Local versus Global Statistical Learning in Language

Luca Onnis¹ and Shimon Edelman²

¹University of Genoa, Italy

²Cornell University, USA

Pre-print

Author Note

Word Count: 6431

Abstract

A generally held assumption about human statistical learning is that learners keep track of the global statistics of the elements of interest across the entire set of stimuli they are exposed to. In naturalistic settings, this assumption is problematic because it requires that the cognitive system keep track of an exponentially growing number of relations while determining which of those relations are relevant and which are not. We investigated a more plausible assumption, namely, that statistical learning proceeds incrementally, using small windows of opportunity in which the relevant relations are assumed to hold over temporally contiguous objects or events. This local statistical learning hypothesis was tested on two learning tasks, one involving non-adjacent structures and the other – novel word-to-world mappings. Our results suggest that human subjects make use of simple general-purpose cognitive heuristics that exploit temporal contiguity and contrast, leading to superior learning in tasks that belong on two different levels of language acquisition.

Keywords: statistical learning non-adjacent dependencies fast mapping
combinatorial explosion sequential learning implicit learning learnability

Local versus Global Statistical Learning in Language

Introduction

A fundamental goal of cognitive science is to understand how learning systems extract and generalize structure from the environment in which they are embedded. Much of learning can be characterized as finding patterns in space and time under conditions of uncertainty — as in inducing categories and rules from observed data (e.g., Feldman 2000; Goodman, Tenenbaum, Feldman, and Griffiths 2008; Tenenbaum and Griffiths 2001), learning word meanings from their co-occurrence with perceived events in the world (e.g., Frank, Goodman, and Tenenbaum 2009), or acquiring different levels of linguistic structure (e.g., Solan, Horn, Ruppin, and Edelman 2005). Any psychologically plausible theory of structure induction must embrace the volume and the “messiness” of real-world perceptual experience (James 1890). In particular, given the extraordinary amount of data available at all times from multiple streams of sensory information, one must ask: how exactly does the cognitive system keep track of an exponentially growing number of relations among various pieces of data, while identifying the significant relations and excluding spurious correlations? This problem is known as the combinatorial explosion of hypotheses. An appeal to combinatorial explosion has been used to argue for the intractability of purely experience-driven learning, for instance the putative unlearnability of human language from experience alone (see Fulop and Chater 2013 for a discussion).

Several approaches to the combinatorial problem have been offered. According to a nativist account, certain representations of the relevant structures are innately available to human learners, acting as a filter that effectively reduces the size of the hypothesis search space. Theorists working in this tradition hold inductive statistical learning to be too slow and incapable of discovering core structural aspects of language (Hauser, Chomsky, & Fitch, 2002), requiring instead that any learning rely on top-down built-in representations (Endress & Bonatti, 2007). For example, according to Peña, Bonatti, Nespor, and Mehler (2002), learning a language requires both statistical computations to identify words in running speech and algebraic-like computations to discover higher

level structure such as grammatical relations between such words.

An alternative view holds child-directed speech to be sufficiently rich in relevant structural information, and inductive learning mechanisms to be sufficiently powerful, so that correct linguistic structures can be learned and generalized without the need for innate structural constraints, and even in the absence of explicit error correction.

Support for this account comes from theoretical analysis (Chater & Vitányi, 2007), computational modeling (Roberts, Onnis, & Chater, 2005), and behavioural experiments (Figueroa & Gerken, 2019; Hsu, Chater, & Vitányi, 2013; Reeder, Newport, & Aslin, 2017). In addition, child-directed speech may contain more indirect negative evidence than generally acknowledged. For instance, Moerk (1991) showed that most parental “expansions” in early childhood involve grammatical corrections.

Notably, the latter approach still needs to explain how the combinatorial explosion problem is resolved. According to the hypothesis that we explore here, learners may never effectively face a full combinatorial explosion, because the learning they engage in is restricted in time and space to small windows of opportunity. Such local statistical learning does not require extracting complex statistics from large batches of linguistic input, yet is sufficiently powerful to acquire information about relevant structures. This account is supported by several lines of evidence. First, empirical studies of caregiver-child social interactions suggest that caregivers provide consistent and temporally contingent forms of social cues that enhance the learning process, thus reducing drastically the hypothesis search space, especially in infants and young children. For example, pre-linguistic infants learn to produce new phonological patterns from speech that is contingent on their babbling, but not from identical but non-contingent speech Goldstein et al. (2010). Further, when adults engage in iterated turn-taking speech with children, they often quickly and contingently respond to and expand upon the learner’s productions, and show fine timing tuning to the children’s vocalisations (Elmlinger, Schwade, & Goldstein, 2019).

In these studies the local temporal structure of social interactions appears crucial, suggesting that when learning proceeds incrementally, the distillation of spatial and

temporal patterns from the stream of experience can occur in the here-and-now.

Indeed, both children and adults are able to make appropriate inferences from data, despite having limited attention spans and working memory capacity. In this sense, sensory experience in the input data stream must be processed as soon as it is encountered, because only a limited amount of information can be retained for longer periods due to processing 'bottlenecks' (Christiansen & Chater, 2016; Goldstein et al., 2010). This constraint calls for learning that is local in the present sense (regardless of whether the memory bottleneck is a problem to be contended with by the learning system, or is itself an evolutionary adaptation, a possibility raised by Lotem, Kolodny, Halpern, Onnis, & Edelman, 2016). Similarly, other contributors have shown learning is constrained by attentional processes (Toro, Sinnett, and Soto-Faraco 2005). In short, any theory of inductive human learning must explain how success is achieved by relying on small windows of opportunity, and under attention and working memory limitations (as opposed to extracting statistics from large batches of data).

