Simultaneous Training on Overlapping Grapheme Phoneme Correspondences Augments Learning and Retention

Tanja C. Roembke
Dept. of Psychological and Brain Sciences
University of Iowa
Tanja-roembke@uiowa.edu

Michael V. Freedberg
Henry M. Jackson Foundation for the Advancement of Military Medicine
Michael.freedberg@nih.gov

Eliot Hazeltine
Dept. of Psychological and Brain Sciences
University of Iowa
Eliot-hazeltine@uiowa.edu

And

Bob McMurray
Dept. of Psychological and Brain Sciences
Dept. of Communication Sciences and Disorders
Dept. of Linguistics
University of Iowa
Bob-mcmurray@uiowa.edu

Running Head: Simultaneous Training on Overlapping Spelling Regularities

Corresponding Author

Tanja Roembke
W311 SSH
Dept. of Psychological and Brain Sciences
University of Iowa
Iowa City, IA 52242
319-335-0692 (voice)
319-335-0191 (fax)
tanja-roembke@uiowa.edu

Abstract

An important component of learning to read is the acquisition of letter-to-sound-mappings. The sheer quantity of mappings and many exceptions suggests that children may use a form of statistical learning to acquire them. However, while statistical models of reading are item-based, reading instruction typically focuses on rule-based approaches involving small sets of regularities. This discrepancy poses the question of how different groupings of regularities, an unexamined factor of most reading curricula, may impact learning. Exploring the interplay between item statistics and rules, this study investigated how consonant variability, an item-level factor, and the degree of overlap among the to-be-trained vowel strings, a group-level factor, influence learning. First graders ($N = 361$) were randomly assigned to be trained on vowel sets with high overlap (e.g., *EA, AI*) or low overlap (*EE, AI*); this was crossed with a manipulation of consonant frame variability. While high vowel overlap led to poorer initial performance, it resulted in more learning when tested immediately and after a two-week-delay. There was little beneficial effect of consonant variability. These findings suggest that online letter/sound processing impacts how new knowledge is integrated into existing information. Moreover, they suggest that vowel overlap should be considered when designing reading curricula.

Keywords: Reading, statistical learning, vowel overlap, variability

Simultaneous Training on Overlapping Grapheme Phoneme Correspondences Augments Learning and Retention

1: Introduction

1.1 The Complexity in Learning to Read

Literacy is correlated with academic success (National Reading Panel, 2000) and life outcomes (e.g., Beck, McKeown, & Kucan, 2002). However, 60% of students lack basic reading skills (U.S. Department of Education, National Center for Education Statistics, 2010). This may be in part due to the complexity of the problem children must solve: Computing the meaning of a written sentence involves skilled word recognition, use of text conventions, and sentence processing abilities including semantics and syntax. Even the most elemental component of reading—reading isolated words—can be complex, given the irregularities in the sound-spelling-system of many languages and the large number of lexical candidates to sort through.

The relationship between spelling and meaning is arbitrary (words that are spelled similarly do not have similar meanings), and as a result acquiring the ability to directly recognize words would appear to require learners to memorize tens of thousands of words. Consequently, during the initial stages of reading acquisition many phonics-based curricula teach children to map letters to sounds, where there are more consistent, generalizable regularities (so-called grapheme-phoneme-correspondence [GPC] regularities). Once these regularities are mastered, a child can access the sound of even an unfamiliar word from its spelling and subsequently access meaning. Thus, one critical aspect of learning to read is the acquisition of GPC regularities.¹

There has been growing interest in applying principles and models of learning from the

¹ We use this term only to describe regularities in the orthographic/phonological system of the language. We do not use this term to imply a specific cognitive representation, for example explicit rules as in models like Coltheart et al. (2001). Such regularities in the world are also expected to be instantiated in a distributed manner in the weights of connectionist models like Harm and Seidenberg (1999).

cognitive sciences to applied problems in education (Apfelbaum, Hazeltine, & McMurray, 2013; Kellman, Massey, & Son, 2010; McDaniel, Anderson, Derbish, & Morrisette, 2007; Rohrer, Dedrick, & Stershic, 2015; Seidenberg, 2016). The present study builds on this effort to examine the acquisition of GPC mappings to that support decoding skills. We examine a disconnect between statistical learning theories of reading acquisition that emphasize exposure to large number of items and common classroom approaches based on teaching rules.

1.2 Statistical Learning

In English and other orthographically deep languages, GPC regularities, specifying for example that *EA* makes the /i/ sound, are not straightforward: They are only quasi-regular. There are many exceptions (e.g., *steak*) and sub-regularities (e.g., *head*, *dead*) (Gates & Yale, 2011; Johnston, 2001; Seidenberg, 2005). For example, *EA* makes the /i/ sound in 127 of 164 monosyllabic words. However, it makes the /ε/ sound in 30 (e.g., *dead*), and the /ei/ sound in 7 (e.g., *steak*). Thus, during early literacy acquisition, even the process of reading a single word can be challenging.

This quasi-regularity has motivated the idea that the mappings are learned via statistical learning mechanisms (Harm, Seidenberg, Macdonald, Thornton, & Zevin, 2004; Plaut, McClelland, & Seidenberg, 1995; Plaut, McClelland, Seidenberg, & Patterson, 1996), which track subtle predictive relationships among letters and between letters and sounds. There is mounting empirical evidence for this: Studies document a correlation between statistical learning ability (in visual and auditory domains) and reading outcomes such as word recognition and phonological processing in both children (Arciuli & Simpson, 2012; Spencer, Kaschak, Jones, & Lonigan, 2015) and adults (Arciuli & Simpson, 2012). Statistical learning is also impaired in dyslexia (Gabay, Thiessen, & Holt, 2015). There is also evidence that the pattern of children's

errors in individual items conforms to the statistical structure of the input: Spelling mistakes in early spellers are more likely to contain letters that children frequently encounter (e.g., the letters of their names) (Pollo, Kessler, & Treiman, 2009, see also Treiman, Gordon, Boada, Peterson, and Pennington, 2014), and more advanced spellers make errors that reflect the statistics of the texts they have been exposed to (Thompson, Fletcher-Flinn, & Cottrell, 1999). Moreover, the way children assign stress to syllables is predicted by the likelihood of those syllables being stressed in age-typical texts (Arciuli, Monaghan, & Ševa, 2010). Together, these data suggest that the acquisition of sound/spelling correspondence may derive from statistical learning.

In such models, learning is fundamentally item-based and GPC regularities are not explicitly stored: Learning a regularity (e.g., that *EA* generally makes the /i/ sound) is conceptualized as emergent from experience with individual items (words). In this case, an individual regularity is not what is learned, but a system of mappings that “works” across all the words that are encountered. Supporting this, Armstrong et al. (2017) showed that after adults are taught a small set of nonwords with irregular pronunciations, they spontaneously generalize this to a new set of untrained nonwords (Kim, Pitt, & Myung, 2013).

If the goal is to harness a learning theory to improve reading instruction, statistical learning suggests the focus should be on manipulations of statistics at the item-level (e.g., the set of words used during training) to achieve better learning of the regularities. That is, a better understanding of the detailed statistics that children attend to could yield principles that might be used to curate sets of items that lead to more robust learning (Armstrong et al., 2017; Rayner, Foorman, Perfetti, Pesetsky, & Seidenberg, 2001).

Apfelbaum, Hazeltine, and McMurray (2013) offer an example of this approach. They examined the acquisition of GPC regularities for vowels, which are most challenging for early

readers as they are less regular than those for consonants (Fowler, Liberman, & Shankweiler, 1977; Näslund, 1999; Shankweiler & Liberman, 1972). A critical problem in statistical learning is determining which statistics to attend to. In this case, the child must learn that the vowel letters are most predictive of the (vowel) pronunciation, and not the consonants. As a result, training on a set of items (words) with similar consonants could impair learning vowel GPCs as children encode the spurious statistical co-occurrence of the target vowel with neighboring consonants, potentially hindering generalization (see also, Juel & Roper-Schneider, 1985, for an analogous discussion in reading). For example, after exposure to *hat*, *cat* and *mat*, the child could learn that the /æ/ sound is predicted by *AT*, not by *A* in a CVC context. The solution to this problem is to train on items that are more variable on irrelevant dimensions (Gómez, 2002; Lively, Logan, & Pisoni, 1993; Rost & McMurray, 2009, 2010).

To test this, Apfelbaum et al. (2013) trained children on six GPC regularities (three digraphs and three short vowels) using a computer-based reading intervention that implemented a variety of sound/spelling tasks (Foundations in Learning, 2010). Target vowels were embedded in training word-lists that were either highly similar across words (e.g., *met*, *pet*, *moat*, *pot*) or variable (e.g., *met*, *leg*, *loaf*, *top*). Children assigned to the variable condition showed larger gains at post-test both for trained and novel words. This supports the idea that manipulations at the item-level—when motivated by a clear learning theory—can benefit instruction.

1.3 Explicit Teaching of Regularities

Statistical learning promotes a focus on items, with regularities as emergent properties. Based on this account, the optimal training set for instruction should mirror the distribution of letters in the language (e.g., in the extreme, all possible items). In contrast, *reading instruction* is typically structured around *regularities*, not items, and typically emphasizes a small set of them

at any one time (e.g., long vowels that end with silent *E* or digraphs like *CH* and *SH*). This may alter the statistics of the input in unintended ways.

They emphasis on rules offers practical benefits: Explicit rules can be easily described, giving students something to remember to bootstrap learning. However, there are too many GPC regularities for a child to receive instruction on all of them, nor could explicit memory processes retain them all. Thus, curricula often focus on the most regular of the GPC regularities (c.f., Fry, 1964). These will generalize to a large number of words, permitting children to learn independently (Share, 1995). While this strategy is insufficient to cover the full scope of the English GPC system, it gets students off the ground, and the success of phonics curricula speak to the broad utility of this approach (Ehri, Nunes, Stahl, & Willows, 2001; National Reading Panel, 2000; Stuebing, Barth, Cirino, & Fletcher, 2008).

However, even within this approach, to avoid overwhelming children, most curricula focus on training a subset of the regularities at a time (e.g., the short vowels), rather than training on all possible regularities. For example, in many phonics curricula, short vowels (*E*, *A*, *O*, *U*, in CVC consonant frames) are introduced first, followed by long vowels (vowels + silent *E*), and then digraphs (e.g., *EA*, *AI*). This approach is not inconsistent with statistical learning, but few statistical learning studies have approached the problem of learning multiple distinct regularities simultaneously and thus may not offer insight into this problem.

1.4 Structuring Items and Regularities to Promote Learning

Little research has examined the types of grouping of GPC regularities that may benefit learning. From a statistical learning perspective, the specific groupings that are commonly used however may not be effective. For example, the short vowels *E* and *A* have no letters in common, and as a result, practicing them together may not offer any benefits. The hard part of learning

may be discriminating a word which includes an *E* and makes the / ε / sound (like *men*) from one that makes the /i/ sound (like *mean*). Thus, groupings like “short vowels” could isolate the wrong statistics (see Carvalho & Goldstone, 2014, for analogous issues in category learning). As this example illustrates, understanding how learning is impacted by structure at the level of the generalizable regularities (in addition to structure at the item-level) may have consequences for the basic understanding of how children learn to read, as well as implications for structuring curricula.

However, there is little basic science on this notion of grouping rules during instruction. And as we describe below, statistical accounts applied to this problem make two competing predictions. Thus, the goal of the present study is to examine the interplay of statistics and rules in shaping learning. To be clear, this question is not about whether statistical learning or a more explicit, rule-based mechanism are at work (e.g., Coltheart & Rastle, 1994; Plaut et al., 1995), but how groupings at the rule-level (i.e., vowel groupings) interact with statistics at the item-level to drive learning.

1.5 The Present Study

In digraph vowels, letters can combine to represent a phoneme or phonemes (as in *EA*, *AI*, *OI*). Unlike the consonant manipulations of Apfelbaum et al, where the consonants need to be ignored, here, the conjunction of letters is critical, and the challenge is to map the *combination* of letters to a sound that may not be associated with either. For example, *OA* and *EA* both contain an *A*, and of course, *A* can also function in isolation. All three strings however map to different phonemes, making the interpretation of *A* contextually dependent. Most curricula do not systematically consider this type of overlap, instead grouping by conceptually related regularities (e.g., long vowels).

