



Full Length Article

Children's narrow learning bottleneck accelerates the emergence of statistical properties of language

Lucie Wolters^{a,*}, Simon Kirby^b, Inbal Arnon^c

^a Department of Cognitive and Brain Sciences, Hebrew University of Jerusalem, Mt Scopus, Israel

^b Centre for Language Evolution, School of Philosophy, Psychology & Language Sciences, University of Edinburgh, UK

^c Department of Psychology, Hebrew University of Jerusalem, Mt Scopus, Israel

ARTICLE INFO

Keywords:

Cultural transmission
Language learning
Transmission bottleneck
Age-related differences
Zipfian frequency distributions

ABSTRACT

Learners are limited in how much they can learn when a language is passed from one generation to the next. This transmission bottleneck is thought to play a crucial role in shaping language by promoting the emergence of systematic structure that enhances learnability. Prior work has investigated the impact of transmission bottlenecks by experimentally manipulating how much input learners receive. However, in natural language learning, bottlenecks can also result from learners' cognitive limitations. Here, we take a naturalistic approach to testing the effects of such learning bottlenecks by comparing cultural transmission in two populations that differ in their cognitive capacities: children and adults. Specifically, we investigate the emergence of two cross-linguistic statistical properties that were shown to facilitate learning: statistically coherent units, and Zipfian frequency distributions. Transmission chains of child and adult learners observed and reproduced sets of colour sequences that were produced by the preceding participant, simulating the process of cultural transmission. For both age groups, initially random sets of sequences evolved to have statistically coherent units with a Zipfian frequency distribution, but crucially, these statistical properties emerged faster with children. In addition, both children and adults learned the final child-generated sequence sets more easily than the adult-generated ones. Together, these findings suggest that children's more limited cognitive capacities create a narrower learning bottleneck, which in turn accelerates the emergence of statistical properties that enhance learnability. These results provide novel evidence that cognitive constraints can modulate the emergence of systematic structure that facilitates learning.

1. Introduction

While the range of linguistic expressions in everyday communication is effectively unbounded, language learners are constrained in how much they can learn as language is passed from one generation to the next. This transmission bottleneck has been proposed to play a crucial role in shaping language by promoting the emergence of systematic structure that enhances learnability (Brighton, 2002; Ferdinand & Zuidema, 2009; Kirby, 2000, 2001; Smith et al., 2003). For example, a productive affix marking past tense (e.g., the English suffix -ed) improves learnability by reducing the need to memorise each verb form separately, facilitating faithful transmission despite limited learning capacity. A substantial body of work has examined how transmission bottlenecks shape linguistic structure by experimentally manipulating their size (Ferdinand & Zuidema, 2009; Kirby et al., 2007; Kirby et al.,

2015; Kirby & Hurford, 2002). This work suggests that narrower bottlenecks accelerate the emergence of systematic structure because learners must generalize from limited input. In contrast, wider bottlenecks lead to more faithful transmission, allowing for the preservation of irregularities and slowing down the emergence of systematic structure. Typically, transmission bottlenecks are imposed on learners by varying *how much* of a behaviour they are exposed to: for example, given an artificial language that has ten novel object-label mappings, studies vary whether learners see all ten meanings (the widest possible bottleneck), or only a subset of them (a more narrow bottleneck). However, in natural language learning, bottlenecks are not only imposed by the quantity of input that learners receive; they are inherent to cognition. Learners' limited memory, attention, and processing capacities inherently restrict how much of the input they can encode and retain (Christiansen & Chater, 2016; Kirby & Tamariz, 2022). In this study, we investigate

* Corresponding author.

E-mail addresses: lucia.wolters@mail.huji.ac.il (L. Wolters), simon.kirby@ed.ac.uk (S. Kirby).

<https://doi.org/10.1016/j.cognition.2026.106622>

Received 25 November 2025; Received in revised form 1 June 2026; Accepted 9 June 2026

Available online 23 June 2026

0010-0277/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

how such learning bottlenecks shape the emergence of systematic structure. Rather than manipulating the quantity of input, we compare two populations that differ in their cognitive capacities: children and adults. By leveraging naturally occurring differences in cognitive capacities, we simulate the effects of narrower and wider learning bottlenecks and examine how they shape the emergence of systematic structure over generations. If cognitive limitations function as transmission bottlenecks, then populations that differ in cognitive capacity should shape language differently over cultural transmission.

Because children's cognition is still developing, they differ from adults in memory, attention, and processing efficiency (Craig & Bialystok, 2006). These limitations impose a narrower learning bottleneck which could accelerate the emergence of systematic structure in cultural transmission with children compared to adults. Several findings from the emerging sign language literature suggest that children indeed introduce more systematic structure than older learners during learning. For example, spatial modulation in Nicaraguan Sign Language (NSL) originated primarily in younger learners (exposed to the language before the age of ten), rather than adults from the same cohort (Senghas & Coppola, 2001). Similarly, younger learners produced more segmented and linearised expressions for complex events compared to older ones (Senghas et al., 2004), a pattern has been linked to children's stronger tendency to segment and linearize their gestures compared to adolescents and adults (Bohn et al., 2019; Clay et al., 2014). Moreover, children have been suggested to be more likely to regularize grammatical inconsistencies during learning, resulting in more systematically structured output: deaf children were found to produce systematic grammatical forms despite exposure to the inconsistent sign language of their parents (Singleton & Newport, 2004), and while adults tend to reproduce probabilistic variation, 5- to 7-year-olds have a tendency to regularize variation and learn categorical rules (Hudson Kam & Newport, 2005, 2009; Newport, 2020; Samara et al., 2017). This difference in regularization behaviour between adult and child learners has been attributed to children's more limited memory, which may prevent them from fully learning and reconstructing the variation (Hudson Kam & Chang, 2009; Hudson Kam & Newport, 2005, 2009). Supporting this idea, adults were found to be better at reproducing variation when their cognitive load was experimentally decreased (Hudson Kam & Chang, 2009), and to regularize more when the input variation was more complex (Hudson Kam & Newport, 2009). Regularization by adults is also impacted by pragmatic assumptions about the domain being learned (Ferdinand & Zuidema, 2009; Montag, 2021; Perfors, 2012), but the contribution of such factors in children is not clear.

Taken together, previous findings suggest that children's more narrow learning bottleneck may accelerate the emergence of systematic structure during cultural transmission. However, direct experimental evidence for this claim remains limited. Cultural transmission can be investigated in the lab using iterated learning paradigms where participants form transmission chains in which each participant learns and produces a behaviour that was produced by the previous participant in the chain (Kirby et al., 2008). Only a small number of studies to date have used this experimental paradigm to test how cultural transmission by children affects the emergence of systematic structure (Kempe et al., 2015, 2019; Raviv & Arnon, 2018). In a non-linguistic iterated learning task where children and adults had to reproduce visio-spatial patterns (dots on a grid), systematic structure emerged more readily with children: children increased the systematicity of the patterns to a level that allowed them to reproduce the patterns as successfully as adults (Kempe et al., 2015). In contrast, studies using artificial languages, where learners acquire mappings between meanings and labels, found that less structure emerged with children (Kempe et al., 2019; Raviv & Arnon, 2018). These contrasting findings may be explained by the higher cognitive load of learning form-meaning mappings, which could have impaired children's generalization and hindered the emergence of systematic structure. Consistent with this interpretation, artificial language learning studies with adults show that generalization depends on

successful learning of the input (Johnson et al., 2020), and computational models found that systematic structure does not emerge when learners only learn a small subset of the language (Kirby & Hurford, 2002). The divergent findings may also be driven by age-related differences in task demands: children may struggle in learning the compositional structure of form-meaning mappings because it requires certain pragmatic insight about communicative expressivity (Kempe et al., 2019; Naigles, 2002). This interpretation aligns with prior work showing that children experience difficulties in referential communication tasks (Kempe et al., 2019; Matthews et al., 2007; Nilsen & Graham, 2009). Taken together, the current findings regarding age-related differences in the emergence of structure are mixed, calling for additional investigation using a task that is simpler and where age-effects from learning are not conflated with pragmatic task demands associated with the learning of form-meaning mappings.

