Language Evolution in the Lab: The Case of Child Learners

Limor Raviv[1] (limor.raviv@mail.huji.ac.il) & Inbal Arnon[2] (inbal.arnon@gmail.com)

[1]Department of Cognitive Science, [2]Department of Psychology, Hebrew University, Jerusalem

Abstract

Recent work suggests that cultural transmission can lead to the emergence of linguistic structure as speakers' weak individual biases become amplified through iterated learning. However, to date, no published study has demonstrated a similar emergence of linguistic structure in children. This gap is problematic given that languages are mainly learned by children and that adults may bring existing linguistic biases to the task. Here, we conduct a large-scale study of iterated language learning in both children and adults, using a novel, child-friendly paradigm. The results show that while children make more mistakes overall, their languages become more learnable and show learnability biases similar to those of adults. Child languages did not show a significant increase in linguistic structure over time, but consistent mappings between meanings and signals did emerge on many occasions, as found with adults. This provides the first demonstration that cultural transmission affects the languages children and adults produce similarly.

Keywords: language evolution; cultural transmission; iterated learning; developmental differences

Introduction

Usage-based theories suggest that the kinds of structures we observe in natural languages arise from general biases and constraints on human capacities, such as learning, memory and processing (Tomasello, 2009). Importantly, these weak individual tendencies can become amplified and fixated over time through the process of *cultural transmission*: the transmission of language over generations through a repeated cycle of use, imitation, observation and induction (Kirby, Griffiths & Smith, 2014). Iterated learning models (ILMs), which simulate the process of cultural transmission of a given behavior over multiple generations, support this theory by showing how the iterative nature of cultural transmission can lead to the creation of strong linguistic universals without the need to assume strong innate biases.

Mathematical and computational simulations of iterated learning show that cultural transmission amplifies weak cognitive biases over time by shaping structural properties to fit agents' existing tendencies and predispositions (Reali & Griffiths, 2009). Moreover, findings from non-linguistic ILM studies with adult participants that examine the transmission of various behaviors (e.g., drawings, whistles, gestures, visual patterns) show that the learned behaviors become significantly more predictable and more structured over generations (Cornish, Smith & Kirby 2013; Verhoef, Kirby & de Boer 2014). In particular, systems become easier to learn, with participants in later generations making considerably fewer mistakes. Notably, systems also become less random and more structured over time, often displaying compositional structure and consistent reuse of smaller building blocks.

Linguistic ILM studies, which are the most relevant to the question of language evolution, show that language learnability increases thanks to an increase in linguistic structure (Kirby, Cornish & Smith, 2008). The languages produced in these studies develop consistent mappings between meanings and signals over time, with similar meanings expressed using similar strings. This is achieved either by the creation of homonyms that mark a shared dimension (e.g., color) but lead to under-specification, or by developing morphological structure, in which different affixes are used to encode different semantic dimensions.

While these findings support the role of cultural transmission in language evolution, they are limited to adult learners. Only one study has used ILM to compare children to adults on a non-linguistic task (Kempe, Gauvrit & Forsyth, 2015), and no published study has looked at the emergence of linguistic structure over time in children. Yet such findings are crucial for making inferences on how learning biases may affect language structure. In particular, adult participants may rely on their extensive and explicit linguistic knowledge when learning an unfamiliar artificial language (Cornish, Tamariz & Kirby, 2009). Consequently, they may have a prior bias in favor of linguistic structure, which (consciously or not) influences their performance, causing structure to emerge. In other words, the structure observed in adult studies may not reflect a cognitive bias *responsible* for the evolution of language over time, but rather a bias that is the *result* of it. If this is true, we cannot draw strong conclusions from these findings on the processes underlying the emergence of linguistic structure in the first place. This criticism can be avoided by looking at children, who have less extensive experience with language (Ramscar & Gitcho, 2007). The lack of evidence from child learners is also problematic because children are the most frequent learners in the actual process of linguistic cultural transmission. Their performance is a test case for the verification of the hypothesis that structure can emerge over time through cultural transmission.

