When Less is Less and When Less is More: Starting Small with Staged Input

Christopher M. Conway (cmc82@cornell.edu)
Department of Psychology; Cornell University; Ithaca, NY 14853, USA

Michelle R. Ellefson (M.Ellefson@warwick.ac.uk)
Department of Psychology; University of Warwick; Coventry CV4 7AL, UK
Morten H. Christiansen (mhc27@cornell.edu)
Department of Psychology; Cornell University; Ithaca, NY 14853, USA

Abstract

It has been suggested that external and/or internal limitations may paradoxically lead to superior learning (i.e., the concepts of "starting small" and "less is more"; Elman, 1993; Newport, 1990). In this paper, we explore what conditions might lead to a starting small effect. We report on three artificial grammar learning experiments with human participants. In Experiment 1, we found an effect of starting small with visual center-embedded, recursive input staged incrementally. Experiment 2 replicated this finding and extended the effect to right-branching recursive structure. Finally, in Experiment 3 we found no effect for starting small with auditory center-embedded input. These results suggest that starting small can confer a learning advantage but perhaps only under certain conditions.

Introduction

Intuitively, learners should learn better when they are unhindered by internal or external limitations, such as those relating to constraints on memory or on the input. However, recent proposals take the somewhat paradoxical stance that cognitive limitations and/or reduced input may confer a computational advantage for learning. These theories, specifically the notion that "*less is more*" (Newport, 1990) and the importance of "*starting small*" (Elman, 1993), are often couched in terms of language acquisition. For analyses involving componential inputs, such as language, limited processing may be advantageous because it acts as a "filter" to reduce the problem space, making it more manageable.

Unfortunately, the evidence related to starting small is far from conclusive. Though it is true that children learn language better than adults, this may be due to any number of factors. Initially, computational work supported the theory of starting small (e.g., Elman, 1993), but more recent simulations appear to contradict those findings (Rohde & Plaut, 1999; in press). Empirical evidence gathered from human participants has also not resolved the issue. Though some of the data support starting small, (Cochran, McDonald, & Parault, 1999; Kareev, Lieberman, & Lev, 1997), other data do not (Fletcher, Maybery, & Bennett, 2000; Ludden & Gupta, 2000; Rohde & Plaut, in press).

This paper attempts to understand under what conditions, if any, starting small might have an effect. First we discuss the inconclusive evidence for starting small. Second, we discuss recursive grammars and why such structures may provide a suitable testbed. Next, we present data from three artificial grammar learning experiments with human participants. Experiment 1 shows that when visual, center-embedded input is staged in a starting small fashion, participants achieve better learning than when the input is presented non-incrementally. Experiment 2 reveals a similar effect of starting small using right-branching recursive structure. Finally, Experiment 3 provides a test of starting small in the auditory modality using center-embedded input. The results of this last experiment suggest that under some conditions, starting small may not be beneficial. Together, this evidence suggests new ways to interpret the starting small hypothesis and the conditions under which less is less and less is more.

Evidence Relating to Starting Small

The "less is more" and "starting small" hypotheses can actually be thought of as two related but separate ideas. Both are similar in that they propose that processing limitations may confer a learning advantage but they differ in terms of the nature of the limitation itself. One possibility is that the processing limitations arise from internal, cognitive (e.g., memory) constraints. A second possibility is that the processing limitations are external in nature, in the form of staged or incremental input. Here we review data related to these two possibilities, starting with the internal constraint viewpoint.

In the context of language acquisition, Newport (1990) proposed that maturational constraints in the form of cognitive limitations are crucial for allowing language to be learned successfully. Elman (1993) tested this idea by training a simple recurrent network (SRN) to learn aspects of an artificial language. Under normal conditions, the network was unable to learn the sequential regularities of the grammar. But when Elman simulated children's working memory limitations by periodically eliminating the network's access to its prior internal states–and allowing the size of this temporal window

to increase over time–the neural network's performance improved. Further support comes from studies with human participants. Cochran, McDonald, and Parault (1999) taught adults portions of a modified version of American Sign Language (ASL). In their first two experiments, they simulated cognitive limitations by supplying a simultaneous capacity-limiting task during training. Cochran et al. found that the participants in the no-load condition displayed more rigid learning and were less adept at using the signs in new contexts. Additionally, Kareev, Lieberman, and Lev (1997) explored the relation between working memory capacity and the detection of correlation. Human participants were tested on their ability to predict the relationship between two binary variables. Participants with lower natural working memory were better at detecting the appropriate correlation and performed better on the task than did high memory capacity participants. This evidence appears to lend direct support to the importance of starting small: in some situations, cognitive limitations appear to confer a learning advantage.