In this study, we will examine the emergence of two cross-linguistic statistical properties that have been shown to emerge through cultural transmission with adult learners but have never been examined in children (Arnon & Kirby, 2024; Wolters, Kirby, & Arnon, 2025). First, that language consists of parts that tend to be statistically coherent. For example, the transitional probabilities between adjacent syllables tend to be higher when both syllables belong to the same word, as opposed to when they cross a word-boundary (Saksida et al., 2017; Stärk et al., 2022). And second, that across languages, word frequencies follow a Zipfian or near-Zipfian distribution when ordered by rank, such that a small number of words occur very frequently while many occur rarely, with frequency declining according to a power law (Ferrerri Cancho, 2005; Kaplan & Yu, 2025; Kimchi et al., 2023; Lavi-Rotbain & Arnon, 2023; Mehri & Jamaati, 2017; Piantadosi, 2014; Zipf, 1949). Infants, children and adults can use the statistical coherence of words to discover word boundaries in continuous speech, a crucial milestone in language learning (for review: Saffran & Kirkham, 2018). In addition, speech segmentation is facilitated when the relevant words follow a Zipfian distribution, as opposed to a uniform one where all words appear equally often (Kurumada et al., 2013; Lavi-Rotbain & Arnon, 2022; Wolters, Ota, & Arnon, 2025). Based on these findings, prior work hypothesized that the learnability advantages of statistically coherent units and Zipfian frequency distributions can drive their emergence in language over the course of cultural transmission, explaining their recurrence across languages (Arnon & Kirby, 2024; Christiansen & Chater, 2016; Cornish et al., 2017; Isbilen & Christiansen, 2020; Shufaniya & Arnon, 2021; Wolters, Kirby, & Arnon, 2025). Recent work tested this hypothesis by having adults reproduce sets of non-linguistic colour sequences in transmission chains (Arnon & Kirby, 2024; Wolters, Kirby, & Arnon, 2025). Sequence sets evolved to have statistically coherent units (sub-sequences) with a frequency distribution approaching a power law. Importantly, sets with more skewed frequency distributions were reproduced with fewer transmission errors, suggesting that the emergence of the statistical properties increased learnability. In the present study, we use the same iterated sequence learning paradigm to investigate the effect of learning bottlenecks on the emergence of statistically coherent units and Zipfian frequency distributions, by comparing the rate of their emergence in transmission chains with children and adults (Study 1). This paradigm is particularly well-suited for comparing the emergence of systematic structure across age groups. Its non-linguistic design substantially reduces task complexity, alleviating the constraints that often make it difficult for children to generalize from limited exposure in short experimental sessions, while also avoiding potential age-related differences in the pragmatic assumptions involved in learning form-meaning mappings (Kempe et al., 2019; Raviv & Arnon, 2018). At the same time, it is complex enough to create a learning bottleneck for both child and adult learners.

We predict that children's more limited cognitive capacities will create a more narrow learning bottleneck in comparison to adults, which in turn will accelerate the emergence of properties that facilitate learning. As an indication of this narrow bottleneck, we expect children

to make more reproduction errors at the start of transmission in comparison to adults. Importantly, we also expect to see the faster emergence of statistically coherent units with a skewed frequency distribution in the transmission chains with children, increasing the learnability of the sequence sets. If alternatively, the more limited cognitive capacities of children hinders their learning, then we expect to see a slower or no emergence of statistically coherent units with a skewed frequency distribution in the transmission chains with children, accompanied by less or no improvement in learnability over transmission. In Study 2, we ask children and adults to reproduce the final sequence sets of the child and adult transmission chains. This allows us to assess whether the statistical properties emerging from transmission chains with children and adults similarly impact learnability across age.

2. Study 1: Iterated sequence learning by children and adults

2.1. Methods

2.1.1. Participants

A total of 100 children and 69 adults participated in the experiment, forming six transmission chains per age-group, with ten generations in each chain. Children ranged in age from 7;5 to 11;0 years (mean age = 8;5). Children in this age-range exhibit more limited cognitive capacities than adults, while being sufficiently old to perform the task, allowing us to investigate the effects of learning bottlenecks on cultural transmission.¹ Children were recruited in the Living Lab at the Bloomfield Science Museum in Jerusalem (<https://www.thelivinglabjerusalem.com/>). Parental consent and spoken assent were obtained for all children. All children were Hebrew speakers² with no known language or learning difficulties. Each child received a small prize for participation. During data collection, 40 children were excluded from the study for having more than nine invalid trials during testing ($n = 33$, see Section 2.1.3), or because they chose not to complete the study ($n = 7$).

Adults were recruited on the crowdsourcing platform Prolific (<https://www.prolific.co/>) and pre-screened to be native English speakers, aged 18–40, have no prior participation in similar experiments run by the research team, and a prolific approval rating above 95%. Participants received monetary compensation at an hourly rate of nine pounds. During data collection, nine participants were excluded due to (1) having more than nine invalid trials during testing ($n = 4$), (2) having more than three incorrect trials during practice ($n = 4$), and (3) using recording devices to complete the study, based on their answer to a question posed at the end of the study ($n = 1$).

2.1.2. Stimuli

The stimuli consisted of sets of 30 colour sequences (red, yellow, green and blue), presented as successive colour flashes on a playing board modelled after the Simon Game (see Fig. 1). Because this was a transmission experiment, participants reproduced sets of sequences that were produced by the preceding participant in the chain (see Procedure, Section 2.1.3). The sequence sets for participants in the initial generation of the transmission chain (generation zero) were created by randomly generating sequences with a length of ten. The sequence length was set to ten, rather than the twelve used in prior studies (Arnon & Kirby, 2024; Wolters, Kirby, & Arnon, 2025), to make the task more suitable for children.



Fig. 1. The playing board (Hebrew version).

2.1.3. Procedure

Children and adults were assigned to transmission chains, forming six independent transmission chains per age group, with ten participants in each chain. Within each chain, each participant represented a single “generation”. To simulate the process of cultural transmission over generations, participants reproduced a sequence set that was produced by the immediately preceding participant. The first participant in each transmission chain saw a randomly generated sequence set (as specified in Section 2.1.2), ensuring that the initial input contained no systematic structure beyond chance. Because each participant only learned from the output of the immediately preceding generation, any systematic changes observed across generations reflect cumulative effects of learning.

The procedures for the children and adults were kept as similar as possible to ensure a valid comparison between groups. The main difference was that adults completed the experiment remotely on a browser-based platform, while children participated in the lab using a tablet, with an experimenter present to ensure they were progressing through the task. Any additional procedural differences were necessary to adapt the adult version for online testing and are described below.

Each participant saw the set of 30 sequences in a randomized order, with a self-timed break after the tenth and twentieth sequence. At the start of the experiment, participants were told they would see and repeat 30 sequences. After seeing the playing board for 400 ms, they saw a sequence of colour-flashes.³ Each colour lit up for 300 ms, with an interval of 300 ms between individual flashes.⁴ At the end of the sequence, the prompt “repeat” appeared. Children reproduced the sequences by tapping on the tablet, whereas adults reproduced the sequences using the q, w, a, and s keys on their keyboard, which respectively mapped onto the top-left, top-right, bottom-left and bottom-right-location of the playing board.⁵ A colour lit up for 300 ms when a key was pressed. After

¹ Informal piloting with an age-range from 5 to 12 years, suggested that children aged over 6 would have no issue to complete the task. We chose to test a more narrow age-range (3.5 years) to justify generalization across age for the child transmission chains. The exact age-range of 7;5–11;0 was decided due to practical reasons, as this was the sample of participants that was expected to be available in our testing location.

² There was a Hebrew and an Arabic version of the experiment. Only one Arabic speaking child participated in the study.

³ Hex colour codes of the colours on the playing board: Green: #005500, red: #550000, blue: #000055, yellow: #D5BE1A. Hex colour codes of the colours on the playing board during the flashes: Green: #00FF00, red: #FF0000, blue: #0000FF, yellow: #FFFC00.

⁴ The presentation rate and inter-stimulus interval were set to match the parameters used in Cornish et al. (2013) and Wolters, Kirby, and Arnon (2025).

⁵ At the start of the experiment, participants were asked to select whether they have a QWERTY or AZERTY keyboard. For participants with an AZERTY keyboard, the a, z, q, and s keys respectively mapped onto the top-left, top-right, bottom-left- and bottom-right-location of the playing board.

each trial, feedback was shown at the centre of the playing board to help maintain participants' engagement throughout the task. For children, this feedback took the form of a smiling emoji randomly chosen from a set of eight, while for adults it was a percentage score indicating their accuracy on that trial.⁶

Before starting the testing phase, participants reproduced randomly generated practice sequences with a length of four. Participants continued to the test phase of the experiment when they reached the maximum of six practice trials, or when they correctly reproduced three sequences in a row. During the testing phase, a trial was considered invalid if the length of the produced sequence fell outside the range of 6 to 14, or if the participant took longer than five seconds to start reproducing the sequence. Invalid trials were repeated at the end of the experiment, with a maximum of nine repeated trials. If participants had more than nine invalid trials, they were removed from the transmission chain and replaced by a new participant. To control the quality of the data of the adult participants, who completed the experiment online, we applied two additional exclusion criteria: participants were excluded (1) when they had more than three incorrect trials during practice, and (2) when using recording devices to complete the study, based on their answer to a question posed at the end of the study. The experiment took approximately ten minutes to complete.

