Statistical language learning: computational, maturational, and linguistic constraints*

ELISSA L. NEWPORT

Georgetown University

ABSTRACT

Our research on statistical language learning shows that infants, young children, and adults can compute, online and with remarkable speed, how consistently sounds co-occur, how frequently words occur in similar contexts, and the like, and can utilize these statistics to find candidate words in a speech stream, discover grammatical categories, and acquire simple syntactic structure in miniature languages. However, statistical learning is not merely learning the patterns presented in the input. When their input is inconsistent, children sharpen these statistics and produce a more systematic language than the one to which they are exposed. When input languages inconsistently violate tendencies that are widespread in human languages, learners shift these languages to be more aligned with language universals, and children do so much more than adults. These processes explain why children acquire language (and other patterns) more effectively than adults, and also may explain how systematic language structures emerge in communities where usages are varied and inconsistent. Most especially, they suggest that usage-based learning approaches must account for differences between adults and children in how usage properties are acquired, and must also account for substantial changes made by adult and child learners in how input usage properties are represented during learning.

KEYWORDS: statistical learning, language universals, child–adult differences, Greenberg word order universals, inconsistent input.

---

1. Introduction

Beginning in research with Jenny Saffran (Saffran, Aslin, & Newport, 1996; Saffran, Newport, & Aslin, 1996), Richard Aslin and I have articulated an approach to first language learning that we have called STATISTICAL LEARNING – a term borrowed from computational linguistics (Charniak, 1993). The central hypothesis in our approach is that language learners – infants and young children, and also adult learners[1] in our experiments in the lab – can use the statistical information derived from the linguistic distribution of elements in the speech stream to determine such things as what sequences of sound form the morphemes and words of the language, in what syntactic contexts these elements can appear, what grammatical categories they form, and what the phrases and sentence structures of the language are. The studies we and others have done, testing learners on their ability to compute such statistics, have shown that language learners are remarkably good at acquiring the information in their linguistic input that signals such facts about the language they are learning.

However, in the present paper I want to make a more complex point about language acquisition in general and about statistical learning in particular. There are two problems in explaining language acquisition. One concerns how learners acquire the details of their particular language from the linguistic input they receive. The second concerns how to account for language universals – that is, principles or common tendencies in the constructions that are found across the languages of the world and that appear as languages change through time. This second problem is most often addressed by positing innate knowledge of linguistic universals. Here, I will suggest, in contrast, that these two problems can be addressed by the same statistical learning mechanism. While learners are indeed adept at acquiring the details of the particular language to which they are exposed, they do not always acquire those details veridically. Under certain circumstances – and particularly when the learners are children – statistical learning results in shifts, sharpening, and regularization of the patterns in the input statistics. In many cases these tendencies within statistical learning can result in changes in the languages acquired and may thus be responsible for the appearance of certain language universals.

---

[1] In this paper – and indeed throughout our work – we use the terms 'adult learners' and 'child learners' rather than 'first language learners' and 'second language learners'. While in many studies these two contrasts are correlated, there are crucial cases in which they are not (cf. Newport, 1990; Mayberry & Eichen, 1992, on child versus adult learners of ASL as a primary language; Mayberry, Lock, & Kazmi, 2002, on child versus adult learners of ASL as a first or a second language). While both variables have an impact on language acquisition, for the points I will make here the important difference is whether the learner is a child or an adult at the time of learning.

448

First I will briefly review our work showing that both adults and children are capable of rapid and adept learning of quite complex distributional patterns in miniature artificial languages, through statistical learning. Then I will turn to our most recent work showing that, when the patterns we present violate regularities that are widespread in natural languages, learners impose these regularities on their input, shifting less natural languages toward patterns that are linguistically more natural. Our findings suggest that statistical learning has inherent constraints or biases in the types of patterns that are most readily learned; that these biases are strongest in young children and weaken substantially with age; and that they may help to explain some of the typologically common patterns found in the languages of the world.

