Amount of Learning and Signal Stability Modulate Emergence of Combinatorial Structure
and Iconicity in Novel Signaling Systems

Vera Kempe ^{*1,}
Nicolas Gauvrit ^{2,}
Nikolay Panayotov ^{3,}
Sheila Cunningham ¹
Monica Tamariz ⁴

¹Division of Psychology, Abertay University, Dundee, Scotland, U.K.

² INSPE, Université de Lille, France

³Division of Games Technology and Mathematics, Abertay University, Dundee, Scotland,
U.K.

⁴Department of Psychology, Heriot Watt University, Edinburgh, Scotland, U.K.

*Corresponding Author: v.kempe@abertay.ac.uk

Keywords: iterated language learning, combinatoriality, compositionality, iconicity

Acknowledgement: We gratefully acknowledge funding from Research Incentive Grant
#70700 of the Carnegie Trust for the Universities of Scotland.

Abstract

Iterated language learning experiments that explore emergence of linguistic structure in the laboratory vary considerably in methodological implementation, limiting generalizability of findings. Most studies also restrict themselves to exploring the emergence of combinatorial and compositional structure in isolation. Here, we use a novel signal space comprising binary auditory and visual sequences and manipulate amount of learning and temporal stability of these signals. Participants had to learn signals for meanings differing in size, shape and brightness; their productions in the test phase were transmitted to the next participant. Across transmission chains of 10 generations each, Experiment 1 varied how much learning of auditory signals took place, and Experiment 2 varied temporal stability of visual signals. We found that combinatorial structure emerged most reliably with greater amount of learning and when signals were temporally stable. Iconicity emerged with reduced amount of learning, as opportunity for rote-memorization appeared to hamper exploration of the iconic affordances of the signal space. However, emergence of combinatoriality and iconicity in these entirely unfamiliar signaling systems was too fragile to allow for compositional signal-meaning mappings to emerge, so learnability did not improve over the course of transmission. These findings underscore the importance of systematically manipulating training conditions and signal characteristics in iterated learning language learning experiments and suggest that combinatorial structure and iconicity may be a prerequisite for emergence of compositional structure in novel signaling systems.

The last two decades have seen a vast expansion of experimental research into the cultural evolution of language. This research aims to understand under what conditions structured and conventionalized communication systems emerge, and how these systems change and diversify over time. At first glance, the incompatibility of language evolution and psychological testing time scales may seem to render any laboratory research that tries to address these questions futile. Yet building on seminal work by Bartlett (1932) and Esper (1966) recent methodological advances have utilized iterated learning in increasingly creative and sophisticated ways to gain insights into the fundamental principles that underpin the emergence of linguistic structure. In these iterated language learning experiments, participants are presented with novel signals for a set of meanings which they have to learn and – in some instances – use in referential communication with a partner. The results of each participant’s learning and use of these novel signals are then presented to the next generation of participants along a diffusion chain. The findings have demonstrated that passing languages through an inter-generational transmission bottleneck while also using them for intra-generational communication leads to the emergence of artificial mini-languages that are increasingly easier to learn and more systematically structured in ways that resemble natural languages (for a review see Tamariz, 2017). It should be noted that compositionality can also emerge without generational transmission, for instance during communication in large social networks (Raviv, Meyer & Lev-Ari, 2019) or when the predictability of referents from the context is low (Winters, Kirby & Smith; 2018).

Although iterated language learning is an increasingly prolific line of research, it so far has focused on a limited set of learning conditions and signal spaces from which conclusions about the emergence of linguistic structure are being drawn. The current study aims to scrutinize some of these conditions to gain a fuller picture of the generalizability of what this line of research can tell us about the cognitive mechanisms that may drive language evolution. Specifically, we ask to what extent the emergence of structural features resembling those in natural languages is affected by amount of learning when controlling for individual differences, and by modality and temporal characteristics of the signals.

Combinatorial and Compositional Structure

Linguistic structure emerges on two levels, which instantiate the duality of patterning considered a hallmark of human language (Hockett, 1960): On the one hand, signal spaces (also sometimes called form spaces [Cuskley, 2019]) are carved up into meaningless units, which can be recombined to produce potentially unlimited sets of discrete signals. In natural languages these units comprise the products of oral, facial or manual gestures (e.g., phonemes or pitch contours) which are recombined in systematic ways (e.g., using phonotactic constraints), to form larger units (e.g. morphemes or words). This level of structure is referred to as *combinatorial* structure and is assumed to arise from cognitive biases that favor reuse and modification of learned building blocks (Roberts & Galantucci, 2012, Verhoef et al., 2014). Iterated language learning experiments have explored a variety of signaling spaces such as sequences of letters (Kirby et al., 2008; 2015, Carr et al., 2017; Cornish et al., 2017), syllables (Raviv & Arnon, 2018), colors (Cornish et al., 2013), tones (Kempe et al., 2019) as well as drawn squiggles (Tamariz & Kirby, 2015), tracings on a continuously moving pad (Galantucci, Kroos & Rhodes; 2010), abstract graphical forms called Ferros (Cuskley, 2019), whistled pitch contours (Verhoef et al., 2014; 2015), and manual gestures (Motamedi et al., 2019). These experiments have shown that both inter-generational transmission (Verhoef et al. 2014) and referential communication (Little et al., 2017), can lead to the emergence of combinatorial structure. The combinatorial units that emerge in these experiments are shaped by the specific

affordances of each signal space. For example, color sequences are transformed into chunks of reusable color combinations (Cornish et al., 2013), unfamiliar squiggles are transformed into assemblies of familiar geometrical or graphical elements (Tamariz & Kirby, 2015), and complex melodies are transformed into concatenations of shorter pitch contours (Verhoef et al., 2014; see Cuskley, 2019, for a recent overview).

On the other hand, when signals refer to meanings the combinatorial units that make up individual signals are mapped onto dimensions of the meaning space in systematic ways. Recombination of such mappings allows for unlimited productivity in the expression of meanings. The mappings are arbitrary in the sense that the nature of a specific link between a set of combinatorial units and a meaning dimension is irrelevant as long as it is used consistently. This level of structure is referred to as *compositional* structure, in analogy to the morpho-syntax of natural languages where sub-lexical and lexical units such as phonemes, morphemes and words systematically map onto syntactic and semantic features. Compositionality can also be positional (Galantucci & Garrod, 2011) such that the ordering of signals conveys dimensions of meaning, often in a hierarchical manner (Saldana et al., 2019). For example, in iterated learning experiments where signals comprise strings of letters certain morpheme-like letter combinations can become associated with a referent's color, while others may map onto a referent's shape or direction of movement, and it does not matter which particular letter combinations map onto each of these meaning dimensions as long as these mappings are consistent (Kirby et al., 2008; 2015).

However, in natural languages, the specific links between features of the signals and dimensions of the meaning space are not entirely arbitrary but exhibit some degree of iconicity. In oral languages, iconicity frequently, but not exclusively, involves *conventional sound symbolism* (Cuskley & Kirby, 2013) where some features of the signal space, e.g. certain sound clusters, are systematically associated with certain meaning dimensions. More generally, such non-arbitrary relationships involve a degree of homeomorphism between the topologies of the signal and meaning spaces (Little et al., 2017b), a relationship that has also been termed *relative iconicity*. For example, both in natural languages and in artificial signaling systems longer signals tend to be associated with bigger referents or more complex meanings (Lewis & Frank, 2016, Little et al., 2017; Kempe et al., 2019). Furthermore, words with similar sounds tend to have similar meanings (Monaghan et al., 2014; Tamariz, 2008), and letters with similar shapes tend to map onto similar sounds (Jee, Tamariz & Shillcock, *under review*). In sign languages, *sensory sound symbolism* is more frequent as the relationship between signals and meanings tends to involve some degree of physical resemblance. For example, in many sign languages signals bear visual similarity with the objects or actions they denote (Perniss et al., 2010); and even in oral languages some sounds can have an onomatopoeic relationship with their referents such as the English sound cluster /gl/, which invokes the impression of shine and smoothness associated with words like 'glow', 'glisten', or 'glide', a phenomenon also termed *absolute iconicity* (Little et al., 2017).