However, there are reasons to be critical of this data. For instance, Rohde and Plaut (1999; in press) conducted neural network simulations that contradicted Elman's (1993) results. Using the same architecture, simulation parameters, and training input, Rohde and Plaut failed to get an advantage for starting small. Rohde and Plaut (in press) also give reasons for questioning Cochran et al.'s (1999) and Kareev et al.'s (1997) conclusions, instead arguing that their data does not support the notion that internal limitations benefit learning. Other studies appear to support this perspective. For example, adult participants in an artificial grammar learning task with a capacity-limiting condition failed to show an effect of starting small (Ludden & Gupta, 2000). Relatedly, children early in development do not surpass more developed children in artificial grammar learning tasks (e.g., Fletcher, Maybery, & Bennett, 2000).

There are fewer experiments testing the external limitation version of starting small. This may be partly because of the widespread belief that the language input that children receive is not substantially different from that of adults. However, as Rohde and Plaut (in press) point out, there is evidence that child-directed speech tends to consist of shorter utterances and less complex sentences than adult-directed speech (e.g., Pine, 1994). Therefore, it may be feasible that starting with simplified input grants a learning advantage in language and other domains.

Elman (1993) provided a test of this version of starting small using neural network simulations. In an incremental input condition, Elman organized the network's input so that it was exposed first only to simple sequences; complex sequences were then introduced to the network gradually. When trained in this way, the networks showed a learning advantage[1].

A recent study with human participants also supports the validity of an external constraints view of starting small (Kersten & Earles, 2001). Adults were exposed to an artificial language consisting of both auditory nonsense sentences and visual, animated events. Some of the participants were exposed to a staged input regiment, in which they received input in three phases: first only single words were presented along with the animated events, then sentences composed of two words, then finally three-word sentences. These participants fared better on tests of their understanding of the language compared to participants who were exposed to a non-staged input presentation. Though Kersten and Earles (2001) view this demonstration as supporting the notion of internal limitations providing a starting small advantage, we agree with Rohde and Plaut (in press) that this conclusion may not be warranted. Instead, we view this data as showing the possible benefits of using a staged input training scheme.

In closing, we note three crucial observations. First, the study by Kersten and Earles (2001), though it may not be an accurate depiction of children's language acquisition, does suggest that staged input may confer a learning advantage. Second, we observe that most of the "successful" tests of starting small have incorporated visual input (e.g., Cochran et al., 1999; Kareev et al., 1997; Kersten & Earlest, 2001) while most of the evidence refuting starting small has relied on auditory input (e.g., Ludden & Gupta, 2000). Third, the input structures that have been used to date in tests of starting small have been relatively simple. However, people are able to learn structures of greater complexity, such as that found in recursion. It is possible that in these more complex learning situations, the effect of starting small may be more pronounced. Based on these three observations, we explore starting small using a staged input scheme, examining both visual and auditory input that is recursively-structured.

Recursive Artificial Grammars

A recursive, grammatical construction is one that is defined by self-reference. Different types of recursion can be found across a variety of linguistic structures. As the amount of self-referencing increases within a recursive construction, the amount of embedding increases. Consider these grammatical English noun-phrases.
a) *The dog* [*on the sidewalk*].
b) *The dog* [*on the sidewalk*] [*near the tree*].
c) *The dog* [*on the sidewalk*] [*near the tree*] [*by the house*].
The above sentences involve *right-branching* recur-

---
[1]However, it should be noted that Rohde and Plaut's (1999; in press) simulations appear to contradict these findings.

sion. In this case, new prepositional phrases can be recursively added onto the right end, creating sentences of potentially infinite length. Sentence (a) comprises 0 levels of embedding, (b), 1 level of embedding, and (c), 2 levels of embedding.