2. Statistical learning from input: words, phrases, and sentence structure

Our earliest work on statistical learning focused on the problem of word segmentation. Since words vary so widely from one language to another (and therefore, of course, could not be known innately), learners must be utilizing cues in the linguistic input to determine which stretches of sound in the speech stream form the words of the language. Structural linguists had long provided suggestions for what these cues might be: the words of the language are those sound sequences that regularly recur within a corpus, in contrast to word boundaries, which vary depending on what words happen to follow one another in a particular sentence (Harris, 1955). In our first statistical learning experiments (Saffran, Aslin, & Newport, 1996; Saffran, Newport, & Aslin, 1996) we exposed learners to streams of speech comprised of four to six words, randomly ordered and produced by a speech synthesizer so that there were no acoustic cues to word boundaries and no meaning or prosodic cues to where the words began and ended; learners could learn the words only if they were capable of utilizing the statistical regularities that arise from the consistency of sounds within a word, as compared with the relative inconsistency of sounds across a word boundary. We hypothesized that learners could segment words within the speech stream by computing the transitional probabilities with which one syllable followed another. The sequences with relatively high transitional probabilities were the words; those with relatively low transitional probabilities were at the boundaries between words. Both infants and adults succeeded at this type of learning, as evidenced by their ability after exposure to select the trisyllabic words over part-words – 3-syllable sequences that spanned a word boundary, consisting of the end of one word and the beginning of another. These results suggested that language learners could indeed compute fairly complex statistics about sound sequences, rapidly and online, and use them to learn the structure of the sound stream.

Since that time, Richard Aslin and I have conducted many studies demonstrating statistical learning in other aspects of language structure, and we have also articulated the specific types of statistics that adult and child learners must be computing in order to acquire them. Our work has delineated the statistical learning mechanisms required to learn adjacent as well as non-adjacent sound patterns in word structure (Newport & Aslin, 2004); syntactic patterns that give rise to form class categories and subcategories (Reeder, Newport, & Aslin, unpublished; Reeder, Newport, & Aslin, 2013; Schuler, Reeder, Newport, & Aslin, 2016); syntactic patterns for different types of verb–argument structure (Wonnacott, Newport, & Tanenhaus, 2008); and syntactic patterns that lead to learning phrase structure (Thompson & Newport, 2007). In some cases we have also shown that the same types of statistical learning can occur in non-linguistic materials as well (Creel, Newport, & Aslin, 2004; Fiser & Aslin, 2002; Gebhart, Newport, & Aslin, 2009; Hunt & Aslin, 2001, 2010; Saffran, Johnson, Aslin, & Newport, 1999). In each of these studies, the statistical cues to linguistic structure have been regular, consistent, and typical of the languages of the world. Under these circumstances, both adults and children are able to acquire these patterns and appear to learn in similar ways. (For studies of children as statistical learners, see Saffran, Aslin, & Newport, 1996; Saffran et al., 1999; Saffran, Newport, Aslin, Tunick, & Barrueco, 1997; Schuler, Reeder, Lukens, Aslin, & Newport, unpublished.)

3. Statistical learning, language universals, and language change

In the experiments just described, learners are exposed to consistent statistical cues to structures that can vary greatly from one language to another and therefore must be learned from the details of linguistic input. Under these circumstances, infants, young children, and adults are all very good at learning these elements and their combinatorial properties from the statistical cues that we provide in the input in our experiments (and that natural languages also provide during real language learning). However, we have also been interested in observing the learning of languages that are not typical – where the input is quite inconsistent, as when it is acquired from models who are late learners or pidgin language speakers; or where the language violates so-called language universals. In these experiments, children and adults look quite different from one another. Most interesting, children do not acquire these languages veridically – they alter the languages, making them more consistent and more aligned with universal tendencies of natural languages. These results reinforce our previous findings on differences between child and adult learners (Johnson & Newport, 1989; Newport, 1990), and also support

a long-standing hypothesis in the linguistics literature that children may play a special role in the emergence and change of languages through time.