Referential communication studies in which participants were asked to memorize or negotiate a set of novel signals showed that for some signaling systems, e.g. drawings or gestures, absolute iconicity can emerge right from the start (Fay et al., 2014a,b; Motamedi et al., 2019), and the iconic signals subsequently become simplified and conventionalized with repeated use (Lister & Fay, 2017). For other signaling systems like Leap Motion (Little et al., 2017) or binary tone sequences (Kempe et al., 2019), where physical resemblance between signal and meaning space is limited, emergence of relative iconicity requires transmission over several generations as it takes time for users to find suitable and well-motivated mappings between signal features and relevant meaning dimensions, e.g. using signal duration as an indicator of referent size.

Manipulating the transmission bottleneck:

The notion of a transmission bottleneck refers to limitations on the amount and kind of input that learners encounter and the constraints on processing the input. In natural language acquisition, children encounter a subset of their ambient language and generalize the structural regularities they are able to extract from it. In iterated language learning experiments, the poverty of the stimulus found in natural language acquisition is sometimes recreated by withholding a proportion of signal-meaning pairings from training, and then testing participants on the full set (e.g. Kirby et al., 2008; Motamedi et al., 2019). In other studies, where learners encounter the entire signaling system designed for a specific experiment (e.g. Kirby et al., 2015; Kempe et al., 2019), memory and processing limitations are assumed to drive the addition of structure to aid learning (Cornish et al., 2017).

Such different operationalizations of the transmission bottleneck in iterated language learning experiments make it difficult to pinpoint specific constraints that drive emergence of structure. First, it is unclear which specific input manipulations best represent input constraints that operate in natural language acquisition. Some studies only present about 50% of signals (e.g. Kirby et al., 2008) to encourage generalization during testing while other studies present the entire signaling system (e.g. Kirby et al., 2015; Carr et al., 2017). This procedural difference may matter because randomly selected sub-sets of signals may or may not contain whatever structural regularities are starting to emerge over the course of transmission. Natural language acquisition often involves exposure to a special register – child-directed speech – that has been shown to emphasize relevant structural features of a language (Kempe & Brooks, 2005; Savickiene et al., 2009) and is therefore likely to facilitate the discovery of structural regularities. Thus, withholding parts of the input in iterated language learning experiments may inadvertently constrain or bias structure discovery in ways that are not well controlled and are different from the way natural input aids language acquisition.

Secondly, iterated language learning experiments show considerable variation with respect to how often individual signals are presented during training. For example, in Kirby et al. (2008; 2015) a relatively small set of signals was presented six times each while in Carr et al. (2017) or Saldana et al. (2019), much larger sets of signals were presented three times each. Again, methodological feasibility seems to determine the overall amount of training but these different training regimens make it difficult to disentangle effects of presentation frequency from effects of structural complexity. Even for identical numbers of categories and meaning dimensions, number of repetitions is likely to affect encoding. For example, frequent exposure to individual signals might encourage rote-memorization thereby preventing learners from extracting structure. If learners subsequently notice discrepancies between training and novel testing items they may ‘invent’ new signals based on pre-existing linguistic knowledge (but see Fedzechkina et al. [2020] for recent evidence that the native language not always exerts an influence on structure addition). Less frequent exposure to individual signals, on the other hand, might encourage structure discovery and generalization.

Finally, structure addition is also affected by individual differences between learners, the effects of which on iterated learning are difficult to disentangle from effects of input selection and training regimen. While some accounts have suggested that poor learners, e.g. children, are more likely to add structure to aid learning (Hudson Kam & Newport, 2005), recent evidence shows that structure addition is driven by strong learners who are better able to extract regularities from the input (Johnson, Siegelman & Arnon, 2020). However, it is not clear whether more structure addition by strong learners results from advantages in encoding

of information or in the representation of structural regularities, and how either of these advantages may impact inter-generational transmission.

The first aim of the present study was therefore to manipulate amount of learning while controlling for individual differences by introducing fixed learning criteria. Participants had to reach either a more stringent or a more lenient accuracy threshold in comprehension and production of the training signals before being allowed to proceed to the testing phase. Rather than simply varying the amount of exposure to individual signals, this manipulation controlled the amount of learning regardless of individual differences. We predicted that if amount of learning facilitates structure extraction then combinatorial and compositional structure should emerge more readily when learners had to reach a more stringent criterion. If, on the other hand, imposing structure serves primarily as an aid to learning, combinatorial and compositional structure should be more likely to emerge in the lenient training condition.

Manipulating modality:

Emergence of structure may be affected not just by constraints that arise from the amount of learning but also by the specific characteristics of signals and meanings. Typically, participants in iterated language learning experiments are presented with printed labels for a set of meanings (but see Carr et al. [2017] for experiments that provided additional synthesized auditory input and Smith & Dediu [2020] for experiments transmitting spoken pseudo-words) and are asked to produce signals by typing letters on a keyboard. The letter strings presented at the outset of learning largely obey phonotactic and graphotactic constraints that are familiar to literate participants. Such a signal space does not permit to study the emergence of duality of patterning because the basic combinatorial units – graphemes – and the rules for how to combine them are already provided. Moreover, a comparison of learners that were either familiar or unfamiliar with musical notation has shown that musical literacy *per se* greatly facilitated the ability to reproduce compositional structure (Tamariz et al., 2010) suggesting that, by analogy, structure emergence in the domain of language may have been over-estimated in studies that use printed word-like letter strings as signals for literate participants.

A number of studies have explored emergence of combinatorial structure using unfamiliar signal spaces such as whistled pitch contours (Verhoef et al., 2014), color sequences (Cornish et al., 2013), or Ferros (Cuskley, 2019). However, without linking such signals to meanings these studies were restricted to exploring combinatorial structure in isolation from compositional structure which is bound to paint an incomplete picture of how linguistic structure can emerge on multiple levels. To study the emergence of duality of patterning novel signaling systems which do not contain pre-specified familiar combinatorial units need to be combined with meanings. Studies that have used unfamiliar signals like Leap Motion generated sounds (Eryilmaz & Little, 2017; Little et al., 2017a,b), whistled pitch contours (Verhof et al., 2015) or tracings on a moving pad (Del Guidice, 2012; Galantucci, Kroos & Rhodes, 2010; Roberts & Galantucci, 2012; Roberts, Lewandowski & Galantucci, 2015) to represent a range of meaning spaces have shown that emergence of structure interacts with the iconic affordances of each signal space: The greater the potential for iconicity of the signal space the longer it takes for combinatorial and compositional structure to emerge. A frequent observation is that visual signals such as manual gestures appear to have greater iconic potential (Little et al., 2017b; Perniss et al., 2010), presumably because, in contrast to vocal gestures, manual gestures possess more degrees of freedom thereby providing greater capacity for homeomorphic mappings onto a wide range of meaning dimensions. Yet when the number of meaning dimensions exceeds the number of signal features – a constellation that is more typical for vocal gestures – combinatorial and compositional structure emerge more readily (Little et al.,

2017b). Thus, when signals are unfamiliar iconic mappings are encouraged and the differences in the gestural repertoire that is used to produce visual or auditory signals can determine modality differences in affordances for iconicity.

