Published in final edited form as:

Lang Cogn. 2012 November ; 4(4): 381–418. doi:10.1515/langcog-2012-0020.

The emergence of duality of patterning through iterated learning: Precursors to phonology in a visual lexicon

Alex Del Giudice

Department of Linguistics, University of California San Diego.

Alex Del Giudice: adelgiud@ucsd.edu

Abstract

Duality of Patterning, one of Hockett's (1960) proposed design features unique to human language, refers in part to the arrangements of a relatively small stock of distinguishable meaningless sounds which are combined to create a potentially infinite set of morphemes. Literature regarding the emergence of this design feature is less abundant than that exploring other levels of structure as focus is more often given to the emergence of syntax. In an effort to explore where combinatorial structure of meaningless elements arises the results of two pilot experiments are presented within which we observe human participants modifying a small lexicon of visual symbols through a process of iterated learning. As this lexicon evolves there is evidence that it becomes simpler and more learnable, more easily transmitted. I argue that these features are a consequence of spontaneous emergence of combinatorial, sub-lexical structure in the lexicon, that the pattern of emergence is more complex than the most widely espoused explanation suggests, and I propose ways in which future work can build on what we learn from these pilot experiments to confirm this hypothesis.

Keywords

Iterated learning; experimental methodologies; sub-lexical structure

1 Introduction

Among research on the evolution of language a greater focus is placed on the emergence of syntactic structure, how morphemes and words are combined, than the emergence of structural properties at the sub-lexical or sub-morphemic level, how meaningless elements combine to form meaningful ones. Of the wealth of computer simulations and mathematical models exploring the evolution of language, only a small proportion focus on the emergence of combinatorial structure at that level (Liljencrants and Lindblom 1972; Nowak et al. 1999; de Boer 2001; Oudeyer 2001, 2006; Zuidema and de Boer 2009). Research of precursors to phonology in animal behavior and animal communication systems abounds (for reviews see Yip 2006 and Fitch 2000), though even in this work, the animal behavior in question is often viewed in terms of how it may relate to complex morphology or syntax, such as in Ouattara et al. (2009), even when that work better speaks to the emergence of structure at the sub-lexical level: from meaningless calls to meaningful components. And we begin to see more experimental work exploring how humans in a laboratory setting create a novel communication system, as in the work of Galantucci (2005), Fay et al. (2008), or Kirby et al. (2008), though this and research like it generally focuses on participant's manipulation of meaningful elements even when those words and morphemes, their sub-lexical components,

are created anew or modified by the participants. The comparative lack of discussion on this topic may in part result from the notion that the main question here has been answered: as the meaning space of the language increases, and more words are needed, sub-lexical combinatorial structure fills that need in the most efficient way (Hockett 1960; Lindblom et al. 1984; Studdert-Kennedy 1998, 2005).

However, recent evidence has emerged suggesting that sub-lexical structure may not necessarily result only from a large meaning space. Where established sign languages do show clear evidence of duality of patterning, researchers propose that a new sign language (Al-Sayyid Bedouin Sign Language, ABSL) lacks the sub-lexical component of this feature despite its speakers' apparent use of a large meaning space (Sandler et al. 2011). It may seem possible that the signal space of a sign language is larger than that of a spoken language, and that it may allow for the production and perception of more distinct signals, such that the pressure for duality of patterning to emerge is weakened. However, established sign languages appear to have a phoneme inventory size that is within the range of those found in spoken languages (Rozelle 2003) and there is little reason to believe that ABSL is different in this regard.

Nevertheless, in an effort to further elucidate our understanding of how the combinatorial, sub-lexical aspects of duality of patterning might emerge I present the following experiments. The pilot work I discuss here intends to test the viability of an experimental paradigm within which we can explore the emergence of sub-lexical structure in a laboratory setting. To do this, I combine the experimental paradigm used by Kirby et al. (2008), which incorporates the concept of iterated learning (Kirby and Hurford 2002), and an element of the paradigm used by Galantucci (2005), which allows for the production of a visual signaling system with speech-like properties. In contrast to previous and existing work, I combine these two paradigms in a specific effort to observe the evolution of this communication system in a controlled environment where

- a. participants use a transmission medium that is influenced by their existing lexicon *as little as possible* and
- b. participants are not explicitly tasked with creating a system, but that any systematic components emerge spontaneously and are retained and expanded by new learners in successive generations.

In these two pilot experiments I establish two diffusion chains of learners, with each learner representing a "generation": each acquires a small visual lexicon and then reproduces it for the next learner to acquire. The state of the lexicon at each generation is itself examined and compared to those of earlier generations to discover if aspects of combinatorial sub-lexical structure emerge.

Despite some methodological problems to be discussed, evidence from these experiments suggest that, via this iterated learning process, meaningless units emerge and spread throughout the lexicon. This emergence of the sub-lexical level of duality of patterning occurs despite the lexicon being relatively small (at fewer than 15 symbols) with much of the available signal space left unexplored. Instead we see that incremental changes at each generation, either intentional or accidental, lead to the re-use of units smaller than the 'word' and this may be reinforced for its ability to make the system more easily learnable.

The paper is structured as follows: In Section 2, I describe the general methodology and major components utilized in the two pilot experiments, as well as some motivations for those choices. In Sections 3 and 4, I present a more detailed explanation of Experiments 1 and 2, respectively, and present the results of those studies. In Section 5 is a qualitative analysis of the data, addressing some aspects of the results that suggest evidence for the

emergence and recombination of subunits. I conclude with Section 6 where I address shortcomings with the methodology, discuss insights gained from these experiments and suggest directions for future work.

2 General Methodology

The goal of this study is to observe how participants modify a small, visual lexicon with little overt sub-lexical structure, and explore whether such structure, or its increased use, emerges in this setting. Laboratory experiments that examine the emergence of properties of communication systems have explored an array of phenomena including morphosyntactic-semantic compositionality (Kirby et al. 2008), sign arbitrariness from initial iconicity (Garrod et al. 2007; Fay et al. 2008), lexical systematicity from semantic structure (Theisen et al. 2010), and systematic structure through increased semantic complexity (Galantucci 2005). Taking cues from this avenue of research, I show how the emergence of a component of duality of patterning may be similarly investigated.