3.1. CHILDREN, ADULTS, AND INCONSISTENT INPUT

An important phenomenon in the deaf signing community is that children learning ASL as their native language, from birth, are often learning that language from parents who are late learners of the language (Fischer, 1978; Newport, 1981, 1982, 1990). While the parents, like late learners of other languages, may use complex constructions inconsistently and with many errors, their children look like other native users, acquiring these constructions without learning their errors (Ross, 2001; Singleton & Newport, 2004). In hearing communities, such improvements among child language learners might arise from input they receive from native speakers outside the family. However, in the deaf community, due to the small numbers of native signers – only 5–10% of deaf signers are native users of the language (Schein & Delk, 1974) – there are many children learning ASL only from their late learning parents, without any exposure to native signers. Singleton and Newport (2004) and Ross (2001) showed that such children make ASL constructions much more consistent, acquiring their parents' regular usages but not their inconsistent errors. Figure 1, taken from Singleton and Newport (2004), shows the production of morphemes in ASL verbs of motion produced by one child, Simon, compared with those of his parents, who provided his only ASL input. While Simon's parents produced each of these ASL morphemes correctly about 70–75% of the time (and produced inconsistent errors in the other 25–30%), Simon produced the same morphemes correctly almost 90% of the time, virtually eliminating their inconsistent errors. We described this finding as REGULARIZATION and suggested that it may be similar to what happens when children acquire young pidgin or early creole languages (sometimes called CREOLIZATION) (cf. also Fischer, 1978; Newport 1988, 1999).

Hudson Kam and Newport (2005, 2009) brought this phenomenon into the laboratory in order to understand the process by which children accomplished this regularization. We created miniature languages in which most properties were very regular; but one construction – 'determiners' (*ka* or *po*) that co-occurred with nonsense nouns – were used very inconsistently. The amount of inconsistency was varied across experimental conditions. Austin, Furlong, Schuler, and Newport (unpublished) replicated these experiments with adult and child learners of different ages and with modifications to make the languages easier to learn. In all of these experiments, we found that adults closely reproduced the inconsistencies of their linguistic input, but young children (ages 5–6) acquired only the most regular and consistently used forms. Older children were in between these two groups.
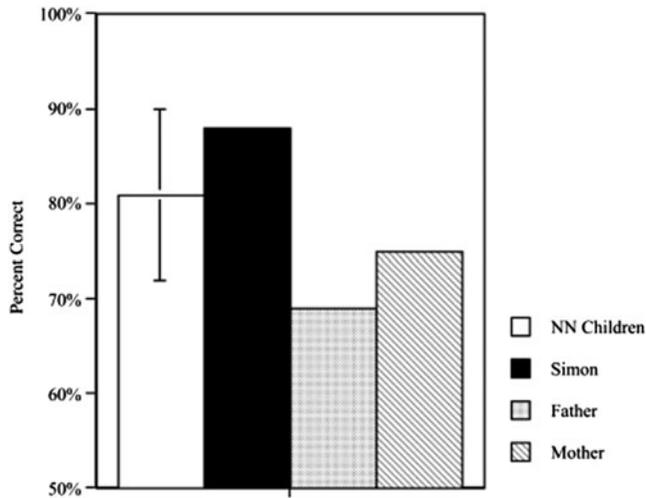
Fig. 1. Simon's use of ASL morphemes, compared to his parents and his native signing peers (from Singleton & Newport, 2004).

Figure 2 shows one example of these results. In this experiment, participants saw short film clips of two puppets interacting, each accompanied by a sentence in a nonsense language with VSO word order and the determiners *ka* or *po* appearing after the nouns. In the input, *ka* appeared after nouns 67% of the time, *po* appeared after nouns 33% of the time. There were no characteristics of the sentences that predicted when *ka* vs. *po* would appear; they simply varied, with 67% *ka* and 33% *po* with every noun and in each sentence position in the language. Adults reproduced this 67/33 variation with amazing precision. Children, in contrast, produced *ka* about 90% of the time and *po* only about 10%.

These results make a number of important points. First, it is important to point out that this is a type of statistical learning: inconsistent variation in these studies follows a set of controlled statistical probabilities, and adult learners indeed reproduced the precise statistics of variation. Children also followed the statistics of their input, but only in the sense that they learned best the form that occurred more frequently or more consistently. However, in contrast to the adults, children reproduced this more consistent form almost all the time – turning probabilistic variation into something more like a rule. In subsequent experiments we have shown that children are capable of learning forms that occur with low frequency: they can learn both high frequency *ka* and low frequency *po* when each is used consistently, in a predictable context. But when each is used unpredictably and inconsistently, they strongly favor producing the more consistent form. This is apparently
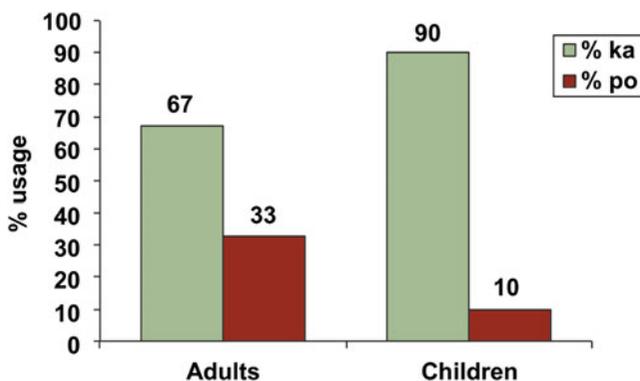
452