Here we seek to offer a perspective on this problem of local learning, and suggest that human learning can efficiently and incrementally extract local statistics that have a positive global effect on learning. The outline of our paper is as follows. We begin by presenting a computational framework, ACCESS, first introduced by Goldstein et al. (2010), which incorporates plausible cognition-general principles of local learning that are psychologically real, and thus usable by human learners. We then show how these cognition-general principles can be used to identify language-like structures without appealing to unrealistic assumptions about the brain. We present three human experiments examining how people succeed in learning from local input in two language-like cognitive tasks. By manipulating the order of presentation of the training data set, we show that people's sensitivity to features of the ACCESS framework produces superior learning compared to a condition that compels subjects into batch learning. In doing so, we contribute to defusing the computational explosion problem and hope to dispel skepticism about the limits of statistical learning as a general-purpose toolkit for human cognition. Finally, in the Discussion section we link

the effect of temporal contiguity to work on language acquisition, which has long recognized the importance of local cues in acquisition, in the form of partial repetitions, variations, and extensions in natural adult conversation as well as in child-directed speech.

Local ACCESS to structure

Our goal is to investigate how learning can be successful while proceeding incrementally, using small windows of opportunity in which the relevant relations are those that hold over spatially and/or temporally neighboring objects or events. This is conveyed by the acronym ACCESS: Align Candidates, Compare, Evaluate Statistical/Social Significance (Goldstein et al. 2010). According to ACCESS, learners integrate, over a restricted time window, prominent statistical regularities with contextual cues such as social interaction and reward. Statistical significance is realized by recognizing patterns of co-occurrences that emerge above background noise.

To provide a practical scenario, you are at a social gathering attended by several international guests, and you interject a snippet of conversation between two speakers of a language unknown to you. Upon hearing the first utterance *kedmalburafuloropesai* you attempt to parse the speech, and wonder whether the speaker speaks a polysynthetic language such as Mohawk or Nahuatl and has just uttered a single long word. For all you know, the utterance could also be composed of a range of possible words, for instance as many words as there are syllables. There is no way to know. Now assume the next utterance in the conversation comes in: *rafuloro*). Here you may be able to instantly *align* the two utterances still present in your working memory, and *compare* the sounds they have in common. A partial repetition occurred such that the initial unsegmented utterance can be now perceived as a set of three candidate lexical units, namely *kedmalbu-rafuloro-pesai*. This process can be repeated as the conversation unfolds. For example, the third utterance in the conversation being *rafupesai*, two new lexical units offer themselves to your ears upon aligning and comparing the third utterance with the second one, namely *rafu* and *loro*. Thus, when aligned, the partially

matching sentences suggest candidate units (here, *kedmalbu*, *rafuloro*, *pesai*) without the need for learners (and eavesdroppers at a social gathering) to entertain all possible candidate matches over a long series of sentences. Local partial repetitions across sentences appears to facilitate learning in this speech segmentation task.

There is reason to believe that language learning proceeds in ways similar to the example sketched above, and is governed by ACCESS principles. Some indications that different windows are at work in different stages of language development come from the work of Vlach and Johnson (2013), although much work remains to be done to clarify the relevant mechanisms (Smith, Suanda, & Yu, 2014). In particular, there has been little work using learning tasks to examine how people find structured patterns based on limited, local information. Behavioral studies in statistical learning (e.g., Gómez and Gerken 2000; Turk-Browne and Scholl 2009) indicate that infants, children, and adults can extract regularities across a set of exemplars distributed in time and/or space. A common and often tacit assumption of these studies is that to extract such regularities, learners collect global statistics across the entire set of stimuli they are exposed to, often over multiple trials or training sessions. We have seen that when applied to naturalistic learning, this ‘global’ assumption becomes problematic.

A study by Onnis, Waterfall, and Edelman (2008) attempted to capture the local aspect of human language learning depicted in the social gathering vignette (but without the glam of the party element). Adult participants were asked to listen to a sequence of novel unsegmented utterances such as *kedmalburafuloropesai* created from miniature artificial language. Utterances were generated by speech synthesis software and removed acoustic and prosodic cues to word boundary, so that sensitivity to distributional statistics alone could be attributed to any learning. Then at test participants were tested on their accuracy in detecting novel words like *rafu* versus part-words like *fulo*, in a forced-choice task. Onnis et al. (2008) found that learners were significantly better at the word segmentation task when consecutive utterances in the training set overlapped in some of their syllables (e.g., *kedmalburafuloropesai* immediately followed by *rafuloro*), compared to a control condition in which the same set of sentences was purposefully

arranged in a ‘scrambled’ order so that no parts of adjacent sentences matched.

Importantly, the study also found evidence for a “trickle-up” effect: not only did learners more reliably preferred word units heard within pairs of partially matching sentences, but also other valid units that never occurred in such pairs (e.g., *gianaber*, *kiciorudanamjeisulcaz*). In a second experiment, Onnis et al. (2008) also found that identifying phrasal groupings of pre-segmented words defined by a phrase-structure grammar improved when the partial lexical overlap between adjacent sentences was consistent with constituency structure. This second finding is important, as it suggests that sensitivity to local statistics might facilitate the learning not only of lexical units, but also of higher-order relations such as syntactic patterns.

In the present work, we aimed to extend the generality of these findings, by testing the local statistical learning hypothesis in two new experiments requiring different types of linguistic induction. The first task called for the learning of predictive relations among non-adjacent pseudowords (Experiment 1); the second one involved the learning of novel arbitrary word-to-world mappings under conditions of uncertainty and fast mapping (Experiment 2 and 3). In the experiments, adult participants were exposed to miniature languages adapted to each task, subject to a crucial manipulation that allowed us to test the effect of incremental, local statistics versus global statistics. Specifically, the participants were exposed to the language in one of two conditions: 1) a Variation Set (Varset for short) condition, in which the majority of training trials were arranged in sequence such that one element was in common between two consecutive trials (elements were specific to each experiment), and 2) a Scrambled condition, in which a smaller proportion of consecutive trials contained element overlap.