Statistical learning approaches offer (at least) two conflicting possibilities. One interpretation of the Apfelbaum et al. (2013) findings is that learners build simple associations between individual letters and phonemes. This is consistent with straightforward statistical learning (e.g., counting co-occurrences), but also with direct associations between letters and sounds. Under this model, variability impedes the formation of incorrect associations between consonant letters and the vowel phonemes (since they never co-occur sufficiently to build up any associative strength), while similarity in the consonant frames reinforces them (see Apfelbaum & McMurray, 2011 for an analogous model in word learning). This naïve associative account is consistent with how statistical learning is typically characterized, i.e., the blind tracking of items' co-occurrence. It is also similar to learning operations posited in the “visual cue” or “paired associate” stage of reading in classic approaches (Ehri, 1995; Gough & Juel, 1991; Juel & Roper-Schneider, 1985). Under this view, digraphs contain a deviation from a one-to-one-mapping between graphemes and phonemes such that the same letter must be mapped to multiple sounds. As a result, digraphs should be harder to learn. Supporting this, vowels—which are more likely to require overlapping strings like digraphs—are particularly difficult for children (e.g., Read, 1975; Treiman, 1993).

Alternatively, if learning is conceptualized in terms of acquiring a system or mapping (e.g., as in a connectionist model), overlap may be advantageous. Novel information that is consistent with established knowledge or schemas may be more easily integrated with existing information and more resistant to interference than inconsistent information (McClelland, 2013). Interleaving inconsistent information may also help learners identify dimensions that separate categories (Carvalho & Goldstone, 2014). Thus, under these views, simultaneous training on overlapping GPC should benefit learning.

Resolving these competing predictions has important implications for the way curricula are structured to promote learning. We thus taught first grade students a small number of GPC regularities for short vowels and digraphs over several days (following Apfelbaum et al., 2013). We asked if overlap among simultaneously taught vowel GPC regularities promotes or impedes learning. Overlap was manipulated by changing the sets of vowels that students were simultaneously trained on to either maximize shared letters across the set or to minimize them. To relate both item-based and regularity-based factors, we further asked if the role of (irrelevant) consonant variability depends on the overlap among the (relevant) vowels. As in Apfelbaum et al. (2013), consonant variability was manipulated by selecting words to instantiate vowel strings that either had highly variable or similar consonant frames.

Building on the naïve associative account (Apfelbaum & McMurray, 2011; Gough & Juel, 1991; Juel & Roper-Schneider, 1985), we predicted that overlapping vowels may be more difficult to learn, as the same vowel phoneme must be associated with two different letters. By this account, variability in consonant frames may highlight the overlapping vowels as a unit, potentially mitigating the effect of overlap. Alternatively, schema driven accounts like McClelland (2013) suggest that overlap may create conditions in which learning is more robust, helping children to embed the relevant knowledge in their preexisting GPC network. It is unclear whether consonant variability will interact with overlap in this case.

We developed two sets of items that included groups of vowels with either low overlap (*AI, OO, EE, OU, I, U*) or high overlap (*EA, OA, OI, AI, E, O*). These were embedded in word lists with either similar or variable consonants for a 2×2 design. As in Apfelbaum et al. (2013), we repurposed an Internet-based reading intervention program (Foundations in Learning, 2010) to deliver training and testing to first graders in their schools. We tested first graders for several

reasons. First, they match the children tested in Apfelbaum et al. (2013), allowing us to follow-up on the effect of variability. Second, they are still in the process of acquiring these GPC mappings, so they were unlikely to perform at ceiling but simultaneously had sufficient skills to perform the tasks.

Students' baseline ability was first assessed with a pre-test, incorporating a variety of simple tasks tapping six GPC regularities. They then completed training (with feedback on each trial) over approximately one week, followed by post-test. As an extension of Apfelbaum et al. (2013), we also included a third testing point one to two weeks later to assess retention. Given that the GPCs differed between high and low overlap, testing could not be identical between overlap conditions (but were identical for variable/similar conditions within an overlap group). Unlike Apfelbaum et al. (2013), GPCs were fully interleaved across training to improve learning (Carvalho & Goldstone, 2014; Kang & Pashler, 2012; Taylor & Rohrer, 2010; Vlach, Sandhofer, & Kornell, 2008).

2. Method

2.1 Participants

Participants were 316 first-grade students (average age: 6.46 years) from eight elementary schools in the [removed for review]. All first-grade students without an Individualized Education Plan for a learning disability were invited to participate. Approximately 50% of first graders in these eight schools participated, yielding a wide distribution of reading abilities. Of the 316 students enrolled in the study, 281 completed pre-test, training and post-test. The remaining 35 withdrew during pre-test or training, or were excluded based on behavioral issues reported by the research staff during testing. Two children did not complete post-test

Table 1

Demographic breakdown of participants across conditions for the final dataset.

		High Overlap		Low Overlap		Total
		Variable	Similar	Variable	Similar	
Gender	Boys	35	33	33	38	139
	Girls	37	32	38	35	142
First	English	67	57	62	63	249
Language	Other	5	8	9	10	32

immediately after training but one week later; for these, post-test results were treated as retention (and these children were excluded from the post-test analysis). Additionally, post-test data from two children were excluded, as they did not complete post-test until five days after end of training (and they did complete retention shortly after). Their retention data remained in the analysis. Thus, data from 277 children were analyzed at post-test, and from 267 at retention.

The final sample included 32 non-native English speakers, roughly equally distributed across conditions. These children's data were included along with the rest of the data, as English language learner status did not interact with experimental factors in prior similar work (Apfelbaum et al., 2013). Table 1 shows the demographic breakdown of the training groups.

In addition to gender, native language and age, we obtained children's reading assessment scores on the Fountas and Pinnell Benchmark Assessment System, a test of children's reading levels, from the school district. Letter scores (A through Y) were converted to an ordinal scale, ranging from 0 to 23 with an average of 6.74 (SD = 4.08).² This was used as a between-subject factor for moderator analyses (see Supplemental Online Materials [SOM] 1, 2).

2.2 Design

2.2.1 Overview. The experiment included four phases: pre-test, training, post-test and retention. Across phases, eight tasks were used. There were four experimental conditions (high

² A score of 6 or 7 (equivalent to letter scores F and G) roughly maps onto the upper level of what is expected for first graders (Fountas & Pinnell, 2011).

Figure 1: Visual depiction of vowel overlap in A) high overlap and B) low overlap condition.

overlap/low overlap GPCs $\times$ similar/variable consonants) which were crossed with two conditions for counterbalancing.

2.2.2 GPC regularities. Children learned one of two sets of six GPC regularities. These GPC regularities captured common digraphs and monographs of English that have high spelling-sound consistency. Each set included two short vowels (e.g. *E* as in *men* or *O* as in *pot*) and four digraphs (e.g. *EA* as in *mean* or *AI* as in *mail*; Table 2, Figure 1). One vowel (*AI*) was present in both the high overlap and low overlap condition. Short vowels and digraphs were interleaved during training and test. Each GPC regularity represented the dominant pronunciation for the given rule (see Supplement S1).

To create overlap conditions, GPC regularities were grouped to maximize or minimize the number of other regularities sharing at least one letter. Given only five vowel letters in English, it was not possible to construct completely non-overlapping sets or to equate the number of overlapping pairs in the overlap set. On average, each GPC regularity in the high overlap set shared a letter with 2.67 others (range 1-4), where each regularity in the low overlap set shared a letter with 1.0 other regularities (range 0-2). The GPC regularities in each condition were balanced on overall difficulty using procedures outlined in Supplement S1.

2.2.3 Tasks. To ensure that learning and testing tapped a robust ability (rather than the

Table 2
GPC regularities used in this study.

Group	Grapheme	Phoneme	Examples	Percentage Consistent	<i>n</i>	<i>N</i> Overlaps
High Overlap	AI+	eɪ	bait, gain, mail	98.7	87	3
	EA*	i	mean, heat, beam	77.4	164	3
	OA*	ou	coal, goat soap	98.1	53	4
	OI	ɔɪ	foil, coin, soil	100	35	3
	E*	ɛ	red, met, pen	100	272	1
	O*	ɑ	pot, hot, rod	82.4	267	2
Low Overlap	AI+	eɪ	bait, gain, mail	98.7	87	1
	OO*	u	food, roof, pool	76.7	133	1
	EE*	i	feel, feed, teen	100	127	0
	OU	aʊ	hour, noun, loud	54.8 ¹	146	2
	I*	ɪ	lid, hit, tin	93.0	491	1
	U*	ʊ	but, mud, run	96.4	305	1

Note. *N* refers to the number of monosyllables in the MRC psycholinguistic database that include that vowel string; regularity is the proportion of monosyllables that include that string, and are pronounced with the target phoneme. + indicates GPC rule that appears in both groups; *'s indicate GPC regularities for which screener data was available.

¹ While this regularity may seem low, there were 6 total pronunciations of OU found, and the next highest only appeared in 19.2% of the items, so this pronunciation is clearly dominant.

ability to perform a single task), and to keep students engaged, eight tasks were used (Table 3).

Students performed each task in multiple blocks of 10 trials, switching frequently. Tasks were grouped into “cycles”; in each cycle, there was a selection screen allowing students to select a task. After completing the tasks in a cycle, they advanced to the next. This gave students some ability to determine task order (to maintain engagement), while not permitting them to complete all the blocks for one task in a row. Spelling tasks required selecting correct letters to complete a word. Recognition tasks required students to select the correct word or word family.³

Training used five cycles of six tasks; testing used four. Two tasks only appeared in testing to assess generalization across tasks. We counterbalanced which tasks were held out for

³ A word family was defined as a group of words that shares the same vowel and ending sounds; this was consistent with students' experience with the concept of word families. The definition of word families was included in the instructions of the *Find the Word Family* and *Find the Nonword Family* task.

Table 3
Tasks used in this study.

Task	Description	Task Type	Stimuli	Role	
				A	B
Find the Word	Hear a word played and find that word among ten displayed alternatives (“Find the word...”).	Recognition	Word	Gen.	Train
Fill in the Blank (Nonword)	Hear a nonword and see a consonant frame and ten vowel options. Choose which vowel completes the played nonword (“Fill the blank to make the word...”).	Spelling	Nonword	Gen.	Train
Make the Word	Hear a word and choose the letters to spell it from ten displayed alternatives for each position (“Make the word...”).	Spelling	Word	Train	Gen.
Nonword Family	Hear the vowel and coda consonant of a nonword, and find the nonword that contains those sounds among ten alternatives (“Find the nonword family...”).	Recognition	Nonword	Train	Gen.
Word Family	Hear the vowel and coda consonant of a word, and find the word that contains those sounds among ten alternatives (“Find the word family...”).	Recognition	Word		
Change the Word (Vowel)	See a consonant frame and ten vowel options. Asked aurally to change one word to another (“Change the vowel in <i>cat</i> to make <i>coat</i> ”).	Spelling	Word		
Find the Nonword	Hear a word played and find that nonword among ten displayed alternatives (“Find the nonword...”).	Recognition	Nonword		
Change the Nonword (Initial)	See vowel and offset consonant and ten onset consonants. Asked to change one nonword to another (“Change <i>poat</i> to make <i>soat</i> ”).	Spelling	Nonword		

Testing

Training

generalization testing with two sets (A, B). In both sets, the same four tasks were used at test; the only difference was which tasks were included in training. Half of the tasks used real words and the other half nonwords. Nonwords allowed a more thorough manipulation of the overall amount of consonant variability and a more direct test of grapheme → phoneme knowledge, since children cannot perform well by “memorizing” the word.

Each trial included ten response options. Foil responses were drawn from other items within that block, unless this would have led to repetitions (e.g., in task like *Change the Vowel*, where 6 vowel strings would have been insufficient). Here additional vowels were added to create ten unique response options.

2.2.4 Random assignment. Students were randomly assigned to one of eight conditions (overlap $\times$ consonant similarity $\times$ task-set A/B). Assignment was done in two phases. Testing items were controlled by the overlap condition, so this assignment was performed at the onset of the experiment, balancing gender, school, and language status. Within an overlap group, similar and variable conditions had the same testing procedures and items. Thus, assignment to the respective variability condition was performed after pre-test, to balance pre-test scores in addition to demographic factors and task set.

