Cognition 181 (2018) 160–173

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

Systematicity, but not compositionality: Examining the emergence of linguistic structure in children and adults using iterated learning

Limor Raviv[a,c,*], Inbal Arnon[b]

[a] Psychology of Language Department, Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands
[b] Department of Psychology, Hebrew University, Jerusalem, Israel
[c] Department of Cognitive Science, Hebrew University, Jerusalem, Israel

ARTICLE INFO

Keywords:
Iterated learning
Cultural transmission
Linguistic structure
Age-related differences
Language learning

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. The lack of evidence from child learners constitutes a problematic gap in the literature: if such learning biases impact the emergence of linguistic structure, they should also be found in children, who are the primary learners in real-life language transmission. However, children may differ from adults in their biases given age-related differences in general cognitive skills. Moreover, adults' performance on iterated learning tasks may reflect existing (and explicit) linguistic biases, partially undermining the generality of the results. Examining children's performance can also help evaluate contrasting predictions about their role in emerging languages: do children play a larger or smaller role than adults in the creation of structure? Here, we report a series of four iterated artificial language learning studies (based on Kirby, Cornish & Smith, 2008) with both children and adults, using a novel child-friendly paradigm. Our results show that linguistic structure does not emerge more readily in children compared to adults, and that adults are overall better in both language learning and in creating linguistic structure. When languages could become underspecified (by allowing homonyms), children and adults were similar in developing consistent mappings between meanings and signals in the form of structured ambiguities. However, when homonimity was not allowed, only adults created compositional structure. This study is a first step in using iterated language learning paradigms to explore child-adult differences. It provides the first demonstration that cultural transmission has a different effect on the languages produced by children and adults: While children were able to develop systematicity, their languages did not show compositionality. We focus on the relation between learning and structure creation as a possible explanation for our findings and discuss implications for children's role in the emergence of linguistic structure.

1. Introduction

How does linguistic structure emerge? Under the classic nativist approach, originally formulated in Chomsky (1965), linguistic structure is driven by a set of abstract and language-specific principles, which are both universal and innate, and impact how languages are shaped. An alternative explanation is offered by usage-based theories, suggesting that the kinds of structures we observe in human languages arise from general biases and constraints on individuals' cognitive capacities, such as learning, memory and processing (Tomasello, 2009). Under this view, languages are shaped through the process of *cultural transmission,* where weak individual tendencies become amplified and fixated over time through a repeated cycle of use, observation, and induction (Kirby, Griffiths & Smith, 2014). This prediction is supported by findings from iterated learning paradigms, which show how the iterative nature of cultural transmission can lead to the creation of linguistic structure over multiple generations without the need to assume strong or language-specific innate biases (Culbertson & Kirby, 2015; Kirby, Cornish & Smith, 2008; Kirby, Smith & Brighton, 2004).

Iterated learning studies simulate the process of cultural transmission by using a diffusion chain paradigm, in which agents (computational or human) are exposed to a target behavior that they need to reproduce. Crucially, the behavior produced by the first agent in the chain becomes the input behavior for the second agent, the behavior of the second agent becomes the input for the third agent, and so on for several "generations" of agents. Mathematical and computational

* Corresponding author at: Max Plank Institute for Psycholinguistics, P.O. Box 310, 6500 AH Nijmegen, The Netherlands.
E-mail address: limor.raviv@mpi.nl (L. Raviv).

https://doi.org/10.1016/j.cognition.2018.08.011
Received 1 September 2017; Received in revised form 12 August 2018; Accepted 17 August 2018

models of iterated learning show that the structural properties of artificial languages can be shaped over time to better-fit agents' existing tendencies and predispositions (e.g., Griffiths & Kalish, 2007; Kirby, Dowman & Griffiths, 2007). For example, agents' weak bias against synonymy was amplified over repeated episodes of learning so that an initial lexicon with multiple labels for each item developed unique object labels (Reali & Griffiths, 2009).

In a seminal study, Kirby et al. (2008) used such a paradigm with adult participants to show that linguistic structure can emerge over time in an artificial language. Participants were exposed to items that varied across three semantic dimensions (shape, color and motion type), and needed to learn and reproduce novel labels describing these items. The first participant was trained on an artificial language without structure (with random mappings between forms and meanings), and their written output was given as input to the next participant in the chain. The resulting languages were evaluated on their learnability (i.e. how faithfully they were transmitted) and on their structure (how systematic they were). Over ten generations of participants, the languages became easier to learn and developed consistent mappings between meanings and signals. In the first study, languages were transmitted without intervention (allowing homonyms). These languages developed systematicity in the form of structured ambiguities, with small and underspecified lexicons in which items sharing a semantic feature were referred to using the same label. For example, all spiraling items were referred to as “poi”, regardless of their shape or color. Although this study resulted in systematic languages, encoding multiple semantic dimensions using holistic labels meant that the artificial languages lost much of their informativity, and differed dramatically from natural languages in their expressivity. In the second study, homonyms were filtered out during transmission to impose an artificial expressivity pressure and prevent underspecification. The result was that languages developed compositional structure – one of the hallmarks of natural languages: sub-strings were systematically reused to express different features. For example, color was marked with a prefix (e.g., “l-” for blue items vs. “n-” for grey items) and motion was marked with a suffix (e.g, “-plo” for bouncing items vs. “-pilu” for spiraling items). Similar increases in compositional structure and in learnability have since been replicated (Beckner, Pierrehumbert & Hay, 2017), and found for a range of linguistic and non-linguistic behaviors (e.g., drawings, whistles, gestures, visual patterns, for review see Tamariz & Kirby, 2016).