2.2. Results

The data was analyzed using linear regression models, implemented using the lme4 library (Bates et al., 2015) in R (R Core Team, 2021). Following Barr (2013), we adopted a maximal random-effects structure. Consistent with standard practice, we restricted maximality to theoretically motivated predictors (referred to as independent variables in the description of the analyses). Analyses included the data of both age groups combined, excluding invalid trials. All data and analyses can be accessed on the Open Science Foundation project page of this study: <https://osf.io/ujfvn7>.

2.2.1. Sets become more learnable over transmission for both children and adults

Transmission error for each sequence was calculated as the by-length normalized Levenshtein edit-distance between a seen and produced sequence. This is calculated as the minimum number of insertions, deletions, and substitutions required to turn one sequence into the other, divided by the length of the longest sequence. Fig. 2 shows that at the start of transmission, generation zero, transmission error was higher for children compared to adults (mean children = 0.47, sd = 0.07; mean adults = 0.34, sd = 0.04). A model including only the data of generation zero, with transmission error on trial as the outcome variable, age group as an independent variable (dummy coded, adults as the reference), trial number (mean-centered) as a control variable, and a random intercept for chain, confirms this difference is significant ($\beta = 0.12$, SE = 0.03, $t = 3.92$, $p < .01$). This indicates that the randomly generated sets were initially harder for children to reproduce, suggesting that they indeed have a more narrow learning bottleneck than adults.

Fig. 2 shows that transmission error decreased over generations in both age groups, with a slightly lower mean transmission error for children. We fitted a model with transmission error on trial as the




Fig. 2. Transmission error in reproducing individual sequences decreases over transmission, showing that sequence sets become easier to reproduce over generations in both age groups. The solid lines with error bars represent the mean error of sequences across chains within each age group (descriptive statistics); error-bars represent the 95% confidence intervals across sequences. The lines without error bars represent the mean transmission error of individual chains.

outcome variable, generation, age group and their interaction as independent variables, trial number as control variable, as well as a by-chain random slope for generation. Generation and trial number were mean-centered for this and following analyses. Transmission error decreases significantly over generations ($\beta = -0.02$, SE = 0.01, $t = -3.26$, $p < .01$), with no significant interaction between generation and age group ($\beta = -0.01$, SE = 0.01, $t = -0.91$, $p = .384$), indicating a similar reduction in transmission error across groups. Transmission error was not affected by trial number ($\beta = 0.00$, SE = 0.00, $t = 1.33$, $p = .184$).

2.2.2. Sequence length impacts learnability for both children and adults

One way in which sequence sets may become more learnable over time, for both children and adults, is through a reduction in sequence length; shorter sequences place lower demands on memory and processing, and are therefore easier to reproduce. In the children's chains, sequence lengths decrease from the starting length of 10 to a mean sequence length of 7.6 (sd by chain = 0.6), compared to a mean of 8.5 in the adult chains (sd by chain = 1.1) (see Fig. 3). This means that by the end of transmission (generation ten), sequences are on average one




Fig. 3. Length of sequences over transmission, showing that sequences decrease in length over generations in both age groups. The solid lines with error bars represent the mean length of sequences across chains within each age group (descriptive statistics); error-bars represent the 95% confidence intervals across sequences. The lines without error bars represent the mean sequence length of individual chains.

⁶ We did not use the smiley feedback with adults since we thought it would not be sufficiently motivating, whereas we expected that simple smiley feedback would not be sufficiently motivating for adult participants. Importantly, both forms of feedback seemed to be similarly effective: the reduction of transmission error was similar for children and adults.

⁷ All reported analyses were repeated after removing the randomly generated sequence sets from generation zero. Excluding generation zero did not lead to any substantive changes in the findings. The results of these analyses are available on the OSF project page.

colour shorter in the children's sets. We fitted a model with sequence length as the outcome variable, generation, age group and their interaction as independent variables, trial number as control variable, and a by-chain random slope for generation. Sequence length decreases significantly over generations ($\beta = -0.17$, $SE = 0.05$, $t = -3.71$, $p < .01$), and is overall shorter in the children's chains compared to the adults' chains ($\beta = -1.60$, $SE = 0.46$, $t = -3.46$, $p < .01$). We do not find a significant interaction between generation and age group ($\beta = 0.01$, $SE = 0.07$, $t = 0.19$, $p = .851$), indicating that there is no significant difference in sequence length reduction between children and adults. To test whether the decrease in sequence length across generations contributes to the increase in learnability of the sequences, we reran the model of transmission error in Section 2.2.1 while including an additional independent variable for sequence length (mean centered) and its interaction with generation and age group. Sequence length significantly predicts transmission error across age groups⁸ ($\beta = 0.02$, $SE = 0.00$, $t = 6.97$, $p < .001$) but we do not find an interaction between sequence length and age group ($\beta = 0.00$, $SE = 0.00$, $t = 0.22$, $p = .828$), suggesting that children and adults are similarly impacted by sequence length during sequence reproduction. Generation is still a significant predictor of transmission error ($\beta = -0.01$, $SE = 0.01$, $t = -2.51$, $p < .05$), suggesting systematic changes beyond reductions in sequence length contribute to the increase in learnability over time.

2.2.3. Segmentation method

To analyze the emergence of statistically coherent units and their statistical properties, we segmented the sequences into units using the method developed by Arnon and Kirby (2024), which has since been used to analyze other culturally transmitted sequential behaviours (Arnon et al., 2025; Wolters, Kirby, & Arnon, 2025). This method identifies unit boundaries based on drops in the transitional probabilities between colours, similar to how language learners can use drops in the transitional probabilities between syllables to discover word boundaries in continuous speech (e.g. Saffran et al., 1996). A colour's transitional probability is estimated as its conditional probability given the two preceding colours in the sequence, based on their occurrence across the entire sequence set. The segmentation method posits a boundary when the probability of a transition is low relative to that of the preceding transition, using the ratio between the two transitional probabilities. A set of non-structured sequence sets are used as a baseline to estimate when a drop in transitional probability is large enough to justify segmentation. To do this for the present study, we generated 500 random sequence sets of 30 sequences with a length of ten, and computed the transitional probability ratios for each set separately. We then created an aggregated distribution of all the transitional probability ratios and took the low five percentile point of this distribution as the segmentation threshold ($= 0.39$). For each sequence set, we applied the segmentation method and then used the posited unit boundaries to extract a lexicon⁹ of unique units that can be analyzed for their statistical properties.¹⁰ Fig. 4 shows an example of the evolution of a single sequence over the ten generations.

⁸ Note that since transmission error is quantified as the length-normalized Levenshtein edit distance, any effect of sequence length on transmission error cannot be attributed solely to learners producing shorter sequences.

⁹ By 'lexicon' we refer to a collection of segmented units extracted from the sequence set, which, unlike the traditional use of the term, do not carry any inherent meaning.

¹⁰ To ensure that any age-related effects we find are not related to the specific use of trigram probabilities (it is theoretically possible that children and adults rely on ngrams of different sizes, so that a certain ngram would be a better fit for capturing one age groups' behaviour), we conducted additional robustness analyses where segmentation is based on different n-gram sizes (bigrams, trigrams, and fourgrams). The overall pattern described in the results sections remains are consistent across n-gram sizes. The full analyses using bigrams and fourgrams can be found in Supplementary Materials A.




Fig. 4. The evolution of a single sequence out of a set of 30 (item 24 from Adult chain C), over ten generations. The black lines between colours indicate the unit boundaries posited by the segmentation method (i.e., points in the sequence where the transitional probability ratio was lower than 0.39).

2.2.4. Statistically coherent units emerge over transmission for both children and adults

To see whether the statistical coherence of units increases over transmission in the two age groups, we asked whether the transitional probabilities of colours within units increase over generations relative to those at unit boundaries. Fig. 5 shows that for both children and adults the within-unit transitional probabilities increase over transmission while the probabilities at unit boundaries remain low, and that this pattern is more pronounced for children. We fitted a model with transitional probability as the outcome variable and generation, age group, transition type (within-unit vs. between-unit, dummy-coded with within-unit as the reference), and their three-way interaction as independent variables, as well as a control variable of sequence length. A by-chain random slope for generation was included. The transitional probability of within-unit transitions is higher than between-unit transitions ($\beta = 0.20$, $SE = 0.01$, $t = 39.45$, $p < .001$), indicating the presence of statistically coherent units across generations in both age groups. As shown in Fig. 5, this pattern is stronger for children than for adults across generations. This is confirmed by a significant interaction between transition type and age group ($\beta = 0.07$, $SE = 0.01$, $t = 9.60$, p




Fig. 5. Mean transitional probability within units and between units over generations (descriptive statistics), for both age groups. The error-bars represent the 95% confidence intervals across individual transitions. The figure shows that for both age groups within-unit transitional probability increases over generations, indicating increasingly statistically coherent units, and that the rate of evolution is faster for children compared to adults.

