, & King, 2002), which presented the item to draw each trial and provided a shared online whiteboard. Participants guessed and corrected their partners' guesses in an MSN Messenger chat window.

Game The items about which the participants communicated were chosen to share salient semantic features; each item can be thought of as one of five entity types (such as person or building) that relates to one of ten themes (such as education or agriculture). There were 26 core items, no more than one of each entity-theme combination. The items appeared with different frequencies, between three and eight times every 126 trials. Additionally, there were 14 filler items, occurring just once per game, intended to prevent participants from assuming that their set of items was closed. The items occurred in random order. Participants knew nothing about the items in advance. In particular, they were never exposed to a list of the items.

Each trial, one participant was the Drawer and other was the Guesser. The Drawer saw an item (such as professor) on his screen and was allowed to draw immediately. The Drawer drew with a mouse, had only black ink, and could not erase anything. The Guesser saw everything the Drawer drew immediately, on her screen. The Guesser did not see the Drawer's mouse movements when he was not drawing, and could not draw herself. When she was ready, the Guesser guessed by typing into a chat window. The Drawer stopped drawing immediately and either confirmed or corrected the guess in the chat window. Players advanced themselves to the next trial. Every six trials, the participants switched Drawer and Guesser roles. The participants played for two hours.

A pair was allowed just one guess per trial. A pair won 1 point for every correct guess but lost 1 point for any incorrect guess or drawing that included a pre-existing symbol or convention (including writing). The goal was to win as many points as possible in the two hours of play. Participants from the three top-scoring pairs were entered into a prize draw for an additional £20.

Observation Participant pairs were organized into four chains of four generations each. The first generation of each of the four chains was chosen at random from the 12 games of the experiment reported in Theisen et al. (2010). These pairs played the same exact game described above but did not observe anyone else’s drawings before playing. Generation 1’s final drawings were observed by Generation 2, Generation 2’s final drawings were observed by Generation 3, and Generation 3’s final drawings were observed by Generation 4.

Pairs observed the previous pair’s signs for 15 minutes. Each participant in a pair saw screenshots of what was drawn for the last 30 trials of the previous generation’s game, in the order they occurred in the original game. In practice, this amounted to observing the signs for 15 – 20 (M = 18.25) of the 26 core items from the previous generation. For each screenshot, the subjects read what the item was, who drew it, and what the Guesser guessed for it. The participants studied each screenshot and its associated information for 20 seconds each.

Procedure Participants read instructions, which described the basics of both the game and the observation phase. Importantly, participants read that observing others’ signs might or might not help them in their own game and that their task when drawing was to do whatever would get their partner to guess correctly and quickly (i.e. not necessarily to copy signs they had observed). Pairs observed the previous generation’s signs for 15 minutes, played the game for 2 hours, and then were debriefed.

Analysis

The experimenter coded each pair’s initial set of signs, their final set of signs, and the set of signs they observed from the previous generation for systematic compositionality.

An initial set of signs or final set of signs consists of one signal for each of the 26 core items. Observed sets consisted of fewer items. When an observed set included more than one drawing for a core item, we coded the last drawing, reasoning that it would be most salient to the pair.

Each set of drawings was printed on a page in a table, organized so that rows and columns contained drawings for similar items. The coder examined each row and each column for any element shared among two or more drawings. The coding instructions stated that the coder should mark an element only if there appeared to be a special understanding between the players to draw certain things certain ways. If there was a shared element, the coder marked which of the drawings in that row or column included it. For the observed sets, if the drawings in a category shared an element that was prohibited (e.g., the cross for medical items), those drawings were excluded from the analysis. This is because the next generation was not allowed to copy a prohibited element. This only happened once. In addition, because the set of items observed was random, sometimes only one item in a

category was observed. Of course, there can be no element shared across drawings in this case, so these drawings were also excluded from the systematic compositionality analysis. This happened just six times, in just four sets of signs.

Each drawing was inspected twice – once as a member of its row and once as a member of its column. Thus, except for a few cases of observed sets (as noted above), each set of drawings could receive a total score of 52. The total score divided by 52 (or the total score possible, in the case of the observed sets) is our systematic compositionality score.

The sets were coded blind and in random order. In addition, the full sets were coded before any of the observed subsets, to prevent the coder from inadvertently looking for sign elements she marked in the subsets while coding the full sets. Reliability of this coding procedure was established in Theisen et al. (2010).