The communication system used in these pilot studies, modified slightly from the implementation found in Galantucci (2005), resembles the transient and linear qualities of speech but is visual in nature. Such a system is used so that learners are less likely to make conscious comparisons to their own language. It allows for some level of iconicity, arguably more than is allowed by speech alone, but limits the kind of iconography possible in a pen-and-paper drawing. Participants create marks with a Wacom Bamboo digitizing pad and stylus. Those marks are translated, in real time, as a black trace onto a clear white panel on a computer screen (presented on a 17-inch MacBook Pro laptop display). Following Galantucci's design, "The horizontal component of the stylus' motions directly controls the horizontal component of the trace's motion on the panel," but "the vertical component of the trace's motion was independent of the stylus' motions, moving with a constant downward drift" (Galantucci 2005: 741). Therefore, irrespective of where on the vertical axis the participant begins, the line on the display panel always begins at the top border of the designated space. When the stylus first makes contact with the digitizing pad, the trace point begins to move downward at approximately 250 pixels per second and if the stylus remains in contact with the digitizing pad, a line is created by the trace point. If the stylus is removed from the pad, the trace point continues to move in a downward direction but no mark is made until the stylus makes contact again. When the trace point reaches the bottom border of the display panel, no more marks can be made (see Figure 1 for examples of this process).

Such a system discourages participants from adopting an orthographic or illustrative strategy for transmitting symbols. Some level of iconicity is still possible: for example a numerical/counting strategy like the one developed by some participants in Galantucci's experiment (see panels 5 and 6 of Figure 1 for an example). The implementation used in the following experiments differs slightly from the original in that the symbols created here do not fade with time. Instead, after participants begin producing a symbol, they have a limited time in which to finish the symbol before the trace point reaches the bottom of the display panel. The choice to not incorporate this aspect was practical as it was better suited to presenting each generation of learners with symbols that could be learned quickly and accurately.

The language creation and transmission tasks instantiate the iterated learning model through a diffusion chain of human participants. Iterated learning is the process by which an individual learns and masters some behavior by observing someone (or the product of someone) who has mastered it in the same way (Kirby and Hurford 2002). In the experimental paradigm, this process is emulated by creating a chain of participants, each of who attempts to learn and reproduce the linguistic behavior that was the output of the previous participant's reproduction. Using this process small changes created by each

participant can accumulate overtime such that no particular individual is solely responsible for creating whatever structure may emerge. Instead, it has been shown in both agent-based computer simulations (Kirby 1999; Kirby et al. 2004) and in laboratory experiments (Kirby et al. 2008) that linguistic compositional structure arises not deliberately but as an effect of the accumulation of factors which are not necessarily under any conscious control by the learners and producers.

In these pilot experiments, the first participant of each chain does not create a symbolic lexicon de novo. Instead a “seed” lexicon forms the input for the first generation of each chain. This seed, used for all chains in both experiments (though the lexicon used in Experiment 2 incorporates an additional symbol) was created using the same communication method described above by a group of 4 naïve participants in a pre-piloting phase. Each of these participants was given approximately 50 attempts to use the communication method to create whatever symbols they could. From the collected set of these 200 “doodles”, 55 were removed for being nearly identical to other symbols in the set¹ and 13 symbols to be used in the experiments were randomly selected from the remaining 145 (see Figure 2). Given the method with which these symbols were selected, it is inevitable that some components making up a symbol may share strong similarities with those making up another. For example, note how symbols 2 and 10 both commence with two perpendicular lines. In examining the output of participants’ learning of these symbols we shall keep in mind whether these components played a role in whatever structure might emerge.

3 Experiment 1

In this first experiment, we explore the effect of iterated learning on the evolution of symbols, which are not assigned any explicit meaning.

3.1 Methods

Participants in this experiment were all undergraduate students or recently graduated students from the University of Edinburgh. All had full use of their dominant writing hand, had normal or corrected vision, and were monolingual speakers of British or American English. Two diffusion chains of participants were run: Chain 1 consisted of seven participants (mean age:20.9, SD:1.6, 6f, 1m, all right-handed) and Chain 2 consisted of five participants (mean age:20.8, SD:3.0, 3f, 2m, 4 right-handed). The difference in numbers of participants is attributed to absenteeism and lack of participation during the small window of time in which the experiment was conducted. All were recruited via an advertisement on an Internet based job advertisement board and were paid 3GBP for approximately 25 minutes of participation.

Participants were told they would be playing a simple *language* game in which they would learn some symbols in an alien language called “Ixwy”. They were informed that the instrument used to create the symbols is not like “pen and paper,” but is intended to mimic the kind of instrument that Ixwy aliens would use to communicate. The seed lexicon for the first participant consisted of the 13 symbols presented in Figure 2.

Each “round” consists of two phases: (1) a learning phase: where participants are exposed to all the symbols; and (2) a recall phase, where participants attempt to recall and produce all the symbols from memory.

¹Participants in this pre-piloting phase would sometimes create two or more identical symbols in a row when they ran out of ideas for possible symbols.

In the learning phase, the participant is presented with a randomly selected symbol for 3 seconds then is immediately instructed, on screen, to copy it and press the space bar to move on. Once all 13 symbols are seen and copied the recall phase begins. In this phase, participants are asked to recall/produce all 13 symbols in any order. They do so by producing each symbol on a blank screen, with the exception of a number in the top right corner. Participants produce one of any symbols they can recall and then press the space bar, which clears the screen and increments the number in the corner so the next recalled symbol can be produced.

Participants generally have difficulty recalling all 13 symbols, but the experiment requires that they produce 13 even if they cannot remember them all. For whatever symbols they cannot recall, participants are prompted to create one that “might fit well with the rest of the symbols of the Ixwy language”. After the 13th symbol is produced this second phase is over and a new instruction screen is displayed informing participants that a new round will begin, identical in form to the previous, and to “prepare to do more copying”. For each new round, the order of the 13 symbols presented in the copy phase is randomized. In the fourth and final round, the 13 symbols created during this recall phase are used as the input for the following participant in the chain. Throughout all four rounds, the symbols presented to a participant always come from that same set of input symbols (the participant never sees her own productions of the symbols again).

The 13 symbols created in the fourth recall phase are processed after completion. The experimenter assigns each symbol a number (1–13) corresponding to the image from that participant’s input lexicon that is the most likely origin. These pairings between input and the final recall symbols are assigned in the following way:

- i. Generally more than half of the symbols participants produced resembled the original input symbols more closely than any other input symbol so that each could be associated with its progenitor with some certainty.
- ii. Each participant’s productions from this final **recall** phase were also compared with her productions from the preceding final **copy** phase (where participants copied symbols that they had just seen). The majority of differences between the input symbols and the final recall symbols could be reconciled this way since differences among the symbols were often produced by participants in the copy phases as well.²
- iii. Any symbols that could not be assigned to an input symbol using the methods above were assigned to the most similar unpaired symbol – as judged by the experimenter – from the input set.