Interestingly, several different predictions can be made regarding children's possible performance in this paradigm: On the one hand, children may be expected to perform like adults, or even better, given their superior language learning abilities in real-world settings (Birdsong, 1999). This prediction is consistent with the claim that children have a special role in the formation of linguistic structure (Bickerton, 1984). Studies with deaf children who are born to hearing parents (and were not exposed to a formal sign language) suggest that children have unique abilities in imposing structure and introducing regularities like word

1

order that are not found in the gestures of their mothers (Goldin-Meadow & Mylander, 1998). Moreover, research on the developing Nicaraguan Sign Language (NSL) suggests that not only do younger learners reach better fluency, but that NSL has evolved to be more learnable and more grammatically structured in the second generation of child learners (Senghas & Coppola, 2001). This prediction is also supported by the single ILM study that compares children to adults on the same non-linguistic task: Kempe et al. (2015) found that in a visual recall task, children created more identifiable and less complex visual patterns in comparison to adults. They conclude that structure (or less random patterns) emerged more readily in child chains, with children reducing complexity to a level that allowed them to reproduce the patterns as successfully as adults – despite having inferior visual working memory. Taken together, these findings suggest that children may impose more structure compared to adults.

On the other hand, children have less mature cognitive resources (e.g., working memory) and are generally worse in artificial language learning tasks in laboratory settings (Ferman & Karni, 2010; Perry, Axelsson & Horst, 2015), suggesting they might show inferior performance overall. There are also reasons to believe that children differ from adults in their learning and processing skills, which may lead to different biases, different preferences and different trends of learning across development (Arcuili & Simpson, 2011; Hudson-Kam & Newport, 2005). Supporting this claim, artificial language learning studies show clear age-related differences in both learning and generalization: children are more conservative in learning a new structures in comparison to adults (Boyd & Goldberg, 2012), and overgeneralize more than adults (Wonnacott, Brown & Nation, 2013). Importantly, if children are guided by biases that are quantitatively and/or qualitatively different than adults', like overgeneralization and eliminating variation, they may differ in their sensitivity to cultural transmission effects and may exhibit different patterns entirely.

The Current Study

Our goal is to contrast these two predictions by conducting the first large-scale study of iterated language learning in children and adults. We use a novel, child-friendly paradigm that closely resembles previous work with adults. Importantly, we will use the same task with both age groups to enable the comparison between them, as was done in Kempe et al. (2015). We examine the changes in the structure and learnability[1] of the languages produced by children and adults over time, with two questions in minds: (1) Is there an overall difference in performance between children and adults? (2) Will children, like adults, produce more learnable and more structured languages over time? Given that skills like statistical learning, explicit learning,

attention and working memory all improve with age, we predict some degree of difference in the overall performance between children and adults. However, we ask whether cultural transmission affects both age groups in a similar way, resulting in similar trends and rate of change. As children may have different biases for regularization, generalization and systemization, linguistic structure may emerge *faster* in one of the age groups.

Importantly, we will directly compare the performance of children and adults by using a more sensitive statistical analysis than used in previous studies. Unlike Kirby et al. (2008), in which the increase in structure and learnability was demonstrated by examining the differences between the first and last generation only, we will examine the course of change over all 10 generations by using regression models. Using regression models has several advantages: we can see changes in linguistic parameters across the course of all generations, rather than just the first and the last; we can examine interactions between our effects of interest, like age group and the effect of time; and we can control for several factors (e.g., gender) and for individual differences using random slopes and intercepts.