If subjects learn better in the Variation Set condition for both language learning tasks, then this finding, together with those of (Onnis et al., 2008), would constitute evidence that learners resort to a common local or incremental strategy when faced with apparently different learning tasks. This in turn would suggest that statistical learning need not accumulate large swaths of data and risk an impasse due to the combinatorial explosion of possible candidate structures.

Experiment 1: ACCESS to non-adjacent dependencies

Non-adjacent dependencies are an important type of structure found in natural languages and difficult to learn. Formal linguistic theory postulates that constituent structure at the core of linguistic knowledge cannot be learned from surface-level associations (e.g., Chomsky 1957; cf. Lashley 1951, for a similar criticism being leveled independently in the area of sequential action). Chomsky pointed out that that key relationships between words and constituents are often conveyed in nonadjacent (or remotely connected) structure as opposed to adjacent associations. In English, for instance, arbitrary amounts of linguistic material may intervene between auxiliaries and inflectional morphemes (e.g., *Jamie Oliver is showing/*showed on TV*, *Mary has long traveled/*travel the country*) or between subject nouns and verbs in number agreement, and may be separated by embedded subordinate clauses (e.g., *those books that sit on your office shelf *is/are so dusty*).

The presence of embedding and nonadjacent dependencies in language thus represents a point of difficulty for distributional approaches that focus on discovering adjacent relations. For instance, a distributional mechanism computing solely neighboring information could end up endorsing the following sentence **those books that sit on your shelf is so dusty* as grammatical because the nearest-neighbor noun (“shelf”) to the verb (“is”) is singular. A similar case, **the book that sits on your shelves are so dusty* suffers from the same problem. Thus, non-adjacencies represent a challenging test case for distributional learning accounts. At the same time, much of the empirical research on statistical learning so far points to a quick and robust ability of learners to track statistical regularities among adjacent elements, e.g., syllables (Saffran, Newport, & Aslin, 1996). These studies suggest that there is a natural bias to make immediate statistical computations among adjacent properties of stimuli (Futrell, Mahowald, & Gibson, 2015).

Given this bias, distributional tracking of nonadjacent probabilities, at least in uncued streams of syllables, has often proven elusive (Newport & Aslin, 2004; Onnis, Monaghan, Richmond, & Chater, 2005; Peña et al., 2002), and several researchers have

documented how non-adjacencies can be learned when supported by additional cues such as semantics (Poletiek & Lai, 2012), prosody (Mueller, Bahlmann, & Friederici, 2010), phonological status (Bonatti, Pena, Nespor, & Mehler, 2005), or similarities in phonological properties of non-adjacent dependent syllables (Frost and Monaghan 2016; Onnis et al. 2005). Other researchers have instead found that in specific cases distributional properties of the input can suffice to learn non-adjacencies, for example with sufficient variation in the intervening word (Gómez, 2002; Onnis, Monaghan, Christiansen, & Chater, 2004). Thus, it is theoretically important to investigate the limits on learnability of non-adjacencies in carefully controlled experiments that manipulate the distributional properties in which such structures occur. Here, we tested whether or not an incremental, local form of statistical learning might support the detection of non-adjacent dependencies.

Method

Participants

Seventy-four university students were recruited and compensated with a \$8. All spoke English as their native language.

Materials

Training Stimuli. We used an artificial grammar to generate sentences composed of three pseudo-words (henceforth words) each. Sentences had the form $A1\ X\ B1$, $A2\ X\ B2$, and $A3\ X\ B3$. The X words were drawn from a pool of 12 different words. The words, recorded by a female voice, were the same that Gómez (2002) used in her study and were originally chosen as tokens among several recorded sample strings in order to eliminate speaker-induced differences in individual strings. The elements $A1$, $A2$, and $A3$ were instantiated as *pel*, *vot*, and *dak* respectively; $B1$, $B2$, and $B3$, were instantiated as *rud*, *jic*, *tood*. The 12 X middle words were: *wadim*, *suleb*, *kicey*, *puser*, *fengle*, *coomo*, *loga*, *gople*, *taspu*, *hifam*, *deecha*, *vamey*. Words were separated by 250-ms pauses and sentences by 750-ms pauses.

Test Stimuli. Test stimuli were eight grammatical and eight ungrammatical sentences. The grammatical sentences conformed to the grammar of the sentences heard during training and had structure $A1\ X\ B1$, $A2\ X\ B2$, and $A3\ X\ B3$; the eight ungrammatical sentences were created by mismatching the non-adjacent dependencies and had the form $*A1\ X\ B2$, $*A2\ X\ B3$, and $*A3\ X\ B2$. To further test for the ability of participants to recognize non-adjacencies in novel contexts (Frost and Monaghan 2016; Onnis et al. 2004, four sentences in each Grammatical and Ungrammatical set contained two X elements heard during training (*wadim*, *suleb*), while another four contained two new unheard X words (*roosa*, *tuceed*). As during training, words were separated by 250-ms pauses and sentences by 750-ms pauses.

Procedure

Learners were asked to listen and pay close attention to sentences of an invented language and were told that there would be a series of simple questions relating to the sentences after the listening phase. They were not informed of the existence of rules to be discovered. During training, participants listened to a total of 288 sentences, wherein each of the three A_B non-adjacencies was repeated 96 times, and each X word appeared 24 times, eight times for each non-adjacency. Training trials were presented in three blocks separated by a 2 min break. At the end of training and before the test, participants were told that the sentences they had heard were generated according to a set of rules, and that they would now hear 16 strings, eight of which violated the rules. They were asked to press “Y” on a keyboard if they thought a sentence followed the rules and to press “N” otherwise.

As pointed out by Gómez (2002), tracking the difference between grammatical and ungrammatical non-adjacencies is a subtle task that cannot rely upon simple adjacent information, as both grammatical and ungrammatical test sequences were made up of the same bigrams (two-word sequences). As an example, compare the grammatical test sentence $*A1\ X\ B1$ and $*A1\ X\ B2$. During training, all bigrams forming these sentences were heard – $A1\ X$, $X\ B1$, and $X\ B2$. Thus, bigram information is useless and

even misleading in this task.