Symbols were organized in this way so as to better compare them across generations, but participants were never influenced by this organization as they always saw a random ordering of the input symbols. Similarities and differences among input and recall symbols noted by the experimenter are not necessarily those noted by the participant during the task. Even an automatic/computational process (such as the one used in the quantitative analysis below) or a set of interrater agreements from a group of human raters would not be free of bias or be expected to approach the judgment of each participant during performance. Therefore, due to the subjective nature of each step of this process care should be taken with any assumptions of symbol lineage across generations.

²The fact that the differences between input and recall symbols also appear during the copy phase, coupled with comments collected during debriefing with participants suggest that several of these differences may not have been a result of poor recall, but a result of either difficulty with manipulating the abstract communication system and/or inaccurate copying stemming from participants’ belief that accurate copying was impossible given the “inaccurate” mapping from stylus to the screen.

In an effort to quantitatively analyze the results, I use a computer algorithm to measure differences among all the symbols within a generation's lexicon and among all the symbols across two lexicons from two successive generations. The first measurement (intra-lexicon distance) provides a rough measure of reuse of symbol sub-units and other similarities among symbols within one participant's learned lexicon (i.e. how similar each symbol is to any other symbol in the lexicon). The second (inter-lexicon distance) provides a measure of each participant's accuracy in learning and reproducing the lexicon (i.e. how similar each generation's lexicon is to the previous generation's). This method of analysis was performed by Verhoef et al. (2011) and is incorporated with small changes here.

To calculate a distance measure between two symbols, each symbol is transformed into a vector representing the x-coordinate (or horizontal coordinate) of the line that makes up that symbol. Only the x-coordinate is necessary because the y-coordinate changes at the same steady rate with respect to time for all participants. The difference between any two such vectors is calculated with the aid of "derivative dynamic time warping" (Keogh and Pazzani 2001). In this process a *derivative vector* is calculated for each symbol vector by calculating the slope between each two successive elements of the symbol vector. The derivative is crucially used in order to reduce the impact of small reproduction errors in the symbols, most often such errors were where small and unintentional breaks in a solid line are introduced.³ The resulting derivative vector is then compared with another symbol's derivative vector by calculating the difference between them (equivalent to a Euclidian distance measure of the two derivatives).

However, as these differences are computed, the time axis (i.e. the y-coordinates) for either one of the symbols might be "warped" (e.g. various segments are compressed or expanded with respect to time) to best match any similarities within the pair of symbols, if any are found. This is done to control for small timing differences between the production of two symbols. This process is similar to elongating or compressing elements within two speech samples, such as equalizing the length of vowels in two production of the same word with the goal of comparing their other qualities (for a simplified example of this, see Table 2 below).

Calculating the distance between every possible pair of symbols in one generation's lexicon and summing the shortest of those distances for each symbol is how intra-lexicon distance is measured. For example: in a simple lexicon of four symbols, we might find the following distances

1. A and B = 5⁴
2. A and C = 6
3. **A and D = 3**
4. **B and C = 5**
5. **B and D = 4**
6. C and D = 7

Thus, the shortest distances for all four symbols (A and D = 3, B and D = 4, C and B = 5, and D and A = 3) give us a minimal intra-lexical distance of 15.

³In cases where there is no line derivatives are calculated as 0.

⁴The magnitude of these values is not representative of the distance values collected from the actual data. Magnitude here is reduced for simplicity.

Measuring inter-lexicon distance requires a more complex approach. In an optimal scenario, each symbol in a generation's lexicon was derived from one matching symbol in the preceding generation's lexicon. In order to find these matches automatically one symbol is randomly selected from generation 'A' and it is compared to each symbol of the other generation 'B' until one with the shortest distance is found. Next, a second symbol is chosen from generation A and compared to all the *unpaired* symbols from generation B. This process continues until all symbols in both generations are paired and then the differences across all pairs are summed. However, there might be more than one optimal set of pairings given this approach, thus finding the global optimum requires searching all possible random order permutations for creating such pairings. In order to make the task more computationally tractable a random search is performed by finding the smallest sum of differences from 50,000 unique potential sets of pairings. The sets of pairings used in the reported results are those with the smallest distance found in these searches.

3.2 Results

The lexicons resulting from the two chains of participants are presented in Figures 3 and 4. The rows are the collection of symbols created at each generation (the lexicon produced by each participant in the fourth and final recall/reproduction phase), with the first row being the initial seed lexicon. The columns represent the experimenter's association of symbols across each generation as determined using the process described in Section 2. Notable patterns seen to emerge in this data will be discussed in Section 5 below.

Intra-lexical distances for both chains are presented in Figure 5. There is a decreasing trend in the distance in both chains suggesting that the symbols within each successive generation's lexicons are becoming more similar. Note that the intra-lexical distance is highest for the seed lexicon (labeled here as Generation "0") suggesting that similarities among words, both holistic or resulting from the potential reuse of subunits, are smaller in the seed lexicon than the lexicons of all other generations.

The inter-lexical distance for both chains is reported in Figure 6. The downward trend notable in these data suggests that each successive generation's symbols are more similar to those created by the previous generation. This further suggests that the language is becoming more easily learnable across the generations as the symbols are more easily reproduced and/or remembered. This trend is also in part due to the symbols becoming simpler to produce. Due to the nature of the symbol creation system, certain symbols require more planning or a better understanding of the communication system. Symbols that tend to be more difficult to reproduce are often altered toward greater ease of production, and as each generation makes these alterations accurate reproduction of the system as a whole can increase.

4 Experiment 2

This second experiment principally differs from the first in that participants were presented symbol-meaning pairs rather than symbols alone. The same communication system was used as in Experiment 1, but the learning task included the implicit mapping of the symbols to a set of meanings (in the form of cartoon images). Meanings were included in this symbol transmission task to determine what role, if any, iconicity or semantic similarity might play in the development of the symbolic system. The presence of meanings could lead to the introduction of changes that favored some kind of iconic representation, making the link to their meanings more transparent. In addition, re-use of sub-lexical pieces, should it emerge, might be found to correlate with the semantic classes of the meanings, such that some or all animate objects or some or all geometric shapes would share some sub-unit (as per the findings of Theisen et al. 2010) providing evidence that the emergence of sub-lexical

structure may be tied to the emergence of what are traditionally believed to be more meaningful levels of structure.