Fig. 2. Adults versus children in Inconsistent 67/33 Condition.

what Simon does in natural language learning, and perhaps what children exposed to young languages or inconsistent language communities do as well.

3.2. CHILDREN, ADULTS, AND LANGUAGE UNIVERSALS

In the cases described above, variation and inconsistencies occur in constructions that are language-specific. The forms that are used inconsistently (e.g., *ka* vs. *po*) do not differ in which is more common or widespread in the languages of the world, and children often favor learning whichever form is used more consistently (though see Hudson Kam & Newport, 2009: children sometimes develop their own rules and do not always learn best the form that is used the most frequently). We have also studied the effects of statistical variation in learning constructions that are in accord with, or that violate, patterns that are widespread or universal in the languages of the world, in order to see whether learners are biased to shift languages toward the typologically common patterns.

3.2.1. *Greenberg word order universals*

Culbertson, Smolensky, and Legendre (2012) adapted the miniature language paradigm described above to ask whether adult learners favor following the word order principles described by Greenberg (1963) as highly common in languages around the world. One widespread typological pattern, captured by Greenberg in his Universal 18 and as the 'head directionality parameter' in the Principles and Parameters framework (Baker, 2001), is a pattern of consistent or harmonic word order: nouns and their various modifiers (adjectives, numerals, relative clauses) tend to be in a consistent order, either with the noun first and all of the modifier phrases after or with the modifier

453

phrases first and the noun after. While not all languages follow these patterns, almost 80% of the world's languages surveyed in the World Atlas of Language Structures do (WALS: Dryer, 2008a, 2008b). Culbertson et al. (2012) presented adult learners with one of four miniature artificial languages in which nominal word order was variable; two of the languages predominantly observed these patterns, while the other two predominantly violated these patterns (though each language contained 30% utterances which had the opposite word order). Importantly, after exposure, adult learners produced languages that were slightly but significantly more harmonic. The results are shown in Figure 3. In contrast to Hudson Kam and Newport (2005, 2009), and to Austin et al. (unpublished), described above, they did not reproduce exactly the same variation as in their input and also did not merely regularize the predominant input pattern; rather, they regularized their input in the direction of the typologically more common (more harmonic) pattern.

3.2.2. *Word order and inflectional morphology*

Fedzechkina, Jaeger, and Newport (2012; Fedzechkina, Newport, & Jaeger, in press) have used a similar paradigm, also adapted from Hudson Kam and Newport (2009), to investigate how adult learners acquire miniature languages that do not exhibit the usual trade-off of natural languages between word order and morphological inflection for marking grammatical case. In natural languages, grammatical case contrasts (who does what to whom – devices marking the subject vs. object of the sentence) are typically marked either by consistent word order (for example, the subject of the sentence is the first NP of the main clause) or by inflection (for example, the subject of the sentence (or its determiner) takes a nominative ending). Languages utilizing inflectional case markers typically have more flexible word order; and languages with optional inflectional case markers will use them more frequently when there is a non-canonical word order or an unexpected subject. In our experiments, however, we created somewhat unnatural languages: languages that had variations in word order or the animacy of its nouns with no corresponding changes in the use of inflectional case markers. For example, in one experiment, two word order variants (SOV and OSV) occurred with 60/40 frequency; nouns could be either animate or inanimate (50/50); and a case marker for direct objects could occur or be absent (60/40); but these variations occurred independently and were not related as they would be in a natural language. Importantly, after several days of learning the languages, learners own productions were slightly but significantly more like natural languages. In Figure 4 we can see that learners used the object case marker more frequently when the direct object was animate (an unexpected type of noun to be the direct object)
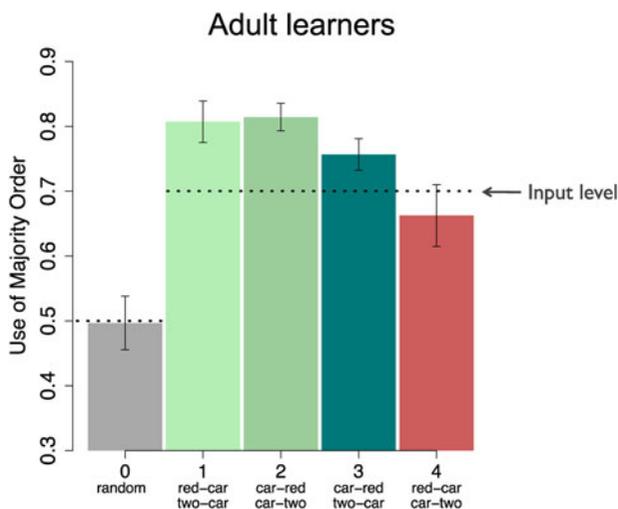
454

