Learning to read with a machine teacher: Discovering efficient procedures for training the orth-to-phon relationships in English

Christopher Cox[1], Matthew Cooper Borkenhagen[2], and Mark Seidenberg[2]

1. Department of Psychology, Louisiana State University
2. Department of Psychology, University of Wisconsin–Madison

Introduction

With limited time and instructional resources, what should a child be taught to maximize their chance of achieving basic reading skills? This aspect of curriculum design can be cast as an optimization problem, to discover sequences of letter-to-sound correspondences that enable the greatest number of English words to be read aloud with the least amount of training. We posit a cognitive model of a **Learner** based on a connectionist architecture of reading development [1] and attempt to directly solve this optimization problem using **Machine Teaching** [2]. This generalizable and scalable approach permits novel investigations of the structured relationships that support reading aloud by exploring idealized learning trajectories of a cognitively plausible Learner.

Machine Teaching

While machine *learning* discovers a solution to a problem, such as the relationship between orthography and phonology, machine *teaching* discovers the best set or sequence of training examples to present a learner so that it solves the problem *efficiently*.

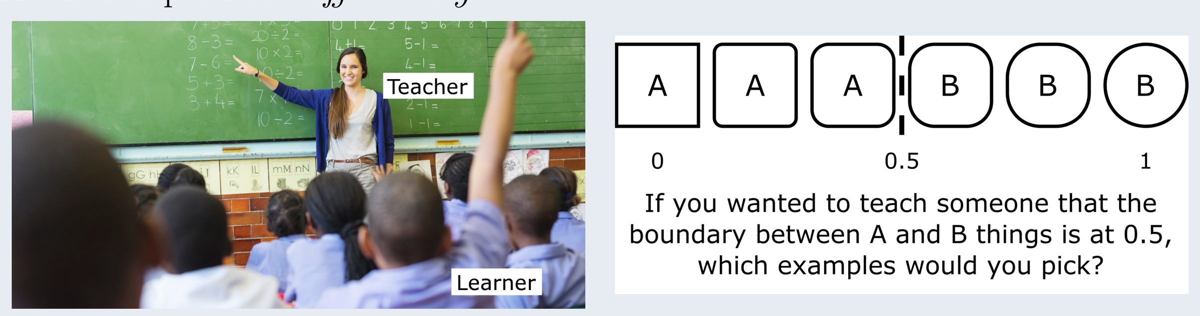

Cognitive Model of a Learner

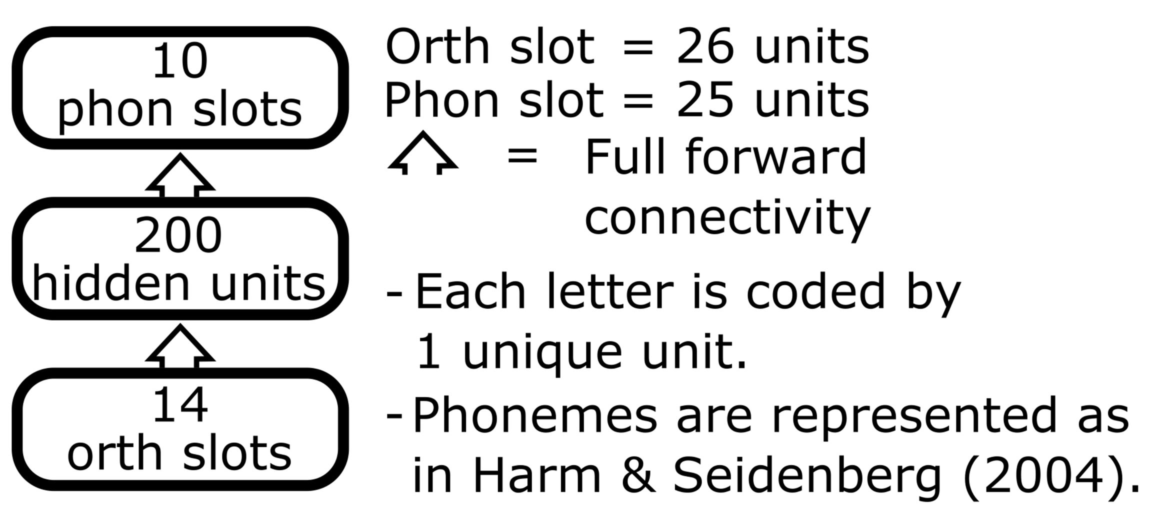

Figure 1: All learners are based on the same model architecture.