The accumulated findings support the prediction that linguistic structure can emerge through cultural transmission. However, they are limited in one interesting way: they are based only on adult learners. To date, only one study has used iterated learning to compare children and adults on a non-linguistic task (Kempe, Gauvrit & Forsyth, 2015; discussed below), and no published study has examined the emergence of linguistic structure over generations of child learners. The lack of evidence from child learners is problematic for several reasons. First, it limits the conclusions that can be drawn about the role of cultural transmission in the emergence of linguistic structure. The process of transmission is complex, and involves several different components (learning, production, and finally transmission to another agent), all of which may impact the resulting behavior. Learning biases in particular have been argued to shape the emergence of structure in the cultural transmission of language (e.g., Kirby et al., 2004; Smith, 2011). If the emergence of compositional structure over generations is influenced by learners’ biases, then similar effects should also be found in children, who are the primary and most prototypical learners of language in real-world transmission. Children’s performance is therefore a test case for the hypothesis that typical cross-generational learning can drive the emergence of linguistic structure. Second, adult participants may rely on their extensive and explicit linguistic knowledge when learning an unfamiliar language: adults may have a stronger prior bias in favor of linguistic structure, which can (consciously or not) influence their performance, causing structure to emerge more readily or more rapidly (Cornish, Tamariz & Kirby, 2009). This criticism is consistent with the wide-spread effects of transfer from individuals’ first language (L1) when learning a second language (e.g., White, 2000), as well as with the effect of L1 knowledge on artificial language learning – for instance, L1 phonotactics impact the segmentation of artificial languages (Finn & Hudson Kam, 2008; Siegelman, Bogaerts, Elazar, Arciuli & Frost, 2018). While strong biases are not *necessary* for structure to emerge (e.g., Kirby, 2001; Kirby et al., 2004, 2007; Smith, 2009), agents’ existing biases could still influence the nature and rapidness of this process. This idea receives some support from computational models: under certain circumstances, changes in bias strength can impact the speed with which transmission fidelity increases (e.g., Ferdinand & Zuidema, 2008b), and even the resulting structural patterns (e.g., Navarro, Perfors, Kary, Brown & Donkin, 2018; Brochhagen, Franke, & van Rooij, 2016; Morgan & Levy, 2016; Smith, 2011). If adults’ experience with their L1 results in a stronger bias for structure, their performance on iterated learning studies might reflect a cognitive bias that is partially the *result* of the evolution of language over time, rather than a bias *responsible* for it. This criticism can be partially overcome by looking at children, who have less explicit meta-linguistic knowledge and are more likely to learn implicitly (Arnon & Ramscar, 2012; Karmiloff-Smith, Grant, Sims, Jones & Cuckle, 1996; Ravid & Malenky, 2001; Ullman, 2001). Finally, examining children’s performance in iterated learning is important given the long-lasting debate about their postulated role in the formation and extension of linguistic structure.

Children’s role in the formation of grammatical structure has been heavily debated in the language emergence literature. On the one hand, children are claimed to play a special role in the formation of linguistic structure in creole languages. Bickerton’s influential Language Bioprogram hypothesis (1984) argues that children, and not adults, are responsible for the formation of grammar in the process of creolization, and that they regularize the language and add structure to it through their reliance on innate linguistic biases. Similar claims have been made in the sign language literature, where children are shown to introduce novel linguistic structures. Studies of deaf children born to hearing parents show that children introduced regularities, like word order, which were not found in their input (e.g., Goldin-Meadow & Mylander, 1998). Research on emerging sign languages further suggests that children have a unique role in making new languages more structured: younger learners (exposed to the developing Nicaraguan Sign Language before the age of ten) produce more structured languages compared to adult learners within the same cohort (Senghas & Coppola, 2001). More generally, children are claimed to add grammatical features (such as linear sequencing) to the language when learning it from a previous cohort (Senghas, Kita & Özyürek, 2004). Additionally, younger children have a stronger tendency to segment and linearize their gestures compared to adolescents and adults (Clay, Pople, Hood & Kita, 2014). Based on these studies, which argue that children create core properties of language, we may predict that children will show similar or even stronger biases for creating structure in linguistic iterated learning. This prediction is supported by the single iterated study that compared children to adults on the same non-linguistic task. Using a visual recall task, Kempe et al. (2015) found that children created more identifiable and less complex visual patterns in comparison to adults, and concluded that structure emerged more readily in child chains. It is also supported by findings showing that children are more likely to regularize compared to adult learners (Hudson Kam & Newport, 2005; 2009; Samara, Smith, Brown & Wonnacott, 2017)

On the other hand, there are reasons to think that structure will emerge less readily in child learners. Children’s postulated role in the process of creolization has been challenged by studies showing that it is a slow and multi-generational process (Arends, 1993; Carden & Stewart, 1988), and that complex grammatical structures emerged long before children were acquiring it as their first language (Arends & Bruyn, 1995; Sankoff & Laberge, 1974). These findings suggest that the main innovators in the process of creolization were adult speakers, and argue that children’s contribution to the process is, if any, in the selection and