< .001), indicating that the salience of the cue to unit boundaries is stronger in the sets produced by children compared to those produced by adults. In addition, a significant interaction between transition type and generation indicates that the probability of within-unit transitions increases more than that of between-unit transitions over generations, suggesting units increase in their statistical coherence over generations across age groups ($\beta = 0.01$, $SE = 0.00$, $t = 3.55$, $p < .001$). Crucially, a significant three-way interaction ($\beta = 0.01$, $SE = 0.00$, $t = 4.73$, $p < .001$) indicates that this effect is stronger in children: the emergence of statistically coherent units—and the corresponding emergence of salient unit boundaries—occurs at a faster rate in children compared to adults. We did not observe a main effect of sequence length, nor did sequence length interact with any of the other predictors in the model.

2.2.5. Unit distributions become more Zipfian over transmission for both children and adults

To measure the statistical structure of the frequency distributions of units over generations, we looked at the entropy of the distributions (Shannon, 1948), which increases with set size (the number of unique units) and decreases with distribution skew. Fig. 6 (left) shows that entropy decreases for both age groups. To test whether frequency distributions become increasingly skewed over generations in the two age groups, we fitted a model with entropy as the outcome variable and generation, set size (mean centered), age group and the two two-way interactions (age group $\times$ generation, and age group $\times$ set size) as predictors. A by-chain random slope for generation was included. The set sizes of lexicons are overall smaller for children, as indicated by a significant interaction between set size and age group ($\beta = 0.01$, $SE = 0.00$, $t = 4.52$, $p < .001$). Since we control for set size, the effect of generation in this model reflects an increase in skew. We find a significant effect of generation on entropy, indicating that the frequency distributions become more skewed over transmission for both age groups ($\beta = -0.01$, $SE = 0.00$, $t = -2.73$, $p < .01$). Crucially, we find that this effect is stronger for children ($\beta = -0.01$, $SE = 0.00$, $t = -3.51$, $p < .001$). This suggests that the sets that were transmitted by children evolved to be more skewed than those transmitted by adults. Fig. 6 (right), shows the change in entropy after normalizing entropy by set size,¹¹ showing the faster increase in distribution skew in the transmission chains of children.

Power law distributions like the Zipfian distribution are characterised by a linear relationship between the logarithm of frequency and the logarithm of rank. Therefore, to evaluate whether the transmitted frequency distributions increasingly resemble a Zipfian power law, we calculated the R^2 of the distributions to a linear regression between frequency and rank on a log-log scale. If the distributions become increasingly Zipfian, then we expect R^2 to increase over transmission. Fig. 7 shows that the R^2 of the frequency distributions of both age groups increase over transmission. A model with R^2 as the outcome variable, generation, age group, and their interaction as independent variables, and a by-chain random slope for generation confirms that this increase is significant ($\beta = 0.04$, $SE = 0.01$, $t = 5.99$, $p < .001$). Additionally, R^2 is overall higher for children ($\beta = 0.20$, $SE = 0.07$, $t = 3.02$, $p < .05$), suggesting that sets transmitted by children exhibit distributions with a better fit to a power law. However, the rate of change towards a Zipfian distribution is similar across both age groups, as indicated by the lack of a significant interaction between generation and age group ($\beta = -0.00$, $SE = 0.01$, $t = -0.20$, $p = .843$).

Several alternative distributions besides the power law can also show an approximate linear relationship between frequency and rank on a

log-log scale, for example the exponential distribution (Clauset et al., 2009). To address this, we conducted additional analyses comparing the goodness-of-fit of the frequency distributions to a power law, uniform and exponential distribution (following Wolters, Kirby, & Arnon, 2025). The uniform distribution was included to test whether the chains start out with a uniformly distributed frequency distribution and gain a better fit to a power law over time. To compare the fit of the observed frequency distributions to the three examined distributions (power law, uniform, exponential), we first estimated the parameters for each distribution using maximum likelihood estimation. Next, we derived the expected frequency distribution for each distribution based on the estimated parameters and the number of units and tokens in the observed distribution. We then quantified the goodness-of-fit of the observed distributions to the different examined distributions by calculating the Akaike Information Criterion (AIC). The AIC values for the different examined distributions were compared by converting them into relative probabilities. This was done by subtracting each AIC from that of the best-fitting distribution (ΔAIC), transforming the differences into relative likelihoods, and normalizing these likelihoods so they summed to one.

Fig. 8 shows that for both children and adults at generation zero, the uniform distribution provides the best fit but over subsequent generations, the relative probability of the power law distribution increases, eventually surpassing the alternatives. Interestingly, the frequency distributions in the children's chains gain a better fit to a power law earlier on in transmission (generation three) compared to those in the adult's chains (generation ten).

2.2.6. The emergent statistical properties reflect the sequential ordering of colours

One potential concern with our results is whether the emergence of the statistical properties, as well as the difference in the rate of emergence in the two age groups, might be an artifact of changes in the distribution of individual colours or the length distribution of the sequences. To confirm this is not the case, we compare the experimental data to a shuffled baseline that maintains the length of the sequences and the unigram distribution of the individual colours but destroys any sequential information in the original data. Specifically we ask whether the emergence of a Zipfian distribution across age groups and the faster emergence with children (as described in Section 2.2.5) stem from the sequential ordering of colours rather than from changes in the unigram distributions or sequence lengths.

To create the baseline, we shuffled the tokens of each set fifty times, generating fifty shuffled sets for each experimental sequence set, and then reapplied the segmentation method to each shuffled set. Like before, we look at the R^2 of the distributions to a linear regression between frequency and rank on a log-log scale. We fitted a model with R^2 per set as the outcome variable and included data set (original: 120 sets; shuffled: 6000 sets, with shuffled as the reference), age group, generation, and their three-way interaction as independent variables, along with a random intercept for chain.¹² We find that the R^2 of the experimental sets is significantly higher than that of the shuffled sets ($\beta = 0.24$, $SE = 0.03$, $t = 9.12$, $p < .001$). Furthermore, the R^2 of the experimental sets increases more over generations, as evidenced by a significant interaction between data set and generation ($\beta = 0.03$, $SE = 0.01$, $t = 4.65$, $p < .001$). This finding suggests that the emergence of statistical properties is indeed driven by changes in the sequential ordering of colours. Additionally, the difference in R^2 between the experimental and shuffled sets is larger for the those transmitted by children, compared to those transmitted by adults, as indicated by a significant interaction between age group and data set ($\beta = 0.19$, $SE = 0.04$, $t = -5.31$, $p < .001$). This suggests that the higher R^2 of sets transmitted by children (as

¹¹ Normalized entropy, also referred to as Efficiency, is calculated by subtracting the actual entropy from the maximum possible entropy for the given set size, and then dividing this difference by the maximum entropy. This results in a number between 0 and 1 that represents the fraction of the maximum possible entropy.

¹² We initially fitted the model with an additional by-chain random slope for generation but the model did not converge.




Fig. 6. *Left:* entropy over transmission; *Right:* normalized entropy over transmission. The lines with error bars represent the mean of chains (descriptive statistics); error-bars represent the 95% confidence intervals across chains. The lines without error bars each represent the mean of individual chains.




Fig. 7. *Left:* The R^2 of the linear regression between log frequency and log rank for the frequency distributions in both age groups, over generations. The lines with error bars represent the mean across chains (descriptive statistics); error-bars represent the 95% confidence intervals across chains. Over transmission, frequency distributions are becoming more Zipfian, reflected by an increase in R^2 . *Right:* The frequency distributions of units of both age groups in generation ten. Error bars represent 95% confidence intervals across chains.




Fig. 8. The relative probability of each model (power law, exponential, uniform) across generations, for each age group. Error-bars represent the 95% confidence intervals across chains. For both age groups, the relative probability of the power law model increases and surpasses the alternatives by the end of transmission.

described in Section 2.2.5) at least partially stems from the sequential ordering of the colours, rather than changes in the unigram distributions of colours or sequence lengths of the children's sets. There is no interaction between generation and age group ($\beta = -0.00$, $SE = 0.00$, $t = -1.88$, $p = .059$), or between generation, age group and data set ($\beta = 0.00$, $SE = 0.01$, $t = -0.08$, $p = .936$).

2.2.7. Lexicons become more learnable over generations for both children and adults

Both children and adults reproduced the sequences with increasing accuracy over generations (Section 3.1). Here, we ask how well children and adults reproduce the lexicons of the sets of sequences. To test this, we measured the similarity of lexicons across generations. Specifically, we created pairings between the units of lexicons from subsequent

generations, and calculated the inverse Levenshtein distance for each pairing, divided by the length of the longest unit, resulting in a similarity measure ranging from zero to one for each pair of units. To find a pairing of units that gives the maximal similarity between subsequent lexicons, we used the Hungarian algorithm (also known as the Munkres assignment algorithm),¹³ a widely used algorithm to solve assignment problems. Crucially, this algorithm ensures that each unit is paired uniquely, preventing multiple assignments of the same unit. Because lexicon sizes differ across sets, some units remained unpaired. These unpaired units were assigned a similarity score of zero.