4.1 Methods

All participants were undergraduate or recently graduated students from the University of Edinburgh and undergraduates from the University of California – San Diego. All had full use of their dominant writing hand, had normal or corrected vision, and were monolingual speakers of British or American English. Two diffusion chains were created, the first consisted of 7 participants (mean age: 21.6, SD: 2.15, 4f, 3m, 6 right-handed) and the second consisted of 9 (mean age: 21.8, SD: 3.70, 6f, 3m, all right-handed). Participants at the University of Edinburgh were recruited via an advertisement on an internet based student job board and were paid 3 GBP for approximately 35 minutes of participation (participants that took longer than 45 minutes were given an additional 1 GBP). Participants at the University of California San Diego were recruited for course credit. Once again, discrepancies between the number of participants resulted from difficulties attracting participation.

The same list of 13 symbols used in Experiment 1 (Figure 2) were again used as the seed lexicon for this experiment with the addition of one more symbol here presented in Figure 7. This addition was made because pre-pilot data suggested that recall accuracy on this task was higher than that of Experiment 1. For the seed lexicon only, these 14 symbols were randomly assigned to a specific set of 14 of 20 images forming the set of ‘meanings’ in this experiment (Figure 8). These images depicted elements from 4 categories: Animate natural objects [man, woman, fish, bird, dog], natural inanimate objects [tree, sun, rock, flower, mushroom], abstract shapes [pentagon, circle, triangle, star, cube], and manufactured inanimate objects [car, house, pencil, apartment, computer].

As before, participants in this experiment are told they will be playing a simple language game in which they will learn some symbols in an alien language called “Ixwy”. Learners are informed that the instrument they will be using to create the symbols won’t feel much like using a pen and paper but is intended to mimic the instrument Ixwy aliens use to communicate.

Experiment 2 involved three distinct trials, a practice trial, a learning trial, and a final trial.

4.1.1 Practice trial—This trial mimics the learning trial from Experiment 1 and is intended to provide the initial exposure to the symbol-meaning pairs in the form of a *guessing game*. Participants are presented with 3 images randomly selected from the set of 20 possible meanings and 1 Ixwy symbol which represents the meaning for one of the 3 images. Each image is labeled with a number (1 for the image on the left, 6 for the image in the center, and 0 for the image on the right). The symbol is presented in the lower 4/5 of the screen and the three “meanings” are equidistant and presented in a row on the top fifth of the screen. Participants use the number keys at the top of the keyboard to select one of the three images that they believe represents the correct meaning for that symbol. Upon making a guess by pressing a key, the symbol disappears leaving a blank canvas in the bottom 4/5 of the screen. If the participant guesses correctly, a green check mark is superimposed over the selected meaning and a message congratulates them on the correct guess and prompts them to reproduce that symbol as they remember it. If participants guess incorrectly, a red cross is superimposed on the meaning they selected and a green check mark is superimposed over the correct meaning. A message reiterating the correct meaning is displayed along with the prompt to reproduce the symbol. The order of the 3 images is random for each symbol of the trial and the foils are selected randomly from any of the remaining 19 images of the total set. This practice trial only occurs once per participant. It continues until all 14 symbol-meaning

pairs have been seen and each of the 14 symbols has been copied, at which point the learning trial begins.

4.1.2 Learning trials—Participants complete four learning trials. These trials are similar to the practice trial except that after each guessing phase, participants no longer simply copy the symbol they have just seen. Now, after feedback about their guess is reported, they are instead shown a randomly selected image from the set of 14 learned images and are prompted to produce the *symbol* corresponding to that given *meaning*. Thus, while the initial practice trial simply exposes participants to the whole lexicon and allows them to copy a symbol that was just observed, the tasks in the learning trials involve alternating between recalling a meaning given a symbol and recalling a symbol given a meaning. The learning trial continues until participants make a guess for the meaning of all 14 symbols and recall/reproduce the symbol for all 14 meanings. For each of the four repetitions of this trial, the order in which symbol-meaning pairs are presented for both guessing and reproducing are randomized.

4.1.3 Final trial—The final trial mimics the recall trials of Experiment 1 except, rather than recall symbols from memory, participants are prompted to recall symbols with their paired meanings. The “guessing-game” component is removed in an effort to eliminate concurrent exposure to input symbols during the final recall. Participants are presented with all 14 meanings one at a time, in random order, and are tasked with reproducing that meaning’s symbol. In addition, interspersed throughout those 14 reproductions, they encounter 6 new meanings and are tasked with creating novel symbols for those meanings. These 6 images may have been presented throughout the trials as foils during the guessing phase, but participants had not been shown “Ixwy” symbols for these meanings in any trial. In the instructions preceding this phase, participants are told to invent some new symbols for these meanings “which would fit in with the rest of the Ixwy Language”.

Rather than transmitting all 20 of the symbols created in the final trial, only a subset is used as the input for the next trial. Of the 14 symbols learned and recreated in the learning trials by each generation, the same set of 10 are transmitted to the next generation so that the development of these 10 can be tracked across all generations. The other 4 symbols transmitted come from the set of 6 created anew during the final trial. These 4 transmitted symbols are alternated across each generation (See Table 3 for further clarification). The purpose of this is to provide an opportunity for innovation to enter into the lexicon. In addition, by asking participants to innovate the six new symbols we create an opportunity to examine whether the method of creating new symbols is influenced by an increased perceived structure in the learned lexicon, providing insight into whether participants generalize any perceived sub-lexical structure within the 14 learned pairs when creating symbols for novel meanings. Transmitting a set of novel symbols to the next generation should ensure that innovation would be allowed to influence the lexicon learned by the following generation.

4.2 Results

The 10 symbol-meaning pairs which were learned across all generations formed one group (Figure 9 and Figure 11 for Chain 1 and 2, respectively), and the remaining 10 symbol-meaning pairs, 4 of which were learned and 6 of which were innovated at each generation, form another (Figure 10 and Figure 12). Notable patterns will be discussed in Section 5. However, it is crucial for upcoming analysis to note a particular result in Generation 2 of Chain 1. Here we note that nearly all of the symbols created during the final recall trial fail to resemble those of the preceding generation. The participant representing this generation had difficulty both in using the novel communication system, but also performed very

poorly in the two recall tasks (recalling symbols given a meaning, and recalling a meaning given a symbol). This led to frustration and a lack of effort such that the participant began to “doodle” whenever tasked with recalling symbols, showing little effort to attempt to memorize symbol-meaning pairs. Consequently, the resulting lexicon was almost entirely created anew and many of the symbols looked very much alike making the task more difficult for following generations.