A learner in the "best of all possible worlds"

We designed a machine teacher to discover a sequence of training examples, sampled with replacement from a predefined pool, such that each subsequent example achieves the maximum possible reduction in error the test set for a given learner. We call this machine teacher **Candide**.

Discover training sequences with a Machine Teacher

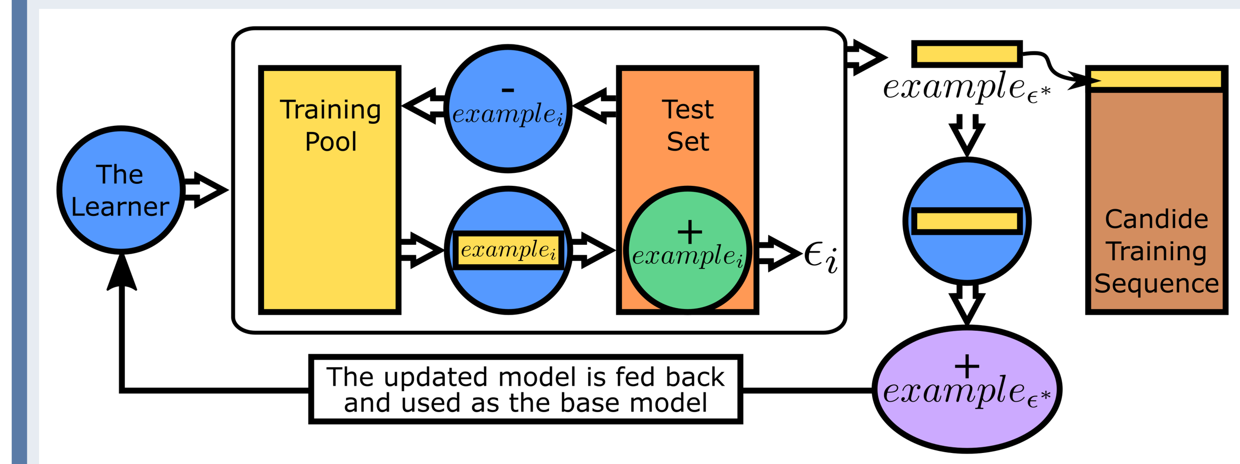

Figure 2: The *Candide Machine Teacher* evaluates all items in the training pool before picking the next single training example. The chosen example will result in the largest error reduction on the full test set. In our experiments, the test set is used as the training pool. Given an initial model state (blue learner at left), we loop over the training pool (yellow box, inner loop) which provides one example at a time to the learner (blue circle w/ yellow bar). This produces a temporary model (green circle), which is then evaluated on the whole test set (orange box). The error is recorded, and the learner is reverted to the initial (blue) state before re-entering the loop with a new example from the training pool. Once exhausting the training pool, the example associated with least error on the test set is added to the sequence (brown box, far right) and provided to the learner as a true training experience. The learner carries this experience into the next pass through the loop with Candide.

Learning Environment

The full corpus has **2882 monosyllabic words**. Sequences were generated based on a training pool of 1000 random words. The remaining 1881 words were assigned to the **Eval Set**. When learning sequences using the Candide machine teacher, the training pool was used as the test set.

Orthographic input patterns and phonological output patterns were shifted so that the first vowel in each word appears in the same slot.