Fig. 9 shows that the similarity of units increases over generations for both age groups. We fitted a model with the similarity of units as the outcome variable, generation, age group, and their interaction as independent variables, and a by-chain random slope for generation. Similarity increases significantly over generations ($\beta = 0.02$, $SE = 0.00$, $t = 4.51$, $p < .001$), suggesting that not only the sequences, but also the lexicons of both children and adults become increasingly learnable over generations. Moreover, similarity is overall higher for children ($\beta = 0.06$, $SE = 0.02$, $t = 2.95$, $p < .05$), indicating that the sequence sets transmitted by children evolve lexicons that are easier to reproduce by its learners, compared to those transmitted by adults. There is no significant interaction between age group and generation ($\beta = 0.01$, $SE = 0.01$, $t = 1.09$, $p = .303$).

2.2.8. Children and adults are sensitive to the transitional probabilities of colours

Previous work has shown that participants' online behaviour reflects the emerging statistical structure of the sequence sets (Wolters, Kirby, & Arnon, 2025). Specifically, the probability of the colour transitions were predictive of the reaction times of adults during reproduction, such that more probable transitions resulted in faster reaction times. This suggests that the transitional probabilities that were used to measure the emergence of the statistical properties are indeed tracked by learners. Here, we test whether this is also true for children. As children reproduced sequences by tapping on a tablet and adults used a keyboard to do so, we cannot directly compare the reaction times of the two age groups. Instead, we will analyze the two age groups separately. For each age

group, reaction times below 100 ms were excluded, to remove extremely short reaction times (children: 0% of the data; adults: 0.3% of the data, 36 data points), as were reaction times outside the range of 2 SD from the participant's mean (children: 4.4% of the data, 448 data points; adults: 4.2% of the data, 556 data points). For children, the mean reaction time decreases from 690 ms in generation one (sd by chain = 80 ms) to 621 ms in generation ten (sd by chain = 90 ms). For adults, the mean reaction time decreases from 589 ms in generation one (sd by chain = 262 ms) to 466 ms in generation ten (sd by chain = 123 ms).

We fitted a model with log reaction time as the outcome variable for each age group. The model included an independent variable for transitional probability, control variables for generation and position in the sequence (mean-centered), and by-chain random slopes for generation. Transitional probability significantly predicts reaction time in both age groups, with faster reaction times for more predictable transitions, as expected if participants' online behaviour reflects the statistical structure of the set (children: $\beta = -0.14$, $SE = 0.02$, $t = -9.63$, $p < .001$; adults: $\beta = -0.12$, $SE = 0.02$, $t = -5.38$, $p < .001$).¹⁴ There was also a significant effect of position (children: $\beta = -0.01$, $SE = 0.0$, $t = -5.62$, $p < .001$; adults: $\beta = -0.02$, $SE = 0.00$, $t = -12.10$, $p < .001$), with reaction time decreasing towards the end of sequences. Reaction time did not decrease significantly with generation (children: $\beta = -0.01$, $SE = 0.01$, $t = -1.60$, $p = .171$; adults: $\beta = -0.00$, $SE = 0.02$, $t = -0.22$, $p = .833$). These findings indicate that both children and adults are sensitive to the transitional probabilities that were used to infer the statistical structure within the sequence sets.

2.2.9. Sets with lower entropy are transmitted with fewer errors

To ask whether sets with more statistical structure were easier to reproduce, we look at the correlation between the entropy and mean error of the sequence sets. We find a clear inverse relationship between the entropy and mean accuracy in the sets produced by children ($r(58) = 0.46$, $CI[0.24, 0.64]$, $p < .001$) and those produced by adults ($r(58) = 0.31$, $CI[0.07, 0.53]$, $p < .05$) (see Fig. 10), suggesting that the increase in learnability over transmission is related to the increase in distribution skew over transmission.

We conducted additional analyses to ensure that the increase in learnability is related to skew, and cannot be fully explained by a general reduction in complexity resulting from cultural transmission (e.g., Kempe et al., 2019; Tamariz and Kirby, 2015). These analyses are described in detail in Supplementary Materials B. We estimated Kolmogorov complexity at both the sequence level (using the *acss* package; Gauvrit et al., 2016) and the set level (using a standard compression algorithm). At the sequence level, complexity decreased over generations, and more so in child chains, suggesting age-related differences in how individual sequences change over time. However, analyses at the set level confirmed that while compression decreases over generations, it does not do so more strongly in child chains, and model comparisons indicate that Zipfian fit better predicts transmission error than compression-based estimates. Moreover, the emergence of the statistical properties is not driven only by changes in sequence-level complexity. Together, these findings suggest that the statistical properties we measure capture aspects of system-level structure that are relevant to learners beyond what compression-based estimates of complexity reflect.

3. Study 2: Are sequence sets transmitted by children easier to learn?

In the previous study, we found that sets transmitted by children




Fig. 9. Lexicon similarity increases over generations, showing that lexicons become easier to reproduce over generations in both age groups. The lines with error bars represent the mean similarity of units in the lexicon across chains within each age group (descriptive statistics); error-bars represent the 95% confidence intervals across units. The lines without error bars represent the mean unit similarity of individual chains.

¹³ We implemented the algorithm using the 'linear_sum_assignment' function in the Scipy Optimize library.

¹⁴ In a separate model we added an interaction between generation and transitional probability to see whether transitional probabilities are more predictive of reaction times in later generations. The model revealed no significant interaction.




Fig. 10. The correlation between entropy and mean transmission error of sets from generation zero to nine for each age groups. The figure shows sets with lower entropy are easier to reproduce.

evolved to have units with higher statistical coherence and a more skewed frequency distribution compared to the sets transmitted by adults. Here, we further investigate the relation between these statistical properties and learnability by asking whether the sequence sets generated by children are more learnable than those generated by adults for both children and adults. To do this, we asked a different sample of children and adults to observe and reproduce a sequence sets taken from the final generations of a child or adult chains.

3.1. Methods

3.1.1. Participants

A total of 49 children and 45 adults participated in the experiment. We used the same recruitment procedure as for the transmission study. (see [Section 2.1.1](#)). During data collection, five adults and nine children were excluded due to having more than nine invalid trials during testing. No participants were excluded based on the other exclusion criteria.

3.1.2. Stimuli

Participants were tested on the sequence sets produced in the last generation of the child and adult chains in the transmission experiment. We had six possible sets to select from for each age group (since there were six transmission chains for each age group). For each age group, we

selected the two sets that were closest to the mean entropy of the six sets (mean child-sets = 5.02, mean adult-sets = 5.28). This was done to ensure that the selected sets were representative of the statistical structure of the chains at the end of transmission. The four sequence sets had similar sequence lengths, with a mean sequence length of eight in both conditions (sd child-sets = 1.6, sd adult-sets = 1.7). We tested 20 children and 20 adults on each sequence set, resulting in 40 children and 40 adults in the two experimental conditions: child-generated sets or adult-generated sets.

3.1.3. Procedure

The procedure for this experiment was identical to that of a single iteration in the transmission experiment, including the procedural differences between the adult and child version of the experiment (see [Section 2.1.3](#)).

3.2. Results

[Fig. 11](#) shows that both children and adults reproduced the child-generated sets with lower transmission error than the adult-generated sets (mean adult error for adult-generated sets = 0.25, sd = 0.2; mean adult error for child-generated sets = 0.07, sd = 0.15; mean child error for adult-generated sets = 0.32, sd = 0.21; mean child error for child-




Fig. 11. Mean transmission error (descriptive statistics) is lower for the child-generated sets than the adult-generated sets for both age groups, indicating the child-generated sets are more learnable. The error-bars represent the 95% confidence intervals across participants. The transparent dots represent the means of individual participants.

generated sets = 0.18, $sd = 0.2$). That is, both children and adults found the child-generated sets easier to reproduce. We fitted a model with error on trial as the outcome variable, condition (child-generated set/adult-generated set), age group (adults/children) and their interaction as independent variables, sequence length (numeric, 6–14) and trial number (numeric, 1–30) as control variables, and a random intercept for participants.¹⁵ Condition and age group were effect-coded (sum-to-zero), so that the model intercept represented the grand mean across conditions age groups. Overall, children made more transmission errors than adults ($\beta = 0.09$, $SE = 0.02$, $t = 5.35$, $p < .001$), in line with the idea that they have a narrower bottleneck. In addition, the child-generated sets were reproduced with fewer errors than the adult-generated sets ($\beta = -0.15$, $SE = 0.02$, $t = -9.48$, $p < .001$), a pattern present in both children ($\beta = -0.14$, $SE = 0.02$, $t = -5.9$, $p < .001$) and adults ($\beta = -0.17$, $SE = 0.02$, $t = -7.5$, $p < .001$). These findings suggest that transmission by children introduces a greater learnability pressure compared to transmission by adults, which in turn increases the learnability of the resulting sequence sets for both child and adult learners.