Increased levels of embedding result in slightly decreased comprehensibility of English right-branching recursive sentences. Decreases in comprehension are even larger for a second type of recursive structure: *center-embedding* (e.g., Bach, Brown, & Marslen-Wilson, 1986). Center-embedded recursion grows a sequence by embedding new material in the center. For example, consider:

d. *The boy likes the dog.*
e. *The boy [the girl loves] likes the dog.*
f. *The boy [the girl [the woman [the man adores] admires] loves] likes the dog.*

Sentence (d) is relatively easy to understand, (e) is harder, but (f) is almost impossible to comprehend.

The difficulty of comprehending and producing deeply center-embedded constructions is well documented (e.g., Bach, et al., 1986). English speakers rarely include them in written or spoken language, despite their conformance to the formal grammatical rules of English. Center-embedded recursive structures might be difficult to comprehend because of the need to learn relationships between non-adjacent elements. With greater levels of embedding, memory is taxed, which may hinder comprehension and learning.

Christiansen and Chater (1999) simulated learning of recursive grammars with connectionist networks. The networks produced fewer errors for right-branching items than for center-embedded items, and produced more errors for items with increased embedding. Ellefson, Young, Christiansen, and Espy (submitted) found that human subjects in an artificial grammar learning task had more difficulty learning center-embedded recursive constructions compared to learning right-branching recursive constructions. Moreover, the presence of structural cues (e.g., spatial orientation, brackets) increased participants' understanding of the underlying structure of recursive visual sequences. The presence of these cues may have highlighted the relation among items in the grammars.

Here, we explore the possibility that starting small may facilitate learning of recursive constructions by focusing learners' attention on the relationship between elements (e.g., as in the number agreement relationship between nouns and verbs). Once this relationship has been learned for simple constructions it can then be generalized to more complex constructions. Thus the purpose of this study was to examine the relative usefulness of starting small when learning recursive structure across two different modalities: vision and audition.

Experiment 1: Visual Learning of Center-Embedded Recursive Structure

In the first experiment, we generated stimuli from an artificial grammar having center-embedded recursion. We created two separate training conditions. In the starting small condition, participants were exposed to three training phases. In the first phase, the input was composed of sentences with 0-level center-embedding. The second phase incorporated sentences with 1-level embedding, while the third phase used sentences with 2-level embedding. In this way, the input "started small" and progressively became more complex.

In the second training condition, participants received the same input though in random order. A third experimental condition received no input; rather, these control participants took part in the testing phase only. The starting small theory predicts that the starting small input group will learn the recursive structure better than both the random and control groups.

Method

Subjects Twenty-four undergraduate subjects were recruited from Psychology classes at Cornell University, earning extra credit.

Materials The stimuli were letter sequences generated from the same artificial grammar as Ellefson et al. (submitted). The sequences were based on the repetition of noun-verb pairs within a recursive structure, in which arbitrary letters designated plural and singular nouns and verbs. As with English sentences, nouns and verbs were paired with each other according to plurality. For example, the singular noun *cat* might be paired with the singular verb *plays*, but it would not be paired with the plural verb *play*. Likewise, the plural noun *cats* would be paired with *play* but not *plays*. In our experiment, the letter S falls in the plural noun category and the letter T falls in the plural verb category. Therefore, S and T comprised a possible singular noun-verb pair. Likewise, the letters M and Z comprise a possible singular noun-verb letter pair.

Twelve consonants, C, Q, M, P, X, S, W, Z, K, H, T, and L represented the singular and plural nouns and verbs. There were three letters assigned to each of the letter roles of singular noun, plural noun, singular verb, and plural verb[2]. The sequences contained 0-, 1-, and 2-level embeddings. To increase embedding, additional noun-verb pairs were added to the sequence, inserted into the middle of the center-embedded sequences to achieve higher levels of embedding. An example of a 0-level sequence is

---

[2]Singular nouns: C, Q, and M. Plural nouns: P, X, and S. Singular verbs: W, Z, and K. Plural verbs: H, T, and V.