Method

The experiment utilizes a diffusion chain paradigm, the most common technique in ILM studies, in which all learners (apart from those in the initial generation) are trained on the output produced by previous learners in the chain. Diffusion chains in this experiment consisted of 10 generations of single participants. Our design is based on Experiment 1 in Kirby et al. (2008) with the following modifications: we used alien figures instead of geometric shapes; the number of items was reduced by half; a human experimenter interactively accompanied the learning; and we used a "syllable bank" instead of free typing. All other conditions were matched to Kirby et al. (2008), including a learning bottleneck, three dimensions of meaning, varying length of words and multiple diffusion chains. In this experiment, there were no limitations on the number of repeated words participants can produce. In Kirby et al. (2008), this design led to a decrease in the number of unique words (improving learnability but creating ambiguities) and to the emergence of linguistic structure, with homonyms assigned according to some semantic dimension.

Participants

90 children (age range: 6.5-12y, mean age: 8:8y, 41 boys and 49 girls) and 40 adults (age range: 21-68y, mean age 33y, 10 men and 30 women), comprising a total of 4 distinct adult chains and 9 distinct child chains. All child participants were visitors at the Bloomfield Science Museum in Jerusalem and were recruited for this study as part of their visit in the Israeli Living Lab in exchange for a small reward. Of the adult participants, half were family members visiting the Living Lab and half were undergraduate students at the Hebrew University, recruited

---

[1] Thanks to Smith & Kirby, who kindly provided us with their code for these analyses, we were able to use the same algorithms to compute structure and learnability as used in the original paper.

2

for this study for credit or a small payment. All participants were native literate speakers of Hebrew.

Materials

At the beginning of each diffusion chain, 12 words were randomly drawn from a closed set of 16 nonsense words[2], all of which did not contain or resemble any existing words in Hebrew (as judged by a separate sample of native speakers). These 12 words were then randomly assigned as labels to 12 different items, creating the initial language on which the first participant was trained. We used different types of alien figures, appearing in different colors, either alone or in a group. Thus, items varied along three semantic dimensions: alien type (A, B or C), color (blue or red) and plurality (single or plural). Stimuli included all possible combinations of these three semantic dimensions. Figure 1 below shows the meaning-space structure used in this task, with an example on either side:
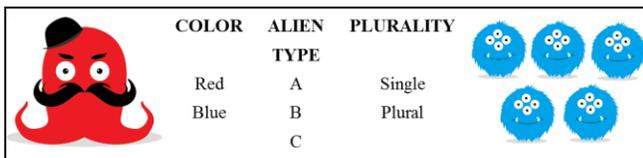
Figure 1: The three semantic dimensions of items in the task

Procedure

Participants were told they are about to learn an alien language that describes many different types of aliens, and that they should try and learn it as best they can. The experiment had three stages: initial exposure, practice and test. Participants were always exposed to a random subset of the target language (SEEN words) during initial exposure and practice, simulating a learning bottleneck. Specifically, participants were trained on just 9 out of 12 words in the language, yet were tested on all items, including UNSEEN words. Note that while adult participants in Kirby et al. (2008) were trained for over 45 minutes, such long sessions are impractical with children. We therefore settled on two sets of exposure to the SEEN words, once during initial exposure (including active verbal production) and once during practice (including active written reconstruction).

The initial exposure phase consists of a random sequence of items from the SEEN subset appearing on the screen together with their label. The experimenter read the label out loud several times and encouraged participants to remember this pairing. Both children and adults were required to reproduce the label aloud before moving on to the next item. During the following practice phase, participants were exposed again to all SEEN items and then had to recreate the labels using a pre-given "syllable bank". Then, participants completed a test phase: they were

presented with a series of items without labels, and were required to provide the correct labels according to what they've learned so far, using the same "syllable bank".

Importantly, transmission was implemented in the following way: for each participant, we took the 12 labels produced by him/her during the test and used them as the input language for the next participant in the chain. Specifically, while the first participant was trained on the random initial language drawn by the computer, the second participant learned the language produced by the first participant, and so on for 10 generations of participants.

Results

We examine the performance of children and adults on the same task by looking at two parameters: (1) language learnability, measured by transmission error (normalized Levinstein distances between input and output strings); and (2) linguistic structure, measured with the z-scores produced by a Monte-Carlo algorithm with 1,000 iterations, in respect to the degree of similarities between signals and meanings in a given language (Pearson correlation between form and meaning distances). For a detailed explanation regarding the coding of these parameters, see Kirby et al. (2008).