Design

To test whether learners could detect the correct non-adjacent $A\ B$ structures, the experiment used a between participants design with two conditions of training type: Varset, and Scrambled. Participants were randomly allocated to one of these conditions, with half of participants receiving each type of training, and the same test. In the Variation set (Varset) condition, 80% of training sentences were arranged in pseudo-random order subject to the constraint that at least two consecutive sentences contain the same non-adjacency. For example, *vot malsig tood* would be followed by *vot fengle tood* (see Table 3, rightmost column).

In the Scrambled condition, the same sentences were used but only 5% percentage of immediately consecutive trials contained an overlap in non-adjacencies. For example, *vot malsig tood* would be followed by *pel fengle jick* (see Table 3, Global column). Our hypothesis was that non-adjacencies could be more easily tracked and learned whenever utterances were heard in consecutive variation sets, because this order arrangement facilitates the immediate alignment and comparison of shared non-adjacencies among consecutive sentences. Tracking non-adjacencies is still possible in the Scrambled condition in principle, but is harder as it requires tracking overriding the bias to process uninformative adjacent bigram information (see Gómez 2002; Onnis et al. 2004).

Results

We ran mixed-effects logistic regressions in R using the package *lme4*, entering Participants and Items as random variables, and Condition (Varset vs Scrambled), Grammaticality (Grammatical vs Ungrammatical items), and Experience (items containing Heard vs Unheard X words) as fixed factors. For our categorical variables we used deviation coding, so as to interpret the regression coefficients as main effects. Coefficient values are expressed in log odds. We performed a backward model selection by first fitting a saturated model with all fixed factors and interactions to predict accuracy scores (correct endorsements of grammatical test sentences and correct

rejections of ungrammatical ones), and progressively removing terms that did not improve the fit of the model, starting from the higher-order interactions.

The most parsimonious model revealed an effect of Condition ($\beta = 0.24$, $z = 2.01$, $p < 0.05$), Grammaticality ($\beta = 1.26$, $z = 10.56$, $p < 0.001$), and a Grammaticality $\times$ Experience interaction ($\beta = 0.81$, $z = 3.39$, $p < 0.001$). Participants in the Varset condition exhibited higher accuracy scores (mean proportion correct $M = 0.57$, $SD = 0.12$, scoring significantly better than chance, $t(33) = 3.63$, $p < 0.001$) compared to the Scrambled condition ($M = 0.51$, $SD = 0.10$, not significantly different from chance, $t(39) = 0.76$, $p = 0.22$). In addition, only 39% of participants in the Scrambled condition had a score higher than the chance baseline of 50%, compared to 62% in the Varset condition ($\chi^2(1) = 6.69$, $p < 0.01$). Overall, participants were more accurate in endorsing grammatical sentences than rejecting ungrammatical ones. The effect of Grammaticality was larger for items containing Heard X words. Overall, the findings support the view that participants learned non-adjacent dependencies only in the condition where variation sets were present, and lend support to our initial incremental statistical learning hypothesis.

Experiment 2: ACCESS to form-meaning mappings

Experiment 1 above together with the Onnis et al. (2008) experiments on speech segmentation and phrase structure induction provide evidence for contingency and proximity as facilitators of structure inference from auditory sequences. However, alignment of linguistic structures also occurs with perceptual stimuli from other modalities such as vision. Language learners do not just learn the sequences of words in a sentence, but also track how to map those words to world referents in the environment. In a realistic unsupervised learning scenario there are many potential word-to-world mappings, and the learner must solve this correspondence problem in the face of uncertainty (Quine, 1960). Can local learning based on the principles of alignment and comparison (Goldstein et al., 2010) reduce the combinatorial explosion of hypotheses arising in the process of matching words to their referents?

To address this question, we modeled Experiment 2 after an established cross-situational word learning paradigm by Yu and Smith (2007) (see also Kachergis, Yu, and Shiffrin 2009 for a similar approach). In each learning trial, participants saw multiple novel pictures and simultaneously heard multiple novel words. This naturally created ambiguity regarding which word to map with which object. Consider trial n in Figure 2, in which four new words are heard in sequence (*jide*, *nerk*, *bim*, *cree*) while four shapes appear simultaneously on the screen. In the absence of prior knowledge of the specific word-referent associations, this situation creates as many as 24 (in mathematical notation, $4!=24$) possible word-referent combinations. Which one is correct? This inductive problem cannot be solved on any single trial, but must rely on collective both positive evidence (repetition of certain pairs) and negative evidence (absence of repetition of candidate pairs) *across* multiple trials, hence the term cross-situational learning. The ability to make cross-situation comparisons when learning novel word-world pairings is found in children (e.g., Akhtar and Montague 1999; Childers and Paik 2009). How exactly are the ambiguities in such learning tasks resolved?

Consider a more abstract learning scenario that consists of 4 successive trials as in Table 1, with each trial ambiguously presenting two words (in lower-case) and two referents (in upper-case). In trial 1, learners could mistakenly link the word *a* to the referent *B*; later they could successfully rectify their mistake and discover the correct *a-A* mapping, if *all* the following conditions applied: 1) they registered that word *a* occurred in trial 4 without the incorrect referent *B*; 2) they remembered the prior word-referent pairing; 3) they registered both co-occurrences and non co-occurrences among all possible pairings; and 4) they calculated the right statistics.