This procedure meant that while pre-test scores were well-matched within a variability conditions, pre-test scores were not matched across overlap condition (Table 4), and our analysis starts by examining this factor in pre-test alone.

2.2.5 Testing. Pre-test, post-test and retention had an identical design. These tests consisted of three cycles of four tasks, resulting in 120 total trials (10 trials/task/cycle). Words were randomly assigned to block in a fixed order for each student, roughly balancing the number of items from each GPC class within a block. Within a block, item order was random. During testing, students did not receive feedback.

Testing trials were evenly divided between words and nonwords, leaving 30 trials (3 cycles) for each task. Within a task, items were evenly divided among the six

Table 4

Percent correct and standard error of the mean for pre-test for all four experimental conditions.

	Variable	Similar
High Overlap	57.01 (2.23)	58.12 (1.81)
Low Overlap	62.46 (2.19)	63.40 (2.26)

GPC regularities. For each regularity there were five types of items tested: 2 trained words, 2 untrained (generalization) words from the same GPC regularity (one with a CVC structure, one with at least one consonant digraph or cluster), and a baseline word that was not from any of the trained GPC regularities. At test, items differed between overlap conditions (to reflect the different GPC regularities) but were identical between the similar and variable conditions.

Testing sessions were run over two days; children were logged out automatically of the program after completing 60 trials (halfway through cycle 2). The following day, they completed the remaining 60 trials. Pre-test generally occurred on Thursday and Friday to allow for the second phase of random assignment over the weekend, though some children missed a day, and pre-test was extended to the following week.

Post-test was conducted on the day immediately after the last day of training (or the following Monday, if training finished on a Friday). Retention was conducted at least a week after post-test (though typically two weeks), depending on the school's schedule.

2.2.6 Training. Training consisted of five cycles of six tasks, for a total of 300 training trials ($10 \text{ trials} \times 6 \text{ tasks} \times 5 \text{ cycles}$). There were an equal number of trials for each GPC regularity, and this was roughly equal within a task. Training items differed as a function of overlap and similarity.

During training, children received feedback on their performance via three mechanisms. First, on each trial they received auditory feedback with a 'bing' for correct answers, and a 'buzz' for incorrect ones. Second, on incorrect trials, students repeated the trial after hearing "Try again!" On these repetitions the child's initial choice was removed from the screen, and for spelling tasks buttons were added to the response letters allowing the child to hear the sound corresponding to each of the responses. If they gave an incorrect response twice, they were told

the correct answer and it was highlighted on the screen. Finally, a score accumulated across trials within each task. If children chose the correct answer on the first attempt, they received ten points. If responded incorrectly the first time, but correctly on the repetition, they received five points; otherwise no points were given.

Participants did not run a fixed number of trials on each training day but ran for a fixed period of 20 minutes. At the end of this period, they were automatically logged out of the system, and would continue where they had left of the following day. Once participants had completed the total number of training trials (e.g., on the last day), they were logged out and were done.

2.3. Items and Stimuli

Items were matched on frequency, imageability (if available) and neighborhood density across conditions (Supplement S2). All words, training instructions, letter cues as well as carrier sentences were recorded by a phonetically trained native speaker of English. She was instructed to talk clearly and slowly. Recordings were conducted in a sound-attenuated room using a Kay Elemetrics CSL 4300b (Kay Elemetrics Corp., Lincoln Park, NJ) at 44,100-Hz sampling rate. For each item, the speaker produced 3-5 exemplars in a neutral carrier phrase (“He said...”). These recordings underwent noise reduction using Audacity (Audacity Team, 2014), to remove background noise. Next, the clearest exemplar of a given item was excised from the phrase. 50msec of silence was then spliced onto the onset and offset of each word, followed by amplitude normalization using Praat (Boersma & Weenink, 2016).

2.4 Procedures

The study was conducted between February and May 2014, with short bursts of activity at each school. Pre-test, training and post-test lasted for about 1.5 weeks at a given school with an additional two days of retention testing. All participants within a given school were tested

over the same period, with students tested in waves depending on the availability of computers. Stimuli were presented over high-quality headphones to minimize disruption from other students taking part in the study at the same time.

After logging in, children saw a task selection screen which displayed icons for all possible tasks. Completed tasks were marked with a checkmark and showed how many points were awarded within that task (during training). After completing a cycle, participants were presented with a new task selection screen in a new color.

Each task began with auditory instructions. Trials started with a carrier sentence that included the target stimulus. For instance, in the *Find the Word* task, students were told to “Find the word *meat*”. This was accompanied by ten written responses on the screen, one of which was the target. Responses were made with a mouse. Students received positive reinforcement (e.g. audio saying, “Thank you for working so hard!” or “Great job!”) during both testing and training; these were given pseudo-randomly, approximately every five trials and were independent of students’ accuracy on a particular trial.

2.5 Analysis

The dependent measure in all analyses was accuracy⁴. Baseline words (untrained words from untrained GPC classes) were not included in the analysis. There was a large effect of vowel overlap at pre-test. As this was of theoretical interest, it prevented the use of ANCOVA and related methods to assess learning (since this would have removed an effect of interest). Thus, to quantify the amount of learning (rather than differences in post-learning performance), we focused on change scores (post-test – pre-test). To investigate the amount of learning in *practical*

⁴ In the *Make the Word* task, children spelled the target item by essentially making three separate responses (they selected the onset consonant(s), the vowel, and then the coda consonant(s)). For this task, we used only the accuracy for the vowel, as this was our primary interest.

terms, we calculated untransformed change scores between pre-test and either post-test or retention. This captures meaningful differences in performance, much as classroom learning would be evaluated. However, this analysis could also reflect ceiling or floor effects such that larger gains are possible from low pre-test scores than from high. Thus, we also investigated the amount of learning using a more theoretically empirical logit scale (which stretches the space near 0 and 1), by computing a change score between logit scaled pre-test and post-test (or retention) accuracy.

Change scores were also analyzed with item ANOVAs (following Clark, 1973). Here, accuracy was not computed for each subject; but averaged for each item (across subjects). This was important as overlap was a between-item manipulation (though variability was within). For some factorial ANOVAs, the assumption for homogeneity of variance was not met (a significant Levene's test). It is not possible to correct for this violation. However, this was not concerning due to our large sample size, the equal distribution of subjects across conditions, and ANOVA's robustness to violations of assumptions.

3. Results

3.1 Pre-test

We first examined pre-test data to detect differences in baseline performance due to the overlap. Initial analyses included variability; however, as these showed no main effects or interactions, analyses reported here collapse across variability.

Figure 2A suggests lower pre-test performance in the high overlap condition was mostly driven by the digraphs. Thus, we performed an ANOVA examining overlap and vowel type (within-subject). We found a significant main effect of overlap ($F(1, 279) = 10.53, p = 0.001, \eta_p^2 = 0.036$) with lower performance in the overlap condition. We also found a significant main

Figure 2: A) Proportion correct at pre-test for both high overlap and low overlap condition separated by vowel type (digraphs vs. monographs). B) Proportion correct at pre-test for both high overlap and low overlap condition for vowel AI only. Both figures include the standard error of the mean as error bars.

effect of vowel type ($F(1, 279) = 258.24, p < 0.001, \eta_p^2 = 0.481$), supporting overall lower performance for digraphs. In addition, there was a significant interaction between overlap and vowel type ($F(1, 279) = 31.38, p < 0.001, \eta_p^2 = 0.101$). This was driven by a significant effect of overlap for digraphs ($F(1, 279) = 34.79, p < 0.001, \eta_p^2 = 0.111$), but not for monographs ($F < 1$).

While GPC regularities were balanced on initial difficulty (Supplement S1), it is possible that differences between groups were due to the GPC regularities in the high and low overlap conditions. We tested this by examining performance on *AI* items, which appeared in both conditions (Figure 2B), though the specific items used in the two conditions differed. *AI* showed a similar pattern with about a 10% decrement in performance for the high overlap group ($F(1, 279) = 17.89, p < 0.001, \eta_p^2 = 0.060$). This suggests that the differences between conditions were unlikely to have derived from the difficulty of the individual GPC regularities, but rather emerged from interrelationships among GPC regularities in the set.⁵

⁵ Finally, to determine if these effects could have arisen from spurious differences between the students, we conducted a 2×2 (overlap $\times$ variability) ANOVA with reading scores (collected by the district) as the dependent variable. As expected, we found no significant differences in reading score as a function of overlap ($F(1, 277) =$

Figure 3: A) Proportion correct at pre-test and post-test across all the 2×2 conditions (overlap $\times$ variability) including standard error of the mean as error bars. B) Untransformed post-test change scores for all conditions. C) Empirical logit transformed proportion correct at pre-test and post-test across all conditions. D) Transformed post-test change scores for all conditions.

3.2 Post-test

We next examined learning gains at post-test as a function of condition (vowel overlap and consonant variability). Figure 3A shows accuracy for each group at pre- and post-test.

First, to investigate the practical value of our manipulations, we entered post-test change scores based on untransformed accuracy into a 2×2 between-subjects (overlap $\times$ variability) ANOVA. We found a significant effect of overlap ($F(1, 273) = 7.05, p = 0.008, \eta_p^2 = 0.025$),

1.81, $p = 0.179, \eta_p^2 = 0.006$) or variability ($F < 1$), and there was no interaction ($F < 1$). This suggests the effect of overlap at pre-test cannot be the consequence of pre-existing differences in ability between groups.

with more learning in the high overlap than the low overlap condition. Variability was not significant ($F(1, 273) = 1.09, p = 0.297, \eta_p^2 = 0.004$), nor was there an interaction of variability and overlap ($F(1, 273) < 1$) (Figures 3A and 3B). In an item analysis, a mixed 2×2 (overlap $\times$ variability) ANOVA was used (with variability as within-item factor and overlap as between-item). Overlap remained significant ($F(1, 94) = 8.09, p = 0.005, \eta_p^2 = 0.079$). There was a significant effect of variability ($F(1, 94) = 9.64, p = 0.003, \eta_p^2 = 0.093$), indicating more learning after variable than similar training. There was no interaction ($F(1, 94) < 1$). Note that variability may have been significant by-item (but not-subject) because in this analysis variability is a within-item comparison (whereas it was between-subject in subject analyses).

We next applied the empirical logit transformation to accuracy and recalculated the difference score (Figure 3C, D). This showed no significant effect of overlap ($F(1, 273) = 2.44, p = 0.120, \eta_p^2 = 0.009$). The effect of variability ($F(1, 273) = 3.35, p = 0.068, \eta_p^2 = 0.012$) was marginally significant, showing slight advantages in variable conditions. Overlap and variability did not interact ($F < 1$). The item analysis with transformed change scores, found significant effects of overlap ($F(1, 94) = 8.18, p = 0.005, \eta_p^2 = 0.080$) and variability ($F(1, 94) = 6.10, p = 0.015, \eta_p^2 = 0.061$), but no interaction ($F < 1$). These results are very similar to the ones observed in the untransformed item analysis: Learning was better for items in the high overlap condition as well as items in the variable condition.

Overall, these analyses indicate two related patterns: Untransformed change scores show that students in the high overlap condition significantly learned more than the ones in the low overlap condition. There was only a small advantage for students in the variable training conditions (only significant by-items). At the same time, however, transformed change scores showed a more modest effect of overlap (only significant by items), though the direction of the

Figure 4: A) Raw proportion correct at pre-test and retention across all the 2×2 conditions (overlap $\times$ variability) including standard error of the mean as error bars (as all other panes). B) Untransformed retention change scores for all conditions. C) Transformed accuracy at pre-test and retention across all conditions. D) Transformed retention change scores across all conditions.

effects was the same—there was no evidence that overlap exerted a cost on learning. This suggests that the benefits observed in the untransformed change score analysis may have been partially due to the lower pre-test range of participants in the high overlap condition. However, this lower pre-test performance in the high overlap condition is likely a consequence of our manipulation, the increased orthographic overlap of the vowels (see our analysis of *AI* at pre-test). That means that although maximizing overlap may generally decrease performance initially, it ultimately leads to either better or at least equivalent learning. In addition, there appears to be a small learning benefit for students in the variable condition (particularly so in the by-item analysis).