4. Discussion

Language learning is constrained by the capacity of our cognition, creating a bottleneck when a language is transmitted from one generation to the next. Prior work has argued that systematic structure emerges as an adaptation to this learning bottleneck by enabling learners to generalize from limited exposure to the language (Brighton, 2002; Ferdinand & Zuidema, 2009; Kirby, 2000, 2001; Smith et al., 2003). Specifically, more narrow bottlenecks are expected to lead to a greater increase in structure compared to wider bottlenecks. Prior work has primarily tested this idea by manipulating the amount of input adult learners receive during learning (Ferdinand & Zuidema, 2009; Kirby et al., 2007; Kirby et al., 2015; Kirby & Hurford, 2002). However, natural language is also shaped by bottlenecks that arise from constraints on our cognitive capacities (Christiansen & Chater, 2016; Kirby and Tamariz, 2022). In this study, we set out to experimentally test the effects of such learning bottlenecks on the emergence of two statistical properties of language: statistically coherent units and Zipfian frequency distributions. We did this by comparing the rate of their emergence in two populations, children and adults, that differ in their cognitive capacities.

We conducted an iterated learning experiment in which chains of child or adult learners observed and reproduced sets of colour sequences. Consistent with prior cultural transmission studies (Arnon & Kirby, 2024; Carr et al., 2017; Kempe et al., 2015; Kirby et al., 2008; Wolters, Kirby, & Arnon, 2025), transmission error decreased over generations in both age groups, indicating that the sequences became more learnable over time. Error rates were significantly lower for children compared to adults in the first generation of the transmission chains, confirming their more limited cognitive capacities and the resulting narrow bottleneck. However, by the end of the transmission, children's error rates had converged with those of adults. How, then, did the sequence sets change such that children, despite their more limited cognitive capacities, achieved adult-level learnability over the course of cultural transmission?

To address this question, we examined the sequence sets for the emergence of two statistical properties that have been shown to facilitate learning (Kurumada et al., 2013; Lavi-Rotbain & Arnon, 2022; Saffran & Kirkham, 2018; Wolters, Ota, & Arnon, 2025). We segmented each sequence set into units based on the transitional probabilities between colours, using dips in these probabilities as indicators of unit boundaries (Arnon & Kirby, 2024). Across both age groups, the initially

random sequence sets evolved to have statistically coherent units with frequency distributions that were best fit by a Zipfian power law. Moreover, we found a significant correlation between distribution entropy and transmission error, suggesting that the emergence of these statistical properties, particularly the skew of the frequency distributions, increased the learnability of the sets for both children and adults. These findings suggest that the same hallmark statistical properties of language, previously only shown to emerge through cultural transmission with adults (Arnon & Kirby, 2024; Wolters, Kirby, & Arnon, 2025), can also reliably emerge under the narrow learning bottlenecks characteristic of child learners, highlighting the potential of these statistical properties as robust learnable patterns in the evolution of linguistic structure.

Crucially, while both children's and adults' sequence sets evolved learnable statistical properties, they evolved faster in transmission chains with children: statistically coherent units, and the corresponding saliency of unit boundaries, emerged at a faster rate in the transmission chains with children. Similarly, the frequency distribution of segmented units became more skewed in the children's chains compared to the adult's chains, as indicated by a larger reduction in entropy and a better Zipfian fit in the children's sets. Looking at the change over generations in more detail reveals an interesting pattern: children introduced high levels of statistical structure in early generations, after which the rate of evolution slowed down, whereas adults showed a more gradual, continuous accumulation of structure over time. This early structuring in children's chains suggests that the more narrow bottleneck creates a stronger pressure towards learnability from the outset. To ensure that the faster evolution rate in the children's chains were not simply driven by differences in the basic properties of the sets produced by children and adults, we compared the sequence sets to a shuffled baseline in which sequence length and the unigram distribution of individual colours were preserved, but all sequential structure was destroyed. This comparison confirmed that the accelerated emergence of statistical structure in the children's chains was driven by changes in the sequential ordering of the colours. Taken together, these findings show that the statistical properties emerge faster in children's chains compared to adults chains, potentially explaining why, despite the initial differences in error rates, the sets generated by children and adults are learned with similar error rates by the end of transmission. In addition, we also find that both children and adults developed increasingly learnable lexicons of units over generations. Lexicons became increasingly similar between successive generations over time, and this effect was stronger with children. This suggests that the structural changes introduced during transmission not only reduced reproduction errors of the sequences, but also resulted in lexicons of units that are easier to reproduce.

To test whether the increased emergence of the statistical properties in the child chains enhances learnability across age groups, we conducted a second experiment in which both adult and child learners were tested on the final sequence sets of the child and adult chains. Both age groups reproduced the child-generated sets with significantly lower transmission error, indicating that the structural changes introduced by children improved learnability across age groups, and reconfirming that sets with increased statistical structure, like those generated by children, are more learnable. These findings suggest that children and adults share similar biases in learning statistical structure, in line with prior research showing a facilitative effect of Zipfian distributions on learning in both children and adults (Lavi-Rotbain & Arnon, 2022).

Similar learning outcomes (in terms of the statistical properties we examined) do not necessarily imply identical learning mechanisms. It is possible that children, in addition to having more limited processing capacities, also have a stronger tendency to regularize statistical patterns (Perfors, 2012, 2016). Relatedly, children may be better implicit learners than adults (under certain conditions), potentially leading to their transmission errors being more systematically shaped by statistical regularities (Janacsek et al., 2012; Juhasz et al., 2019). More generally, our operationalization of the learning bottleneck collapses across

¹⁵ We initially ran a model including a random intercept for the sequence set but this model did not converge due to low variance in the random effect structure.

multiple cognitive capacities (memory, attention, implicit learning), each of which may contribute differently to the observed bottleneck effect (Ferdinand & Zuidema, 2009; MacDonald, 1999; MacDonald, 2013; Smalle et al., 2021). It therefore remains an open question to what extent the differences observed between child and adult chains reflect differences in memory, attention, production biases, or their interaction. In addition, the use of visual stimuli may also interact with these mechanisms, as several studies have revealed modality differences in statistical learning (Conway & Christiansen, 2005; Conway & Christiansen, 2006; Frost et al., 2015). Given that language is primarily processed in the auditory modality, future work should test whether our findings reflect general learning dynamics rather than modality-specific effects. A further question for future work, relevant not only to the present study but also to other iterated sequence learning studies that use non-linguistic stimuli (Arnon & Kirby, 2024; Wolters, Kirby, & Arnon, 2025), is how pressures from sequence learning express themselves when signals are used to express meaning. Although learning the sequential structure of language, including the task of word segmentation, constitutes to some extent an independent learning problem, naturalistic language acquisition involves the simultaneous acquisition of multiple aspects of language, including the learning of form–meaning mappings. An interesting direction for future research, therefore, is to understand how pressures from sequence learning interact with pressures introduced by the learning of form–meaning mappings, and how this in turn shapes the statistical properties of language.

Although numerous iterated learning experiments have been conducted with adults, this is only the second study to show that cultural transmission by children leads to the emergence of more systematic structure and increased learnability, and the first to demonstrate that transmission by children can drive the emergence of cross-linguistic statistical properties. The only other study reporting similar findings is Kempe et al. (2015), in which children and adults reproduced visuo-spatial patterns (dots on a grid). Consistent with our results, they found that structure emerged more readily in children's chains, with the trajectories of change similar to those in our study: children introduced high levels of structure early, while adults exhibited a more gradual, incremental emergence over time. These findings diverge from cultural transmission studies that used an artificial language learning paradigm, where structure did not emerge with children. This divergence may be explained by the different experimental tasks: whereas the present study and Kempe et al. (2015) used non-linguistic stimuli without meaning (sequential in our case, spatio-visual in theirs), Raviv and Arnon (2018) and Kempe et al. (2019) used artificial languages, involving the learning of compositional structure, which may pose higher cognitive and pragmatic demands (Kempe et al., 2019; Naigles, 2002).