Thus, successful mapping in this scenario requires tracking multiple statistical relations both locally (over adjacent trials) and globally (over the entire learning experience). For the purposes of this study, we refer to this mechanism as a global statistical learner, because disambiguation can only occur with certainty after collecting and remembering a sizeable sample of trials. We can compare it directly with the incremental learning scenario of Table 2 that presents partial repetitions/variation sets of word-referent

mappings. Here the pair a - A occurs in trial 2 immediately after trial 1. In the Words column, there is a single word a that is presented across two successive trials; likewise, in the Scenes column a single referent A is presented across two successive scenes. Thus, by simultaneously keeping track of the two constant elements in Words and Scenes (the auditory and visual modality), the learner can in principle immediately identify a unique word-mapping that remains constant over any two successive trials, without the need to hold in memory other potential relations present and absent, as the global learning strategy would require. Mathematically, the combinatorial problem reduces drastically and progressively each time a new candidate pair is found. For example

These considerations lead to the following predictions. If, on the one hand, learners are able to implicitly sample global cross-trial statistics accumulated over all learning trials, then learning in scenarios like 1 and 2 above should be equally successful, because the global statistics are the same — the only difference is the order of presentation of the trials. Conversely, if local alignment and comparison are partly responsible for learning, the scenario depicted in Table 2 should lead to significantly better learning than the scenario depicted in Table 1.

A related question of interest to us was whether or not local learning is equally useful in tasks with varying intrinsic complexity. Yu and Smith (2007) made the cogent case that having to learn more word-referent pairs led counter-intuitively to better overall learning. In their Experiment 2, they found that subjects exposed to 18 word-reference pairs during training actually learned more pairs on average than subjects exposed to only 9 pairs. One explanation for their finding is that a global statistical learner in the 18-pair condition would experience fewer spurious associations that might be expected to make learning from these individually more ambiguous trials difficult. Spurious correlations are co-occurrences between a given word and an incorrect referent given that multiple words and referents were presented on each trial. For a global statistical learner, Yu and Smith (2007) argued that smaller data sets may not be as good as large ones because spurious correlations are more likely to occur.

Thus, one might posit that local statistical learning constitutes a more useful heuristic

when experiencing smaller data sets and when spurious correlations are more likely and can be winnowed out immediately by partial repetitions. Variation set cues may not benefit learners exploiting higher global systematicity across trials with more complex data sets. Alternatively, local learning based on alignment and comparison of successive trials might be operative regardless of the complexity of data set, suggesting that these heuristics flexibly adapt to different data sets. In our design, crossing presence or absence of variation sets with the complexity of the mapping problem allowed us to test whether or not learning heuristics change from local to global.

Method

Participants. Forty-eight undergraduate students who had not participated in Experiment 1 were recruited and compensated with a small sum. All spoke English as their native language.

Materials and Procedure. Materials and design were partly adapted from Experiment 2 of Yu and Smith (2007). Each training trial contained 4 spoken words and 4 pictures of individual objects, with no information about specific word-picture correspondence. This 4 (auditory) x 4 (visual) learning scenario yields a large number of possible word-referent associations: 16 per trial. Because the referent of a word cannot be unambiguously determined during any single trial, disambiguation must rely on information from multiple trials.

Participants were assigned to one of four conditions, crossing two 2-factor manipulations. The first was the manipulation of learning condition as in Experiment 1. Participants in the Varset condition experienced 80% of trials in which one word-referent pair was shared between two consecutive trials. Participants in the Scrambled condition experienced only 5% of consecutive trials containing auditory and visual overlap. An example of two trials with overlapping form-meaning pair is shown in Figure 2. The second crossed manipulation concerned the complexity of the data set: In the 9-words/8-repetitions (9x8) condition, subjects had to discover a total of 9 word-referent pairs each repeated 8 times over the course of training. In the the

18-words/6-repetitions (18x6) condition, subjects had to discover 18 word-referent pairs and were given 6 repetitions of each word-referent pair.

Test trials included each a single word from among those heard during training, for which participants had to choose the correct referent among four possibilities. As in Experiment 1, because the two learning conditions differed only in the order of trials, the global statistics of word/picture co-occurrence were identical across conditions. This allowed us to make differential predictions, as follows. On the global statistical learning account, learners solve the indeterminacy problem by keeping track of multiple word-referent statistics across many individually ambiguous words and scenes, across the several trials in the experiment (Yu & Smith, 2007). Thus, learning should not differ across our two learning conditions, but is expected to be better with a more complex data set, in accordance with the previous findings by Yu and Smith (2007). Thus, the global statistical learner account would predict a main effect of complexity and no effect of learning condition. Conversely, on the local statistical learning account, learners benefit from the contiguous arrangement of partially overlapping trials and should therefore learn better in the Varset condition, regardless of complexity of the data set. A learning condition x complexity interaction can be expected where the effect of local learning is evident only in the less complex data set condition. As discussed above, this outcome would suggest that local statistical learning is more effective for inductive problems with small data set. More importantly for a theory of inductive learning, perhaps learners adopt different heuristics depending on the perceived complexity of the data set at hand.

Results

We fitted a mixed-effects model with Participant and Item as random effects (R procedure *glmer*, with a binomial link function), and Learning Condition, Complexity, and their interaction as fixed factors. We used deviation coding, so as to interpret the regression coefficients of the fixed factors as main effects. Coefficient values are expressed in log odds. A model selection procedure ran using the *glmulti* function

yielded an effect of Complexity (level: 18x6, $\beta = 0.58$, $z = 3.00$, $p < 0.01$), with more word-referent pairs discovered correctly in the 18-pair condition. This finding replicates that of Yu and Smith (2007), suggesting that within-trial uncertainty is a more critical factor in learning than is the number of pairs in the learning set. We also found an effect of Learning Condition (level: Varset, $\beta = 0.42$, $z = 2.26$, $p = 0.026$), with more word-referent pairs discovered correctly in the Varset condition compared to the Scrambled condition. This pattern of findings, coupled with a non-significant Variation set x Complexity interaction, suggests that contrary to an exclusively global account of statistical learning, learning proceeded on a more local trial-by-trial basis, exploiting temporally contiguous relations. Furthermore, the main effect suggests that local statistical learning is as useful a heuristic in small as in larger data sets, and it may thus scale up to more realistic language learning situations.