Fig. 3. Harmony bias in adult learners (from Culbertson et al., 2012).

than when it was inanimate (the more prototypical direct object), and more frequently when the word order was the uncommon OSV than when it was the more common SOV word order.

3.2.3. *Greenberg word order universals: comparing children and adults*

While Culbertson et al. (2012) found that even adult learners tended to shift variable languages toward Greenberg's typologically more common harmonic patterns, Jenny Culbertson and I wanted to ask what children would do when faced with the same variation – after all, in our studies of inconsistent *ka/po* variation, children regularized much more extensively and under a much broader range of circumstances than did adults. We exposed six- to seven-year-old children to the nominal word order patterns that Culbertson et al. had used with adults; we adapted the task for children by making the languages slightly simpler and extending exposure over two days. As in Culbertson et al., participants were exposed to one of four languages: two whose word orders were 75% harmonic / 25% non-harmonic and two whose noun phrase word orders were 25% harmonic / 75% non-harmonic. While adult learners had made each of the languages about 10% more harmonic than the ones to which they were exposed (see Figure 3), children increased harmony in their word orders much more dramatically. Because children did not merely regularize the word order patterns to which they were exposed, we determined a preferred word order pattern for each child (the noun–modifier and noun–numeral pattern that the child produced most often). Virtually all the children
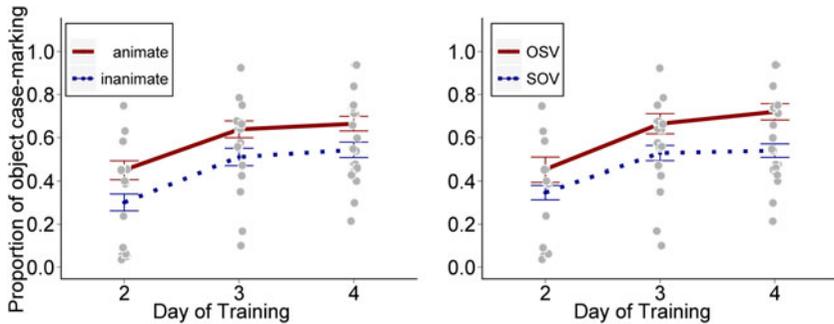
455

Fig. 4. Differential object case marking (Fedzechkina, Jaeger, & Newport, 2012).

had a preferred word order pattern that they used more than 75% of the time. Figure 5 shows the percentage of children who preferred each of the possible word order patterns. The two harmonic patterns – Adj+N and Num+N or N+Adj and N+Num – together were the preferred word order patterns for a whopping 85% of the children. This result shows that the bias toward using harmonic word orders, which was significant but small in adult learners, is even stronger in young children.

In Culbertson et al. (2012) with adults, and in this study with children, we presented learners with languages that each contained a variety of word order patterns, some harmonic and some not. This type of design allows learners to display learning biases in terms of which of these input patterns they reproduce more frequently. One final study, still in progress, with children ages four to seven compared with adults, asks what learners will do if they are exposed to languages that are CONSISTENTLY NON-HARMONIC – where adjectives are always before their nouns and numerals are always after, or the reverse. Under these circumstances, where languages are internally perfectly consistent but do not follow typologically common patterns, will learners still display any biases toward harmonic Greenberg word orders? Our results thus far are shown in Figure 6.