The transmission error reflects participants' accuracy in reproducing the language, with easier languages eliciting fewer mistakes. Thus, an increase in learnability should be accompanied by a decrease in transmission error. As for structure, the z-score (or structure score) of a given language indicates how likely it is that its structure is created by chance. The higher the z-score is, the smaller the chances that the mapping between words and meanings in this language is random. Thus, an increase in linguistic structure should result in a significant increase in structure score. If the structure score for a given language is higher than 1.96, the language has significantly consistent mappings between words and meanings which is less than 5% likely to have been created by chance.

Language Learnability

Figure 2 shows the changes in mean transmission error as a function of time in both child and adult chains. A decrease in error over generations indicates an increase in language learnability. As can be seen, transmission error generally decreased over time. Following Kirby et al. (2008), we examined the difference in error between the first and final generations. This analysis confirmed a significant reduction in error for both children (mean error at generation 1 =0.75, mean error at generation 10 =0.43, t(8.2)=3, p<0.05) and adults (mean error at generation 1 =0.69, mean error at generation 10 =0.15, t(5.1)=9.91, p<0.01).

We used a mixed-effect linear regression model to predict mean transmission error in each generation (Table 1). The fixed effects were gender, generation number, age group and the interaction between the latter two. The model had the maximal random effects structure justified by the data that would converge, including random intercepts for different chains.

---

[2] All words in the initial and later languages were limited to the same 8 syllables, chosen based on Hebrew phonology: "šu", "gu", "di", "ki", "so", "mo", "bal" and "taz". We included CVC syllables with open vowels, which are common used in Hebrew.

3
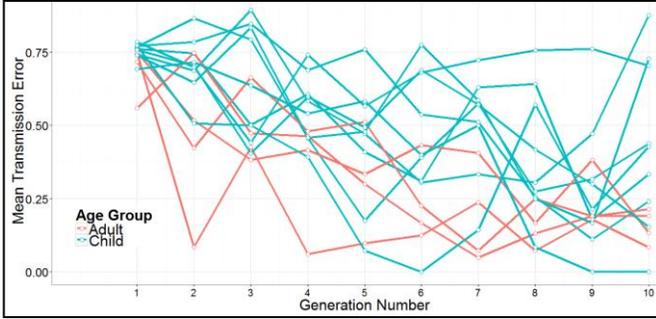
Figure 2: Mean transmission error by generation number and age group.

Table 1: Learnability model

	Estimate	Std. Error	z value	*p*-value
(Intercept)	0.35943	0.0416	8.63893	< .001 ***
Generation Number	**-0.05099**	**0.00933**	**-5.46028**	**< .005 ****
Age Group (Children vs. Adults)	**0.19643**	**0.04973**	**3.94921**	**< .005 ****
Gender (Male vs. Female)	-0.02982	0.03213	-0.92797	> .1
Age Group X Generation Number	0.00549	0.01136	0.48344	> .1

The model showed that generation number has a strong negative effect on transmission error, with errors significantly decreasing over generations (β=-0.05, SE=0.009, t=-5.46, p<0.005). That is, the languages of both children and adults become easier to learn over time. We found a significant difference between children and adults, with the mean transmission error being significantly higher in children (β=0.19, SE=0.04, t=3.94, p<0.005). This suggests that children make more mistakes than adults overall. Crucially, the interaction between age group and generation number was not significant (β=0.005, SE=0.01, t=0.48, p>0.1), so though children are inferior learners, the effect of time on learnability is similar across age groups.