In Figure 13, the mean score for the “guessing game” across the 4 learning trials is plotted for both Chains 1 and 2. Here we note the dip in performance at Generation 2 of Chain 1 but also note that performance in this task recovered in later generations suggesting that, while the resulting lexicon of Generation 2 may have been difficult to learn, changes made to it in later generations may have aided the recovery seen here.

Intra-lexical distance was calculated over all the symbols within a generation, those that were both learned and created anew. The results of the intra-lexical distance measure for both chains of Experiment 2 are presented in Figure 14. Here, as in Experiment 1, we find a reduction in Intra lexical distance (distance between symbols within a single generation), though the pattern in Chain 1 is noticeably different. The steep drop and plateau notable there is related to the steep drop in symbol complexity that occurred at Generation 2. Due to a lack of motivation, the participant representing Generation 2 created many symbols that were either nearly identical or variations on a specific theme, resulting in very small distance measures.

The inter-lexical difference is presented in Figure 15. In calculating inter-lexical distances, distance measures are calculated only for those 14 symbols that were learned from the previous generation and not for the additional 6 symbols created *de novo*. Here, the comparably high value notable at Generation pair 2 (comparing Generation 2’s lexicon to Generation 1’s) is related to the nearly completely re-invented lexicon at Generation 2.

5 Qualitative analysis

Close observation of the symbols created in each chain by each generation provides hints that sub-lexical units are emerging. Sometimes these units appear for two or three generations only to disappear almost as quickly. Other times the units survive, they remain in the lexicon for several generations and seem to spread to new words. The abbreviated discussion that follows is an attempt to highlight some of the emergent aspects that seemed most striking to the author. This discussion is not exhaustive, however it is also subjective as I am not aware of any useful methodology for the *unbiased* and direct assessment of the existence sub-lexical units in this kind of symbolic system. I leave it to the reader to formulate judgments about the validity of the following observations. In the closing Section I discuss design methods to minimize such problems in future studies.

5.1 Use of space/location

In Chain 1 of Experiment 1 the first generation is seen to recall and reproduce 12 of the 13 symbols comparatively accurately, though symbol 3 appears to be a new invention that closely resembles Symbol 7⁵ (see Figure 16). Two features distinguish symbols 3 and 7: The high frequency wave patterns at the top and bottom are linked by a straight line in 7 but are separated in 3, and symbol 7 appears to the right of center and 3 to the left.

The location distinction between these two symbols propagated through the next 6 generations and by Generation 7 it appears to have been incorporated into other signs (e.g.

⁵In debriefing, the participant representing this generation reported they believed to have perfectly recalled all 13 symbols.

see Generation 3, Symbols 6, 11, and 13). By that final generation, as we will see is often the case in all diffusion chains in both experiments, symbols become less complex, having lost many features that previously distinguished them. However, the left/right location distinction of the symbols remains throughout one of the more salient characteristics by which signals may have been differentiated by participants.

This use of left versus right space reappears in Chain 2-Experiment 2 to distinguish symbols 4 and 6 (representing the meanings “car” and “dog”, see Figure 11 reprinted below). At Generation 5, symbol 6 was changed or reinterpreted to resemble that of symbol 4. But while symbol 6 was produced mostly near the center of the space, symbol 4 was produced to the left of center. The following generation, in reproducing these symbols, eliminated the extra segment in Symbol 4 thus making the two symbols more similar, but the distance between them is exaggerated by symbol 6 having shifted to the right of center. This distinction was retained in the lexicon of the next 3 generations.

5.2 Similarities among segments embellished

In Generation 4 of Chain 1-Experiment 1, we see that symbol 2 has been slightly altered from that of the previous generation: the two relatively straight lines are substituted by two high amplitude “peaks” (see Figure 17 for more detail). There is no way to know for certain why this symbol was changed in this way, but these two peaks are reminiscent of elements that also appear in symbols 8 and 11 in the same generation suggesting a potential source of this change. The next generation appears to amplify the similarity of these high amplitude peaks across all three symbols. Here we see a path toward the emergence of sub-lexical structure: Symbols which already share *perceived* similar structure with others may be modified to better match that structure. The resulting similarities and differences may then be more clearly noted by successive generations. However, this effect was short-lived: production or recall errors in Generation 6 caused the similarities among these symbols to be less striking leading Generation 7 to lose them entirely. Such loss is commonplace in this incarnation of the iterated learning paradigm since each generation depends on the efforts and skills (not to mention errors or lack of motivation) of a single participant.

We see a similar sequence of events in close examination of Chain 2 of Experiment 2 (Figures 11 and 12 reproduced below). The “s” shaped stroke (S) that appears first in Symbols 3, 8, 9, 14 of Generation 3 is incorporated into symbols 4, 6 and 13 in Generation 4, and then in symbols 2, 7, 15, 16, and 17. By Generation 6 nearly every item in the lexicon is made up of this unit. Furthermore, we note that beginning with Generation 5 and continuing with all the following generations, for all those symbols that participants were tasked with creating anew (see Figure 12), this S unit was incorporated in most or all of these. This suggests that, consciously or not, the participants in Generation 5 and beyond have extracted a generalization about the components of the lexicon and how to best create symbols that “fit well with the language”.

5.3 “mirroring”

The results of Chain 2-Experiment 1 (Figure 4, reproduced below) we note how space can be used for more than just left-right distinctions. In Chain 2 we can find several symbols that are near-mirror images. For example, symbols 9 and 10 of Generation 5 are reflections of one another, as are 11 and 13, 5 and 8, etc. This mirror strategy is innovated in Generation 1 (see symbols 6 and 11) and transmitted within some lexical pair or pairs across each generation until Generation 4.

A similar strategy also makes an appearance in the last few generations of Chain 2-Experiment 1, first notable in Generation 2 where symbols 8 and 9 (representing the

meanings “circle” and “rock”) are distinguishable by their inversion across a central axis. This element survives through the next 4 generations even while the shape of the mirrored elements slowly changes. Then, beginning in Generation 4, symbols 5 and 10 (representing the meanings “sun” and “bird”) begin to share components that are inverted along the center axis (see Figure 11).