When languages are perfectly consistent – even if they do not follow typologically common patterns – adults acquire the language to which they are exposed. Here they do not display ANY tendency to shift their language toward Greenberg harmonic patterns. Apparently a bias toward language universal patterns is weak enough in adults that it will only be visible in our experiments when there is some support for these patterns in the input language, allowing them to amplify the frequency of these word orders in the language to which they are exposed. But when the input language does not contain any harmonic patterns, adults acquire precisely the patterns to which they are exposed. In striking contrast, young children strongly prefer
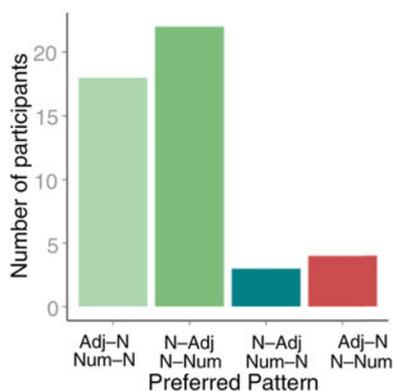
456

Fig. 5. Harmony bias in children (Culbertson & Newport, 2015a).

harmonic patterns, regardless of whether their input language consistently displays another pattern. As in the earlier experiment, children each follow a preferred pattern of their own; here 94% of the youngest children (ages 4–5) prefer a harmonic pattern. Older children (ages 6–7) still significantly favor a harmonic order, but less strongly, showing response patterns between those of young children and those of adults.

What are these biases in children, and why do they appear more strongly in children than adults? One possibility is that these are INNATE LINGUISTIC PRINCIPLES (e.g., the head directionality parameter) that strongly constrain young children's acquisition of language, and weaken, relative to the tendency to learn input patterns, over age (see Culbertson & Smolensky, 2012, and Culbertson, Smolensky, & Wilson, 2013, for a probabilistic formulation of such constraints). Another possibility – which I favor – is that these are BIASES TOWARD PATTERN CONSISTENCY THAT ARISE FROM MORE GENERAL CONSTRAINTS ON LEARNING AND PROCESSING and that are strongest in young children due to their more limited cognitive processing capacities (Culbertson & Newport, 2015a; Hudson Kam & Newport, 2009; Newport, 1990). On this latter formulation, learners more easily acquire linguistic structures (and perhaps other types of patterns) that are more consistent; in young children this may regularize and change the language to which they are exposed. Indeed, in many recent experiments on learning rules and exceptions in our lab, a bias toward consistency and regularization is very strong in young children – a consistent and productive use of morphemes that are inconsistently used or not fully productive in the input, and here a consistent ordering of words across a form class category (Austin et al., in preparation; Culbertson & Newport, 2015a, 2015b; Hudson Kam & Newport, 2009; Schuler, Yang, & Newport, 2016). In contrast, older children

457
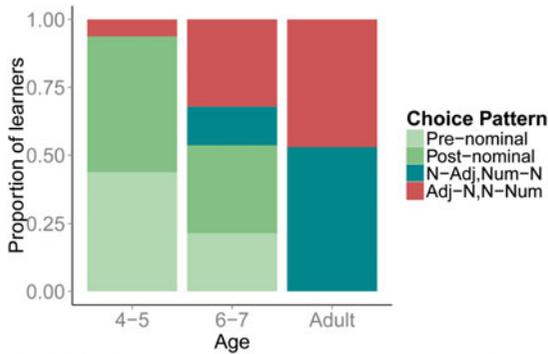
Younger children: 94% harmonic; Older children: 54%: Adults 0%

Fig. 6. Harmony bias in children and adults when input is consistently non-harmonic (Culbertson & Newport, 2015b).

and adults reproduce inconsistencies more faithfully, though they too will shift languages toward consistency and word order harmony to some degree.

4. Conclusions: statistical learning of input regularities in children and adults

Overall, these studies make some strong suggestions about language learning in general and statistical learning in particular. First, as Richard Aslin and I have shown in many studies, language learners show an extraordinary ability to acquire many different types of linguistic structures through the statistical regularities that these patterns create in their input. Adults, children, and even young infants can segment potential words from a continuous stream of speech, assign words to grammatical categories, and group these categories into hierarchical phrase structures. When these statistical patterns are strongly cued and consistent in the input, and when they are in accord with the typological patterns of natural languages, learners of all ages succeed in acquiring them veridically.