But why are languages becoming more learnable over time? Similar to Kirby et al (2008), the languages of both children and adults in this experiment were characterized by a rapid decrease in the number of distinct words. Here, the number of unique words dropped to as few as only two words in certain chains. Because there were less unique words to memorize overall as chains progress, participants in later generations had smaller chances to make a mistake, which naturally increases learnability. Confirming this claim, lower rates of transmission error were strongly associated with a smaller number of distinct words for both age groups (t(128)=11.2, *r*=0.7, p<0.01)

When participants use this strategy of underspecification, multiple semantic dimensions are encoded using holistic labels, losing much of their informativity. Using just a hand-full of words is functionally useless and indeed varies from natural language, which tend to be expressive. Yet while it is underspecified, such ambiguous languages can still be structured and systematically encode some dimensions of meaning. We now turn to examine language structure.

Language Structure

Figure 3 below shows the changes in structure score as a function of time in both child and adult chains. An increase in structure score over generations reflects an increase in linguistic structure. Dots that fall above the black line represent languages that have consistent and non-random signal-to-meaning mapping at p<0.5.

Importantly, the emergence of consistent mappings between meanings and signals was accomplished by both children and adults multiple times during this experiment, with approximately 20% of languages having more structure than randomly structured languages. Most of these languages were produced in later generations, supporting the role of cultural transmission in the emergence of linguistic structure. Looking at these significantly structured languages confirmed that homonyms were not assigned at random: children and adults created languages in which homonymity was structured along some sematic dimension, closely resembling the result of Kirby et al. (2008).

For example, in one child chain the final language converged to 3 distinct words representing each alien type regardless of color and quantity: all aliens of type A were called "*didi*", all aliens of type B were called "*balgu*" and all aliens of type C were called "*šuki*" (Figure 4). Similar structure emerged in adults' chains. In a different child chain, systematic structure emerged already in generation 8, and was transmitted flawlessly to the last two participants. This language converged to just 2 distinct words representing alien color, regardless of type and quantity: "*ditaz*" for all red aliens and "*balšu*" for all blue aliens.

Like Kirby et al. (2008), we found that adult languages did show a significant difference in structure score between first and final generation (mean structure at generation 0 =0.6, mean score at generation 10 =1.89, t(5.3)=-3.14, p<0.05). Interestingly, children's languages did not show such a change (mean structure at generation 0 =0.54, mean structure at generation 10 =1.05, t(10.7)=-0.78, p=0.45).

We used a mixed-effect linear regression model to predict the structure score in each generation (Table 2), with similar fixed and random effects structure as in the previous model. The model showed a significant difference between children and adults (β=-0.6, SE=0.2, t=-2.3, p<0.05), with adult languages being significantly more structured than those of children. However, though an increase in generation number was associated with higher structure scores, this positive effect was unfortunately not strong enough to reach significance[3] (β=0.1, SE=0.08, t=1.22, p>0.1). There was also no significant interaction between age group and generation number (β=0.02, SE=0.1, t=0.28, p>0.1), indicating that though adults' languages were more structured overall, time affects structure similarly across age groups: both children and adults showed the same non-significant trend of increase in structure over generations.

---

[3] We believe the source of this null-effect is the Monte-Carlo algorithm's poor approximation in case the input sample is uniformly distributed and/or has little variation, as in the case of languages with a small number of homonyms.

4
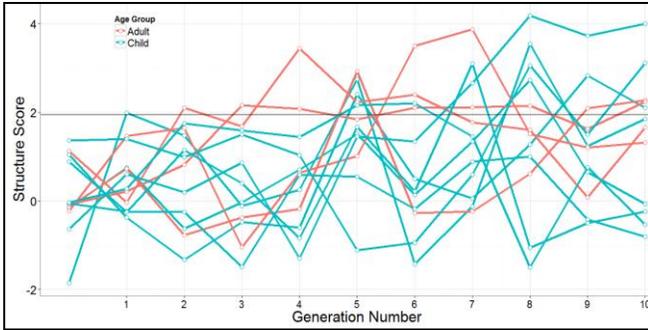
Figure 3: Mean structure score by generation number and age group.