The mixed findings with regard to age-effects in iterated learning studies relate to a larger debate on the role of children in driving language emergence and change. On the one hand, children have been argued to play a significant role in the emergence and systematization of linguistic structure in several settings. For example, in the formation of grammatical structures in the process of creolization (Bickerton, 1984),¹⁶ the emergence of compositional and linearized expressions in sign languages (Senghas et al., 2004; Senghas & Coppola, 2001; Vogt, 2005), and the spontaneous emergence of grammatical structure in the home sign systems of deaf children (Goldin-Meadow, 2003). Moreover, it has been argued that in the presence of unpredictable variation and a lack of normative mechanisms (e.g. in places of intense language contact or a mixture of L1 and L2 speakers) children's learning biases will systematize linguistic structure (Dale & Lupyan, 2012; Kirby & Tamariz, 2022; Sanz-Sánchez & Moyna, 2023). On the other hand, children's contribution to language change has been challenged (see

Raviv et al., 2025 for an in-depth argument). Children do not innovate novel linguistic forms until later in childhood, and agents of innovation are generally considered to have some social influence in society, which young children typically lack (Kerswill & Williams, 2000; Labov, 2007; Nettle, 1999; Raviv et al., 2025). More generally, there is little empirical evidence of child-generated innovations that have spread through language communities. These apparent discrepancy between these findings may be reconciled by two observations: first, that the effects of children's systemization will be most pronounced in cases of language emergence where children are the agents of change, for example, in the emergence of deaf community sign language where there is peer-to-peer transmission or in home-sign systems. In such scenarios, children's narrow bottleneck may accelerate the emergence of systematic structure compared to scenarios of language change where children learn from more experienced individuals (Kirby & Tamariz, 2022). And second, that when children have a role in language emergence or change, it will mostly concern the systematization of language, including linear sequencing, segmentation, or regularization. This means that while such processes may introduce novel linguistic forms through, for example, overregularizations, children likely play a smaller role in the introduction and diffusion of entirely novel grammatical constructions (Raviv and Arnon, 2018). Our findings align with these observations by showing that children systematize statistical structure, and that the transmission by children alone accelerates the emergence of systemic statistical properties in comparison to transmission by adults.

In conclusion, we show here that two core statistical properties of language—statistically coherent units and Zipfian frequency distributions—emerge reliably over cultural transmission with adult and child learners, and that the rate of this emergence is faster under the more limited cognitive capacities characteristic of child learners. With that, our findings extend previous work examining the effects of transmission bottlenecks by exposing learners to only a subset of a language during training. Specifically, they show that transmission bottlenecks arising from cognitive limitations similarly modulate the emergence of structural properties that facilitate learning.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the author(s) used ChatGPT in order to improve language and readability. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

CRediT authorship contribution statement

Lucie Wolters: Writing – original draft, Methodology, Formal analysis, Conceptualization. **Simon Kirby:** Writing – review & editing, Supervision, Conceptualization. **Inbal Arnon:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare no conflicts of interests.

Acknowledgements

We wish to thank the Living Lab staff for their help with data collection, with special thanks to Omer Benacot, Tmami Abo Mokh, Ariel Yaffe, Tamar Matas, and Lihi Troyansky. We also want to thank the Bloomfield Science Museum in Jerusalem, as well as the parents and children who participated. This research was funded by a research grant from the Israel Science Foundation awarded to the third author.

¹⁶ Although note that empirical data from creole languages suggests that complex grammatical structures emerged long before children were acquiring it as their first language (Arends, 1996; Sankoff & Labov, 1979).

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.cognition.2026.106622>.

Data availability

All data and analysis code reported in this article are publicly available on the Open Science Framework (OSF) project page: <https://osf.io/ujfvn>.

References

- Arends, J. (1996). *The early stages of creolization* (pp. 1–313).
- Arnon, I., & Kirby, S. (2024). Cultural evolution creates the statistical structure of language. *Scientific Reports*, 14(1), 5255.
- Arnon, I., Kirby, S., Allen, J. A., Garrigue, C., Carroll, E. L., & Garland, E. C. (2025). Whale song shows language-like statistical structure. *Science*, 387(6734), 649–653. <https://doi.org/10.1126/science.adq7055>
- Barr, D. J. (2013). Random effects structure for testing interactions in linear mixed-effects models. *Frontiers in Psychology*, 4, 328.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). *Fitting Linear Mixed-Effects Models using lme4* (arXiv:1406.5823). arXiv. <https://doi.org/10.48550/arXiv.1406.5823>.
- Bickerton, D. (1984). The language bioprogram hypothesis. *The Behavioral and Brain Sciences*, 7(2), 173–188. <https://doi.org/10.1017/S0140525X00044149>
- 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. <https://doi.org/10.1073/pnas.1904871116>
- Brighton, H. (2002). Compositional syntax from cultural transmission. *Art&Life*, 8(1), 25–54. <https://doi.org/10.1162/106454602753694756>
- 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. <https://doi.org/10.1111/cogs.12371>
- Christiansen, M. H., & Chater, N. (2016). The now-or-never bottleneck: A fundamental constraint on language. *The Behavioral and Brain Sciences*, 39, Article e62. <https://doi.org/10.1017/S0140525X1500031X>
- Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review*, 51(4), 661–703. <https://doi.org/10.1137/070710111>
- Clay, Z., Pople, S., Hood, B., & Kita, S. (2014). Young children make their gestural communication systems more language-like: Segmentation and linearization of semantic elements in motion events. *Psychological Science*, 25(8), 1518–1525. <https://doi.org/10.1177/0956797614533967>
- Conway, C. M., & Christiansen, M. H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(1), 24.
- Conway, C. M., & Christiansen, M. H. (2006). Statistical learning within and between modalities: Pitting abstract against stimulus-specific representations. *Psychological Science*, 17(10), 905–912.
- 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)*.
- 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), Article e0168532. <https://doi.org/10.1371/journal.pone.0168532>
- Craik, F. I., & Bialystok, E. (2006). Cognition through the lifespan: Mechanisms of change. *Trends in Cognitive Sciences*, 10(3), 131–138. <https://doi.org/10.1016/j.tics.2006.01.007>
- Dale, R., & Luppian, G. (2012). Understanding the origins of morphological diversity: The linguistic niche hypothesis. *Advances in Complex Systems*, 15(03n04), 1150017. <https://doi.org/10.1142/S0219525911500172>
- Ferdinand, V., & Zuidema, W. (2009). Thomas' theorem meets Bayes' rule: A model of the iterated learning of language. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 31, 1069–7977.
- Ferreri Cancho, R. (2005). The variation of Zipf's law in human language. *The European Physical Journal B - Condensed Matter and Complex Systems*, 44, 249–257. <https://doi.org/10.1140/epjb/e2005-00121-8>
- 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.
- Gauvrit, N., Singmann, H., Soler-Toscano, F., & Zenil, H. (2016). Algorithmic complexity for psychology: A user-friendly implementation of the coding theorem method. *Behavior Research Methods*, 48(1), 314–329.
- Goldin-Meadow, S. (2003). *The resilience of language: What gesture creation in deaf children can tell us about how all children learn language*. Psychology Press.
- Hudson Kam, C. L., & Chang, A. (2009). Investigating the cause of language regularization in adults: Memory constraints or learning effects? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(3), 815–821. <https://doi.org/10.1037/a0015097>
- 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. <https://doi.org/10.1080/15475441.2005.9684215>
- Hudson Kam, C. L., & Newport, E. L. (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology*, 59(1), 30–66. <https://doi.org/10.1016/j.cogpsych.2009.01.001>
- Isbilen, E. S., & Christiansen, M. H. (2020). Chunk-based memory constraints on the cultural evolution of language. *Topics in Cognitive Science*, 12(2), 713–726. <https://doi.org/10.1111/tops.12376>
- Janacek, K., Fiser, J., & Nemeth, D. (2012). The best time to acquire new skills: Age-related differences in implicit sequence learning across human life span. *Developmental Science*, 15(4), 496–505. <https://doi.org/10.1111/j.1467-7687.2012.01150.x>
- 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), Article e12877.
- Juhász, D., Nemeth, D., & Janacek, K. (2019). Is there more room to improve? The lifespan trajectory of procedural learning and its relationship to the between- and within-group differences in average response times. *PLoS One*, 14(7), Article e0215116.
- Kaplan, B. E., & Yu, C. (2025). Near-Zipfian distribution is prevalent in infant input. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 47. <https://escholarship.org/uc/item/1q63787s>
- Kempe, V., Gauvrit, N., & Forsyth, D. (2015). Structure emerges faster during cultural transmission in children than in adults. *Cognition*, 136, 247–254. <https://doi.org/10.1016/j.cognition.2014.11.038>
- 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. <https://doi.org/10.1093/jole/lzy012>
- Kerswill, P., & Williams, A. (2000). Creating a new town koine: Children and language change in Milton Keynes. *Language et Societe*, 29(1), 65–115. <https://doi.org/10.1017/S0047404500001020>
- Kimchi, I., Wolters, L., Stamp, R., & Arnon, I. (2023). Evidence of Zipfian distributions in three sign languages. *Gesture*, 22(2), 154–188. <https://doi.org/10.1075/gest.23014.kim>
- Kirby, S. (2000). Syntax without natural selection: How compositionality emerges from vocabulary in a population of learners. *In The evolutionary emergence of language: Social function and the origins of linguistic form* (pp. 303–323).
- Kirby, S. (2001). Spontaneous evolution of linguistic structure-an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2), 102–110. <https://doi.org/10.1109/4235.918430>
- 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. <https://doi.org/10.1073/pnas.0707835105>
- Kirby, S., Dowman, M., & Griffiths, T. L. (2007). Innateness and culture in the evolution of language. *Proceedings of the National Academy of Sciences*, 104(12), 5241–5245.
- Kirby, S., & Hurford, J. (2002). The emergence of linguistic structure: An overview of the iterated learning model. *In A. Cangelosi, & D. Parisi (Eds.), Simulating the evolution of language*. Springer.
- Kirby, S., & Tamariz, M. (2022). Cumulative cultural evolution, population structure and the origin of combinatoriality in human language. *Philosophical Transactions of the Royal Society, B: Biological Sciences*, 377(1843).
- Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141, 87–102.
- Kurumada, C., Meylan, S. C., & Frank, M. C. (2013). Zipfian frequency distributions facilitate word segmentation in context. *Cognition*, 127(3), 439–453. <https://doi.org/10.1016/j.cognition.2013.02.002>
- Labov, W. (2007). Transmission and diffusion. *Language*, 83(2), 344–387.
- Lavi-Rotbain, O., & Arnon, I. (2022). The learnability consequences of Zipfian distributions in language. *Cognition*, 223. <https://doi.org/10.1016/j.cognition.2022.105038>
- Lavi-Rotbain, O., & Arnon, I. (2023). Zipfian distributions in child-directed speech. *Open Mind*, 7, 1–30.
- MacDonald, M. C. (1999). Distributional information in language comprehension, production, and acquisition: Three puzzles and a moral. *In B. MacWhinney (Ed.), The emergence of language* (pp. 177–196). Mahwah, NJ: Lawrence Erlbaum Associates Publishers.
- MacDonald, M. C. (2013). How language production shapes language form and comprehension. *Frontiers in Psychology*, 4, 226.
- Matthews, D., Lieven, E., & Tomasello, M. (2007). How toddlers and preschoolers learn to uniquely identify referents for others: A training study. *Child Development*, 78(6), 1744–1759.
- Mehri, A., & Jamaati, M. (2017). Variation of Zipf's exponent in one hundred live languages: A study of the holy bible translations. *Physics Letters A*, 381(31), 2470–2477. <https://doi.org/10.1016/j.physleta.2017.05.061>
- Montag, J. L. (2021). Limited evidence for probability matching as a strategy in probability learning tasks. *In , 74. Psychology of learning and motivation* (pp. 233–273). Academic Press.
- Naigles, L. R. (2002). Form is easy, meaning is hard: Resolving a paradox in early child language. *Cognition*, 86(2), 157–199. [https://doi.org/10.1016/S0010-0277\(02\)00177-4](https://doi.org/10.1016/S0010-0277(02)00177-4)
- Nettle, D. (1999). Using social impact theory to simulate language change. *Lingua*, 108(2), 95–117. [https://doi.org/10.1016/S0024-3841\(98\)00046-1](https://doi.org/10.1016/S0024-3841(98)00046-1)
- Newport, E. L. (2020). Children and adults as language learners: Rules, variation, and maturational change. *Topics in Cognitive Science*, 12(1), 153–169. <https://doi.org/10.1111/tops.12416>
- Nilsen, E. S., & Graham, S. A. (2009). The relations between children's communicative perspective-taking and executive functioning. *Cognitive Psychology*, 58(2), 220–249.