In addition, signal modality is often confounded with temporal stability. Many visual signals, such as letter strings or drawings, remain accessible over much longer time periods than auditory signals. Even manual gestures used in natural sign languages comprise temporal changes on a slower time scale than vocal gestures where relevant contrasts rely on changes in the tens of milliseconds. Differential sensitivity to temporal characteristics may be responsible for the attested modality-specificity in statistical learning (Frost, Armstrong, Siegelman & Christiansen, 2014): Regularities extracted in one modality do not reliably transfer to another (Redington & Chater, 1996; Tunney & Altman, 1999). Moreover, within individuals, there is no correlation across modalities in the capacity to extract regularities (Siegelman & Frost, 2015). Finally, developmental trajectories of statistical learning abilities also seem to be modality-specific with continuous improvement with age found in the visual but not the auditory modality (Raviv & Arnon, 2018). Such modality-specific differences in statistical learning arise because similar computational principles operate over different neural networks which are characterized by somewhat differing processing and maturation constraints: the auditory cortex displays greater sensitivity to temporal information while the visual cortex is more sensitive to spatial information and less capable of temporal accumulation of information over extended periods of time (Chen & Vroomen, 2013). Accordingly, different temporal presentation rates have opposite effects on statistical learning in the auditory and the visual domain (Emberson et al., 2011).

It is thus unclear to what extent findings from iterated language learning experiments that use temporally stable visual signals generalize to oral languages which comprise rapidly fading auditory signals. On the one hand, the ‘Now-or-Never’-bottleneck created by the temporal transience of auditory input imposes the need for rapid chunking to pass information on for further processing at various levels of abstraction (Chater & Christiansen, 2016), a process that underpins the need for structure creation. Indeed, when rapid fading was imposed on visual signals in an iterated language learning experiment, combinatorial structure emerged more readily than when the same visual signals were temporally stable (Galantucci et al., 2010). On the other hand, rapidly fading auditory signals may be more taxing on working memory capacity than temporally stable visual signals and may thus restrict capacity for structure discovery. In line with this idea iterated language learning experiments that use spoken pseudo-words show greater variability in structure emergence than what has been observed for written pseudo-words (Smith & Dediu, 2020). Moreover, for children, who have limited working memory capacity, iterated learning experiments using transient auditory signals have so far been unable to find evidence for the emergence of structure (Raviv & Arnon, 2018; Kempe et al., 2019) while use of temporally more stable gestural signs did provide some evidence for emergence of structure (Bohn, Kachel & Tomasello, 2019).

Thus, in addition to examining the role of amount of learning in iterated language learning experiments in Experiment 1, Experiment 2 aimed to disentangle effects of signal modality from effects of temporal stability on structure emergence by presenting learners with stable vs. fading visual signals. Moreover, because the visual signals in Experiment 2 had the same information-carrying capacity as the auditory signals in Experiment 1, we were able to perform a direct modality comparison of fading signals from both experiments to see whether structure emerges more readily in the auditory domain when controlling for temporal characteristics of the signals.

The present study:

Below we report on two iterated language learning experiments in which participants were arranged into 12 chains of 10 generations each, and had to learn to associate 4-6-bit binary sequences with eight meanings that differed in shape (spiky or rotund), size (large vs. small) and brightness (dark vs. light). The signal space was designed to be unfamiliar to participants to control for effects of prior knowledge and pre-existing biases. Training comprised both production and comprehension tasks. In the comprehension task, participants were presented with one signal at a time and had to select the associated meaning from the set of eight referents. In the production task, participants were presented with one referent at a time and had to produce the associated signal. In both experiments, training lasted as long as it took participants to either comprehend or produce a pre-specified number of signals accurately. During testing, participants had to produce all signals again, but without feedback. Finally, we also gave participants a self-comprehension task where half of the signals they themselves had produced during testing were presented for comprehension.

In Experiment 1 we manipulated the amount of learning. Participants learned binary tone sequences comprising low and high tones and had to complete as many rounds of alternating comprehension and production training as it took them to use either three (*Short Training condition*) or five (*Long Training condition*) signal-meaning mappings correctly over the course of two comprehension-production blocks. In Experiment 2 we manipulated the temporal dynamics by compared iterated learning of binary color sequences in a *Fading condition* where individual blue or orange circles either appeared and faded one by one at the same rate as the tones in Experiment 1 with a *Stable condition* where all circles were presented simultaneously and remained visible for the same duration it had taken to present a fading sequence of corresponding length. To ensure comparability with the auditory signals, Experiment 2 employed the same training regimen as the Short Training condition of Experiment 1, i.e. participants had to use three signal-meaning mappings correctly in two comprehension-production blocks before being allowed to proceed to testing. Finally, we provide a direct comparison between modalities by comparing the *Short training condition* using binary tone sequences from Experiment 1 with the *Fading condition* using binary visual sequences from Experiment 2. This comparison allows us to disentangle effects of modality from effects of temporal stability on emergence of structure.

Both experiments received ethical clearance from Abertay University and were conducted online. We performed a manipulation check for the learning criterion of Experiment 1 to see whether the more stringent criterion indeed resulted in longer learning by comparing the number of trials participants completed until they reached the criterion. In both Experiments, we measured learnability of signals, combinatorial structure, compositional structure and iconicity as the link between sequence length and referent size as previous research with binary signals had shown that length-size association is the preferred form of iconicity for this signal space (Kempe et al., 2019).

Experiment 1: Binary Tone Sequences

Method

Participants:

One hundred and twenty student participants above the age of 18 years were recruited via social media and through the Abertay SONA system. We had hoped for the task to be sufficiently engaging to spread on social media but were able to recruit only about 10 participants using this method. As a result, most participants received course credit for participation.

Materials:

The meaning space comprised eight visual images of abstract shapes differing in shape (spiky *kiki*-type vs. rotund *bouba*-type), size (small shapes occupied half the size of large shapes) and brightness (light grey vs. dark grey). All referents also had unique properties based on their particular shapes (see Fig. 1). For the signal space, two 500 ms sine-wave tones (high: 440 Hz = musical note *a*; low: 293.7 Hz = musical note *d*) were synthesised and concatenated with an interval of 500 ms. A set of eight sequences was created by randomly combining the two tones into sequences half of which were four, and the other half six tones long to introduce variability in length. Variability in signal duration was introduced to encourage variability in the duration of participants' reproductions and to maintain compatibility with previous studies.

Figure 1: Meanings differing in size (large vs. small), shape (spiky vs. rotund) and brightness (light vs. dark gray)

Procedure:

The experiment was implemented for web browsers using HTML5 technologies to be compatible with a variety of devices, including desktops, laptops, smartphones and tablets. After accessing the weblink participants viewed an instruction screen explaining that they would learn a two-tone language to label a number of unfamiliar shapes, that correct comprehension and production of the two-tone 'words' would earn them points, that they had to collect a certain number of points (three in the short learning condition; five in the long learning condition) per training block to obtain a star and that they needed two stars to proceed to the final round. Next, participants saw a screen with two buttons marked with an 'up' and a

‘down’ arrow arranged next to each other and were invited to press them to familiarise themselves with the high and low tones. Position (left vs. right) of the buttons was counterbalanced across participants.

Training always started with a Comprehension Training block. On each trial, participant first saw a display of all eight meanings dispersed on random locations across the screen. Participants were prompted to press a button labelled ‘Listen’ to hear a signal and to select the object they thought it referred to. If they chose the correct object that object was then surrounded by a green circle and participants earned a point, which was displayed on their scoreboard on the screen. If they chose a wrong object it was surrounded by a red circle while the correct object was surrounded by a green circle, and their score remained unchanged. Immediately after that all objects except the correct one encircled in green faded away. The correct object remained visible on the screen and the associated sound sequence was played to provide feedback. For the next trial, all objects reappeared on the screen in a different spatial configuration and participants listened to the next signal trying to identify its meaning. After completing eight trials presenting all eight signals, participants moved on to the Production block. If participants had correctly identified the pre-specified number of meanings they received a star.