Table 2: Structure model

	Estimate	Std. Error	*t* value	*p* value
(Intercept)	1.261367	0.229747	5.490242	<.001 ***
Generation Number	0.106558	0.08679	1.227778	> .1
Age Group (Children vs. Adults)	**-0.65355**	**0.273864**	**-2.3864**	**< .05 ***
Gender (Male vs. Female)	0.325006	0.220094	1.47667	> .1
Age Group X Generation Number	0.029634	0.105019	0.28218	> .1

	Alien A	Alien B	Alien C	
Red	didi	balgu	šuki	Single
	didi	balgu	šuki	Plural
Blue	didi	balgu	šuki	Single
	didi	balgu	šuki	Plural

Figure 4: A significantly structured child language in generation 10

	Alien A	Alien B	Alien C	
Red	ditaz	ditaz	ditaz	Single
	ditaz	ditaz	ditaz	Plural
Blue	balšu	balšu	balšu	Single
	balšu	balšu	balšu	Plural

Figure 5: A significantly structured child language in generations 8 through 10

Discussion

We found that when underspecification was possible, children and adults behaved similarly: their languages became easier to learn over time at a similar pace and by using the same strategy of reduced words and structured homonymity. Significantly structured languages with non-random signal-to-meanings mapping emerged in child and adult chains on many occasions. However, there was no evidence of significant increase in structure over time for children on either analysis. Importantly, despite adults' overall superiority in this experiment (making less mistakes and creating more structured languages), the effects of cultural transmission on languages' structure and learnability were similar for both age groups. Taken together, this study includes several novel findings: (a) Adults significantly outperform children in this paradigm; (b) Children's

languages become significantly more learnable over time in the same manner and pace as adult languages; (c) Children, like adults, can create significantly structured languages. In light of these findings, we can now discuss the questions introduced in the beginning of this paper.

We found that adults were overall better than children on both measured parameters (question 1): they were better learners in general (reflected by lower transmission errors overall) and created more structured languages from the very beginning. This result is in line with previous artificial language learning studies showing that children are inferior learners in laboratory settings despite their optimal acquisition of natural language (Ferman & Karni, 2010). This finding may be driven by developmental differences in key cognitive functions: children have a more limited memory capacity, more problems in sustaining attention, less mature problem-solving strategies and more difficulties in making object-label associations in artificial language tasks, all of which are relevant skills in this paradigm.

With regard to the emergence of more learnable and structured languages (question 2), children and adults showed a somewhat different trend. Language learnability significantly increased over time for both children and adults in a similar fashion, with the same reduction in error as chains progress. Importantly, despite making more mistakes in general, children developed easier and more learnable languages in the same pace as adults and by using the same strategies (i.e., introducing under-specification), suggesting that both age groups are basically guided by the same learnability biases. Similarly, Kempe et al. (2015) reported no differences in the learnability biases of children and adults. Nevertheless, we found no evidence for a significant increase in structure over time: generation number had the same positive (yet *not* significant) effect on structure for both children and adults using mixed-effects models. In other words, while we predicted that children's languages would become more learnable and more structured over repeated iterations, our results support this prediction only partially. Note that if we examine adults' performance using the less-subtle comparison between initial and final generations, our results mirror Kirby et al. (2008): adults did show a significant increase in structure between the first and last generations, while children did not. This discrepancy raises the question of the reliability of comparing only the first and last generations: a different pattern may be seen in previous studies when including information from all 10 generations. Interestingly, both analyses do not align with the predictions drawn from NSL studies and Kempe et al. (2015): children did not have a stronger bias for linguistic structure in comparison to adults.

