velocity™

RESEARCH FOUNDATION

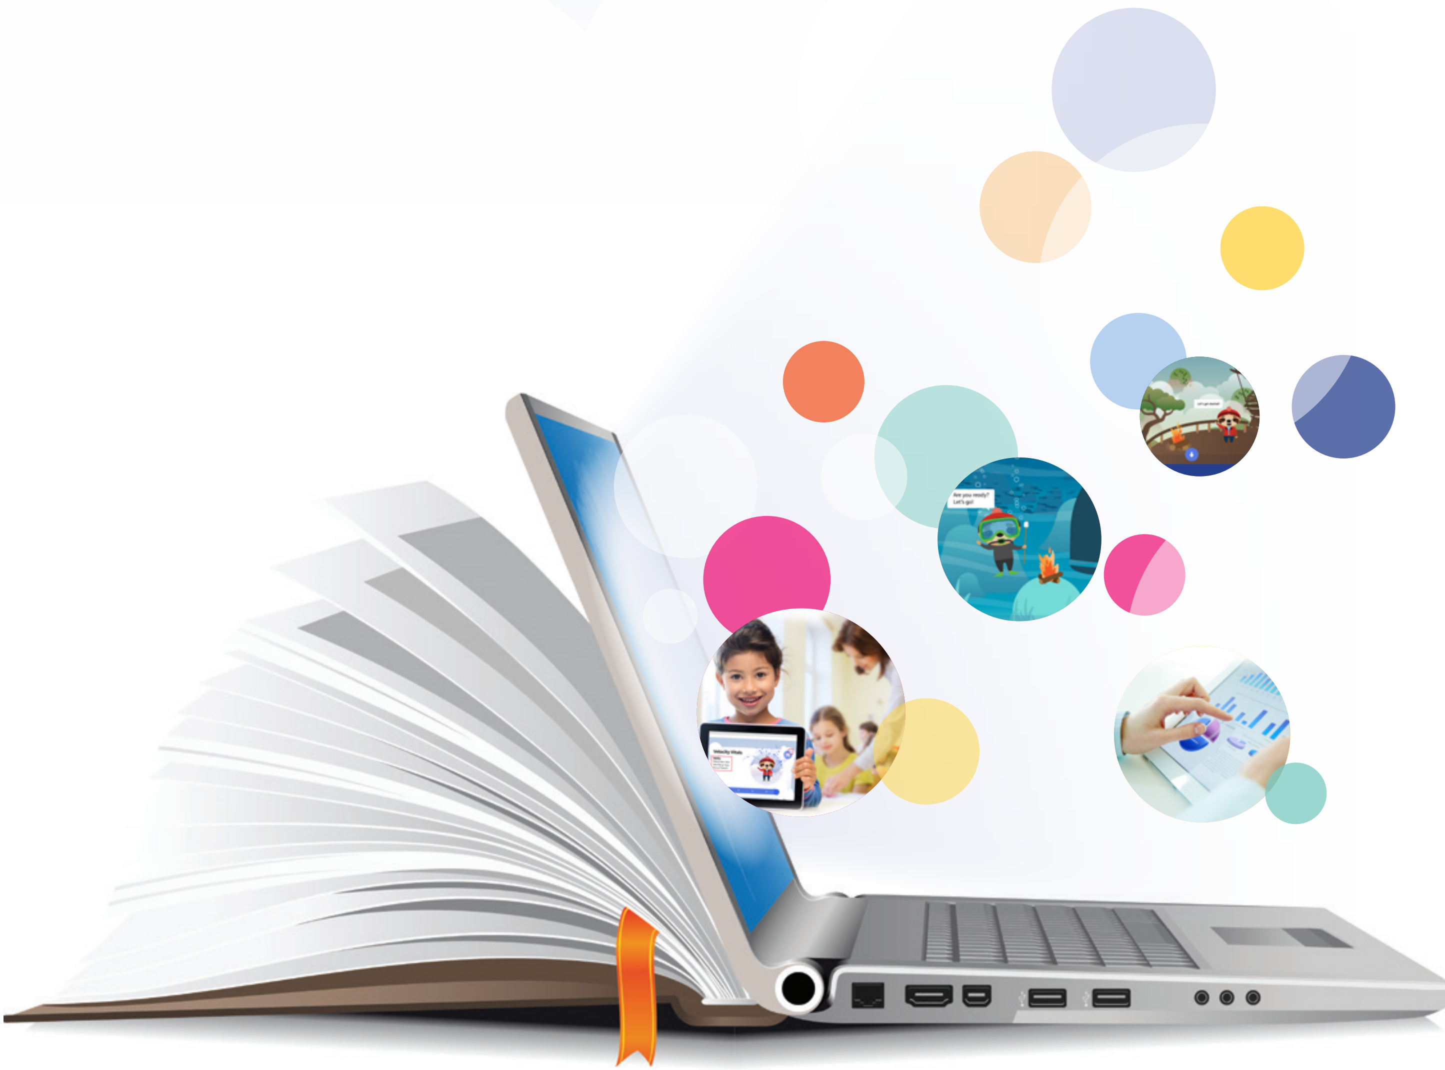

DIGITAL ENGLISH LANGUAGE ARTS
INSTRUCTIONAL SOLUTION FOR GRADES K–5

VOYAGER SOPRIS LEARNING®

Introduction

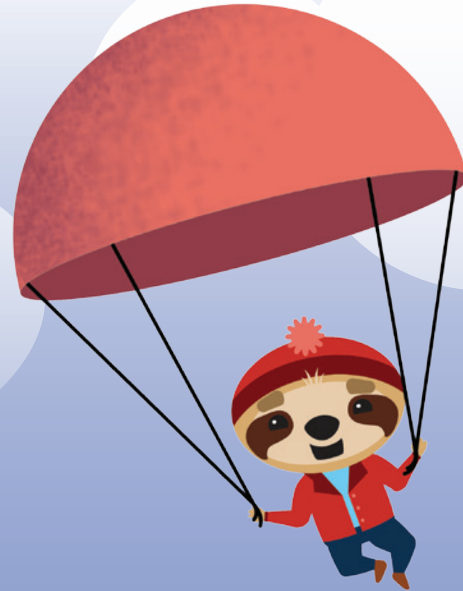

Velocity is defined as the speed of something in a specific direction. *Velocity*™ also is a supplemental digital English Language Arts (ELA) instructional solution for students in kindergarten to fifth grades. Incorporating the use of machine learning, *Velocity* identifies the best learning paths to achieve the best possible results more quickly, helping students move toward and achieve grade-level goals and standards. *Velocity*, designed from the ground up, incorporates pedagogically-sound content with teacher-led lessons and teacher dashboards, providing up-to-the-minute information about student progress, including student needs only a teacher can satisfy.

Enlearn, a Seattle nonprofit, partnered with Voyager Sopris Learning®, one of the largest U.S. educational publishers of intervention and supplementary materials in grades K–12, to create *Velocity*. Each organization is contributing what we do best, Enlearn with the learning platform for adapting and personalizing digital educational materials, and Voyager Sopris Learning with sound, targeted ELA content based on years of successful creation of programs for students struggling to learn to read and do math.

This report elaborates on how *Velocity* works. Since *Velocity* is so different, we start with a machine learning primer and a brief review of the Enlearn Learning Platform that powers the machine learning in *Velocity*. Next, the ELA content will be discussed since it was developed to support *Velocity*, not repurposed from a previous program. The teacher and student experience will be briefly discussed. Briefly because seeing is necessary to get the real flavor of why teachers and students like what they see. This report will close with a summary of why *Velocity* is the right solution for most elementary classrooms.