However, statistical learning is not always veridical or invariant over age. When linguistic input is variable and inconsistent, or when the input patterns conflict with tendencies that are widespread in natural languages, children look quite different from adults. In these circumstances, adults still learn what is presented to them, or may alter the language a bit, particularly if the input contains some evidence for patterns that follow common linguistic tendencies. In contrast, children dramatically change variable, inconsistent, and uncommonly structured languages to be consistent and rule-governed and to follow structural principles that are widespread in natural languages. These tendencies change gradually over age, with older children performing

458

between young children and adults. While most of our studies involve linguistic materials, we have also done some with non-linguistic patterns and see the same results. We therefore believe that these tendencies are cognitive and not specific to language, though much more research is needed to distinguish these possibilities.

Perhaps most significant, our results suggest that statistical learning is strikingly powerful and precise, but also – particularly in children – involves sharpening input statistics and making patterns more consistent across the language. These processes potentially explain why children acquire language (and other patterns) more effectively than adults, and also may explain how systematic language structures emerge in communities where usages are varied and inconsistent. Most especially, they suggest that input- and usage-based learning approaches must account for differences between adults and children in how usage properties are acquired, and must also account for the substantial changes made by child learners in how input usage properties are represented during learning.

REFERENCES

Aslin, R. N., & Newport, E. L. (2012). Statistical learning: Ffrom learning items to generalizing rules. *Current Directions in Psychological Science*, **21**, 170–176.

Baker, M. (2001). *The atoms of language: the mind's hidden rules of grammar*. New York: Basic Books.

Charniak, E. (1993). *Statistical language learning*. Cambridge, MA: MIT Press.

Creel, S. C., Newport, E. L., & Aslin, R. N. (2004). Distant melodies: statistical learning of non-adjacent dependencies in tone sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **30**, 1119–1130.

Culbertson, J., & Newport, E. L. (2015a). Harmonic biases in child learners: in support of language universals. *Cognition*, **139**, 71–82.

Culbertson, J., & Newport, E. L. (2015b). Language learning and word order regularities: children's errors reflect a typological preference for harmonic patterns. Paper presented at Boston University Conference on Language Development, November.

Culbertson, J., & Smolensky, P. (2012). A Bayesian model of biases in artificial language learning: the case of a word-order universal. *Cognitive Science*, **36**, 1468–1498.

Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition*, **122**, 306–329.

Culbertson, J., Smolensky, P., & Wilson, C. (2013). Cognitive biases, linguistic universals, and constraint-based grammar learning. *Topics in Cognitive Science*, **5**, 392–424.

Dryer, M. (2008a). Order of adjective and noun. In M. Haspelmath, M. S. Dryer, D. Gil, & B. Comrie (Eds.), *The world atlas of language structures online* (chapter 87). Munich: Max Planck Digital Library.

Dryer, M. (2008b). Order of numeral and noun. In M. Haspelmath, M. S. Dryer, D. Gil, & B. Comrie (Eds.), *The world atlas of language structures online* (chapter 89). Munich: Max Planck Digital Library.

Fedzechkina, M., Jaeger, T. F., & Newport, E. L. (2012). Language learners restructure their input to facilitate efficient communication. *Proceedings of the National Academy of Sciences*, **109**, 17897–17902.

Fedzechkina, M., Newport, E. L., & Jaeger, T. F. (in press). Balancing effort and information transmission during language acquisition: evidence from word order and case-marking. *Cognitive Science*.

459

Fischer, S. D. (1978). Sign language and creoles. In P. Siple (Ed.), *Understanding language through sign language research* (pp. 309–331). New York: Academic Press.

Fiser, J., & Aslin, R. N. (2002). Statistical learning of higher order temporal structure from visual shape-sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **28**, 458–467.

Gebhart, A. L., Newport, E. L., & Aslin, R. N. (2009). Statistical learning of adjacent and non-adjacent dependencies among non-linguistic sounds. *Psychonomic Bulletin & Review*, **16**, 486–490.

Greenberg, J. (1963). Some universals of grammar with particular reference to the order of meaningful elements. In J. Greenberg (Ed.), *Universals of language* (pp. 73–113). Cambridge, MA: MIT Press.

Harris, Z. S. (1955). From phoneme to morpheme. *Language*, **31**, 190–222.