Table 5

Planned comparisons of post-test change scores to test learning.

Overlap Group	Similarity Group	Mean $\pm$ SD	Statistics
High Overlap	Similar	0.112 $\pm$ 0.119	$t(64) = 7.57, p < 0.001$
	Variable	0.129 $\pm$ 0.133	$t(68) = 8.07, p < 0.001$
Low Overlap	Similar	0.065 $\pm$ 0.176	$t(72) = 3.16, p = 0.002$
	Variable	0.084 $\pm$ 0.142	$t(69) = 4.90, p < 0.001$

We followed up on these analyses by asking whether a variety of moderating factors influenced any of these factors (Supplement S3). These included task and item (trained vs. generalization), word/nonword, gender and reading ability. No evidence of any moderation was found suggesting this pattern holds across a wide variety of conditions.

Given the relatively small amount of change in the low overlap group, untransformed change scores were used to investigate whether there was a clear, measurable benefit for each manipulation (that would be helpful in a teaching context). We compared the untransformed change scores against zero using a one-sample t-test (Table 5). As predicted, post-test change scores were significantly different from zero in all groups, indicating improvement at the end of training: Learning gains ranged from 6.5 to 12.9%, considerable given the relatively short amount of training students received.

3.3 Retention

Finally, we examined retention by comparing accuracy at the one-to-two-week interval to pre-test scores (Figure 4). As with post-test, we first calculated untransformed retention change scores (retention – pre-test) and entered these into a 2 $\times$ 2 ANOVA. A significant effect of overlap was observed ($F(1, 263) = 9.67, p = 0.002, \eta^2 = 0.035$), with better retention in the high overlap than low overlap conditions. No other effect reached significance (variability: $F < 1$; variability $\times$ overlap: $F(1,263) = 1.20, p = 0.275 \eta^2 = 0.005$). The item analysis replicated the significant

effect of overlap ($F(1, 94) = 18.87, p < 0.001, \eta^2 = 0.167$). As with post-test, the item analysis also showed evidence for a variability effect ($F(1, 94) = 7.51, p = 0.007, \eta^2 = 0.074$), but this time there was a significant interaction ($F(1, 94) = 16.38, p < 0.001, \eta^2 = 0.148$): We examined this by separating the data by overlap condition and comparing similar and variable condition within each (with a factorial ANOVA adjusted using the Bonferroni correction, $\alpha = 0.025$). For the low overlap condition, there was no effect of variability ($F < 1$). However, for the high overlap condition, there was a significant effect of variability ($F(1, 47) = 28.37, p < 0.001, \eta^2 = 0.376$). Unlike the small learning benefits observed for the variable condition at post-test, retention was better with similar than variable consonants for high overlap students.

We next evaluated retention with the more theoretically motivated transformed retention scores. We again found a significant effect of overlap ($F(1, 263) = 4.17, p = 0.042, \eta^2 = 0.016$), but no effect of variability ($F < 1$) or interaction ($F < 1$). An item analysis using the transformed change scores also found a significant effect of overlap ($F(1, 94) = 22.97, p < 0.001, \eta^2 = 0.196$). There was also a significant effect of variability ($F(1, 94) = 9.86, p = 0.002, \eta^2 = 0.095$) and a significant interaction ($F(1, 94) = 15.58, p < 0.001, \eta^2 = 0.142$). We again split the data by overlap condition, and adjusted for repeated tests using the Bonferroni correction, $\alpha = 0.025$. Again, we found a significant effect of variability in the overlap condition only (high overlap: $F(1, 47) = 34.29, p < 0.001, \eta^2 = 0.422$; low overlap: ($F(1, 47) < 1$)) with better retention for similar consonants in the high overlap condition.

Table 6
Planned comparisons of retention change scores.

Overlap Group	Similarity Group	Mean $\pm$ SD	Statistics
High Overlap	Similar	0.114 $\pm$ 0.143	$t(61) = 6.27, p < 0.001$
	Variable	0.072 $\pm$ 0.160	$t(66) = 3.71, p < 0.001$
Low Overlap	Similar	0.018 $\pm$ 0.226	$t(69) = 0.68, p = 0.500$
	Variable	0.026 $\pm$ 0.199	$t(67) = 1.09, p = 0.278$

Table 7

*Summary of major findings: For factors (overlap/variability), a capital letter indicates significant effect, lower case is marginal; O: overlap>no overlap; V: variable>similar; S: similar>variable; * indicates presence of interaction with overlap; + indicates marginal interaction with moderator (for moderator analyses, see SOM2 and 3).*

		Post		Retention	
		Subject	Item	Subject	Item
Overlap	Raw	O	O	O	O
	Logit	-	O	O	O
	Moderators	Digraphs/ Monographs Reading ability ⁺		Word/NW	
Var.	Raw	-	V	-	S*
	Logit	v	V	-	S*
	Moderators	Spelling/Rec. Task ⁺		-	
Groups showing Learning / Retention		All		O only	

Overall, children learning high overlap GPCs retained their skills better than children in the low overlap condition. This was true for both untransformed and transformed scores, suggesting advantages for long-term retention with overlapping materials. There was also an overall benefit for similarity, though item analyses indicate this benefit may be restricted to the high overlap condition. Moderator analyses were conducted (Supplement S3) and found no moderation by any of the task, item or participant factors.

Finally, we asked whether untransformed change scores in each group were significantly different from zero at retention (Table 6). Only students assigned to either of the high overlap groups exhibited significant retention with similarly high change scores as at post-test, suggesting only minimal loss over the testing delay.

4. Discussion

4.1 Overview

A marked effect of overlap was evident before training: At pre-test, children performed poorer on overlapping vowels than low-overlapping vowels. That pattern however was more complex with training (Table 7 for a summary). At immediate post-test, in raw change scores

greater learning was observed with high overlap; however, with the empirical logit transformation, evidence for this advantage was less robust (though significant by items). More importantly, across all the analyses retention was consistently better for students in the high overlap condition than in the low overlap one. This overlap benefit is notable given that initial performance in the high overlap condition was much lower. A clear take-away from these findings is that there is no cost to training overlapping items simultaneously, and in fact students may learn more.

Evidence for the variability effect was mixed. At post-test, there was a benefit for variability that was significant by items for both analyses. By subjects it was not significant with raw scores, but marginally significant with transformed scores. This suggests a small learning advantage for variable consonants (at least by-items), though this is not as strong as evidence from prior studies. At retention, the effect was significant but reversed. However, this similarity benefit was largely restricted to the high overlap condition.

4.2 Limitations

Before we discuss our findings, some limitations are worth noting. Many of these are unavoidable consequences of the structure of English and the items available for first graders.

First, overlap was manipulated by selecting distinct sets of regularities on which to train. This makes it difficult to eliminate the possibility that effects were driven by the particular GPC regularities selected for each condition. Pilot data (with a large sample, see Online Supplement S1) do not suggest numerically large differences in four of the six GPC regularities, and one regularity was shared between conditions. Moreover, items were balanced on frequency and neighborhood density between conditions. Overall, this makes it unlikely that differences in the inherent difficulty of the items or the GPC regularities were a large contributor to performance.

Nevertheless, future studies may match overlap conditions, so that some proportion of items is used in both conditions (thus making comparisons across subjects easier), or could use more extensive (multi-day) piloting of items and GPC classes to equate difficulty.

Second, we tested only a small subset of the vowel GPCs of English (not to mention the consonants). Digraphs were selected because they offer a clear way to instantiate and manipulate overlap. While there is no reason to believe that they are not representative, it is important to replicate these findings with other classes of GPC regularities (such as secondary pronunciations, e.g., *EA* as / ϵ / as in *bread* and *dead*) to test their general applicability.

Finally, this is a laboratory study of learning, not a test of a curriculum. These results may not translate directly to classroom learning, and such translation should consider several factors. Classroom learning may not always include immediate feedback (or any at all), which has been shown to be important for some forms of implicit learning (Freedberg, Glass, Filoteo, Hazeltine, & Maddox, 2016; Maddox, Ashby, & Bohil, 2003). Further, it may interact with the statistical structure of the input features and classification problem to drive learning (Carvalho & Goldstone, 2015). Classroom instruction also typically employs different media (e.g. worksheets) and often a much smaller set of items. Given these differences (and others), our work does not directly speak to the efficacy of a real curriculum embodying these principles. Nevertheless, as we discuss shortly, it may have important implications for teaching.

4.3 The Effect of Consonant Variability

The effect of variability was small (Table 7). At immediate post-test (the closest comparison to Apfelbaum et al., 2013), effects were consistent with a variability benefit, but statistical evidence was not robust, with significant effects by items and a marginal effect by subject (with logit transformed scores). Nonetheless, numerically the effect took the form of a

variability benefit. Thus, while we did not observe evidence *against* an immediate variability benefit, the observed variability benefit was small. The fact that these results were statistically weaker than expected was surprising given both Apfelbaum et al.'s (2013) results, and the large body of evidence from other areas of learning that variability in irrelevant elements leads to strong and robust learning benefits (Gómez, 2002; Huet et al., 2011; Kerr & Booth, 1978; Lively et al., 1993; Rost & McMurray, 2009, 2010).

However, there are several differences between Apfelbaum et al. (2013) and the present study that could explain the difference in the size of variability benefit. First is statistical power: Apfelbaum et al. (2013) split the sample along only one factor (variability), not two (variability and overlap) as here. This may explain why some by-subject analyses were not significant, but by-item analyses were (since variability was manipulated between subjects but within-items). However, it does not account for the relatively small numerical effect.

Second, our manipulation of variability may have been inadvertently weaker than that of Apfelbaum et al. (2013): In that study's similar condition, each item shared its onset consonants with 21.8 items (variable=10.4); in contrast in the similar condition of this study, each item shared an onset with only 6.23 items (variable=3.2). Along these lines, similar items in Apfelbaum et al. (2013) shared the full consonant frame with 2.4 other items (variable=0.2); compared with only 1.99 here (variable=0.31). To ensure overlap between vowels, we had to use vowels from a small pool of appropriate words, affording less flexibility to maximize variability. Thus, the smaller difference between the variability conditions may have resulted in the smaller size of the variability benefit observed here.

Potentially more important theoretically, the Apfelbaum et al. (2013) training regime blocked trials by vowel: Student underwent a block of trials using monographs only, then a block

of trials of digraphs only, and then finally a block on which they were interleaved. In contrast, our design interleaved all vowels during training. This was done based on findings that interleaving items may maximize learning (Carvalho & Goldstone, 2014; Kang & Pashler, 2012; Taylor & Rohrer, 2010; Vlach et al., 2008).

However, it is possible that by interleaving words students could isolate the vowel as a unit, in the absence of consonant variability: Fully interleaved GPC regularities might have helped eliminate spurious associations between graphemes and phonemes to isolate the respective vowels without requiring consonant variability. This was investigated in a follow-up study by McMurray, Roembke, and Hazeltine (submitted), which simultaneously manipulated consonant variability and how vowels were blocked during training. Six vowel GPCs were trained and blocked in pairs that either overlapped (EA with OA) or did not (EE with OA), or were all completely interleaved. They showed a variability benefit when blocking highlighted the overlapping strings, but no variability benefit for interleaved training. These findings offer a unifying explanation as to why the observed variability effect was small and less statistically robust than in Apfelbaum et al. (2013) and suggest that how vowels are blocked impacts whether variability benefits learning or not.

It remains unclear why the variability effect was reversed at retention and only significant in the high overlap condition. Apfelbaum et al.'s (2013) study did not include a retention test; however, the previously described follow-up study by McMurray et al. (submitted) observed a similarity benefit at retention under blocked/no-contrast conditions. Further studies are needed to better understand under what circumstances similarity may benefit learning, and how this process is impacted by consolidation.