CW, a 1-level sequence is CPTW, and a 2-level sequence is CPQMTW.

Unique sequences were created for the training and testing sessions. Fifty sequences comprised the training session. Of these 50 training sequences, 10 were 0- level embedding, 20 were 1-level embedding, and 20 were 2-level embedding.

Fifty sequences comprised the test session. Of these testing sequences, 25 were generated from the same grammar as the training sequences (Grammatical) and 25 did not follow the grammar (Ungrammatical). Ungrammatical sequences were created by changing one letter of a grammatical test sequence. The substituted letter was one that was of the proper noun-verb category but with an incorrect plurality. The positions in which the substituted letters occurred in the sequences were distributed evenly across all items. The test session comprised 16 sequences of 0-level embedding, 16 of 1-level embedding, and 18 of 2-level embedding, with each level of embedding having half grammatical and half ungrammatical structures.

Procedure The experiments were run using the E-Prime presentation software with stimuli presented on a computer monitor. Participants were randomly assigned to one of three conditions: starting small, random, or control. The starting small and random subjects were instructed that they were participating in a memory experiment. They were told that in the first part of the experiment they would see sequences of letters displayed on the screen and that they would be tested later on what they observed. Each sequence in its entirely was presented individually, for a duration of four seconds each. Each of the 50 training items was presented 3 times, for a total of 150 input exposures. The starting small participants received the input staged by level embedding. The 0-level sequences were presented first; next the 1-level sequences, and last the 2-level sequences. Sequences were randomized within levels. The random group received all the sequences across all levels of embedding intermixed with one another, in random order. Thus, both the starting small and the random groups received the same training input but in different orders of presentation. The control group participants did not take part in the training phase.

After the training phase, the starting small and random participants were then told that the items they had just seen had been generated by a complex set of rules which determined the order of the letters. They were instructed that they would now see new letter strings, some of which followed the rules of the grammar, and some of which did not. Their task was to classify whether each letter string followed the same rules as the training sequences or not, by pressing a button marked "YES" or "NO". Both the starting small and random groups received the same test instructions and the same set of 50 test items were presented in random order for each participant. The control group participants received the test phase only.

Results and Discussion

The mean number of correct endorsements on the 50 test items was 31.5 (63.0%) for the starting small group, 26.4 (52.8%) for the non-starting small group, and 26.5 (53.0%) for the no-training control group. We conducted single group t-tests and found that only the starting small group performed significantly above chance levels ($t(7) = 4.08; p < 0.005$). We also compared performance between each of the three groups. The starting small group performed significantly better than both the non-starting small group ($t(7) = 2.88; p < 0.05$) and the control group ($t(7) = 3.88; p < 0.05$).

These analyses show that only when the input was presented in a staged fashion were participants able to successfully learn aspects of the recursive structure of the artificial grammar. Crucially, the starting small group out-performed the non-starting small group, lending empirical support to the starting small hypothesis.

Experiment 2: Visual Learning of Right-Branching Recursive Structure

Experiment 1 demonstrated that staging center-embedded, visual recursive input can lead to better learning compared to when the same input is presented non-incrementally. In Experiment 2, we test whether a similar result occurs for the learning of another type of recursive structure: right-branching.

Method

Subjects Fourteen undergraduate subjects were recruited from Psychology classes at Cornell University, earning extra credit.

Materials The materials were identical to those in Experiment 1 except that the sequences were formed using a right-branching structure. That is, additional embedding was increased by adding new noun-verb pairs to the end of a sequence. For example, CW is a 0-level sequence, CWPT is a 1-level sequence, and CWPTQM is a 2-level sequence.

Procedure The procedure was identical to Experiment 1, though a control group was not included as the results from the previous experiment indicated that such a control was not necessary.