In terms of the proportion of word-referent pairs discovered, subjects discovered more pairs than expected by chance in all conditions except the Scrambled-9x8 condition, as shown in Figure 3, $t(7) = 0.9$, $p > 0.1$. The fact that some learning did occur in the Scrambled-18x6 condition suggests that learners were able to register both co-occurrences and non co-occurrences among all possible pairings in a global fashion. This suggests that both local and global computations contribute to learning even under conditions of limited exposure, as in the artificial grammar learning task in this experiment.

Experiment 3: a recency effect?

Experiment 2 suggests that in the absence of local statistics in the form of cross-trial partial self-repetitions learners found it at least somewhat harder to discover the correct word-reference pairings. However, the specifics of how presentation order affected learning in the Varset condition remained unknown. For instance, it is possible that learning occurred in the last two trials, driving this effect in its entirety (i.e., learners could just be remembering the last few trials, and if they happened to contain a repetition, that is what was learned). The motivation for Experiment 3 is to eliminate

the possibility that the differences across the Scrambled and Varset condition in Experiment 2 could be due to participants in the varset condition gaining an advantage from the last few learning trials if they contained a repetition. We thus divided the word-referent items to be learned into two sets; the first set was presented in variation sets in the first half of training, while the second group in the second half. If fast mapping is indeed driven by a ‘recency’ effect, items in the second half should be learned better than those in the first half.

Method

Participants. We recruited 36 undergraduate students who had not participated in Experiment 1 and 2. All spoke English as their native language, and received a small sum for participation.

Materials. The same training and test materials as Experiment 2 with the 9x8 data set were used. The only difference was that 4 of the word-referent pairings occurred across successive trials in the first half of the training set (i.e., in the first 9 trials, Early Pairs), while another four occurred in the last 9 trials (Late Pairs).

Procedure. A similar procedure to that of Experiment 2 was used. Given that the focus in this experiment was on better understanding the Varset condition, and that in Experiment 2 learners in the Scrambled condition failed to learn above chance, we did not assign participants to a Scrambled condition. At test, participants received 8 test trials overall; each trial was a single word from among those heard during training, for which they had to choose the correct referent among four possibilities. Among the 8 test trials, 4 were drawn from the Early Pairs (words presented in the first half of the training) and 4 from Late Pairs (words presented in the second half of the training), in random order. Thus, the dependent variable was the number of correctly endorsed test trials, and the independent variable was the Order of trial (Early versus Late).

Results

A mixed-effects model with Order of pairs (Early versus Late) as fixed factor and Participant and Item as random factors indicated no effect of Order between Early and

Late Pairs, ($\beta = -0.17$, $z = -0.54$, $p = 0.59$). This was confirmed by a Kruskal-Wallis test. Separate t-tests showed that both Early and Late Pairs were learned better than chance (Early Pairs, $M = 0.49$, $SD = .50$, $t(17) = 4.27$, $p < .001$; Late Pairs, $M = 0.44$, $SD = 0.50$, $t(15) = 3.05$, $p < .01$). If anything, the mean for Early Pairs was numerically higher (Figure 4). Thus, there is little support for a ‘recency’ effect: participants learned both items presented in either half of training.

General discussion

What principles guide learning from multiple parallel streams of sensory information? How can humans find structure in sustained exposure to auditory and visual stimuli? Under naturalistic learning circumstances, detecting structure in time and space in the environment presents a problem of combinatorial explosion: there are simply too many statistical relations to track, so that detecting the ones that are relevant for any given task would involve collecting large numbers of statistics over time before confidence is reached. For this reason, distributional approaches to structure induction have been criticized as too weak to be of use. Here we showed the potential benefits of temporally local learning in such situations of uncertainty. We made use of two established experimental paradigms with miniature languages, and established that local learning yields global benefits that extend beyond the relations encountered locally, thus helping the learner manage the computational complexity of structural inference. Taken together, our results suggest that experiencing data in a temporally contiguous manner facilitates learning, compared to a non-contiguous presentation of the very same set of items. How exactly do variation sets helped learners in our experiments? A crucial aspect of the variation set is a repetition. But not just any repetition will do. For example, a full repetition of a sentence (e.g., *vot malsig tood ... vot malsig tood*) does not yield candidate sub-units for analysis when submitted to the heuristics of Alignment and Comparison. Likewise, repeating trial #1 in a form-referent mapping task (see Table 4) *a b A B ... b a A B* would not allow disambiguation of which words map to which referents. Thus, repetitions across time have to be *partial-repetitions* to

be useful to ACCESS. In addition, these partial repetitions have to involve the structural elements to be learned. For instance, the repetition *vot malsig tood ... pel malsig jick* may be less informative as it highlights the middle element, and not the non-adjacent element. It may also indicate that the X word *malsig* can appear in multiple contexts, appears. Follow-up experiments should peruse whether repeating what changes (e.g., the X elements in Experiment 1) versus what stays constant (e.g., the A B elements) obtains different learning outcomes.

We believe the local repetition is aiding learning in a similar way across different tasks. For example, in Experiment 1 having the same non-adjacent dependency repeat twice in time may cue the listeners in to what structure is important and invariant across sentences. It also reduces ambiguity about all other possible relations available in the input, thus allowing a faster detection of structures. Similarly, in Experiment 2 and 3, having the same word-referent pairing appear on the next training trial reduces ambiguity for that individual word-referent pairing, but also reduces the number of spurious correlations. For example, if the initial combinatorics of pairings for 4 words and 4 references can be calculated as $4!=24$ possible word-referent combinations, identifying one such pair across two consecutive trials reduces the combinatorial problem to $3!=6$ on the second trial. Thus, fast identification of pairs through immediate repetition may assist in identifying, by inference, what some of the other word-referent pairings are present on that second trial (Romberg and Yu 2013).

Although it is possible that separate mechanisms underlie learning in our experimental tasks, it is more parsimonious to assume that each of these tasks provides a particular perspective on a single processing mechanism. In this article, the ACCESS model is offered as a candidate for such a mechanism.