4.4 The Effect of Vowel Overlap

4.4.1 Overlap at pre-test. Performance at pre-test was lower in the high overlap than the low overlap condition. This likely was a consequence of two factors. First, on a trial-to-trial basis simply alternating between similar GPC strings could add difficulty, analogous to phenomena like cumulative semantic interference (Howard, Nickels, Coltheart, & Cole-Virtue, 2006; Oppenheim, Dell, & Schwartz, 2010), where naming a sequence of semantically related (conceptually overlapping) images leads to slower performance. Second, items or letters that were assigned to the same task served as each other's foils. Consequently, foil items were more likely to overlap with the target item in the high overlap condition than in the low overlap condition.⁶ Thus, accurately selecting the target in the high overlap condition requires both adequate understanding of the relevant GPC regularities but also the ability to inhibit competitors in the moment (McMurray, Horst, & Samuelson, 2012 for an analogy in word learning). Under this view, the difference in accuracy between the two conditions stems from difficulty inhibiting overlapping competitors in the high overlap condition. This has clear analogues in the well-known effects in word recognition, where words with many overlapping neighbors are harder to recognize (Luce & Pisoni, 1998; Magnuson, Dixon, Tanenhaus, & Aslin, 2007) and also in work on overlapping motor patterns (Wifall, McMurray, & Hazeltine, 2014).

The pre-test decrement associated with overlap was only observed in the digraphs (e.g., Figure 2A). One explanation for this is that the digraphs were more difficult. Monographs are often introduced first in reading curricula and thus have been practiced the most. This early, and often isolated, exposure may lead to more robust mappings that are more resilient to interference from overlapping strings later. It is also possible that digraphs are not perceived as a unit: The

⁶ When teaching a small set of vowels, the only way to avoid this overlap in the foils would be to introduce new vowel strings that only appear as foils. These of course could be learned to be eliminated by students since they are never correct, making the task unintentionally easy.

two individual letters may therefore compete with a larger number of foils on each trial than a monograph. This could also explain why accuracy for digraph vowels was particularly low in the high overlap condition, in which—as described earlier—foils were more likely to be closer competitors to the target.

4.4.2 The effect of overlap on learning and retention. Even though the high overlap condition was markedly harder at pre-test, participants' performance improved significantly more with high overlap training than with the low overlap. This difference was particularly pronounced at retention where the low overlap condition essentially showed no retention (see Figure 4B). Training with minimal overlap—while it did show immediate learning—did not result in longer-term benefits. In contrast, overlapping vowels, which introduce initial difficulties during pre-test (and presumably during training), may be helpful in the long run.

Our moderation analyses (SOM2, 3) suggest that overlap did not interact with item novelty or task novelty at post-test or retention. Although caution is warranted when interpreting these sorts of null effects, it suggests that overlap not only improved learning, but also led to robust generalization to novel words and tasks. Thus, potential learning gains due to overlap in the vowel sets are not fragile or specific to the items/tasks that were used.

The naïve associative account does not predict these findings. The fact that individual letters must be mapped to more than one sound was expected to impede learning. While this difficulty was observed at pre-test, the predictions of this account for learning were not supported. So why may vowel overlap be helpful for learning and retention?

First, recent connectionist instantiations of statistical learning (McClelland, 2013) show that new materials can be more rapidly and robustly integrated into an existing network if they share schematic structure with already learned items than if they are inconsistent with what is

known. In this case, the learning system is not characterized as isolated associations between elements like letters and sounds, but rather between sounds, letters and most importantly, complex intermediate representations. In situations in which new material can be better integrated into these representations (which are based on similarity structure in the training materials), learning is more efficient

This is also consistent with a recent study on visual category learning. Carvalho and Goldstone (2014) manipulated the amount of within- and between-category similarity and temporal arrangement (interleaved/blocked). Their results showed that if the goal is to discriminate items (e.g., if items are highly similar across categories), rather than to find characteristic features, interleaving is beneficial to learning. In learning GPC regularities, a similar process may be at play, in which one is required to learn that *A* makes one sound if presented with an *O*, but a different one if presented with an *E*, forcing the comparison of overlapping representations. Thus, Carvalho and Goldstone's data (2014) are consistent with the idea that simultaneous training on overlapping GPC strings may be a form of schema-driven learning.

This view of statistical learning as embedded in a network raises the possibility of a second mechanism for the overlap effect: eliminative learning. In complex networks, correct connections must be built, but incorrect connections may also be eliminated or pruned. Because of the intrinsic overlap in (English) vowel GPC regularities, spurious associations may exist between letters and phonemes early in reading acquisition (Gough & Juel, 1991; Juel & Roper-Schneider, 1985). For instance, the strong connection between the letter *O* and the phoneme /a/ (as in short vowels) may generalize to a (weaker) association between the letters *OA* and /a/ (even though *OA* most often maps to the phoneme /ou/). If this is the case, then a critical goal for

learning is to prune these connections in favor of a sparser network of just the relevant connections. Under this view, the high overlap condition included more opportunity for students to inhibit or prune letter or item competitors (as discussed in relation to our pre-test findings), yielding an opportunity to tune multiple associations on any given trial (Roembke, Wasserman, & McMurray, 2016; also see McMurray et al., 2012 for evidence for this in a connectionist model of word learning). Moreover, given some evidence for enhanced pruning as part of consolidation (McMurray, Kapnoula, & Gaskell, 2016 for a model in human word learning; Tononi & Cirelli, 2006), this may explain why the overlap effect was pronounced at retention.

At a broader level, our work offers clear evidence that statistical learning in reading must be considered as part of a complex network (e.g., Seidenberg, 2005): Students are not learning an entirely novel set of associations from scratch. Rather, they are embedding new associations in a rich network of existing knowledge. When we consider statistical or associative accounts in this light—as part of a developmental process, not a laboratory learning exercise—the nature of learning may change.

4.5 Broader Theoretical Implications

At a broad theoretical level, our findings offer further evidence that statistical learning facilitates literacy acquisition (e.g., Apfelbaum et al., 2013; Arciuli & Simpson, 2012; Spencer et al., 2015). We would not have observed an effect of overlap if children were not sensitive to the statistics and associative make-up of the training input. Moreover, the study gives further insight into how one may want to conceptualize statistical learning: Our findings are not consistent with a simple accumulation of statistics or simple associations between letter and sounds. In contrast, they illustrate how the statistics at the item level may not fully describe learning when multiple regularities are learned simultaneously.

When we consider how item-level statistics intersect with groupings at the level of the regularities being taught, a similarly complex story emerged. The naïve associative account was an extrapolation from item-level statistics to the problem of learning multiple regularities. Here, the pure item-level statistics should have impeded learning as the same letter is now correlated with two different sounds (in different contexts). This was not supported. Instead, the effect of overlap (particularly at retention) is more consistent with the idea of integrating regularities into a complex network, depending on one's task goals, than blindly accumulating statistics. All told, models closer to connectionist instantiations than pure statistical learners may better describe the early phases of reading (Seidenberg, 2005).

At the same time, our pre-test results suggest that participants' performance does not only depend on their underlying knowledge, but also the in-the-moment characteristics of a trial: The presence of greater overlap between targets and foils in the high overlap condition likely caused differences in pre-test scores (between the overlap conditions) because of increased competition. Of course, these differences in real-time processing—accumulated over many trials—cascade into differences in long-term learning and retention, an idea that is gaining traction in spoken language learning (Apfelbaum & McMurray, 2016; Bion, Borovsky, & Fernald, 2013; Fernald & Marchman, 2012; McMurray et al., 2012; Roembke & McMurray, 2016).

4.6 Implications for Teaching

This study encompasses an infrequently sampled point in the space of reading research: Our goal was to neither directly test opposing theoretical models of reading, as would be typical of basic science work, nor do we offer a direct test of different types of curricula. Instead, our goal was to identify learning principles by which materials can be structured to improve learning, considering both implications from statistical learning and from typical reading instruction. In

particular, as we describe in the introduction, this work was motivated by the fact that when confronted with the real-world of grouping GPC regularities for instruction, statistical learning appeared to make competing predictions for what should benefit learning.

Such studies may offer an important addition to our understanding of instructional practice, given the dearth of work on the microstructure of curricula (e.g., the specific groupings of items or regularities in a curriculum). For example, much research (e.g., Gates & Yale, 2011)—without citing support for this idea—suggests that vowels should be initially trained with word families or pattern words (e.g., *A* should be introduced by training children on words such as *bat*, *fat*, *cat*) to promote learning. This assumes a specific class of statistical regularities (analogous to the effect of consonant similarity) that may or may not be optimal for learning (though children indeed learn from such items). Short term studies like ours and Apfelbaum et al. (2013) can help explore the space of the microstructural choices.

In that light, these findings offer additional insights into why vowels may be particularly hard for early readers (Read, 1975; Treiman, 1993). While clearly part of the difficulty lies in the fact that most digraph rules have several alternative pronunciations (e.g., *EA* can be /i/ in *meat*, /ɛ/ in *threat*, and /eɪ/ in *steak*), our work suggests that the overlap may also contribute to their (in-the-moment) difficulty (even as focused training on the overlapping sets may ultimately lead to more robust learning).

Despite this, our study suggest that it may not be effective to avoid overlap in reading curricula that target subsets of the GPC system (at least if children already have some knowledge on GPC regularities). At immediate post-test, high overlap training led to at least as much learning, if not more, than low overlap training, and retention was more robust. This principle can easily be incorporated into curricula by reworking the way GPC strings are introduced,

though clearly such curricula require randomized control testing for efficacy. As discussed previously, the typical groupings used in most curricula (short vowels, long vowels, etc.) do not embrace this overlap, and may implicitly avoid it. Teaching vowels in blocked “units” of unrelated or non-overlapping vowels (like short vowels, digraphs, etc.) may leave children (at least initially) with difficulties when they later encounter overlapping vowels in natural text. Instead, it may be useful to group vowels by overlap or to include a given vowel string in multiple blocks paired with different sorts of overlapping strings. This may lead to more robust learning and retention.

With regard to items, the findings confirm that variability in irrelevant letters may be important (as in Apfelbaum et al., 2013), but may also be impacted by how vowels are presented (e.g., interleaved with other vowels or not; McMurray et al., submitted). Specifically, our study suggests that a benefit of adding variability is not as ubiquitous as previously thought; however, as we described there are several parameters of training that may account for this, and recent follow-up data suggest that how vowels are grouped/contrasted may be an important factor in determining whether variability is beneficial or not (McMurray et al., submitted).

More broadly, the findings highlight the power of systematically manipulating how items or regularities are grouped during training, and the value of implementing a small set of tasks over a large number of items. Computer-based training may be helpful in enacting these principles, as they can efficiently instantiate controlled manipulations of items; they can stress different forms of overlap at different times, and they can provide immediate feedback (which benefits procedural learning). While this cannot replace real teachers and texts, computer-based training may supplement the curriculum by leveraging learning principles for better outcomes. Thus, it is essential for future research to explore further interactions of factors (such as overlap,

variability and blocking) to develop successful supplemental learning tools.

If one takes statistical learning seriously as a mechanism of reading acquisition (as supported by Apfelbaum et al., 2013; Arciuli & Simpson, 2012; Spencer et al., 2015; Treiman & Kessler, 2006), it must be considered in the context of how reading instruction is frequently structured. Our findings offer clear recommendations for how curricula could be designed, both on the item level as well as at the level of grouping regularities that are taught simultaneously, while simultaneously providing evidence for a network view of statistical learning. In this way, basic research suggests alternative ways of structuring materials and outlines a platform for testing them.

5. References

- Apfelbaum, K. S., Hazeltine, E., & McMurray, B. (2013). Statistical learning in reading: Variability in irrelevant letters helps children learn phonics skills. *Developmental Psychology*, 49(7), 1348–1365.
- Apfelbaum, K. S., & McMurray, B. (2011). Using variability to guide dimensional weighting: associative mechanisms in early word learning. *Cognitive Science*, 35(6), 1105–1138.
- Apfelbaum, K. S., & McMurray, B. (2016). Learning during processing: Word learning doesn't wait for word recognition to finish. *Cognitive Science*.
- Arciuli, J., Monaghan, P., & Ševa, N. (2010). Learning to assign lexical stress during reading aloud: Corpus, behavioral, and computational investigations. *Journal of Memory and Language*, 63(2), 180–196.
- Arciuli, J., & Simpson, I. C. (2012). Statistical learning is related to reading ability in children and adults. *Cognitive Science*, 36(2), 286–304.
- Armstrong, B. C., Dumay, N., Kim, W., & Pitt, M. A. (2017). Generalization from newly learned words reveals structural properties of the human reading system. *Journal of Experimental Psychology: General*, 146(2), 227–249.
- Audacity Team. (2014). Audacity(R): Free Audio Editor and Recorder. Version 2.0.6.
- Beck, I. L., McKeown, M. G., & Kucan, L. (2002). *Bringing words to life, second edition: Robust vocabulary instruction. Solving problems in the teaching of literacy*. New York: Guilford Publications.
- Bion, R. A. H., Borovsky, A., & Fernald, A. (2013). Fast mapping, slow learning: disambiguation of novel word-object mappings in relation to vocabulary learning at 18, 24, and 30months. *Cognition*, 126(1), 39–53.