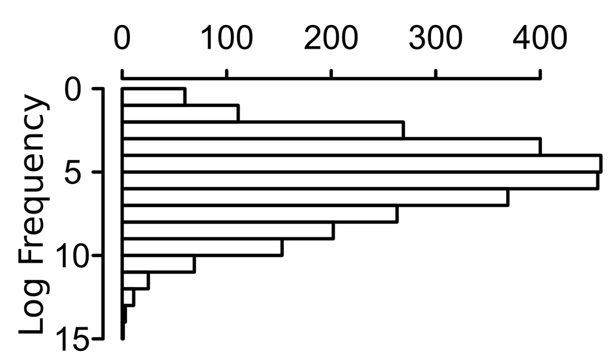

letters	2	3	4	5	6	7	8
n	5	440	1249	937	225	23	2

Experiments

We trained **15 Learners** (Figure 1) to read **1000 random words** (see "Learning Environment"). Each learner was initialized with random weights between layers, but otherwise were identical. Learners updated their knowledge about the environment after each training example (batch size = 1). The learners were randomly assigned to **3 curriculum conditions**:

1. Random sampling with replacement. The same example might be sampled consecutively.
2. Permutation sampling. An example can only be repeated after the training pool is exhausted, but each pass is a random sequence.
3. Machine taught with a sequence generated as depicted in Figure 2 in a prior step. The length of these 5 sequences was determined by allowing the Candide optimization to run for a fixed amount of time (not a fixed number of iterations), which leads to different sequence lengths.

The 15 unique, pre-defined, fixed-length sequences were based on 10 training pools, such that a random, permuted, and machine taught sequence would all be based on the same pool.

The trained models were evaluated on the **1881 items** that were not in the training pool (the **Eval Set**).