The Production Training block always followed the Comprehension Training block. A trial started with one of the eight objects being presented in the middle of the screen. Participants were asked to use the on-screen arrow buttons located underneath to produce a sound sequence to label the object. Each button press elicited a high or low 500 ms buzzing tone depending on the arrow, and participants could string as many tones together as they wanted. Once participants had decided that they had finished creating a tone sequence, they clicked on the ‘Next’ button which triggered a ‘Correct’ or ‘Wrong’ prompt to indicate whether their input had matched the learned signal or not. They then heard the correct sound sequence before proceeding to the next trial. Participants completed all eight trials in random order, gaining a point every time they correctly produced the trained sequence. If they had correctly reproduced the pre-specified number of signals a star appeared on the scoreboard.

If during the Comprehension and Production Training blocks participants had not gained two stars, they underwent further Comprehension and Production training. This set-up made it possible to reach the learning criterion exclusively for correct comprehension, exclusively for correct production or for correct trials in both tasks combined. Successfully collecting two stars completed the training phase.

The Production Testing block started after a short chime to refocus the participants’ attention. It followed the same procedure as the Production Training block, except no feedback was given as the ‘correct’ sequences were not played back to the participants. Unbeknownst to the participants, their input in this block was used as the training input for the next participant in the chain. To prevent degeneration of the language into a set of ambiguous signals participants were alerted when they had entered a previously produced sound sequence and were told to retry until they had created a unique signal for each meaning. The Self-Comprehension Testing block followed immediately after the Production testing block. It was identical to the Comprehension Training block except that participants were told that they would hear the sound sequences that they themselves had created in the previous block to test how well they were able to ‘understand themselves’. No item-level feedback was given about the accuracy of their individual choices and only four randomly selected objects were presented to prevent a strategy of selecting objects by elimination of previously chosen objects. Upon completion participants were provided with an overall accuracy score in the Self-Comprehension task.

As we had initially hoped for the experiment to spread on social media we ended it with a brief questionnaire asking how participants had found out about it and asking them to rate their enjoyment of participation on a 5-point Likert scale. We also provided a comment box to allow participants to include any feedback they would like to give before being debriefed. The subsequent debriefing screen explained the rationale behind the experiment and showed participants the percentage of correct trials they had achieved in the Self-Comprehension Testing block. Average duration of the entire experiment was 27.5 min, ranging between 8 and 128 minutes depending on how long it took participants to reach the learning criterion. A detailed description of the web implementation is provided in the Supplementary Materials at <https://osf.io/2uqma/>.

Results

All data were subjected to Growth Curve analysis with R (Version 3.6.3; R Core Team 2020) using the *lme4* R-package (Bates et al., 2015). Statistical significance of each term was evaluated using the Satterthwaite approximation implemented in the *lmerTest* R-package (Kuznetsova et al., 2017). We fitted linear mixed-effects models with centered fixed effects of Generation and Condition as well as their interaction. For system-level dependent variables (Number of Trials to Criterion, Combinatorial Structure and Compositional Structure), we included random intercepts and slopes of Generation by Chain. For the item-level dependent variables (Edit Distance, Hamming Distance and Sequence Length), we additionally included random effects of individual meanings and implemented the maximal random effect structure that would allow for model convergence (Barr et al., 2013). All data and analyses are available in the Supplementary Materials at <https://osf.io/2uqma/>. Because of the similarity between the experiments in terms of design and methodology we present the output of the statistical analyses in the same table (Table 1) and the graphs in the same figure (Fig. 2) to facilitate direct comparison.

Manipulation Check: To see whether a more stringent criterion resulted in an overall longer learning phase we the two conditions respect to how many trials it took participants to reach criterion. The Growth Curve Analysis confirmed a main effect of Condition (see Table 1) showing that, as expected, participants completed fewer trials to reach the more lenient learning criterion. Number of trials is an indicator for learnability and should decrease if the signaling systems become easier to learn yet the results showed no evidence for this (see Fig. 2, Panel A).

Combinatorial Structure: Determining combinatorial structure for a set of binary sequences that comprise a signaling system is challenging because the specific sub-components that participants might be recombining when producing different sequences are not known. For example, two sequences like 0010 and 0100 can be comprised of the bigrams 01, 10 and 00 or of the single 0 and the trigram 010, and it is not clear whether the bigram 00 has been reused across sequences or the trigram 010. We measured combinatorial structure using a combinatorial coefficient (CC) inspired by the *form recombination index* (Galantucci et al., 2010) and the minimum description length (MDL) principle (Barron et al., 1998): For a given language (i.e., a set of signals) such as $L = \{0010, 0100, 0011, 0101\}$, we determine all possible n-grams and check for the shortest one that would give the shortest overall description. For instance, the bigram $A = 01$ leads to the recoding of L as $\{A:01, 0A0, A00, 0A1, AA\}$. While the original language L was 16 elements long, the recoded language is reduced to a length of

15 elements, counting only alphanumeric symbols. Thus, compared with all other possible n-grams, 01 is the shortest one that leads to the best reduction. The MDL principle therefore suggest that 01 is an elementary unit. Scanning all possible n-grams again, we find that no other n-gram would lead to a shorter recoding of L, so that we keep {A:01, 0A0, A00, 0A1, AA} as the best compressed version of L. Depending on the language, the elementary units determined in this way may be used in several signals. To quantify how much the elementary units tend to be reused, we use the *form recombination index* on the best compressed version of L, returning CC. Hence, CC(L) as our measure of compositionality evaluates the propensity of elementary units to be used across different signals in L. The R-function used to compute CC is provided in the Supplementary Materials available at <https://osf.io/2uqma/>.

The Growth Curve analysis showed a main effect of Generation providing evidence for the emergence of combinatoriality over the course of transmission (see Table 1). Neither Condition nor the interaction showed a significant effect which would suggest that combinatorial structure emerged regardless of the amount of learning (see Fig. 2, Panel B).

Compositional Structure: We determined compositional structure for the sets of signals produced at testing by correlating similarities between all pairs of signals with similarities between the corresponding pairs of meanings and determined the z-score of this correlation within a distribution of 1000 random permutations of signal-meaning pairings (Mantel, 1967). The Growth Curve Analysis provided no evidence for an increase in compositional structure nor for differences between conditions (see Table 1 and Fig. 2, Panel C).

Learnability: We measured learnability of signals by computing length-normalized Levenshtein edit distances between signals produced for the same meaning during testing in consecutive generations. This measure takes account of the number of insertions, deletions and substitutions required to transform one sequence into another, divided by the length of the longer sequence. The higher the values, which can range from 0 to 1, the less accurately a sequence is reproduced. If learnability increases this measure should decrease over the course of transmission. The Growth Curve Analysis showed a main effect of Condition, indicating that reproduction was overall more accurate in the Long Training Condition where participants had to achieve a more stringent learning criterion before testing (see Table 1). As for number of trials to criterion, there was no indication that learnability improved over the course of transmission (see Fig. 2, Panel D).

Self-comprehension: While edit distances indicate accuracy of signal reproduction we also checked whether there was improvement in self-comprehension, i.e. whether participants were able to correctly identify the meaning of their own signals produced during testing. To this end, we computed Hamming distances which indicate whether a selected referent differed from the target referent in one, two or all three dimensions. The Growth Curve analysis showed no indication for a difference between conditions nor for improvement in self-comprehension over the course of transmission (see Table 1 and Fig. 2, Panel E).