Results and Discussion

The mean number of correct endorsements on the 50 test items was 35.0 (70.0%) for the starting small group and 27.43 (54.9%) for the non-starting small group. Only the starting small group performed significantly above chance levels ($t(6) = 6.99; p <$

0.001). The starting small group also performed significantly better than the non-starting small group ($t(6) = 3.86; p < 0.01$).

Like Experiment 1, these results show that only the starting small group was able to successfully learn aspects of the recursive structure of the artificial grammar. Besides serving as a replication of the general effect of starting small with staged input, it extends its applicability to right-branching structures.

Experiment 3: Auditory Learning of Center-Embedded Recursive Structure

The first two experiments reveal that starting small is applicable for visual recursive input. Next we investigate whether starting small also extends to the auditory domain, using the same center-embedded sequences from Experiment 1.

Method

Subjects Eighteen undergraduate subjects were recruited from introductory Psychology classes at Cornell University. Subjects earned extra credit for their participation.

Materials The same center-embedded input sequences were used from Experiment 1 except that each letter was mapped onto a consonant-vowel-consonant syllable: C = "biff"; Q = "rud"; M = "sig"; P = "vot"; X = "mib"; S = "jux"; W = "nep"; Z = "dak"; K = "tood"; H = "jic"; T = "cav"; L = "dup". An example of a 0-level sequence is "biff-nep", a 1-level sequence is "biff-vot-cav-nep", and a 2-level sequence is "biff-vot-rud-sig-cav-nep".

Auditory sequences were generated using the Festival speech synthesizer, which converts written text to synthesized speech (Black, Taylor, & Caley, 1998).

Procedure The procedure was the same as the previous experiments except that the auditory sequences were presented over headphones at a sound level of 70 dB.

Results and Discussion

The mean number of correct endorsements on the 50 test items was 26.4 (52.8%) for the starting small group and 26.1 (52.2%) for the non-starting small group. Neither group performed significantly better than chance levels ($p's > 0.1$). In addition, there was not a difference in performance between the two groups ($t(8) = -0.23; p = 0.82$). Thus, the learning of auditory, center-embedded structure does not appear to be facilitated with a staged input scheme, at least not with the present stimuli and experimental design.


Figure 1: Test performance for starting small (shaded) and random (unshaded) conditions in Experiments 1, 2 and 3.

General Discussion

These three experiments provide insight as to when less is more and when less is less. Experiments 1 and 2 revealed an advantage for participants learning visual, recursively structured input. When the input was staged in an incremental fashion, learners performed better at classifying novel input than did participants exposed to the same input in a non-staged fashion. Experiment 3 showed that at least for the present stimuli and procedure, there was no effect of starting small for auditory, center-embedded structure. The data is summarized in Figure 1.

Under what conditions is less more? We suggest that there may be multiple factors that determine whether there will be a learning advantage, including: whether starting small is implemented in terms of external or internal limitations; which sensory modality receives the input; and the input's level of complexity. Our results reveal that starting small may be most advantageous when the input is staged incrementally, is presented in a visual, spatially-distributed fashion, and is relatively complex (i.e., recursive).

One might wonder why an effect of starting small would be found for visual but not auditory input. One possibility is that modality constraints exist that differentially affect learning in each sensory modality. Such modality differences in sequential learning tasks have been previously observed (Conway & Christiansen, 2002; Saffran, 2002). These studies suggest that humans may be better at encoding and processing serial input when it is presented in the auditory rather than the visual modality; in addition, visual learning occurs best for spatially-distributed input. The current results add to this picture by showing that under certain conditions,

learning in the visual and auditory modalities is heavily constrained. Specifically, auditory center-embedded recursive structure was difficult for participants to learn, even when the input was staged in an incremental fashion. This likely is due to the memory constraints imposed on learning non-adjacent relationships between elements separated over long distances. On the other hand, when the recursive input was presented in a spatially-distributed, visual manner (i.e., in Experiments 1 and 2), learning occurred–but only if that input was staged incrementally.