- Perfors, A. (2012). When do memory limitations lead to regularization? An experimental and computational investigation. *Journal of Memory and Language*, 67(4), 486–506. <https://doi.org/10.1016/j.jml.2012.07.009>
- Perfors, A. (2016). Adult regularization of inconsistent input depends on pragmatic factors. *Language Learning and Development*. <https://doi.org/10.1080/15475441.2015.1052449>
- Piantadosi, S. T. (2014). Zipf's word frequency law in natural language: A critical review and future directions. *Psychonomic Bulletin & Review*, 21(5), 1112–1130. <https://doi.org/10.3758/s13423-014-0585-6>
- R Core Team. (2021). *R: A language and environment for statistical computing [computer software]*. R Foundation for Statistical Computing. <https://www.R-project.org/>.
- Raviv, L., & Arnon, I. (2018). Systematicity, but not compositionality: Examining the emergence of linguistic structure in children and adults using iterated learning. *Cognition*, 181, 160–173. <https://doi.org/10.1016/j.cognition.2018.08.011>
- Raviv, L., Blasi, D., & Kempe, V. (2025). Children are not the main agents of language change. *Psychological Review*. <https://doi.org/10.1037/rev0000580>
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928. <https://doi.org/10.1126/science.274.5294.1926>
- Saffran, J. R., & Kirkham, N. Z. (2018). Infant statistical learning. *Annual Review of Psychology*, 69(1), 181–203.
- Saksida, A., Langus, A., & Nespore, M. (2017). Co-occurrence statistics as a language-dependent cue for speech segmentation. *Developmental Science*, 20(3), Article e12390. <https://doi.org/10.1111/desc.12390>
- Samara, A., Smith, K., Brown, H., & Wonnacott, E. (2017). Acquiring variation in an artificial language: Children and adults are sensitive to socially conditioned linguistic variation. *Cognitive Psychology*, 94, 85–114. <https://doi.org/10.1016/j.cogpsych.2017.02.004>
- Sankoff, D., & Labov, W. (1979). On the uses of variable rules. *Language et Societe*, 8(2–3), 189–222. <https://doi.org/10.1017/S0047404500007430>
- Sanz-Sánchez, I., & Moyna, M. I. (2023). Children as agents of language change: Diachronic evidence from Latin American Spanish phonology. *Journal of Historical Linguistics*, 13(3), 327–374. <https://doi.org/10.1075/jhl.21033.san>
- Senghas, A., & Coppola, M. (2001). Children creating language: How Nicaraguan sign language acquired a spatial grammar. *Psychological Science*, 12(4), 323–328. <https://doi.org/10.1111/1467-9280.00359>
- Senghas, A., Kita, S., & Özyürek, A. (2004). Children creating Core properties of language: Evidence from an emerging sign language in Nicaragua. *Science*, 305(5691), 1779–1782. <https://doi.org/10.1126/science.1100199>
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
- Shufaniya, A., & Arnon, I. (2021). A cognitive bias for Zipfian distributions? Uniform distributions become more skewed via cultural transmission. In *Proceedings of the annual meeting of the 43th cognitive science society*. <https://escholarship.org/uc/item/66c40451>.
- Singleton, J. L., & Newport, E. L. (2004). When learners surpass their models: The acquisition of American sign language from inconsistent input. *Cognitive Psychology*, 49(4), 370–407. <https://doi.org/10.1016/j.cogpsych.2004.05.001>
- Smalle, E. H., Muijle, M., Duyck, W., & Szmalec, A. (2021). Less is more: Depleting cognitive resources enhances language learning abilities in adults. *Journal of Experimental Psychology. General*, 150(12), 2423.
- Smith, K., Brighton, H., & Kirby, S. (2003). Complex systems in language evolution: The cultural emergence of compositional structure. *Advances in Complex Systems*, 06(04), 537–558. <https://doi.org/10.1142/S0219525903001055>
- Stärk, K., Kidd, E., & Frost, R. L. A. (2022). Word segmentation cues in German child-directed speech: A Corpus analysis. *Language and Speech*, 65(1), 3–27. <https://doi.org/10.1177/0023830920979016>
- Tamariz, M., & Kirby, S. (2015). Culture: Copying, compression, and conventionality. *Cognitive Science*, 39(1), 171–183. <https://doi.org/10.1111/cogs.12144>
- Vogt, P. (2005). On the acquisition and evolution of compositional languages: Sparse input and the productive creativity of children. *Adaptive Behavior*, 13(4), 325–346. <https://doi.org/10.1177/105971230501300403>
- Wolters, L., Kirby, S., & Arnon, I. (2025). Cultural transmission promotes the emergence of statistical properties that support language learning. *Cognitive Science*, 49(12), Article e70153.
- Wolters, L., Ota, M., & Arnon, I. (2025). Skewed distributions facilitate infants' word segmentation. *Cognition*, 263, Article 106221. <https://doi.org/10.1016/j.cognition.2025.106221>
- Zipf, G. K. (1949). *Human behavior and the principle of least effort: An introduction to human ecology*. Addison-Wesley Press.