- Boersma, P., & Weenink, D. (2016). Praat: doing phonetics by computer.
- Carvalho, P. F., & Goldstone, R. L. (2014). Putting category learning in order: Category structure and temporal arrangement affect the benefit of interleaved over blocked study. *Memory & Cognition*, 42, 481–495.
- Carvalho, P. F., & Goldstone, R. L. (2015). The benefits of interleaved and blocked study: Different tasks benefit from different schedules of study. *Psychonomic Bulletin & Review*.
- Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 12(4), 335–359.
- Coltheart, M., & Rastle, K. (1994). Serial Processing in Reading Aloud: Evidence for Dual-Route Models of Reading. *Journal of Experimental Psychology: Human Perception and Performance*.
- Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: a dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108(1), 204–256.
- Ehri, L. C. (1995). Phases of development in learning to read words by sight. *Journal of Research in Reading*, 18(2), 116–125.
- Ehri, L. C., Nunes, S. R., Stahl, S. a., & Willows, D. M. (2001). Systematic phonics instruction helps students learn to read: Evidence from the national reading panel’s meta-analysis. *Reading Research Quarterly*, 36(3), 250–287.
- Fernald, A., & Marchman, V. A. (2012). Individual differences in lexical processing at 18 months predict vocabulary growth in typically developing and late-talking toddlers. *Child Development*, 83(1), 203–222.
- Foundations in Learning, I. (2010). Access Code. Iowa City, IA: Foundations in Learning.

- Fountas, I. C., & Pinnell, G. S. (2011). *Fountas & Pinnell Benchmark Assessment System 1*. Portsmouth, NH: Heinemann.
- Fowler, C. A., Liberman, I. Y., & Shankweiler, D. P. (1977). On interpreting the error pattern in beginning reading. *Language and Speech*, 20, 162–173.
- Freedberg, M., Glass, B., Filoteo, J. V., Hazeltine, E., & Maddox, W. T. (2016). Comparing the effects of positive and negative feedback in information-integration category learning. *Memory & Cognition*, 1–14.
- Fry, E. (1964). A frequency approach to phonic. *Elementary English*, 41(7), 759–765.
- Gabay, Y., Thiessen, E. D., & Holt, L. L. (2015). Impaired statistical learning in developmental dyslexia. *Journal of Speech Language and Hearing Research*, 58(3), 934.
- Gates, L., & Yale, I. (2011). A logical letter–sound system in five phonic generalizations. *The Reading Teacher*, 64(5), 330–339.
- Gómez, R. L. (2002). Variability and detection of invariant structure. *Psychological Science : A Journal of the American Psychological Society / APS*, 13(5), 431–436.
- Gough, P. B., & Juel, C. (1991). The first stages of word recognition. In *Learning to read: Basic research and its implications* (pp. 47–56).
- Harm, M. W., & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: insights from connectionist models. *Psychological Review*, 106(3), 491–528.
- Harm, M. W., Seidenberg, M. S., Macdonald, M., Thornton, R., & Zevin, J. (2004). Computing the meanings of words in reading : cooperative division of labor between visual and phonological processes. *Psychological Review*, 111(3), 662–720.
- Howard, D., Nickels, L., Coltheart, M., & Cole-Virtue, J. (2006). Cumulative semantic inhibition in picture naming: experimental and computational studies. *Cognition*, 100(3), 464–482.

- Huet, M., Jacobs, D. M., Camachon, C., Missenard, O., Gray, R., & Montagne, G. (2011). The education of attention as explanation of variability of practice effects: Learning the final approach phase in a flight simulator. *Journal of Experimental Psychology: Human Perception and Performance*, 37(6), 1841–1854.
- Johnston, F. P. (2001). The utility of phonic generalizations: Let's take another look at Clymer's conclusions. *The Reading Teacher*, 55(2), 132–143.
- Juel, C., & Roper-Schneider, D. (1985). The influence of basal readers on first grade reading. *Reading Research Quarterly*, 20, 134–152.
- Kang, S. H. K., & Pashler, H. (2012). Learning painting styles: Spacing is advantageous when it promotes discriminative contrast. *Applied Cognitive Psychology*, 26(1), 97–103.
- Kellman, P. J., Massey, C. M., & Son, J. Y. (2010). Perceptual learning modules in mathematics: Enhancing students' pattern recognition, structure extraction, and fluency. *Topics in Cognitive Science*.
- Kerr, R., & Booth, B. (1978). Specific and varied practice of motor skill. *Perceptual and Motor Skills*, 46(2), 395–401.
- Kim, W., Pitt, M. A., & Myung, J. I. (2013). How do PDP models learn quasiregularity? *Psychological Review*, 120(4), 903–916.
- Lively, S. E., Logan, J. S., & Pisoni, D. B. (1993). Training Japanese listeners to identify English /r/ and /l/. II: The role of phonetic environment and talker variability in learning new perceptual categories. *The Journal of the Acoustical Society of America*, 94(3 Pt 1), 1242–1255.
- Luce, P. A., & Pisoni, D. B. (1998). Recognizing spoken words: The neighborhood activation model. *Ear and Hearing*, 19(1), 1–36.

- Maddox, W. T., Ashby, F. G., & Bohil, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(4), 650–662.
- Magnuson, J. S., Dixon, J. a, Tanenhaus, M. K., & Aslin, R. N. (2007). The dynamics of lexical competition during spoken word recognition. *Cognitive Science*, 31(1), 133–156.
- McClelland, J. L. (2013). Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. *Journal of Experimental Psychology: General*, 142(4), 1190–1210.
- McDaniel, M. A., Anderson, J. L., Derbish, M. H., & Morrisette, N. (2007). Testing the testing effect in the classroom. *European Journal of Cognitive Psychology*.
- McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review*, 119(4), 831–877.
- McMurray, B., Kapnoula, E., & Gaskell, M. G. (2016). Learning and integration of new word-forms: Consolidation, pruning and the emergence of automaticity. In M. G. Gaskell & J. Mirković (Eds.), *Speech Perception and Spoken Word Recognition*. London, UK: Taylor and Francis.
- McMurray, B., Roembke, T. C., & Hazeltine, E. (submitted). Field tests of learning principles learning can support pedagogy: Overlap and variability jointly affect sound/letter acquisition in first graders. *Journal of Cognition and Development*.
- Näslund, J. C. (1999). Phonemic and graphemic consistency: Effects on decoding for German and American children. *Reading and Writing*, 11(2), 129–152.
- National Reading Panel. (2000). *Teaching children to read: An evidence-based assessment of the*

- scientific research literature on reading and its implications for reading instruction. NIH Publication No. 00-4769 (Vol. 7). <https://doi.org/10.1002/ppul.1950070418>*
- Oppenheim, G. M., Dell, G. S., & Schwartz, M. F. (2010). The dark side of incremental learning: A model of cumulative semantic interference during lexical access in speech production. *Cognition, 114*(2), 227–252.
- Plaut, D. C., McClelland, J. L., & Seidenberg, M. S. (1995). Reading exception words and pseudowords: Are two routes really necessary. *Connectionist Models of Memory and Language*, 145–159.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: computational principles in quasi-regular domains. *Psychological Review, 103*(1), 56–115.
- Pollo, T. C., Kessler, B., & Treiman, R. (2009). Statistical patterns in children's early writing. *Journal of Experimental Child Psychology, 104*(4), 410–426.
- Rayner, K., Foorman, B. R., Perfetti, C. A., Pesetsky, D., & Seidenberg, M. S. (2001). How psychological science informs the teaching of reading, 2(2).
- Read, C. (1975). Children's categorization of speech sounds in English. In *NCTE Research Report No. 17*.
- Roembke, T. C., & McMurray, B. (2016). Observational word learning: Beyond propose-but-verify and associative bean counting. *Journal of Memory and Language, 87*, 105–127.
- Roembke, T. C., Wasserman, E. A., & McMurray, B. (2016). Learning in Rich Networks Involves Both Positive and Negative Associations. *Journal of Experimental Psychology: General, 145*(8), 1062–1074.
- Rohrer, D., Dedrick, R. F., & Stershic, S. (2015). Interleaved practice improves mathematics

- learning. *Journal of Educational Psychology*.
- Rost, G. C., & McMurray, B. (2009). Speaker variability augments phonological processing in early word learning. *Developmental Science*, 12(2), 339–349.
- Rost, G. C., & McMurray, B. (2010). Finding the signal by adding noise: The role of noncontrastive phonetic variability in early word learning. *Infancy*, 15(6), 608–635.
- Seidenberg, M. S. (2005). Connectionist models of word reading. *Current Directions in Psychological Science*, 14(5), 238–242.
- Seidenberg, M. S. (2016). *Language at the Speed of Sight: How We Read, Why So Many Can't, and What Can Be Done About It*. Basic Books.
- Shankweiler, D. P., & Liberman, I. Y. (1972). Misreading: A search for causes. In J. F. Kavanagh & I. G. Mattingley (Eds.), *Language by ear and by eye: The relationships between speech and reading* (pp. 293–317). Cambridge, MA: MIT Press.
- Share, D. L. (1995). Phonological recoding and self-teaching: sine qua non of reading acquisition. *Cognition*, 55(2), 151–218.
- Spencer, M., Kaschak, M. P., Jones, J. L., & Lonigan, C. J. (2015). Statistical learning is related to early literacy-related skills. *Reading and Writing*, 28(4), 467–490.
- Stuebing, K., Barth, A. E., Cirino, P. T., & Fletcher, J. M. (2008). A response to recent reanalyses of the National Reading Panel Report: Effects of systematic phonics instruction are practically significant. *Journal of Educational Psychology*, 100(1), 123–134.
- Taylor, K., & Rohrer, D. (2010). The effects of interleaved practice. *Applied Cognitive Psychology*, 24(6), 837–848.
- Thompson, G. B., Fletcher-Flinn, C. M., & Cottrell, D. S. (1999). Learning correspondences between letters and phonemes without explicit instruction. *Applied Psycholinguistics*,

20(01), 21–50.

Tononi, G., & Cirelli, C. (2006). Sleep function and synaptic homeostasis. *Sleep Medicine Reviews, 10*(1), 49–62.

Treiman, R. (1993). *Beginning to spell: A study of first-grade children*. Oxford University Press.

Treiman, R., Gordon, J., Boada, R., Peterson, R. L., & Pennington, B. F. (2014). Statistical learning, letter reversals, and reading. *Scientific Studies of Reading, 18*(6), 383–394.

Treiman, R., & Kessler, B. (2006). Spelling as statistical learning: Using consonantal context to spell vowels. *Journal of Educational Psychology, 98*(3), 642–652.

U.S. Department of Education, National Center for Education Statistics, I. of E. S. (2010). *The nation's report card: Reading, 2009*. Washington, D.C.

Vlach, H. A., Sandhofer, C. M., & Kornell, N. (2008). The spacing effect in children's memory and category induction. *Cognition, 109*(1), 163–167.

Wifall, T., McMurray, B., & Hazeltine, E. (2014). Perceptual similarity affects the learning curve (but not necessarily learning). *Journal of Experimental Psychology: General, 143*(1), No-Specified.

Supplemental Online Material:

Simultaneous Training on Overlapping Grapheme Phoneme Correspondences

Augments Learning and Retention

S1: Items Used

S1.1: GPC Selection

Our design was guided in part on a control experiment that was conducted several years prior to the present experiment (in 2011). This experiment tested a set of 140 first grade students from West Des Moines on a set of 19 different tasks and a range of different items. 190 trials were presented without feedback, and children selected the tasks in whatever order they wished. These data were used to balance tasks and GPC classes for difficulty across conditions.