Interestingly, center-embedded constructions are infrequent in the world's languages. However, right-branching structure is fairly common, most likely because learning this type of structure does not require the learning of long-distance, non-adjacent dependencies. Though we found no effect of starting small for center-embedded auditory structure, it is possible that other types of recursive auditory structure may be ammenable to starting small. A new experiment is currently underway to explore this idea further by testing whether staged input may benefit the learning of auditory right-branching structure. The outcome of this experiment will further elucidate the nature of starting small, helping us to determine whether the crucial factor is one of sensory modality or of non-adjacent learning.

Conclusion

Whether less is more or less is less appears to depend on a number of factors. We have found that an incrementally-staged training scheme improves the learning of visual, recursive, spatially-distributed input. It remains to be seen whether learning in the auditory modality benefits from staged input presentation, as well. For center-embedded structures at least, this does not appear to be the case. Although the lack of a starting small effect for auditory stimuli (Experiment 3) suggests starting small may not play a major role in spoken language acquisition, it is also likely that starting small can be considered as one cue out of many used in the service of learning language. If that is the case, then starting small may have a more noticeable effect when it is combined and integrated with other cues (Christiansen & Dale, 2001).

The present results also may point to differences in the way that spoken vs. visual-based languages (such as ASL) are acquired. If starting small aids visual learning, as Experiments 1 and 2 show, then the learning of complex structure in sign languages may also benefit from a staged input training regiment. Future experiments may help verify this hypothesis, as well as uncover what role starting small plays in spoken language acquisition and other complex learning domains.

References

Bach, E., Brown, C., & Marslen-Wilson, W. (1986). Crossed and nested dependencies in German and Dutch: A psycholinguistic study. *Language and Cognitive Processes*, *1*, 249-262.

Black, A., Taylor, P., & Caley, R. (1998). *The Festival speech synthesis system* [On-line]. Available: http://festvox.org/festival

Christiansen, M. H., & Chater, N. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science. 23*, 157-205.

Christiansen, M.H. & Dale, R.A.C. (2001). Integrating distributional, prosodic and phonological information in a connectionist model of language acquisition. In *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 220-225). Mahwah, NJ: Lawrence Erlbaum.

Cochran, B.P., McDonald, J.L., & Parault, S.J. (1999). Too smart for their own good: The disadvantage of a superior processing capacity for adult language learners. *Journal of Memory and Language*, *41*, 30-58.

Conway, C.M. & Christiansen, M.H. (2002). Sequential learning by touch, vision, and audition. In *Proceedings of the 24th Annual Conference of the Cognitive Science Society* (pp. 220-225). Mahwah, N.J.; Lawrence Erlbaum.

Ellefson, M. R., Young., M. E., Christiansen, M. H., & Espy, K. A. (submitted). Toward a Better Understanding of School Learning Using a Sequential Learning Task. Submitted manuscript.

Elman, J.L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, *48*, 71–99.

Fletcher, J., Maybery, M.T., & Bennett, S. (2000). Implicit learning differences: A question of developmental level? *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *26*, 246-252.

Kareev, Y., Lieberman, I., & Lev, M. (1997). Through a narrow window: Sample size and the perception of correlation. *Journal of Experimental Psychology: General*, *126*, 278-287.

Kersten, A.W. & Earles, J.L. (2001). Less really is more for adults learning a miniature artificial language. *Journal of Memory and Language*, *44*, 25-273.

Ludden, D. & Gupta, P. (2000). Zen in the art of language acquisition: Statistical learning and the less is more hypothesis. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (pp. 812-817), Mahwah, NJ: Lawrence Erlbaum Associates.

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

Pine, J.M. (1994). The language of primary caregivers. In C. Gallaway & B.J. Richards (Eds.), *Input and interaction in language acquisition* (pp. 109-149)

Rohde, D.L.T. & Plaut, D.C. (1999). Language acquisition in the absence of explicit negative evidence: How important is starting small? *Cognition*, *72*, 67-109.

Rohde, D.L.T. & Plaut, D.C. (in press). Less is less in language acquisition. To appear in Quinlin, P. (Ed.), *Connectionist modeling of cognitive development.* Hove, UK: Psychology Press.

Saffran, J.R. (2002). Constraints on statistical language learning. *Journal of Memory and Language*, *47*, 172-196.