Dependent Variable	Intercept	Generation	Condition	Generation x Condition
--------------------	-----------	------------	-----------	---------------------------

Experiment 1: long vs. short training of auditory signals (tone sequences)

Trials to criterion ¹	66.568***	-3.604	-14.929**	-1.661
Combinatorial structure ¹	0.252***	0.014*	0.001	0.001
Compositional structure ¹	0.283	0.061	0.136	0.104
Edit distance ²	0.199***	0.024	0.048*	0.015
Hamming distance ²	0.625***	0.007	0.049	0.007

Experiment 2: stable vs fading visual signals (color circle sequences)

Trials to criterion ¹	58.510***	-0.858	-0.830	3.413
Combinatorial structure ¹	0.257***	0.012 [†]	0.008	0.006
Compositional structure ¹	0.265*	0.141	0.041	-0.002
Edit distance ^{2, a}	0.280***	-0.018	0.031*	0.015
Hamming distance ^{2, a}	0.857***	0.019	-0.024	-0.011

Modality Comparison: short training auditory signals vs. fading visual signals

Trials to criterion ¹	54.016***	-4.758	-5.341*	-0.502
Combinatorial structure ¹	0.251***	0.010	0.001	0.004
Compositional structure ¹	0.322	0.154	0.097	0.010
Edit distance ²	0.248***	0.003	-0.001	0.036 [†]
Hamming distance ²	0.778***	0.021	-0.104 [†]	-0.003

Table 1: Parameter estimates of Growth Curve Analyses fitted to the different dependent variables in Experiments 1 and 2 as well as for the direct comparison between short training auditory sequences from Experiment 1 and fading visual sequences from Experiment 2, ¹ model: Dependent Variable ~ Generation * Condition + (Generation | Chain), ² model: Dependent Variable ~ Generation * Condition + (Generation | Chain) + (Generation * Condition | Meaning), ^a Model did not converge so we removed the random slope of the interaction by Meaning, *** p < .001, * p < .05, [†] p < .1.

Figure 2: Generational change in number of trials (panels A, F, K), combinatorial structure (panels B, G, L), compositional structure (panels C, H, M), edit distance of produced signal to previous generation signal (panels D, I, N) and Hamming distance of selected meaning to target meaning in self-comprehension (panels E, J, O) for

auditory tone sequences in Experiment 1 (left panels), visual color sequences in Experiment 2 (middle panels) and modality comparison (right panels). Error bars indicate ± 1 S.E.M.

Iconicity: To determine whether participants formed an iconic association between meaning size and signal length as in previous studies with the same signal space (Kempe et al., 2019) we fitted Growth Curve Analyses using a full model of fixed effects of Generation, Condition and Size and random effects of chains and individual meanings using a full random effect structure. The parameter estimates and associated p-values are given in column 2 of Table 2. The results show that signals became shorter over the course of transmission and that signals associated with small meanings were consistently shorter than those associated with large meanings (see Fig. 3, Panel A).

Fixed effect	audio (tones) Experiment 1	visual (color circles) Experiment 2	audio vs. visual Exp. 1 vs. Exp. 2
Intercept	4.589***	4.613***	4.455***
Generation	-0.209*	-0.245**	-0.226*
Size	-0.165	-0.158	-0.269*
Condition	-0.197	0.096	-0.063
Generation x Size	-0.107	0.028	-0.107
Generation x Condition	-0.013	-0.014	0.004
Size x Condition	-0.296*	-0.081	-0.191 [†]
Generation x Size x Condition	-0.092	0.043	-0.092

Table 2: Parameter estimates of Growth Curve Analyses fitted to signal length in Experiments 1 and 2 as well as for the direct comparison between short training auditory sequences from Experiment 1 and fading visual sequences from Experiment 2 using the model: Signal Length $\sim$ Generation * Size * Condition + (Generation * Size | Chain) + (Generation * Condition | Meaning), *** $p < .001$, * $p < .05$, [†] $p < .1$.

Figure 3: Length of signals for large and small meanings as a function of training condition for auditory sequences in Experiment 1 (panel A) and of presentation condition for visual color sequences in Experiment 2 (panel B). Error bars indicate ± 1 S.E.M.

Discussion

We had hypothesized that amount of learning would be related to emergence of structure. First, we confirmed that the stringent criterion required an overall greater amount of learning as indicated by number of trials to criterion. Participants in that condition also reproduced the signals more accurately overall than participants who were given a more lenient criterion, as indicated by edit distances. However, no benefit of the long training condition was found in self-comprehension. An asymmetry between production and comprehension performance is in line with findings from Little et al. (2017b) where Leap Motion was used to signal geometrical shapes, and performance in signal recognition was lower than in signal reproduction. However, here the difference between conditions in accuracy of signal production might have in part been an artefact of the difference in signal length: while the overall length difference between conditions did not reach statistical significance, signals for large referents were longer in the Long Training condition resulting in lower length-normalized edit distances for the same number of deviating elements (i.e., length-normalized edit distances between sequences differing in the same number of elements are lower the longer the sequences). Given that dimensionality of the meaning space was identical between conditions similar artefacts in the computing of Hamming distances can be ruled out for comprehension suggesting that there is little evidence for learnability differences as a function of amount of learning.

Crucially, in contrast to many previous iterated language learning studies, learnability did not improve over the course of transmission: when presented with a specific meaning it did not become easier for participants to produce the corresponding signal nor did it become easier for them to identify the meaning when presented with a signal they had just produced. Even though signals became progressively shorter the increase in production efficiency was not associated with commensurate gains in learning signals for specific meanings. Without emerging compositional structure, participants were unable to create stable links between signal features and meaning dimensions and relied instead on rote-memorization to pass the learning criterion thereby creating unique holistic languages throughout the entire course of transmission.

Iconicity can provide an initial solution to the problem of linking unfamiliar signals to meanings by exploiting salient and transparent associations that are easier to remember, and – in the context of communicative situations – constitute shared knowledge (Sulik & Lupyan, 2018). While the present experiment did not include communication, the pressure to memorize unique signal-meaning links was sufficient to develop an association between length and size such that signals for small referents became reliably shorter over the course of transmission while signals for big referents did not. The interesting finding of this experiment is that this iconic association emerged only in the Short Training condition. It appears that when learners are given enough training to allow for rote-memorization as the dominant learning strategy the potential for establishing any kind of systematicity in signal-meaning links, be it iconic or compositional, remains unexplored. In other words, in the absence of communicative pressure constraints on exposure can drive the emergence of iconicity.

It is noteworthy that while iconicity only emerged when exposure and learning were limited combinatorial structure, i.e., the reuse of sub-sequences as combinatorial units or building blocks, emerged in both conditions. Thus, emergence of combinatoriality seems to be a robust phenomenon in the transmission of unfamiliar sequences regardless of meaning

(Cornish et al., 2013; 2017) and can occur even for holistic languages. In contrast to Verhoef et al. (2016) we did not find a delay of emergence of combinatoriality in the condition which favored iconicity. However, while Verhoef et al. (2016) varied iconic affordances, we varied amount of learning which only affected emergence of iconicity but not emergence of combinatoriality. The finding reported here suggests that rote-learning seems to restrict exploration of the signal space with respect to its affordances for potential cross-modal mappings. We therefore used the Short Training condition with its somewhat greater potential for the emergence of iconicity to explore the role of signal stability in transmission of a novel signaling system.

Experiment 2: Binary Color Sequences

Method

Participants:

One hundred and twenty participants above the age of 18 were recruited through the crowd-sourcing website Prolific Academic and paid £4.00 for participation. The change in recruitment strategy was necessary as Experiment 1 had exhausted the available naïve participant pool of this small university.

Materials:

The meanings were identical to Experiment 1. The visual signals comprised sequences of blue and orange dots. In the Fading condition, dots were revealed one by one from left to right with each dot fading away once the next dot to the right of it had appeared. The inter-dot onset interval was 1000 ms mimicking the interval between the onset of tones in Experiment 1. The fading animation lasted 400 ms and started simultaneously with the onset of the next dot in the sequence. In the Static condition, all the dots in a sequence were revealed simultaneously and remained on screen for a total time determined by multiplying the number of dots in the sequence by 500 ms.