GPC regularity was calculated as described: We identified all of the mono-syllabic words that included a given vowel string, and categorized them by pronunciation. From this we computed the GPC's *regularity* as the proportion of words containing the target vowel string that used the intended phoneme. This was computed for the each of several pronunciations, and the highest was selected. These data are reported in Table 2 of the main manuscript.

To ensure that the two sets of GPC regularities were not individually more difficult, we consulted the pilot data. Pilot data were available for 8 of the GPC regularities that were used (marked with a * in Table 2 of the main manuscript). These averaged 78.0% correct (SD=15.9%) in the high overlap group and 79.7% (SD=14.9%) in the low overlap group. In order to fill out the conditions, *AI* was added to both groups, and two new vowels (*OI* and *OU*) were added for which pilot data were not available. Note that the less regular (more difficult) *OU* rule was added to the low overlap group to account for the slightly higher

performance of the other rules based on the pilot data.

S.1.2: Item Selection

Separate lists of items were used for the high overlap and low overlap conditions for training and test. Within an overlap group, items for test were the same across both variability conditions; items during training differed. It should be noted that for OI and OU, there are only 13 monosyllabic, non-cluster items each (of which only a subset is appropriate for first graders). Thus, our selection was considerably limited when selecting words that are similar in consonant structure to other GPC items.

With the exception of the cluster items used in testing (see below), all words had a CVC form and were selected to be words that first graders are likely to know. Word lists were reviewed by a person with educational experience with this age group.

Testing items. For each of the six GPC regularities, we selected five words and five nonwords. For both words and nonwords, four were items that were also used in training for both variable and similar list (within each overlap condition); four followed the particular GPC regularity, but would not be included in training (two of these had a CVC structure, two included a consonant cluster or digraph) to test generalization to novel items; and the last two items were *baseline* items, which used vowels that children had not been trained on. We included clusters in the generalization items to test how easily children would be able to generalize their knowledge to more complex items. For instance, the vowel OI was used in the testing stimuli FOIL and LOIN (shared words that had also been used in training), VOID and MOIST (novel items, one CVC and one cluster)¹, and LAP (baseline). Items were not repeated during a testing phase.

¹ We also completed analyses in which we distinguished between novel-CVC and novel-cluster words; however, there was no difference in performance in this study. We did not include these analyses here.

Training items. As in testing, training items were balanced among the six respective GPC regularities for the participant's overlap condition. For each GPC regularity, five words and five nonwords were selected for a total of 60 words. Each was repeated 5 times (in different tasks) during training. The two item lists within overlap condition differed in the amount of variability in the consonants, resulting in a total of four training lists. Similar item lists included words that shared one or both consonants (e.g. MEAL, MOAN, PEN, MEN); words in the variable group were selected to minimize the number of shared consonants.

As described, two items for each vowel appeared in both the similar and variable conditions (of each overlap condition) during training as well as testing. The only exception to this was for the vowel OU in the low overlap condition, for which insufficient words were available to identify six unique words with the CVC frame that were likely to be known by first graders. Thus, instead of just two words that were shared between the similar and variable conditions, four words were used in both conditions (see Appendix for the full list of items used both in training and testing).

Lexical statistics. We conducted a number of quantitative analyses of the items to determine if the conditions instantiated the intended manipulations, and if any properties of the items were confounded with the between-subject variable overlap.

Training lists. In the similar word list, items shared both the onset and coda consonant with an average of 1.97 other items in the high overlap group, and 2.0 in the low overlap group. In contrast, the variable word list minimized this: Each item shared a frame with 0.33 (high overlap) and 0.30 (low overlap) items. When we considered the onset and coda consonants individually, items in the similar group shared onsets with an average of 6.23 (high overlap) and 6.23 (low overlap) other items, and shared codas with 13.30 and 9.47 items (respectively). In the variable group, each item shared onsets with an average of 3.33 (high overlap) or 3.07 (low overlap) items, and shared offsets with an average of 8.90 (high

overlap) and 9.07 (low overlap) items.

We next examined the potentially confounding factors of frequency, imageability and neighborhood density for the training items. Frequency was estimated from the Kucera-Francis written frequency, and imageability from the Medical Research Council [MRC] Psycholinguistic Database; Coltheart, 1981). Both estimates were only available for words, and not available for all words. The number of words contributing to the analysis is indicated. Neighborhood density was estimated with the Irvine Phonotactic Online Dictionary using the unstressed log10 estimate (IPhOD, Version 2.0; Vaden, Halpin, & Hickok, 2009) for both words and nonwords.

We investigated each of these factors (frequency, imageability, and log neighborhood density) using a 2×2 (overlap $\times$ variability) ANOVA. For frequency, there was no main effect of overlap ($F(1, 112) < 1$), variability ($F(1, 112) < 1$) nor an interaction between the two ($F(1, 112) < 1$), suggesting that frequency was well-matched across all conditions. For imageability, there was a main effect of overlap ($F(1, 88) = 8.97, p = 0.004, \eta_p^2 = 0.092$), indicating items with lower imageability were included in the low overlap ($M = 470.00, N = 42/60$) than the high overlap condition ($M = 525.66, N = 50/60$). However, there was no main effect of variability ($F(1, 88) < 1$) nor an interaction between overlap and variability ($F(1, 88) < 1$). Finally, we compared lists on log neighborhood density. We found no main effect of overlap ($F(1, 236) < 1$) nor did it interact with variability ($F(1, 236) < 1$), suggesting that items in both overlap conditions had similar amounts of phonological neighbors. However, variability reached marginal significance (Similar: $M = 35.96, N = 120/120$; Variable: $M = 32.30, N = 120/120$; $F(1, 236) = 3.44, p = 0.065, \eta_p^2 = 0.014$). As in Apfelbaum et al. (2012), both the variable list in the high overlap and low overlap condition appeared to be more likely to include less frequent consonants than the similar ones, leading to a lower number of phonological neighbors. This was an inevitable consequence of trying to find

words with much similarity in consonant frames (or not), forcing us to oversample more frequent letters for the similar word lists. We do not believe that this influenced our results (but see Apfelbaum et al. for a discussion of this factor), in part because these differences were only present during learning.

Testing lists. To evaluate the overlap manipulation, we estimated these values for the testing lists (testing lists were the same for both the variable and similar conditions). We found no main effect of overlap for frequency ($F(1, 44) = 1.06, p = 0.308, \eta_p^2 = 0.024$), imageability ($F(1, 31) < 1$) nor phonological neighborhood density ($F(1, 94) = 1.94, p = 0.167, \eta_p^2 = 0.020$), suggesting that the testing lists were well matched across overlap conditions. This makes it unlikely that differences in testing performance are due to the particular items included.

S2: Moderator Analyzes at Post-test

We asked whether post-test results held when a variety of moderators were considered. This includes experimentally manipulated factors, like whether a task was trained or untrained, as well as subject level factors like gender. We used the untransformed post-test change scores for additional between-subject (e.g., gender) analyses; in addition, we calculated untransformed post-test change scores for within-subject (e.g., item type) factors. These factors were entered one-by-one into ANOVAs along with overlap and variability. The results are summarized in Table S2.1. While we briefly discuss each class of moderators below, there were a few points in common. The main effect of overlap remained significant ($p < 0.05$) throughout all of the analyses and will not be further reported; similarly, variability ($p > 0.20$) and the three-way interaction of overlap, variability, and moderating factors ($p > 0.30$) never reached significance, and thus will not be included in our description either.

We started by examining task role, i.e. whether a task appeared in training or not.

There was no effect of task role, indicating children exhibited similar learning in tasks they had been trained on and tasks that were only performed at test. That is, learning appeared to generalize across tasks. One of the tasks for each counterbalanced set A or B had been used in word/ nonword version during training (e.g., students in set A completed the *Find the Nonword* task during training and the *Find the Word* task as one of their generalization tasks during testing). At the same time, the remaining generalization task in each set was unique (*Fill in the Blank*; *Make the Word*); thus, these high levels of generalization cannot be due simply to task similarity only.

Next, we examined the different types of task. Performance on spelling tasks (e.g. *Make the Word*; $M = 0.14$) improved significantly more over training than performance on recognition tasks ($M = 0.05$). Variability interacted marginally with task type, suggesting that while students' performance in the similar and variable condition was comparable for recognition tasks, there may have been a slight advantage for variable consonant frames for spelling tasks. This may reflect that variability is particularly valuable for more demanding tasks.

For item factors, there was a significant effect of word role: Training gains differed for trained and novel items. Specifically, students' performance improved most for trained items ($M = 0.11$), and less for novel (generalization) items ($M = 0.08$). However, this did not interact with overlap or similarity. We also observed a main effect of item type, as performance on nonwords ($M = 0.13$) recorded greater gains than the ones on words ($M = 0.07$) across all conditions. Again, this did not interact with either experimental factor.

In contrast to these analyses, when we considered vowel type (monograph or digraph, Figure S2.1A), we found a significant interaction with overlap. As follow-up tests, we split the data by vowel type and compared the overlap conditions within each vowel type using a univariate ANOVA with overlap as the only between-subject factor (with an adjusted $\alpha =$

0.025). We found that there was a significant effect of overlap for digraphs, indicating higher change scores for children assigned to the high overlap condition than the low overlap one ($F(1, 275) = 13.658, p < 0.001, \eta_p^2 = 0.047$). For monographs, however, there was no significant effect of overlap ($p = 0.566$) as a result of similar change scores in both overlap conditions. This indicates that overlap training mattered more for digraphs than monographs (much as was observed for pre-test).

For subject factors, we found a main effect of gender: While girls and boys started off at similar pre-test levels, girls showed greater learning gains ($M = 0.12$) gained more than boys ($M = 0.08$). However, this did not interact with overlap or variability. We did not observe a main effect of reading ability; this shows that children of all skill levels benefited similarly from the training. However, we observed a marginal interaction of overlap with reading ability. Inspecting the data (in Figure S2.1B), there appears to be some evidence that overlap does not play a role for children with low reading scores, but that participants with high reading scores benefitted more from high overlap.

In summary, untransformed change scores were significantly higher for nonwords than words and spelling tasks than recognition tasks, independently of overlap or variability condition. This suggests that students benefitted most from training on items and tasks that required deep understanding of the underlying GPC regularities. Whilst generalization was high across tasks with similar learning rates in both trained and generalization tasks, students performed better on trained items than novel ones. Most importantly, we found that the overlap effect may be moderated by at least two variables – the type of GPC items (monograph or digraph) and the abilities of the learner.

Figure S2.1: A) Post-test change scores for overlap and vowel types (monographs and digraphs). B) Change scores for students with high and low reading scores (split by median) and overlap (averaged across variability condition).

Table S2.1:

Summary of ANOVAs examining moderating factors on the effects of overlap and similarity factors at post-test. Marginal significance is indicated with a star in brackets (<0.1), significance with one star (<0.05), and highly significant results are marked with two stars (<0.01).

Class	Factor	Subcategories	Results	Sig.
Task	Task Role	Trained, Generalization	$F(1, 273) = 2.32, p = 0.129$	
	× Overlap		$F(1, 273) < 1$	
	× Variability		$F(1, 273) < 1$	
	Task Type	Recognition, Spelling	$F(1, 273) = 82.28, p < 0.001$	**
	× Overlap		$F(1, 273) < 1$	
	× Variability		$F(1, 273) = 3.06, p = 0.081$	(*)
Item	Word Role	Trained, Novel	$F(1, 273) = 22.10, p < 0.001$	**
	× Overlap		$F(1, 273) = 1.87, p = 0.173$	
	× Variability		$F(1, 273) < 1$	
	Item Type	Word, Nonword	$F(1, 273) = 60.44, p < 0.001$	**
	× Overlap		$F(1, 273) < 1$	
	× Variability		$F(1, 273) < 1$	
Subject	Vowel Type	Monograph, Digraph	$F(1, 273) = 81.11, p < 0.001$	**
	× Overlap		$F(1, 273) = 10.01, p = 0.002$	**
	× Variability		$F(1, 273) < 1$	
	Gender	Female, Male	$F(1, 269) = 4.25, p = 0.040$	*
	× Overlap		$F(1, 269) = 2.20, p = 0.139$	
	× Variability		$F(1, 269) < 1$	
	Reading Score	Median Split of District Reading Assessment	$F(1, 269) = 1.18, p > 0.2$	
	× Overlap		$F(1, 269) = 3.47, p = 0.064$	(*)
	× Variability		$F(1, 269) < 1$	

S3: Moderator Analyzes at Retention

As with the post-test, we followed up these analyses by investigating how overlap and variability interacted with other factors at retention (Table S3.1). We again used

untransformed change scores for these analyses. The main effect of overlap remained significant throughout ($p < 0.05$); the effect of variability never reached significance.