Results

Figure 3 shows the final drawings for fire engine and fire station for the four generations in one chain. Generation 1’s signs for fire engine and fire station do not share an element with each other. However, generation 2 changed the signal for fire station to include an element from fire engine: what looks like a ladder. Generations 3 and 4 maintained this. Changes like these occurred across pairs’ sets of signs.

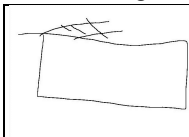
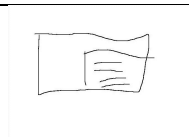




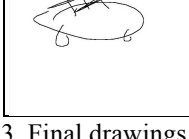
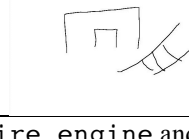
	fire engine	fire station
Gen. 1		
Gen. 2		
Gen. 3		
Gen. 4		

Figure 3. Final drawings of fire engine and fire station from the four generations of one chain. Generation 2 changed the signal for fire station to include an element from fire engine. Changes like these occurred across the pairs’ sets of signs, leading to an increase in systematic compositionality over generations.

Figure 4 shows the systematic compositionality in each pair’s final set of signs, organized by chain and generation.

The graph suggests that systematic compositionality is increasing over generations within a chain. Page's Trend Test confirmed this ($L = 111$, $p < 0.05$).

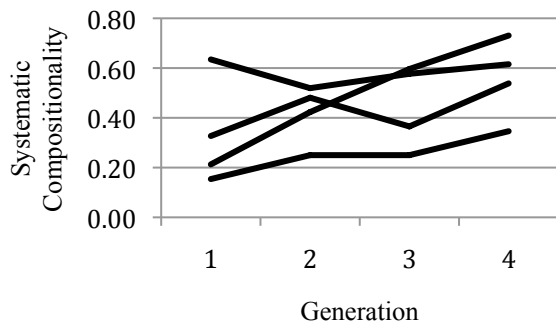


Figure 4. Systematic compositionality of final sets of signs by generation. Each line represents a chain. There is a trend for systematic compositionality to increase over generations.

There are a few possibilities for where the increase in systematic compositionality originated. It could happen that pairs' final sets of signs were less systematically compositional than the subsets of them observed by the next generation. Specifically, by chance, the idiosyncratic signs (signs that do not share elements with other signs) might not have been observed. Since our systematic compositionality measure is a proportion, this would make the systematic compositionality of the observed subset greater than that of the final set. This turns out not to be the case. A Wilcoxon Signed-Rank Test found no significant difference between the systematic compositionality scores of final sets of signs and that of their associated observed subsets ($M_{\text{Final}} = .415$, $SD = .165$; $M_{\text{Observed}} = .364$, $SD = .196$, $p = 0.945$).

It could also happen that systematic compositionality increased during each generation's game, but this is not the case either. As in the experiment reported in Theisen et al. (2010), pairs did not increase the systematic compositionality from their initial signs to their final signs – a Wilcoxon Signed-Rank Test found no significant difference between the systematic compositionality of pairs' initial sets of signs and that of their final sets of signs. ($M_{\text{Initial}} = .441$, $SD = .128$; $M_{\text{Final}} = .450^1$, $SD = .171$, $p = 0.40$)

The third candidate for the increase in systematic compositionality is from the set of signs a pair observed to the initial signs the pair produced. A Wilcoxon Signed-Rank Test showed that this too was below significance ($M_{\text{Observed}} = .364$, $SD = .196$; $M_{\text{Initial}} = .441$, $SD = .130$, $p = 0.065$). However, comparing this with the other two possibilities

¹ This mean is slightly different than the mean reported in the previous paragraph because that was the mean systematic compositionality of the final sets of signs produced by pairs who had observed another pair's set. The current mean is the mean systematic compositionality of the final sets of signs produced by all pairs (including the first generation, who did not observe anyone's signs).

suggests that this is where the increase in systematic compositionality occurs.

We already knew that systematic compositionality spontaneously arises through horizontal transmission in pairs in the game. (Theisen, et al., 2010) These new results show that vertical cultural transmission then increased the systematic compositionality. Thus, the communication systems were shaped by *both* interaction between members of a closed group and the transmission of a communication system from some individuals to others.