Hudson Kam, C., & Newport, E. L. (2005). Regularizing unpredictable variation. *Language Learning and Development*, **1**, 151–195.

Hudson Kam, C., & Newport, E. L. (2009). Getting it right by getting it wrong: when learners change languages. *Cognitive Psychology*, **59**, 30–66.

Hunt, R. H., & Aslin, R. N. (2001). Statistical learning in a serial reaction time task: simultaneous extraction of multiple statistics. *Journal of Experimental Psychology: General*, **130**, 658–680.

Hunt, R. H., & Aslin, R. N. (2010). Category induction via distributional analysis: evidence from a serial reaction time task. *Journal of Memory and Language*, **62**, 98–112.

Johnson, J. S., & Newport, E. L. (1989). Critical period effects in second language learning: the influence of maturational state on the acquisition of English as a second language. *Cognitive Psychology*, **21**, 60–99.

Mayberry, R., & Eichen, E. (1992). The long-lasting advantage of learning sign language in childhood: another look at the critical period for language acquisition. *Journal of Memory and Language*, **30**, 486–512.

Mayberry, R. I., Lock, E., & Kazmi, H. (2002). Linguistic ability and early language exposure. *Nature*, **417**, 38.

Newport, E. L. (1981). Constraints on structure: evidence from American Sign Language and language learning. In W. A. Collins (Ed.), *Aspects of the development of competence: Minnesota symposium on child psychology* (pp. 93–124). Hillsdale, NJ: Lawrence Erlbaum Associates.

Newport, E. L. (1982). Task-specificity in language learning? Evidence from speech perception and American Sign Language. In E. Wanner & L. R. Gleitman (Eds.), *Language acquisition: the state of the art* (pp. 450–486). New York: Cambridge University Press.

Newport, E. L. (1988). Constraints on learning and their role in language acquisition. *Language Sciences*, **10**, 147–172.

Newport, E. L. (1990). Maturational constraints on language learning. *Cognitive Science*, **14**, 11–28.

Newport, E. L. (1999). Reduced input in the acquisition of signed languages: contributions to the study of creolization. In M. DeGraff (Ed.), *Language creation and language change: creolization, diachrony, and development* (pp. 161°178). Cambridge, MA: MIT Press.

Newport, E. L., & Aslin, R. N. (2004). Learning at a distance: I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, **48**, 127–162.

Reeder, P. A., Newport, E. L., & Aslin, R. N. (2013). From shared contexts to syntactic categories: the role of distributional information in learning linguistic form-classes. *Cognitive Psychology*, **66**, 30–54.

Ross, D. S. (2001). *Disentangling the nature–nurture interaction in the language acquisition process: evidence from deaf children of hearing parents exposed to non-native input*. Unpublished doctoral dissertation, Department of Brain & Cognitive Sciences, University of Rochester.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month old infants. *Science*, **274**, 1926–1928.

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tonal sequences by human infants and adults. *Cognition*, **70**, 27–52.

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: the role of distributional cues. *Journal of Memory and Language*, **35**, 606–621.

Saffran, J. R., Newport, E. L., Aslin, R. N., Tunick, R. A., & Barrueco, S. (1997). Incidental language learning: listening (and learning) out of the corner of your ear. *Psychological Science*, **8**, 101–105.

Schein, J., & Delk, M. (1974). *The deaf population of the United States*. Silver Spring, MD: National Association of the Deaf.

Schuler, K. D., Reeder, P. A., Newport, E. L., & Aslin, R. N. (2016). The effect of Zipfian frequency variations on category formation in adult artificial language learning. *Language Learning and Development* (submitted).

Schuler, K. D., Yang, C., & Newport, E. L. (2016). Testing the Tolerance Principle: children form productive rules when it is computationally more efficient to do so. *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.

Singleton, J. L., & Newport, E. L. (2004). When learners surpass their models: the acquisition of American Sign Language from inconsistent input. *Cognitive Psychology*, **49**, 370–407.

Thompson, S. P., & Newport, E. L. (2007). Statistical learning of syntax: the role of transitional probability. *Language Learning and Development*, **3**, 1–42.

Wonnacott, E., Newport, E. L., & Tanenhaus, M. K. (2008). Acquiring and processing verb argument structure: distributional learning in a miniature language. *Cognitive Psychology*, **56**, 165–209.

461