While a number of child languages did have significant systematic structure with non-random signal-to-meaning mapping, such languages were only a small subset of child languages. There are several possible interpretations of this result: it could indicate that children are less likely to introduce structure during cultural transmission, a finding that is in line with accounts that view adults as the major

5

agents of linguistic change (Labov, 2007). Alternatively, it could reflect children's difficulty with the additional cognitive demands posed by artificial language learning tasks. In our study, the discovery of the semantic features of items (such as color, type and plurality) was crucial for the emergence of linguistic structure. If children failed to remember all the features of the aliens they saw, this would dramatically reduce their chances of creating non-random languages with consistent mappings between meanings and signals. Since children have trouble in remembering all the features of novel items presented to them (Perry et al., 2015), this could explain why we didn't find an increase in linguistic structure for child chains. Interestingly, the regression analyses revealed parallels between children and adults, with the effects of cultural transmission being similar in pace and magnitude for both age groups.

Finally, the languages that emerged in our study were degenerated in terms of expressivity (as in Kirby et al., 2008). In their paper, a second experiment was conducted where homonyms were filtered out of the language before transmission to the next participant. Under this condition, morphology-like structure emerged in adults. Future work with children should also include a similar experiment in which ambiguities are not allowed to examine whether children also create compositional structure under such conditions. Another option is to introduce a communicative pressure, which serves as the natural equivalent to disfavoring underspecification (Kirby, Tamariz, Cornish & Smith, 2015). Such work is important for evaluating children's ability to create compositional linguistic structure, a crucial feature of natural languages.

Conclusions

In sum, the results of this study suggest that iterated learning models operate in the same way on both children and adults. This finding strengthens the claim that cultural transmission can truly shape languages to be more compatible with learners' limitations and needs. Yet, since children failed to show an increase in structure, more work is required in order to verify this theory.

References

Arciuli, J., & Simpson, I. C. (2011). Statistical learning in typically developing children: the role of age and speed of stimulus presentation. Developmental Science, 14(3), 464-473.

Bickerton, D. (1984). The language bioprogram hypothesis. Behavioral and brain sciences, 7(02), 173-188.

Birdsong, D. (Ed.). (1999). Second language acquisition and the critical period hypothesis. Routledge.

Boyd, J., & Goldberg, A. (2012). Young children fail to fully generalize a novel argument structure construction when exposed to the same input as older learners. Journal of Child Language, 39, 457–481.

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 Proc.

35th Annual Conference of the Cognitive Science Society (pp. 340-345).

Cornish, H., Tamariz, M., & Kirby, S. (2009). Complex adaptive systems and the origins of adaptive structure: What experiments can tell us. Language Learning, 59(s1), 187-205.

Ferman, S., & Karni, A. (2010). No childhood advantage in the acquisition of skill in using an artificial language rule. PloS one, 5(10).

Goldin-Meadow, S., & Mylander, C. (1998). Spontaneous sign systems created by deaf children in two cultures. Nature, 391(6664), 279-281.

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.

Kempe, V., Gauvrit, N., & Forsyth, D. (2015). Structure emerges faster during cultural transmission in children than in adults. *Cognition*, *136*, 247-254.

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., Griffiths, T., & Smith, K. (2014). Iterated learning and the evolution of language. Current opinion in neurobiology, 28, 108-114.

Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. Cognition, 141, 87-102.

Labov, W. (2007). Transmission and diffusion. *Language*, 344-387.

Perry, L. K., Axelsson, E. L., & Horst, J. S. (2015). Learning What to Remember: Vocabulary Knowledge and Children's Memory for Object Names and Features. Infant and Child Development.

Ramscar, M., & Gitcho, N. (2007). Developmental change and the nature of learning in childhood. Trends in cognitive sciences, 11(7), 274-279.

Reali, F., & Griffiths, T. L. (2009). The evolution of frequency distributions: Relating regularization to inductive biases through iterated learning. Cognition, 111(3), 317-328.

Senghas, A., & Coppola, M. (2001). Children creating language: How Nicaraguan Sign Language acquired a spatial grammar. Psychological Science, 12(4), 323-328.

Tomasello, M. (2009). Constructing a language: A usage-based theory of language acquisition. Harvard University Press.

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.

Wonnacott, E., Brown, H. E., & Nation, K. (2013). Comparing generalization in children and adults learning an artificial language. Poster from Child language seminar. Manchester, UK.