5.4 Segmentation

Segmentation of symbols plays an important role in the resulting sub-lexical structure of the lexicon in all chains. There is a clear trend toward segmentation of continuous symbols. As they are transmitted, discontinuous symbols seldom become more continuous and their discontinuity is often embellished, while continuous symbols can more often be seen to become discontinuous. For example, see the evolutions of the following symbols:

Experiment 1 – Chain 1 symbols 6, 7, 11, and 12

Experiment 1 – Chain 2 symbol 12

Experiment 2 – Chain 1 symbol 6

Experiment 2 – Chain 2 symbols 1, 7, 8, and 9

The initial emergence of segmentation may have often been the result of production error. Accidentally removing the stylus from the pad during production caused a small “break” to appear in the symbol. However, breaks in the signal were generally magnified by all generations. This preference appears to be an additional foundation within which emerging sub-lexical structure can gain a foothold, at least in this kind of communication system. Debriefing interviews suggest that a common strategy used in recalling signals involved participants making references to the differences in the numbers of peaks or of ‘S’ units among symbols. Such an overt strategy might result in the observed trend toward segmentation.

5.5 How units emerge from meaning

In Figure 18, we compare the symbol for “man” (symbol 2) and the symbol for “woman” (symbol 18). Here, below each of the darkened line we see the behavior of a participant tasked with creating the symbol referring to “woman”. It would appear that in all three such generations (1, 3, and 5), the strategy for creating this symbol involved modifying the symbol learned for ‘man’ by adding or subtracting a component to it. There is no clear evidence that the component being subtracted or added is itself meaningful (referring to maleness or femaleness). Thus we note here an additional avenue for the emergence of meaningless units, one in which perceived semantic links might play a role.

6 Discussion

In these pilot experiments I have provided some initial evidence of how elements of sub-lexical structure emerge in a small lexicon of visual symbols that participants learn and produce in a process of iterated learning. The sub-units making up this structure are introduced spontaneously by some generation, then are noted to appear in other items as the lexicon is transmitted, and then to pervade the lexicon as predicted in the iterated learning model.

These pilot experiments have flaws that need to be remedied for more empirically sound, future iterations. Below I discuss some of these problems and potential solutions, and conclude with how the data from these pilot experiments contributes to our understanding of the emergence of duality of patterning.

6.1 Methodological shortcomings

These pilot experiments were designed as a proof of concept in utilizing existing methods with slight modifications toward a new purpose. Therefore, whatever conclusions might be drawn from the data need to be tempered with understanding of the potential flaws and drawbacks of these novel experiments.

First, despite an inclusion of two computational, quantitative measurements of certain aspects of the data, the conclusions largely rely on a subjective analysis. To the extent that we must rely on human judgments for identifying elements like “sub-lexical units”, in future incarnations, an inter-rater reliability task would serve to more accurately identify such units.

However, even if multiple raters of these symbols agree that, for example, the left-right or mirrored characteristics, noted ex post facto, are a real and salient characteristic of the data it should be independently verified that they are actually useful to participants’ learning or production. Future attempts using this paradigm should test whether lexicons where these elements are most salient, such as is generally the case in the final generations of the diffusion chains, are more easily learned than lexicons where they are not so salient, such as in the first generation (for an example of such a test see Verhoef 2012). Should we find this to be the case, then we have better evidence that the lexicons are being modified toward increased learnability, and that sub-lexical units of these kinds may be playing a role.

An additional difficulty, which has been noted in other manifestations of the iterated learning experimental paradigm (Kirby et al. 2008), is related to the loss of expressivity that often occurs within these lexicons. It is apparent in the data from both experiments that symbols across the whole lexicon can become less distinct over time, and that this is not always a result of emerging structure. In certain cases, this is attributable to errors made by a single participant who either failed to recall or reproduce many or all of the symbols accurately. Given that the lexicon is always passed through the mind of a single individual, each successive lexicon is only as “good” as that participant is able or willing to make it, and in a paradigm where there is little external pressure for accuracy there is less motivation to learn and reproduce symbols accurately. There are several methods for correcting these concerns, such as having each generation represented by multiple participants who may or may not interact with each other through the learning process.

Lastly, two additional issues were of concern to an anonymous reviewer. First, the instructions given to participants with regards to creating novel symbols (i.e. when participants were tasked with creating a new symbol for a given meaning, or when they were forced to create a symbol that they could not recall), namely to create a symbol that “fits well with the rest of the Ixwy language” was potentially biasing participants to create structure. This is certainly a possibility, and future experiments might be better served from less explicit instructions, such as simply “create a symbol”. And, lastly, there is a discrepancy in the size of each chain, resulting from the small pool of participants available at the place and time when this experiment was run, and a related concern is that only two chains of participants create the whole of the data being analyzed in each experiment. Comparisons across chains would certainly be better founded if those chains had equal participants, and better and more accurate generalizations could be made from a larger number of chains. Future work will be more balanced in this regard. However, it should be noted that many of the observations made with this data were corroborated in Verhoef (2012) where these problems were addressed with a more balanced set of diffusion chains.

6.2 Contributions

Despite the drawbacks of these pilot experiments, the data gathered with this new methodology shows much promise. The main result of these experiments is that the emergence of this level of duality of patterning exhibits a more complex pattern than is predicted by the most common assumptions, those encapsulated by Hockett (1960: 12) himself: “one can find little if any reason why a communicative system should have [duality of patterning] unless it is highly complicated. If a vocal-auditory system comes to have a larger and larger number of distinct meaningful elements, those elements inevitably come to be more and more similar to one another in sound. There is a practical limit for any species or any machine, to the number of distinct stimuli that can be discriminated, especially when the discriminations typically have to be made in noisy conditions.”

So, given such assumptions, why should we find hints of the emergence of sub-lexical structure in a closed system of symbols? If the optimum strategy in developing symbols is to differentiate them in order to increase transmission fidelity, then we would expect that the symbols in the lexicon would most often change such that differentiation would be maximized. Granted, only Experiment 2 includes a pressure to keep symbols discriminable. Yet we find that as symbols, which already begin in the seed lexicon as fairly discriminable, move toward indiscriminability, maximization of the distinctiveness of the symbols can occur as a strategy to counteract this move. In *both* experiments we find that near-identical symbols did not often survive longer than 2 generations. A possible interpretation is that participants focus on slight differences between the two similar symbols and then embellished both the similarities and the differences.