Procedure:

The procedure was similar to Experiment 1 aside from the following adjustments to the visual modality of the signals: Participants were told they would learn a two-color language. Instead of buttons with ‘up’ and ‘down’ arrows participants pressed an orange or a blue button arranged in the same way as the arrow buttons in Experiment 1. In the Comprehension block, participants were prompted to press a button labelled ‘Show Word’ to reveal the sequence of color circles. In the Production block, participants pressed the colored buttons successively to create the color sequences. In the Fading condition, any previous color circle disappeared over the course of 400 ms as a new one was entered by button click. In the Stable condition, all circles remained visible until participants pressed the ‘Next’ button. As in Experiment 1, to receive a star, participants had to score three correct trials per Comprehension or Production block. The Production and Self-Comprehension Testing blocks were presented once participants had scored two stars. Average duration of the entire experiment was 22.3

min, ranging between 6 and 70 minutes depending on how long it took participants to reach the learning criterion.

Results

Data were analysed using the same model structure as described for Experiment 1.

Trials to criterion: As the learning criterion was identical in both conditions number of training trials would only differ between conditions if one presentation mode was inherently more difficult than the other. The analysis revealed that this was not the case (see Table 1 and Fig. 2, Panel F).

Combinatorial Structure: This was determined as described for Experiment 1. Growth Curve Analysis showed a trend towards increased combinatoriality which fell short of statistical significance (see Table 1 and Fig. 2, Panel G)

Compositional Structure: This was determined as described for Experiment 1. As in Experiment 1, Growth Curve Analysis provided no evidence for an increase in compositional structure nor for differences between conditions (see Table 1 and Fig. 2 Panel H).

Learnability: Edit distances were computed as in Experiment 1. Growth Curve Analysis showed a main effect of Condition, indicating that reproduction was overall more accurate in the Fading Condition (see Table 1). As in Experiment 1, the effect of Generation was not significant suggesting that learnability did not improve over the course of transmission (see Fig. 2, Panel I).

Self-comprehension: Hamming distances computed as in Experiment 1 showed no difference between conditions nor improvement in comprehension over the course of transmission (see Table 1 and Fig. 2, Panel J).

Iconicity: We fitted the same model as in Experiment 1 to signal length. The parameter estimates and associated p-values, given in column 3 of Table 2, show that signals became shorter over the course of transmission (see Fig. 3, Panel B).

Discussion

We had hypothesized that temporal features of signal presentation can affect emergence of structure. Controlling for amount of learning, we found that during testing fading signals were produced more accurately than stable signals. As in Experiment 1, learnability did not improve even though signals became shorter over the course of transmission confirming again that it is not signal complexity *per se* but ability to link signals to meanings that determines learning success. As in Experiment 1, the novelty of the signaling system precluded emergence of compositional structure which is necessary for the establishment of stable links between signals and meanings.

Unlike for the auditory signals in Experiment 1, iconicity did not emerge for visual signals although there was a trend towards shorter signals for smaller meanings in the stable signals which did not reach statistical significance (see Fig. 3, Panel B). Similarly, the trend towards emergence of combinatorial structure also fell short of significance. Taken together, the findings of Experiment 2 suggest that while fading color sequences showed some processing advantage over stable sequences, no further reliable gains in terms of learnability, structure or iconicity were made for transmission of these visual signals.

Modality Comparison

We also performed a direct comparison of the auditory signals in the Short Training condition and the visual signals in the Fading condition (depicted in the right column of Fig. 2) using the same model structure as described for Experiments 1 and 2 to disentangle temporal effects from modality effects.

Trials to criterion: The main effect of Condition indicated that participants required fewer trials overall with auditory compared to visual signals to reach the learning criterion of three correct responses in two training rounds (see Table 1 and Fig. 2, Panel K).

Combinatorial Structure: This analysis showed no effects providing no evidence for emergence of combinatorial structure (see Table 1 and Fig. 2, Panel L)

Compositional Structure: This analysis showed no effects providing no evidence for emergence of compositional structure (see Table 1 and Fig. 2, Panel M).

Learnability: The analysis of edit distances showed a trend towards an interaction between Condition and Generation which fell short of statistical significance (see Table 1) suggesting that compared to visual signals, learning auditory signals appeared to be somewhat easier in earlier generations and more difficult in later generations (see Fig. 2, Panel N).

Self-comprehension: The analysis of Hamming distances showed a trend towards lower distances, and, hence, better self-comprehension for auditory signals but this trend fell short of significance (see Table 1 and Fig. 2, Panel O).

Iconicity: The parameter estimates and associated p-values for the same model as used in Experiments 1 and 2 are provided in column 4 of Table 2. The main effect of generation indicates that signals became shorter over the course of transmission (compare the right graph in Panel A and the left graph in panel B of Fig. 3). The main effect of Size indicates that signals for small referents were overall shorter than signals for big referents. The interaction between Condition and Size fell short of significance: while there was a trend for the emergence of iconicity to be more pronounced in the auditory signals this modality difference was not statistically reliable.

Discussion

The comparison of fading auditory and visual binary sequences in the same training condition showed that fewer trials to criterion were necessary for the auditory signals. This suggests that, all things being equal, auditory sequences were easier to learn than visual signals. There was no evidence for emergence of combinatorial or compositional structure, despite the fact that emergence of combinatorial structure had been observed in Experiment 1 for auditory sequences when the Long and Short Training condition were analyzed jointly. This suggests that the reduced amount of training that allowed participants to reach the more lenient criterion was insufficient for combinatorial structure to emerge and that the emergence of combinatoriality observed for the auditory sequences in Experiment 1 must have been mainly carried by the Long Training condition.

No modality difference was found with respect to the emergence of iconicity: in both conditions shorter sequences were used to signal smaller referents, although the effect appeared to be more pronounced in the auditory signals.

General Discussion

In this study, we asked whether amount of learning, temporal characteristics and modality of unfamiliar signals affect emergence of combinatorial and compositional structure as well as of iconicity in iterated learning experiments. Experiment 1 manipulated amount of learning of binary auditory sequences using a stringent vs. lenient accuracy criterion resulting in longer vs. shorter training phases. Experiment 2 manipulated temporal decay by presenting binary visual sequences in either a stable or a fading condition, with the latter matching the temporal characteristics of the auditory sequences of Experiment 1. Finally, we performed a direct modality comparison controlling for amount of learning and temporal characteristics.

Our findings showed emergence of combinatorial structure for the auditory sequences in Experiment 1 and a similar trend for the visual sequences in Experiment 2. However, when fading auditory and visual signals were analyzed jointly the generational increase was no longer significant suggesting that emergence of combinatoriality was more prominent in the other two conditions: the long training condition for auditory signals in Experiment 1 and the stable condition for visual signals in Experiment 2. This suggests that for combinatorial structure to emerge in an unfamiliar signaling system individual signals need to be presented sufficiently frequently or be sufficiently stable in time to allow for extraction of chunks that can serve as reusable building blocks. If training is too brief or signals fade too quickly the Now-or-Never-bottleneck (Christiansen & Chater, 2016) may be too narrow for such chunking to occur.