ACCESS to natural language

Our results highlight the effectiveness of temporal contiguity and of contrast across four distinct language learning tasks, namely: a) speech segmentation, b) induction of phrase structure, c) non-adjacency learning, and d) multimodal learning under conditions of

uncertainty. The findings thus support the view that local statistical learning may be operating at different levels of language induction, thus strengthening the notion that language acquisition is subserved by a small set of general-purpose inductive mechanisms. In addition, the importance of order of presentation of learning materials has intriguing implications for various practical learning and teaching situations.

Artificial grammars that generate partially lexically overlapping consecutive sentences mimic the manner in which caregivers spontaneously produce successive utterances with partial self-repetitions, or ‘variation sets’ Küntay and Slobin (1996). The findings of Onnis et al. (2008) lend experimental support to a body of work pointing to child-directed speech containing rich adaptations that help the learner discover the structure of language, on multiple levels. More generally, those findings are compatible with a recently proposed theoretical framework, ACCESS, which aims to explain the learning of structure in space and time in terms of general principles of cognitive computation (Goldstein et al., 2010).

Conclusion

Contiguity and contrast were first invoked by Aristotle (*De mem. et rem.*) as fundamental laws of association. Following the early insights of Hume, researchers have come to appreciate the crucial role of statistical inference in ensuring the reliability of experiential learning. A recently proposed theoretical framework, ACCESS (Align Candidates, Compare, Evaluate Statistical/Social Significance), aims to explain the learning of structure in space and time in terms of general principles of cognitive computation (Goldstein et al., 2010). In agreement with those principles, our results suggest the effectiveness of temporal contiguity and contrast in sequential and multimodal learning under conditions of uncertainty, and the importance of order of presentation of learning materials. Gómez (2002) proposed that “learners are driven to seek invariant structure in the stimulus array”; the ACCESS framework offers a mechanistic computational account of how invariant structure may be detected.

Acknowledgements

Luca Onnis was supported by Singapore Ministry of Education's Tier 1 grants #M4011320 and #M4011750.

Pre-print

References

- Akhtar, N., & Montague, L. (1999). Early lexical acquisition: The role of cross-situational learning. *First Language*, 19(57), 347–358.
- Bonatti, L. L., Pena, M., Nespor, M., & Mehler, J. (2005). Linguistic constraints on statistical computations the role of consonants and vowels in continuous speech processing. *Psychological Science*, 16(6), 451–459.
- Chater, N., & Vitányi, P. (2007). ‘ideal learning’ of natural language: Positive results about learning from positive evidence. *Journal of Mathematical psychology*, 51(3), 135–163.
- Childers, J. B., & Paik, J. H. (2009). Korean-and english-speaking children use cross-situational information to learn novel predicate terms. *Journal of Child Language*, 36(1), 201.
- Chomsky, N. (1957). *Syntactic structures*. Walter de Gruyter.
- Christiansen, M. H., & Chater, N. (2016). The now-or-never bottleneck: A fundamental constraint on language. *Behavioral and Brain Sciences*, 39. doi: 10.1017/S0140525X1500031X
- Elmlinger, S. L., Schwade, J. A., & Goldstein, M. H. (2019). The ecology of prelinguistic vocal learning: parents simplify the structure of their speech in response to babbling. *Journal of child language*, 46(5), 998–1011.
- Endress, A. D., & Bonatti, L. L. (2007). Rapid learning of syllable classes from a perceptually continuous speech stream. *Cognition*, 105(2), 247–299.
- Feldman, J. (2000). Minimization of boolean complexity in human concept learning. *Nature*, 407(6804), 630.
- Figueroa, M., & Gerken, L. (2019). Experience with morphosyntactic paradigms allows toddlers to tacitly anticipate overregularized verb forms months before they produce them. *Cognition*, 191, 103977.
- Frank, M. C., Goodman, N. D., & Tenenbaum, J. B. (2009). Using speakers’ referential intentions to model early cross-situational word learning. *Psychological science*, 20(5), 578–585.

- Frost, R. L., & Monaghan, P. (2016). Simultaneous segmentation and generalisation of non-adjacent dependencies from continuous speech. *Cognition*, 147, 70–74.
- Fulop, S. A., & Chater, N. (2013). Learnability theory. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(3), 299–306.
- Futrell, R., Mahowald, K., & Gibson, E. (2015). Large-scale evidence of dependency length minimization in 37 languages. , 112, 10336-10341.
- Goldstein, M. H., Waterfall, H. R., Lotem, A., Halpern, J. Y., Schwade, J. A., Onnis, L., & Edelman, S. (2010). General cognitive principles for learning structure in time and space. *Trends in cognitive sciences*, 14(6), 249–258.
- Gómez, R. L. (2002). Variability and detection of invariant structure. *Psychological Science*, 13(5), 431–436.
- Gómez, R. L., & Gerken, L. (2000). Infant artificial language learning and language acquisition. *Trends in cognitive sciences*, 4(5), 178–186.
- Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive science*, 32(1), 108–154.
- Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: what is it, who has it, and how did it evolve? *science*, 298(5598), 1569–1579.
- Hsu, A. S., Chater, N., & Vitányi, P. (2013). Language learning from positive evidence, reconsidered: A simplicity-based approach. *Topics in cognitive science*, 5(1), 35–55.
- James, W. (1890). *The principles of psychology, vol. 2. ny, us: Henry holt and company*.
- Kachergis, G., Yu, C., & Shiffrin, R. M. (2009). Temporal contiguity in cross-situational statistical learning. In *Proceedings of cogsci* (Vol. 31).
- Küntay, A., & Slobin, D. I. (1996). Listening to a turkish mother: Some puzzles for acquisition.
- Lashley, K. S. (1951). The problem of serial order in behavior. In *Cerebral mechanisms in behavior* (pp. 112–136).
- Lotem, A., Kolodny, O., Halpern, J. Y., Onnis, L., & Edelman, S. (2016). The