If Hockett’s assumptions hold we would expect an influx of symbols to fill the signal space, yet we note here that in participants’ novel creations the entirety of the signal space was never used. Those participants who felt comfortable creating new symbols with the communication system tended to be conservative⁶. So rather than explore some new dimension of the signal space (i.e. modifying the length of each segment, the frequency of each curve, the length of each segment break), participants generally utilized dimensions that already played a role in the lexicon as they perceived it (i.e. left vs. right space, mirror image). Furthermore, those symbols that were strikingly distinct from many of the others within a generation, those that best and most clearly exploited the limits of the signal space by incorporating some new dimension, often failed to survive to the next generation. Such symbols were either forgotten or changed slightly to gain elements from other symbols. When participants invented new lexical items, they inclined to use elements that already appear in other familiar items, especially if those elements appear in more than one such item as we noted with the emergence of the “S” unit in Chain 2-Experiment 2.

As the lexicon is constructed by more and more of such units it may also be passed on to the next generation with higher accuracy. This is because items with familiar sub-units are learned more quickly by association with other items made up of those sub-units, and production of these items is more accurate since the increased frequency of the sub-units requires participants to learn fewer complex motions.

There is some evidence that the emergence of re-used units among symbols improved learning: when two symbols shared a sub-unit participants tended to recall them as a pair. In Experiment 2 participant debriefings suggest that this resulted in an association of the symbols’ meanings. For example, one participant in Chain 2 reported that they recalled the similar symbols for “dogs and cars because they are both things that live in the garage”,

⁶Though it is equally important to remember that participants were given instructions to invent symbols which they felt “fit well in the Ixwy language.”

while another suggested that the similarities between the symbols for “bird and sun” arose from the fact that they are both things you find in the sky. But there is evidence for the link between symbol reuse and memory also in Experiment 1 in the order that participants recalled symbols in the recall phase. For example, in Chain 1, Symbols 2 and 10 each begin with two horizontal lines in Generations 1, 2, and 3. During the final recall phase in each those generations, these two symbols were recalled and produced one immediately after the other. Similarly, in that chain, Symbols 3 and 7 (those exhibiting the left-right distinction throughout all 7 generations) were recalled one immediately after the other in 4 of the 7 generations. In Chain 2, this was also true of symbols 6 and 11 (mirrored pairs) for Generations 2, 3 and 4. Thus, the use of combinatorial sub-lexical structure may emerge not as a necessity of an overburdened signal space but as a main strategy for remembering or storing words, a strategy that is compatible with many modern models of lexical access (Marlsen-Wilson 1987; Goldinger et al. 1989).

The results of this research are preliminary as they lead us to formulate more questions than they resolve. I have argued that when human participants learn a set of word-like visual symbols then produce them, with imperfections, for other human participants to learn from, combinatorial sub-lexical structure can emerge. This emergence of a component of duality of patterning cannot be said to result entirely from a large meaning space taxing a small signal space. The data suggests that new explanations are needed, and experiments such as these provide a novel method for testing those explanations.

Acknowledgments

I would like to thank two anonymous reviewers whose astute observations and comments greatly increased the quality of this paper. In addition special thanks go to Carol Padden and Simon Kirby who were instrumental at all stages of formulating the questions relevant to this research, designing the pilot experiments, and interpreting the results. I would like to thank Tessa Verhoef for insightful discussion, and help applying the methodology behind the quantitative analysis. This work was supported by a grant from the National Institutes of Health (NIDCD 0006473).

References

- de Boer, B. The origins of vowel systems. Oxford: Oxford University Press; 2001.
- De Boer B, Sandler W, Kirby S. New perspectives on duality of patterning: Introduction to the special issue. *Language and Cognition*. 2012; 4(4)
- Fay N, Garrod S, Roberts L. The fitness and functionality of culturally evolved communication systems. *Philosophical Transactions of the Royal Society: Biological Sciences*. 2008; 363:3553–3561. [PubMed: 18799421]
- Fitch WT. The evolution of speech: A comparative review. *Trends in Cognitive Science*. 2000; 4:258–267.
- Galantucci B. An experimental study of the emergence of human communication system. *Cognitive Science*. 2005; 29:737–767. [PubMed: 21702792]
- Garrod S, Fay N, Lee J, Oberlander J, MacLeod T. Foundations of representation: Where might graphical symbol systems come from? *Cognitive Science: A Multidisciplinary Journal*. 2007; 31(6): 961–987.
- Goldinger SD, Luce PA, Pisoni DB. Priming lexical neighbors of spoken words: Effects of competition and inhibition. *Journal of Memory and Language*. 1989; 28:501–518.
- Hockett CF. The origin of speech. *Scientific American*. 1960; 203:88–96. [PubMed: 13683472]
- Keogh E, Pazzani M. Derivative dynamic time warping. *Proceedings of the First International SIAM Conference on Data Mining*. 2001:1–11.
- Kirby, S. *Function, selection, and innateness: The emergence of language universals*. Oxford: Oxford University Press; 1999.