Retention was higher for trained tasks ($M = 0.06$) than tasks that only appeared in testing ($M = 0.04$); this means that whilst there was high generalization across tasks at post-test, an advantage for having experienced a task during training emerged at retention. At the same time, retention on spelling ($M = 0.10$) tasks stayed better than performance on recognition tasks ($M = 0.00$), with recognition performance being identical to baseline accuracy at pre-test. Neither of these factors interacted with overlap or variability.

For item factors, retention was similar for almost all of the item-level factors that we investigated ($F < 1$). Whereas an advantage for trained tasks was observed at retention (but not post-test), trained items (the effect of word role) showed the opposite pattern with better learning at post-test but no difference at

retention, suggesting that it is easier to exhibit immediate generalization of information across tasks than particular items, but longer term, item differences appear to wash out. We also found a significant effect of vowel type, indicating that participants retained more

Figure S3.1: A) Change scores for both overlap and variability for nonwords only. B) Change scores for both overlap and variability for words only. Whereas no retention was present for words in the low overlap condition, participants had retained information for words in the high overlap condition. Learning for nonwords appeared in both overlap conditions.

information about the digraph target items ($M = 0.07$) than the monograph target items ($M = 0.00$), but this did not interact with condition. This is only partially consistent with our post-test findings (which observed the main effect but not the interaction).

Retention was higher for nonwords than for words; more specifically, learning for words approached zero ($M = 0.01$), whereas students' retention for nonwords was very high ($M = 0.09$). Additionally, item type (word/nonword) also interacted with overlap at retention. More specifically (as seen in Figure S3.1A and S3.1B), training benefits were retained for nonwords in both the high overlap and low overlap condition, but only children assigned to the high overlap “curriculum” retained information for words. As a post-hoc test, we split the data by item type and compared change scores for high overlap and low overlap for words and nonwords separately. For words, we found a significant effect of overlap ($F(1, 265) = 9.594, p = 0.002, \eta_p^2 = 0.035$) as a result of higher change scores in the high overlap condition. For nonwords, the difference between the two curricula was not significant ($F(1, 265) = 2.991, p = 0.085, \eta_p^2 = 0.011$). Together, this means that students retained novel information about nonwords (no matter what condition they were in), but only those in the high overlap curriculum showed any training benefit for words. There was no significant main effect for reading score or gender. This stands in contrast to higher post-test change scores for girls.

In summary, we found that higher learning scores at post-test translated into better retention in the majority of cases (e.g. nonwords, digraphs, spelling tasks), but there were exceptions (e.g. for trained items or girls).

Table S3.1:

Effects of between-subject and within-subject factors at retention. Marginal significance is indicated with a star in brackets (<0.1), significance with one star (<0.05), and highly significant results are marked with two stars (<0.01).

Class	Factor	Subcategories	Results	Sig.
Task	Task Role	Trained, Generalization	$F(1, 263) = 6.170, p = 0.014$	*
	× Overlap		$F(1, 263) < 1$	
	× Variability		$F(1, 263) < 1$	
	Task Type	Recognition, Spelling	$F(1, 263) = 110.144, p < 0.001$	**
	× Overlap		$F(1, 263) < 1$	
	× Variability		$F(1, 263) = 1.363, p = 0.244$	
Item	Word Role	Trained, Novel	$F(1, 263) < 1$	
	× Overlap		$F(1, 263) < 1$	
	× Variability		$F(1, 263) < 1$	
	Item Type	Word, Nonword	$F(1, 263) = 110.696, p < 0.001$	**
	× Overlap		$F(1, 263) = 4.446, p = 0.036$	*
	× Variability		$F(1, 263) < 1$	
Subject	Vowel Type	Monograph, Digraph	$F(1, 263) = 75.659, p < 0.001$	**
	× Overlap		$F(1, 263) < 1$	
	× Variability		$F(1, 263) = 1.617, p = 0.205$	
	Gender	Female, Male	$F(1, 259) = 2.044, p = 0.154$	
	× Overlap		$F(1, 259) < 1$	
	× Variability		$F(1, 259) < 1$	
	Reading Score	Median Split of District Reading Assessment	$F(1, 259) < 1$	
	× Overlap		$F(1, 259) < 1$	
	× Variability		$F(1, 259) < 1$	

S4: All Items

Table S4.1a:

Training Words for High Overlap.

Items Similar	Items Variable	Role	Vowel	Item Type
goat	Goat	Test and Train	OA	W
coal	Coal	Test and Train		W
loan	Loaf	Train		W
moat	soap	Train		W
moan	boat	Train		W
main	main	Test and Train	AI	W
sail	sail	Test and Train		W
fail	rail	Train		W
mail	rain	Train		W
tail	paid	Train		W
mean	mean	Test and Train	EA	W
heat	heat	Test and Train		W
seal	fear	Train		W
meal	beak	Train		W
beam	neat	Train		W
loin	loin	Test and Train	OI	W
foil	foil	Test and Train		W
soil	boil	Train		W
toil	coil	Train		W
coin	join	Train		W
met	met	Test and Train	E	W
pen	pen	Test and Train		W
pet	bed	Train		W
men	leg	Train		W
red	jet	Train		W
tot	tot	Test and Train	O	W
pot	pot	Test and Train		W
hot	bog	Train		W
rod	top	Train		W
con	lot	Train		W

Table S4.2a:

Training Nonwords for High Overlap.

Items Similar	Items Variable	Role	Vowel	Item Type
poat	poat	Test and Train		NW
foat	foat	Test and Train		NW
boam	zoam	Train	OA	NW
hoat	doan	Train		NW
poan	koag	Train		NW
paim	paim	Test and Train		NW
tait	tait	Test and Train		NW
baim	zaip	Train	AI	NW
pait	naid	Train		NW
cail	baip	Train		NW
pean	pean	Test and Train		NW
fean	fean	Test and Train		NW
peam	neap	Train	EA	NW
geat	jeal	Train		NW
feam	veat	Train		NW
foin	foin	Test and Train		NW
goit	goit	Test and Train		NW
poin	woip	Train	OI	NW
foim	joik	Train		NW
foit	poil	Train		NW
het	het	Test and Train		NW
wef	wef	Test and Train		NW
bem	fem	Train	E	NW
pel	cen	Train		NW
tet	jed	Train		NW
mot	mot	Test and Train		NW
pom	pom	Test and Train		NW
fot	rom	Train	O	NW
wof	vop	Train		NW
pol	zon	Train		NW

Table S4.3a:

Training Words for Low Overlap.

Items Similar	Items Variable	Role	Vowel	Item Type
bait	bait	Test and Train		W
gain	gain	Test and Train		W
bail	raid	Train	AI	W
mail	pail	Train		W
fail	fair	Train		W
food	food	Test and Train		W
cool	cool	Test and Train		W
fool	loop	Train	OO	W
roof	room	Train		W
pool	moon	Train		W
feel	feel	Test and Train		W
teen	teen	Test and Train		W
peel	peer	Train	EE	W
reef	weed	Train		W
feed	deep	Train		W
foul	foul	Test and Train		W
loud	loud	Test and Train		W
sour	sour	Test and Train	OU	W
noun	noun	Test and Train		W
pout	hour	Train		W
lid	lid	Test and Train		W
hit	hit	Test and Train		W
bit	fix	Train	I	W
tin	dim	Train		W
pit	kin	Train		W
hut	hut	Test and Train		W
rut	rut	Test and Train		W
but	jug	Train	U	W
gun	run	Train		W
mud	bun	Train		W

Table S4.4a:

Training Nonwords for Low Overlap.

Items Similar	Items Variable	Role	Vowel	Item Type
faip	faip	Test and Train	AI	NW
raif	raif	Test and Train		NW
hain	zaip	Train		NW
sair	naid	Train		NW
tain	fain	Train		NW
mool	mool	Test and Train	OO	NW
moop	moop	Test and Train		NW
hoon	boof	Train		NW
foon	joor	Train		NW
poom	koob	Train		NW
beel	beel	Test and Train	EE	NW
peem	peem	Test and Train		NW
geen	veen	Train		NW
neen	veek	Train		NW
meed	teep	Train		NW
foud	foud	Test and Train	OU	NW
rouf	rouf	Test and Train		NW
foun	pouf	Train		NW
coul	woub	Train		NW
moul	zoun	Train		NW
hin	hin	Test and Train	I	NW
fip	fip	Test and Train		NW
pim	bix	Train		NW
rit	zib	Train		NW
fid	lif	Train		NW
pum	pum	Test and Train	U	NW
mup	mup	Test and Train		NW
fud	sug	Train		NW
fup	bup	Train		NW
lud	rud	Train		NW

Table S4.5a:

Testing Items for High Overlap.

Items Word	Items Nonword	Role	Vowel	Condition
coal	poat	Trained		Overlap
goat	foat	Trained	OA	Overlap
load	voan	Novel-CVC		Overlap
toast	woast	Novel-Cluster		Overlap
main	paim	Trained		Overlap
sail	tait	Trained	AI	Overlap
wait	maip	Novel-CVC		Overlap
paint	baird	Novel-Cluster		Overlap
mean	pean	Trained		Overlap
heat	fean	Trained	EA	Overlap
deal	beal	Novel-CVC		Overlap
beard	teant	Novel-Cluster		Overlap
foil	foin	Trained		Overlap
loin	goit	Trained	OI	Overlap
void	boin	Novel-CVC		Overlap
moist	koist	Novel-Cluster		Overlap
met	het	Trained		Overlap
pen	wef	Trained	E	Overlap
fed	neb	Novel-CVC		Overlap
belt	dest	Novel-Cluster		Overlap
pot	mot	Trained		Overlap
tot	pom	Trained	O	Overlap
not	zot	Novel-CVC		Overlap
bond	bost	Novel-Cluster		Overlap
mat		Baseline	A	Overlap
mood		Baseline	OO	Overlap
haunt		Baseline	AU	Overlap
lap		Baseline	A	Overlap
look		Baseline	OO	Overlap
faun		Baseline	AU	Overlap
	maut	Baseline	AU	Overlap
	mool	Baseline	OO	Overlap
	kaun	Baseline	AU	Overlap
	gar	Baseline	AR	Overlap
	kal	Baseline	A	Overlap
	moop	Baseline	OO	Overlap

Table S4.6a:

Testing Items for Low Overlap.

Items Word	Items Nonword	Role	Vowel	Condition
bait	faip	Trained		Overlap
gain	raif	Trained	AI	Overlap
pain	maip	Novel-CVC		Overlap
waist	baint	Novel-Cluster		Overlap
cool	mool	Trained		Overlap
food	moop	Trained	OO	Overlap
toot	bool	Novel-CVC		Overlap
booth	post	Novel-Cluster		Overlap
feel	beel	Trained		Overlap
teen	peem	Trained	EE	Overlap
keen	veem	Novel-CVC		Overlap
teeth	neest	Novel-Cluster		Overlap
loud	rouf	Trained		Overlap
foul	foud	Trained	OU	Overlap
rout	zoom	Novel-CVC		Overlap
couch	foust	Novel-Cluster		Overlap
hit	hin	Trained		Overlap
lid	fip	Trained	I	Overlap
fit	mif	Novel-CVC		Overlap
hint	bist	Novel-Cluster		Overlap
hut	pum	Trained		Overlap
rut	mup	Trained	U	Overlap
cup	vud	Novel-CVC		Overlap
bunk	tust	Novel-Cluster		Overlap
met		Baseline	E	Overlap
haunt		Baseline	AU	Overlap
foil		Baseline	OI	Overlap
pet		Baseline	E	Overlap
coin		Baseline	OI	Overlap
faun		Baseline	AU	Overlap
	nep	Baseline	E	Overlap
	kaun	Baseline	AU	Overlap
	fem	Baseline	E	Overlap
	doil	Baseline	OI	Overlap
	gaut	Baseline	AU	Overlap
	goit	Baseline	OI	Overlap