- bottleneck may be the solution, not the problem. *Behavioral and Brain Sciences*, 39. doi: 10.1017/S0140525X15000886
- Moerk, E. L. (1991). Positive evidence for negative evidence. *First Language*, 11, 219-251.
- Mueller, J. L., Bahlmann, J., & Friederici, A. D. (2010). Learnability of embedded syntactic structures depends on prosodic cues. *Cognitive science*, 34(2), 338-349.
- Newport, E. L., & Aslin, R. N. (2004). Learning at a distance i. statistical learning of non-adjacent dependencies. *Cognitive psychology*, 48(2), 127-162.
- Onnis, L., Monaghan, P., Christiansen, M. H., & Chater, N. (2004). Variability is the spice of learning, and a crucial ingredient for detecting and generalizing in nonadjacent dependencies. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 26).
- Onnis, L., Monaghan, P., Richmond, K., & Chater, N. (2005). Phonology impacts segmentation in online speech processing. *Journal of Memory and Language*, 53(2), 225-237.
- Onnis, L., Waterfall, H. R., & Edelman, S. (2008). Learn locally, act globally: Learning language from variation set cues. *Cognition*, 109(3), 423-430.
- Peña, M., Bonatti, L. L., Nespor, M., & Mehler, J. (2002). Signal-driven computations in speech processing. *Science*, 298(5593), 604-607.
- Poletiek, F. H., & Lai, J. (2012). How semantic biases in simple adjacencies affect learning a complex structure with non-adjacencies in agl: a statistical account. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 367(1598), 2046-2054.
- Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.
- Reeder, P. A., Newport, E. L., & Aslin, R. N. (2017). Distributional learning of subcategories in an artificial grammar: Category generalization and subcategory restrictions. *Journal of Memory and Language*, 97, 17 - 29. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0749596X17300529>
doi: <https://doi.org/10.1016/j.jml.2017.07.006>

- Roberts, M., Onnis, L., & Chater, N. (2005). Acquisition and evolution of quasiregular languages: Two puzzles for the price of one. *Evolutionary prerequisites for language*, 334–356.
- Romberg, A., & Yu, C. (2013). Integration and inference: Cross-situational word learning involves more than simple co-occurrences. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 35).
- Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: The role of distributional cues. *Journal of memory and language*, 35(4), 606–621.
- Smith, L. B., Suanda, S. H., & Yu, C. (2014). The unrealized promise of infant statistical word–referent learning. *Trends in Cognitive Sciences*, 18, 251–257.
- Solan, Z., Horn, D., Ruppin, E., & Edelman, S. (2005). Unsupervised learning of natural languages. *Proceedings of the National Academy of Sciences of the United States of America*, 102(33), 11629–11634.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and bayesian inference. *Behavioral and brain sciences*, 24(04), 629–640.
- Toro, J. M., Sinnett, S., & Soto-Faraco, S. (2005). Speech segmentation by statistical learning depends on attention. *Cognition*, 97(2), B25–B34.
- Turk-Browne, N. B., & Scholl, B. J. (2009). Flexible visual statistical learning: transfer across space and time. *Journal of Experimental Psychology: Human Perception and Performance*, 35, 195–202.
- Vlach, H. A., & Johnson, S. P. (2013). Memory constraints on infants’ cross-situational statistical learning. *Cognition*, 127, 375–382.
- Yu, C., & Smith, L. B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, 18(5), 414–420.

Table 1

Speech segmentation (Onnis et al. 2008, Exp.1)

Structure	Trial	Global	Local
$S \rightarrow ABC$	1	kosifamapiu	da per aguklo
$A \rightarrow \{da, kosi, spino\}$	2	spinoperaguklo	kosi per apiu
$B \rightarrow \{pera, kadro, fama, zupa\}$	3	kosikadropiu	spin oper aguklo
$C \rightarrow \{piu, prati, guklo\}$	4	daperaguklo	da per apрати

Table 2

Phrase structure discovery (Onnis et al. 2008, Exp.2)

Structure	Trial	Global	Local
$S1 \rightarrow Phrase1Phrase2Phrase3$	1	arv bim skiv	arv bim skiv
$S2 \rightarrow Phrase1Phrase2$	2	dro goz heeb irg kuhl larp tood bim	dro goz bim skiv kuhl larp bim skiv
$S3 \rightarrow Phrase3Phrase1Phrase2$	3	plam roo	plam roo
$Phrase1 \rightarrow AB G$	4	nerk tiv dro cree heeb irg	quive tiv dro cree bim skiv
$Phrase2 \rightarrow CD$			
$Phrase3 \rightarrow EF$			

Table 3

Syntactic dependencies (Exp. 1)

Structure	Trial	Global	Local
$S1 \rightarrow A1XB1$	1	vot malsig tood	vot malsig tood
$S2 \rightarrow A2XB2$	2	pel fengle jick	vot fengle tood
$S3 \rightarrow A3XB3$	3	dak feenam rud	pel suleb jick
$A1 \rightarrow vot, A2 \rightarrow pel, A3 \rightarrow dak$	4	vot fengle tood	pel fengle jick
$B1 \rightarrow tood, B2 \rightarrow rud, B3 \rightarrow jik$	5	pel suleb jick	dak feenam rud
$X \rightarrow \{malsig, fengle, \dots\}$	6	dak deecha rud	dak deecha rud

Table 4

Form-referent mapping (Exp. 2-3)

Structure	Trial	Global	Local
		Words Referents	Words Referents
Pairs $\rightarrow \{a - A, b - B,$ $c - C, d - D,$ $f - F, g - G\}$	1	a b B A	a b B A
	2	c d D C	g a A G
	3	e f E F	c d D C
	4	g a G A	e f E F

Figure 1. Test results for Experiment 1. Error bars reflect standard errors.

Figure 2. Example of trials experienced by subjects in the Varset condition of Experiment 2. Only the word *bim* is repeated on two consecutive trials with its associated referent.

Figure 3. Test results for Experiment 2. Error bars reflect standard errors.

Figure 4. Test results for Experiment 3. Error bars reflect standard errors.