- Kirby, S.; Hurford, JR. Cangelosi, A.; Parisi, D. Simulating the evolution of language. London: Springer-Verlag; 2002. The emergence of linguistic structure: An overview of the iterated learning model; p. 121-148.
- Kirby S, Cornish H, Smith K. Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*. 2008; 105(31):10681–10686.
- Kirby S, Smith K, Brighton H. From UG to universals: Linguistic adaptation through iterated learning. *Studies in Language*. 2004; 28:587–607.
- Liljencrants J, Lindblom B. Numerical simulations of vowel quality systems: The role of perceptual contrast. *Language*. 1972; 48:839–862.
- Lindblom, B.; MacNeilage, PF.; Studdert-Kennedy, M. Self-organizing processes and the explanation of phonological universals. In: Butterworth, B.; Comrie, B.; Dahl, O., editors. *Explanations for language universals*. Berlin: Mouton de Gruyter; 1984. p. 181-203.
- Marslen-Wilson WD. Functional parallelism in spoken word-recognition. *Cognition*. 1987; 25:71–102. [PubMed: 3581730]
- Nowak MA, Krakauer D, Dress A. An error limit for the evolution of language. *Proceedings of the Royal Society of London, Series B, Biological Sciences*. 1999; 266(1433):2131–2136.
- Ouattara K, Lemasson A, Zuberbühler K. Campbell's monkeys concatenate vocalizations into context-specific call sequences. *Proceedings of the National Academy of Sciences*. 2009; 106(51):22026–22031.
- Oudeyer, PY. Coupled neural maps for the origins of vowel systems. In: Dorffner, G.; Bischof, H.; Hornik, K., editors. *Papers from the parasessions on language and behavior*. Chicago: Chicago Linguistics Society; 2001. p. 178-203.
- Oudeyer, PY. *Self-organization in the evolution of speech*. Oxford: Oxford University Press; 2006.
- Rozelle, L. *The structure of sign language lexicons: Inventory and distribution of handshape and location*. University of Washington doctoral dissertation. 2003.
- Sandler W, Arono M, Meir I, Padden C. The gradual emergence of phonological form in a new language. *Natural Language and Linguistic Theory*. 2011; 29(2):503–543. [PubMed: 22223927]
- Studdert-Kennedy, M. Evolutionary implications of the particulate principle: Imitation and the dissociation of phonetic form from semantic function. In: Knight, C.; Studdert-Kennedy, M.; Hurford, J., editors. *The evolutionary emergence of language*. Cambridge: Cambridge University Press; 1998. p. 161-176.
- Studdert-Kennedy, M. How did language go discrete. In: Tallerman, M., editor. *Language origins – Perspectives on evolution*. Oxford: Oxford University Press; 2005. p. 48-67.
- Theisen CA, Oberlander J, Kirby S. Systematicity and arbitrariness in novel communication systems. *Interaction Studies*. 2010; 11(1):14–32.
- Verhoef T. The origins of duality of patterning in artificial whistled languages. *Language and Cognition*. 2012; 4(4)
- Verhoef, T.; Kirby, S.; Padden, C. Carlson, L.; Hölscher, C.; Shipley, T. *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society; 2011. Cultural emergence of combinatorial structure in an artificial whistled language; p. 483-488.
- Yip MJ. The search for phonology in other species. *Trends in Cognitive Science*. 2006; 10(10):442–445.
- Zuidema W, de Boer B. The evolution of combinatorial phonology. *Journal of Phonetics*. 2009; 37(2): 125–144.

Fig. 1.

In these nine images the symbol on the left side of each box represents the marks being made with the stylus on the digital pad. The figure on the right is the image that is translated to the screen. In all nine images, the marks made with the stylus begin at the top left (for example, the production of the tree in panel 9 began with the leaves while producing the face in panel 8 began with the left eye).

Fig. 2.
13 symbols making up the training set of the first participant in each chain.

Fig. 3.

The results from Chain 1: Each row is made up of the symbols output from each generation of participants, with the first row being the experimenter selected input or “seed” lexicon.

Fig. 4.
The results from Chain 2: Each row is made up of the symbols output from each generation of participants, with the first row being the experimenter selected input or “seed” lexicon.

Fig. 5.
Intra-lexical distance for Chains 1 and 2 of Experiment 1 (Generation 0 = seed lexicon).

Fig. 6. Inter-lexical distance for Chains 1 and 2 of Experiment 1. The value at each generation represents the distance calculated between it and its predecessor. (Value at Generation 1 = shortest distance between Generation 1 symbols and seed lexicon symbols).

Fig. 7.
14th symbol, used in the seed lexicon in Experiment 2.

Fig. 8.
The “meaning space” arranged in order of its assigned symbol (as numbered in Figure 2).

Fig. 9.

The results from the first 10 symbols created in Chain 1: Each row is made up of the symbols output from each generation of participants, with the first row being the experimenter selected input or “seed” lexicon.

Fig. 10.

The results from the second 10 symbols created in Chain 1: Each row is made up of the symbols output from each generation of participants, with the first row being the experimenter selected input or “seed” lexicon. Darkened lines between generations represent “breaks in the transmission”. The symbols above these lines were not transmitted to the next generation. The symbols below this line were created anew by that generation.

Fig. 11.

The results from the first 10 symbols created in Chain 2: Each row is made up of the symbols output from each generation of participants, with the first row being the experimenter selected input or “seed” lexicon.

Fig. 12.

The results from the second 10 symbols created in Chain 2: Each row is made up of the symbols output from each generation of participants, with the first row being the experimenter selected input or “seed” lexicon. Darkened lines between generations represent “breaks in the transmission”. The symbols above these lines were not transmitted to the next generation. The symbols below this line were created anew by that generation.

Fig. 13. Mean recall performance during recall of meaning given symbol across 4 learning trials. The maximum score for a single trial is 14.

Fig. 14.
Intra-lexical distance for Chains 1 and 2 of Experiment 2. (Generation 0 = seed lexicon).

Fig. 15.

Inter-lexical distance for Chains 1 and 2 of Experiment 2. The value at each generation represents the distance calculated between it and its predecessor. (For example, the value at Generation 1 = shortest distance between Generation 1 symbols and seed lexicon symbols).

Fig. 16.

Example of a newly created symbol. Symbol 3 from the seed lexicon appears to be forgotten by Generation 1, which created a novel symbol resembling symbol 7.

Fig. 17.
The emergence of a recurring sub-unit.

Fig. 18.

Symbols 2 and 18 of Generations 1–5 from Chain 2-Experiment 2. The symbols for meanings 'man' and 'woman' juxtaposed. The darker lines between two generations represent breaks in the transmission.

\$watermark-text

\$watermark-text

\$watermark-text

Table 1

An example of how a simple Euclidian distance would be calculated between two derivative vectors.

Distance calculation example	Pixel 1	Pixel 2	Pixel 3	Pixel 4	Pixel 5	...	Total Distance
Derivative vector A	10	9	8	13	13	x	
Derivative vector B	11	11	9	9	13	y	
Distance at each pixel	1	2	1	4	0	$ x - y $	$1 + 2 + 1 + \dots = \text{distance}$

Table 2

An example of how a distance between two vectors is computed using “time warping” .

Time warped distance calc. example	Pixel 1	Pixel 2	Pixel 3	Pixel 4	Pixel 5	...	Total Distance
Derivative vector A	10	9	8	13	13	x	
Derivative vector B	11	11	9	9	13	y	
Distance at each pixel	$10 - 11 = 1$	$10 - 11 = 1$	$9 - 9 = 0$	$8 - 9 = 1$	$13 - 13 = 0$	...	$1 + 1 + 0 + 1 + \dots = \text{distance}$

Segments with the same shading are those whose distances are measured by the algorithm.

Table 3

Items in bold represent symbol-meaning pairs that are learned from the previous generation, items not in bold represent pairs that are newly created during the final phase at that generation.

Seed	1	2	3	4	5	6	7	8	9	10	11	12	13	14				
Gen 1	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Gen 2	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Gen 3	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Gen 4	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18