Machine Learning Primer[1]

Machine learning is a branch of computer science and statistics devoted to predicting the future given the past. Machine learning is all around us and is responsible for showing targeted advertisements, predicting which movies you'll like, automatically identifying credit card fraud, protecting your email inbox from spam, reading zip codes on mail, and more. In other words, machine learning is about doing better in the future based on what was learned or experienced in the past.

In the case of personalized learning, machine learning is continuously trying to figure out what educational intervention to give a student that is personalized to their specific abilities, needs, and learning context. The end result is that through machine learning, students are given the precise content they need, at the right time, so they can learn more and learn faster. This is the aspect of machine learning *Velocity* is using.

Alternatives to Machine Learning

People have been making advancements for a long time without the help of machine learning. What are the pros and cons of other possible approaches?

Rely on expert knowledge:

If a subject, topic, or concept is well-understood, there might not be a need to run an experiment because a solution already is known. The system would tell the student what to do. The downside is, if there is something better, the student will not have the opportunity to find it. Rule-based or heuristic adaptive learning systems fall into this category. The expert might declare students should learn one concept until they can answer three different problems in a row without errors, then move on to the next concept. However, if the expert was wrong and students should really learn one concept until they are able to answer five problems in a row without errors, the system will never learn this fact.

Through machine learning, students are given the precise content they need, at the right time, so they can learn more and learn faster.