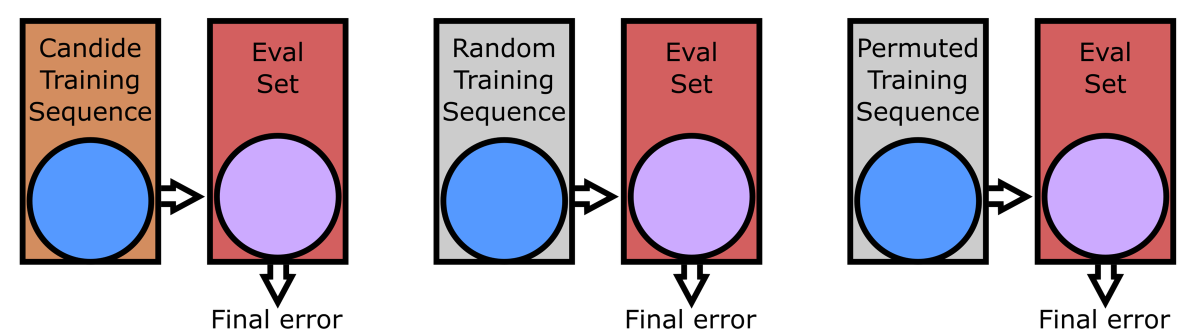

Results

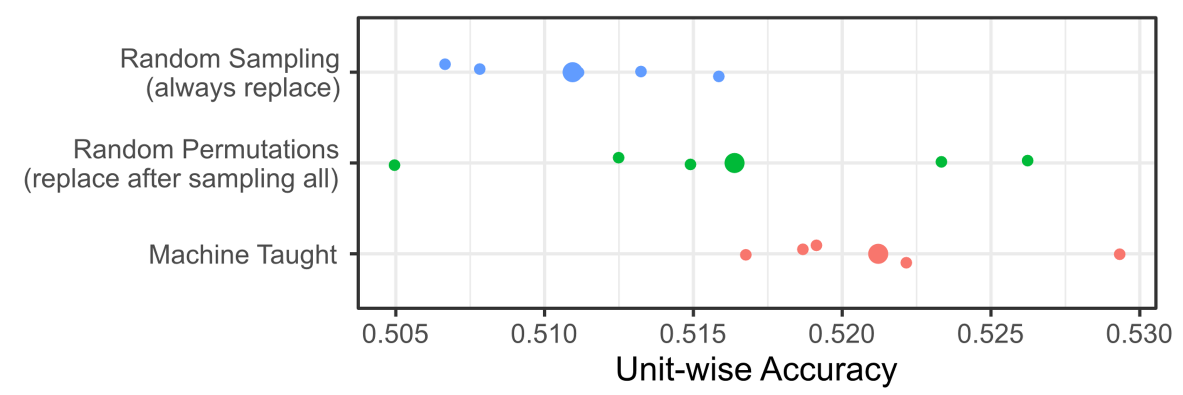

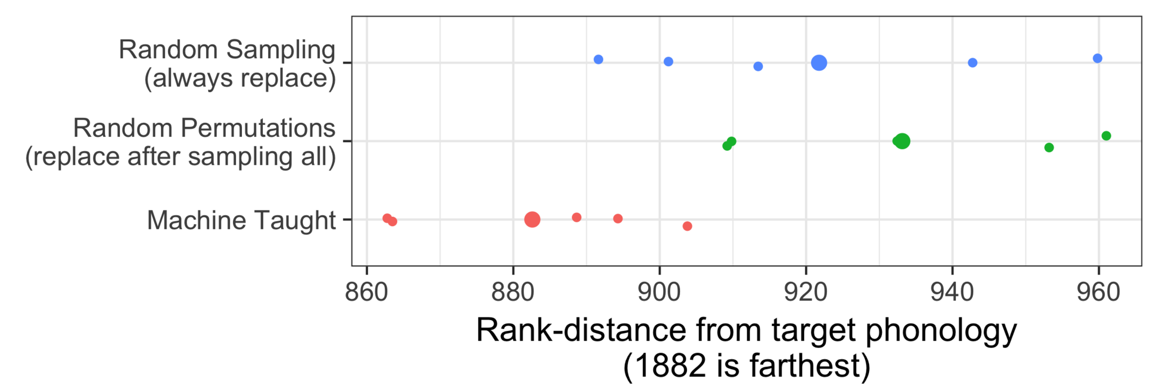

Why is Candide a successful teacher?

Investigations into what structure Candide is exploring at these early learning stages is ongoing. However, there seems to be a systematic emphasis on syllable structure. Considering **30 repetitions** of the Candide machine teaching optimization, the following figure shows the 25 most common words among the first 50 taught. The number indicates the number of repetitions where the word appeared in the first 50.

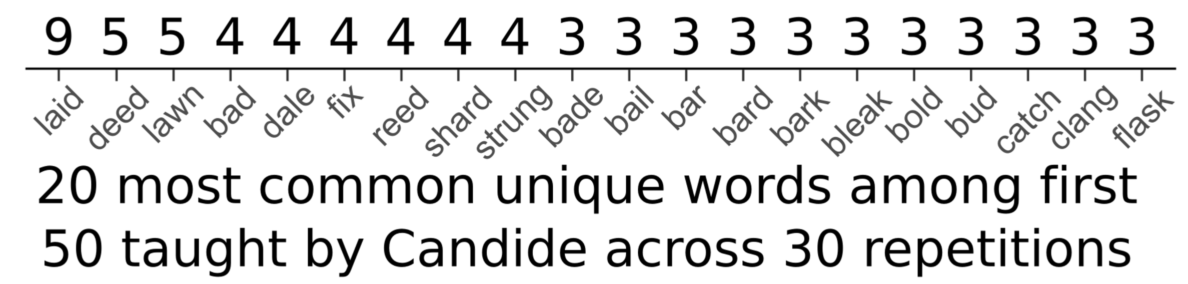

Conclusion

Using **Machine Teaching**, we identified sequences that promote generalizable word-reading as it emerges in a simulated **Learner** better than random and permuted sequences of the same length. It is particularly interesting that Learners generalize better after exposure to a biased sequence obtained through Machine Teaching relative to a permuted sequence that promotes maximal exposure to the training pool. The implication is that highly-efficient curricula can be designed in a data-driven way, which achieve generalizable reading skills with less classroom time.

References

[1] Mark S Seidenberg and James L McClelland. "A distributed, developmental model of word recognition and naming.". In: *Psychological review* 96.4 (1989), p. 523.

[2] Xiaojin Zhu. "Machine Teaching: An Inverse Problem to Machine Learning and an Approach Toward Optimal Education.". In: *AAAI*. 2015, pp. 4083–4087.

Acknowledgements

This work has been supported by the Vilas Trust at UW-Madison and by the Institute of Education Sciences, US Department of Education, through Award R305B150003 to UW-Madison. The opinions expressed are those of the authors and do not represent views of the US Department of Education.

Contact Information

- Web: http://lcnl.wisc.edu/
- Email: chriscox@lsu.edu, cooperborken@wisc.edu