Even when combinatorial structure emerged it may not have supplied enough morpheme-like chunks that could be linked systematically to the different meaning dimensions. Indeed, our results showed no evidence whatsoever for emergence of compositional structure. There is strong evidence that compositionality requires communicative pressure: signal-meaning links need to be systematic enough so that they can be recovered by an interlocutor who is trying to identify a meaning communicated to them (Kirby et al., 2015; for a review see Smith, 2018). It could be argued that without a communicative task there was insufficient pressure for compositionality in these experiments. However, some iterated language learning experiments have demonstrated that compositionality can emerge without communication when ambiguity filters are used, just as we had done in these experiments: Whenever participants attempted to produce an ambiguous signal, i.e. a signal they had already used, they were alerted to this and asked to produce another signal. Under these filtering conditions one would expect compositionality to emerge. The fact that it did not suggests that in addition to

expressivity pressures signaling systems need to have acquired sufficient combinatorial structure to enable compositionality. We would argue that combinatorial structure serves as a pre-requisite for compositional structure: Only when enough combinatorial units are available can these be linked to various dimensions of the meaning space in a consistent and systematic way. One could argue that in these experiments the obvious solution would have been to create a perfect compositional system that maps 3-bit-sequences onto a $2 \times 2 \times 2$ meaning space. Yet such a minimalist compositional system never emerged in any of the 240 signaling systems produced in both experiments. One possibility is that single tones and single color circles did not have enough perceptual salience to serve as morpheme-like compositional units. Another possibility is that such a minimalist system violates duality of patterning as the constituent components would cease to be meaningless building blocks. Future research should explore the criteria that guide the inter-relation between combinatorial and compositional units.

Many previous iterated language learning experiments were not designed to explore the relationship between emergence of combinatoriality and of compositionality because they either supplied combinatorial units from the outset, like graphemes and associated phonemes as the building blocks of pseudo-words (e.g. Kirby et al., 2008; 2015), or omitted meanings altogether when presenting color sequences (Cornish et al., 2013), consonant letter sequences (Cornish et al., 2017), slide whistles (Verhoef, 2012; Verhoef et al., 2014) or Ferros (Cuskley, 2019). As we argued above, such signal spaces are unlikely to capture the entirety of constraints that govern the emergence of duality of patterning. The relationship between emergence of combinatorial and compositional structure can only be explored when unfamiliar signals are paired with unfamiliar meanings. One problem with the unfamiliar signaling spaces that have been used so far to express meaning, e.g. physical or digital slide whistles (Verhoef et al., 2015, 2016), tracings on a moving digitizing pad (Del Giudice, 2012) or Leap Motion sound controllers (Eryilmaz, & Little, 2017; Little et al., 2017b), is that it is difficult to quantify combinatorial and compositional structure as the signals comprise continuous change over time on a range of parameters, e.g. pitch or spatial location, and it is unclear how to segment continuous signals into constituent building blocks. Although interesting inroads into computing combinatorial structure have been made using Hidden Markov Models to analyze either hand coordinates or amplitude and frequency values of auditory signals created with Leap Motion controllers (Little et al., 2017b), we suggest that the simple binary sequences used here constitute a promising signaling system that offers an opportunity to explore and quantify combinatorial and compositional structure in tandem.

While emergence of combinatorial structure required a sufficient amount of learning or temporally stable presentation of signals, emergence of relative iconicity – here instantiated through an association between sequence length and referent size – appeared to benefit from less learning, as evidenced by the interaction between training condition and referent size in the auditory signals where a robust length-size association emerged only with the lenient learning criterion. Although iconicity can serve as an initial solution to the problem of forging transparent and recoverable signal-meaning links in unfamiliar signaling systems (Lister & Fay, 2017), our findings suggest that the opportunity for rote-memorizing afforded by longer training precluded exploration of the iconic affordances of the auditory signaling space. In the visual modality, however, the trend towards emergence of iconicity did not reach statistical significance, despite the lenient learning criterion. We should note that the slightly different recruitment strategies in the experiments (predominantly students gaining course credit in Experiment 1 vs. online participants gaining payment in Experiment 2) may limit comparability in subtle ways. Still, taken together with the finding that a length-size association was evident in both modalities when signals were fading it would appear that emergence of iconicity benefitted from less learning and less temporal stability of the signals. While future

research should explore further why rote-learning limits exploration of iconic affordances our results demonstrate that subtle variations in learning conditions and signal presentation rate can affect the emergence of iconic mappings.

Although there was evidence for emergence of combinatoriality and iconicity in some conditions, our experiments found no increase in learnability over the course of transmission despite signals generally getting shorter. This shows that learnability is driven by the transparency and systematicity of signal-meaning links which in the absence of prior experience with a signaling system are difficult to forge and remain fragile. Even the emerging length-size association was insufficient to allow for reliable signal reproduction because it linked the signals only to one out of three meaning dimensions. Specifically, when attempting to label a small referent, participants may have produced shorter sequences but shortening alone did not sufficiently specify all features of the target sequence. Similarly, when being presented with their own shorter sequences in the self-comprehension task signal length *per se* would not have provided participants with sufficient information to be able to select the correct referent among the four small ones. Thus, our findings suggest that both combinatoriality as a way to induce internal structure of signals and iconicity as a way to establish initial signal-meaning links may constitute the pre-requisites for emergence of combinatorial structure.

Finally, it is noteworthy that when auditory and visual signals were equated in terms of training regimen and temporal characteristics participants needed fewer trials to reach the learning criterion with auditory compared to visual signals. Although this preliminary finding needs further corroboration through careful manipulation of temporal characteristics of the signals using a design that segregates participant variability from modality-induced variability it is in line with evidence for modality-specificity in statistical learning that reflects the predisposition of the auditory cortex for rapid temporal integration (Chen & Vroomen, 2013; Siegelman & Frost, 2015).

In summary, our results underscore the importance of examining emergence of combinatorial and compositional structure in tandem, and to systematically explore the parameter space that is defined by signal familiarity, signal modality, temporal characteristics of signals and the specifics of the training regimen, to better understand how the various processing constraints and their interaction can affect the evolution of language.

References:

- Barron, A., Rissanen, J., & Yu, B. (1998). The minimum description length principle in coding and modeling. *IEEE Transactions on Information Theory*, 44(6), 2743-2760.
- Bartlett, F. C. (1932). *Remembering*. Oxford: Macmillan.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48.
- Bohn, M., Kachel, G., & Tomasello, M. (2019). Young children spontaneously recreate core properties of language in a new modality. *Proceedings of the National Academy of Sciences*, 116(51), 26072-26077.
- Carr, J. W., Smith, K., Cornish, H., & Kirby, S. (2017). The cultural evolution of structured languages in an open-ended, continuous world. *Cognitive Science*, 41(4), 892-923.
- Chen, L., & Vroomen, J. (2013). Intersensory binding across space and time: a tutorial review. *Attention, Perception, & Psychophysics*, 75(5), 790-811.
- Christiansen, M. H., & Chater, N. (2016). The now-or-never bottleneck: A fundamental constraint on language. *Behavioral and Brain Sciences*, 39.
- Cornish, H., Dale, R., Kirby, S., & Christiansen, M. H. (2017). Sequence memory constraints give rise to language-like structure through iterated learning. *PloS one*, 12(1), e0168532.
- Cornish, H., Smith, K., & Kirby, S. (2013). Systems from sequences: An iterated learning account of the emergence of systematic structure in a non-linguistic task. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 35, No. 35).
- Cuskley, C. (2019). Alien forms for alien language: investigating novel form spaces in cultural evolution. *Palgrave Communications*, 5(1), 1-15.
- Del Giudice, A. (2012). The emergence of duality of patterning through iterated learning: Precursors to phonology in a visual lexicon. *Language and Cognition*, 4(4), 381-418.
- Emberson, L. L., Conway, C. M., & Christiansen, M. H. (2011). Timing is everything: Changes in presentation rate have opposite effects on auditory and visual implicit statistical learning. *Quarterly Journal of Experimental Psychology*, 64(5), 1021-1040.
- Eryilmaz, K., & Little, H. (2017). Using leap motion to investigate the emergence of structure in speech and language. *Behavior Research Methods*, 49(5), 1748-1768.
- Esper, E. (1966). Social transmission of an artificial language. *Language*, 42(3), 575-580.
- Fay, N., Ellison, M., & Garrod, S. (2014a). Iconicity: From sign to system in human communication and language. *Pragmatics & Cognition*, 22(2), 244-263.
- Fay, N., Lister, C. J., Ellison, T. M., & Goldin-Meadow, S. (2014b). Creating a communication system from scratch: gesture beats vocalization hands down. *Frontiers in Psychology*, 5, 354.
- Frost, R., Armstrong, B. C., Siegelman, N., & Christiansen, M. H. (2015). Domain generality versus modality specificity: the paradox of statistical learning. *Trends in Cognitive Sciences*, 19(3), 117-125.
- Galantucci, B., Kroos, C., & Rhodes, T. (2010). The effects of rapidity of fading on communication systems. *Interaction Studies*, 11(1), 100-111.