Discussion

It is important to emphasize that it is highly unlikely that, if the first generation of pairs had played for four times as long, the systematic compositionality would have increased as much as it did over four generations. Recall that systematic compositionality did not increase over the course of the pairs' games – the systematic compositionality of the pairs' initial signs was not significantly different from their final signs. This held both for the pairs reported in Theisen et al. (2010), from which the generation 1 pairs were drawn, and the pairs reported here.

It is interesting to notice that the trend of an increase in systematic compositionality over generations is significant but the increase in systematic compositionality from the set of signs a pair observed to the initial set of signs they produce is not quite significant. In fact, this can be understood as a feature of cumulative cultural evolution, in which possibly very small changes at the generational level are amplified. (Kirby, Dowman, & Griffiths, 2007)

What particular features of each type of transmission were responsible for the systematic compositionality? As discussed above and in Theisen et al. (2010), for the pairs in generation 1, the shared interaction history between participants in a pair seemed to play a large role in the appearance of systematic compositionality (which was immediate) – the new signs that partners created for use with each other tended to re-use elements from signs they already shared. As for vertical transmission, we noted above that the increase appears to originate from the set of signs a pair observed to the initial signs the pair produces. The point of transmission gives each subsequent generation the opportunity to replace some idiosyncratic signs with signs that make the set more systematically compositional. This could be because the observing pair did not observe a signal for a particular item or because they did observe it but did not adopt it (because the participants couldn't remember it, it didn't make sense to the participants, or the participants rejected it for some other reason). Interestingly, then, one could argue that it is the same thing responsible for the systematic compositionality in both the horizontal transmission and the vertical transmission cases: a demand for novel signs.

We set out in this paper to integrate horizontal and vertical transmission in one experiment, answering the call of several recent reviews. We did this, and found that both horizontal transmission and vertical transmission had an

effect on systematic compositionality. This result lends strong support to our argument above that the full history of a communication system must include both its horizontal transmission between members of a closed set of individuals and its vertical transmission beyond them through generations of individuals. We hope that future work will not ignore this.

Acknowledgments

This work was supported by the European Union Integrated Project JAST (Joint Action Science and Technology), grant FP6-IST2-003747. The third author also gratefully acknowledges the support of the JAMES project (Joint Action for Multimodal Embodied Social Systems), grant FP7-ICT-270435.

References

- Fay, N., Garrod, S., Roberts, L., & Swoboda, N. (2010). The Interactive Evolution of Human Communication Systems. *Cognitive Science*, 34(3), 351-386.
- Galantucci, B. (2009). Experimental semiotics: A new approach for studying communication as a form of joint action. *Topics in Cognitive Science*, 1(2), 393-410.
- Galantucci, B., & Garrod, S. (2010). Experimental semiotics: A new approach for studying the emergence and the evolution of human communication. *Interaction Studies*, 11(1), 1-13.
- Galantucci, B., & Garrod, S. (2011). Experimental Semiotics: A Review. *Frontiers in Human Neuroscience*, 5(Article 11).
- Galantucci, B., Theisen, C., Gutierrez, D., Kroos, C., & Rhodes, T. (in press). The Diffusion of Novel Signs Beyond the Dyad. *Language Sciences*.
- 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.
- Garrod, S., Fay, N., Rogers, S., Walker, B., & Swoboda, N. (2010). Can iterated learning explain the emergence of graphical symbols? *Interaction Studies*, 11(1), 33-50.
- Healey, P., Swoboda, N., & King, J. (2002). A tool for performing and analysing experiments on graphical communication. *People and Computers*, 55-68.
- Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, 105(31), 10681-10686.
- Kirby, S., Dowman, M., & Griffiths, T. (2007). Innateness and culture in the evolution of language. *Proceedings of the National Academy of Sciences*, 104(12), 5241-5245.
- Kirby, S., & Hurford, J. (2002). The emergence of linguistic structure: an overview of the iterated learning model. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (pp. 121-148): Springer-Verlag New York, Inc.
- Scott-Phillips, T., & Kirby, S. (2010). Language evolution in the laboratory. *Trends in Cognitive Sciences*, 14(9), 411-417.
- Theisen, C., Oberlander, J., & Kirby, S. (2010). Systematicity and arbitrariness in novel communication systems. *Interaction Studies*, 11(1), 14-32.