- Gauvrit, N., Soler Toscano, F., Zenil, H., & Delahaye, J.-P. (2014). Algorithmic complexity for short binary strings applied to psychology: A primer. *Behavior Research Methods*, 46(3), 732-744.
- Hartley, L. H., & Fedzechkina, M. (2020). Learners' bias to balance production effort against message uncertainty is independent of their native language. *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*.
- Hockett, C. F. (1960). The origin of speech. *Scientific American*, 203, 88-111.
- Hudson Kam, C. L., & Newport, E. L. (2005). Regularizing unpredictable variation: The roles of adult and child learners in language formation and change. *Language Learning and Development*, 1(2), 151-195.
- Jee, H., Tamariz, M., & Shillcock, R. How systematically can phonology be intentionally represented in orthography? Quantifying phonographic systematicity in Korean Hangeul, Pitman's shorthand, the Shavian alphabet, and English. Submitted to *Behavior Research Methods*.
- Johnson, T., Siegelman, N., & Arnon, I. (2020). Individual differences in learning abilities impact structure addition: Better learners create more structured languages. *Cognitive Science*, 44(8), e12877.
- Kempe, V. & Brooks, P. J. (2005). The role of diminutives in the acquisition of Russian gender: Can elements of child-directed speech aid in learning morphology? *Language Learning*, 55, Supplement: The Best of Language Learning, 139-176.
- Kempe, V., Gauvrit, N., Gibson, A., & Jamieson, M. (2019). Adults are more efficient in creating and transmitting novel signalling systems than children. *Journal of Language Evolution*, 4(1), 44-70.
- 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., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141, 87-102.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13), 1-26.
- Lewis, M. L., & Frank, M. C. (2016). The length of words reflects their conceptual complexity. *Cognition*, 153, 182-195.
- Lister, C. J., & Fay, N. (2017). How to create a human communication system: A theoretical model. *Interaction Studies*, 18(3), 314-329.
- Little, H., Eryilmaz, K., & De Boer, B. (2017a). Conventionalisation and discrimination as competing pressures on continuous speech-like signals. *Interaction studies*, 18(3), 352-375.
- Little, H., Eryilmaz, K., & de Boer, B. (2015). Linguistic modality affects the creation of structure and iconicity in signals. In *37th Annual Meeting of the Cognitive Science Society (CogSci 2015)* (pp. 1392-1398). Cognitive Science Society.
- Little, H., Eryilmaz, K., & De Boer, B. (2017b). Signal dimensionality and the emergence of combinatorial structure. *Cognition*, 168, 1-15.

- Monaghan, P., Lupyan, G., & Christiansen, M. (2014). The systematicity of the sign: Modeling activation of semantic attributes from nonwords. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 36, No. 36).
- Motamedi, Y., Schouwstra, M., Smith, K., Culbertson, J., & Kirby, S. (2019). Evolving artificial sign languages in the lab: From improvised gesture to systematic sign. *Cognition*, 192, 103964.
- Perniss, P., Thompson, R., & Vigliocco, G. (2010). Iconicity as a general property of language: evidence from spoken and signed languages. *Frontiers in Psychology*, 1, 227.
- Raviv, L., & Arnon, I. (2018). The developmental trajectory of children's auditory and visual statistical learning abilities: modality-based differences in the effect of age. *Developmental Science*, 21(4), e12593.
- Raviv, L., & Arnon, I. (2018). The developmental trajectory of children's auditory and visual statistical learning abilities: Modality-based differences in the effect of age. *Developmental Science*, 21(4), e12593.
- Raviv, L., Meyer, A., & Lev-Ari, S. (2019). Compositional structure can emerge without generational transmission. *Cognition*, 182, 151-164.
- Redington, M., & Chater, N. (1996). Transfer in artificial grammar learning: A reevaluation. *Journal of Experimental Psychology: General*, 125(2), 123.
- Roberts, G., & Galantucci, B. (2012). The emergence of duality of patterning: Insights from the laboratory. *Language and Cognition*, 4(4), 297-318.
- Savickiene, I., Kempe, V. & Brooks, P. J. (2009). Acquisition of gender agreement in Lithuanian: Exploring the effect of diminutive usage in an elicited production task. *Journal of Child Language*, 36, 477-494.
- Siegelman, N., & Frost, R. (2015). Statistical learning as an individual ability: Theoretical perspectives and empirical evidence. *Journal of Memory and Language*, 81, 105-120.
- Smith, K. (2018). The cognitive prerequisites for language: insights from iterated learning. *Current Opinion in Behavioral Sciences*, 21, 154-160.
- Smith, K., & Dediu, D. (2020). Did partial hearing loss shape Australian phonology? An empirical test using spoken iterated learning. In Ravignani et al. (Eds.) *The evolution of language: Proceedings of the 13th International Conference*. Nijmegen, The Netherlands.
- Tamariz, M. (2008). Exploring systematicity between phonological and context-cooccurrence representations of the mental lexicon. *The Mental Lexicon*, 3(2), 259-278.
- Tamariz, M. (2017). Experimental studies on the cultural evolution of language. *Annual Review of Linguistics*, 3, 389-407.
- Tamariz, M., & Kirby, S. (2015). Culture: copying, compression, and conventionality. *Cognitive Science*, 39(1), 171-183.
- Tamariz, M., Brown, J. E., & Murray, K. M. (2010). The role of practice and literacy in the evolution of linguistic structure. In A. D. M. Smith, M. Schouwstra, B. de Boer, & K. Smith (Eds.), *The evolution of language: Proceedings of the 8th International Conference*. Singapore: World Scientific.
- Tunney, R. J., & Altmann, G. (1999). The transfer effect in artificial grammar learning: Reappraising the evidence on the transfer of sequential dependencies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(5), 1322.

- Verhoef, T. (2012). The origins of duality of patterning in artificial whistled languages. *Language and Cognition*, 4(4), 357-380.
- Verhoef, T., Kirby, S., & De Boer, B. (2014). Emergence of combinatorial structure and economy through iterated learning with continuous acoustic signals. *Journal of Phonetics*, 43, 57-68.
- Verhoef, T., Kirby, S., & de Boer, B. (2016). Iconicity and the emergence of combinatorial structure in language. *Cognitive Science*, 40(8), 1969-1994.
- Verhoef, T., Roberts, S. G., & Dingemanse, M. (2015). Emergence of systematic iconicity: transmission, interaction and analogy. In *37th Annual Meeting of the Cognitive Science Society (CogSci 2015)* (pp. 2481-2486). Cognitive Science Society.
- Winters, J., Kirby, S., & Smith, K. (2018). Contextual predictability shapes signal autonomy. *Cognition*, 176, 15-30.
- Zenil, H., Soler-Toscano, F., Delahaye, J.-P., & Gauvrit, N. (2015). Two-dimensional Kolmogorov complexity and an empirical validation of the coding theorem method by compressibility. *PeerJ Computer Science*, 1:e23