Direct evaluation:

Designers commonly brainstorm multiple, very different approaches to solving a problem, then test them on users and pick the one that gets the best response. This would be similar to giving a lab technician several recipes for producing a medicine. She would run a series of experiments, select the best one based on some set of criteria, and use that recipe forever after.

A/B tested learning systems fall under this category. For example, we might have five possible rules about how a system should adapt, and three different sets of content. We could run a fifteen-way (5x3) experiment with the different rules and content, either in the lab or with students in a classroom, then pick the one rule

[1] This portion of the document and the next, The Enlearn Learning Platform, were adapted from the paper, "Machine Learning: A Primer" by Yun-En Liu, Ph.D., Data Scientist at Enlearn.

and set of content that seemed to lead to the greatest learning. These experiments can be expensive and time consuming because of the number of different ways to combine content, algorithms, and interfaces.

Generalization, a.k.a. scientific method:

Scientists try to find rules or principles that are generalizable, so they can predict what will happen in the future when circumstances are different than they are now. For example, if we know students praised in the classroom for being smart then prefer to avoid difficult problems that challenge their intelligence (Dweck, 2006), then we might suspect a digital system should avoid rewarding high performance, and instead, reward effort. By finding these general rules or principles, we can run fewer experiments than is possible with the direct approach. For instance, we could find the best of five adaptive algorithms, the best of three sets of content, then assume that we should just pick the best of both (i.e., 5+3=8 experiments instead of 5x3=15 experiments).

Adaptive algorithms based on educational or psychological research fit into this category. For instance, expanded spaced-retrieval algorithms for memorizing flashcards are based on a large body of cognitive psychology evidence that suggests students remember materials better when there are increasing intervals of time between subsequent reviews or tests (cf. Karpicke & Bauernschmidt, 2011; Karpicke & Roediger, 2010). The potential downside of the scientific method is the generalization might not be valid. The particular group of students participating in a study might prefer to randomly guess answers rather than try to recall the right answer when in the classroom. In that case, the expanded spaced retrieval might not work at all for this group.

Machine learning:

Machine learning is actually a mix of all the different components mentioned above. It relies on expert knowledge in the initial setup of the algorithms, including the policies that govern the kinds of experimentation allowed and the tests that can be used to evaluate outcomes. It then collects data and uses statistical techniques to directly evaluate different options, determining, perhaps, content A is best when students answered the last question wrong, while content B is best when students answered the last question correctly. Finally, the choice of algorithm and what we think makes students similar or different is used to generalize the knowledge gained, determining, perhaps, we now serve content A for all students who struggle with a certain concept and content B for all students who do not struggle with that concept. The best machine learning algorithms will do all of this automatically: (a) make efficient use of data, (b) are guaranteed to converge to the best option, and (c) do all of this in the least amount of time possible. That is why machine learning is such a promising technology.

The Enlearn Learning Platform

Enlearn's core machine learning platform consists of three levels of adaptation. The top layer, skills adaptation, keeps track of which skills a student has learned and chooses skills to improve on that lie just beyond the student's reach—the so-called proximal zone of development. The bottom layer, step adaptation, keeps track of errors a student makes within the many steps of an individual problem, giving different amounts of hints and scaffolding, depending on the errors students have made on similar steps in the past. Through step adaptation, the Enlearn Platform is able to detect student misconceptions or learning obstacles within a problem prior to the final answer, and then provide scaffolding to help course correct immediately during the problem, rather than waiting for a final answer as is typical of most adaptive applications in education.

The middle layer, problem adaptation, lies at the heart of the system and is responsible for choosing a specific problem to give once skills adaptation has decided on the topic that students should work on. It relies on a machine learning technique called reinforcement learning (Sutton & Barto, 1998), which is designed specifically for environments that are known to require adaptive behavior. Reinforcement learning is most commonly used in the field of robotics to plan how the robot should move to achieve its goals (Kober, Bagnell, & Peters, 2013). Using an example with a robot provides a great illustration of reinforcement learning.

At a high level, reinforcement learning works by trying different actions in different states, seeing what happens in response, and measuring whether anything good has occurred. A robot coffee server, for example, might try moving different limbs (actions) depending on its orientation and location (state), observing what happens, such as where it moves or whether it falls over (observation), and eventually discovers that something good has happened, such as picking up a cup of coffee (reward). Over time, the algorithm learns how actions affect the state (rolling one wheel only causes the robot to turn but not move), and what states are good (being physically closer to a coffee cup or having an arm out with fingers open). With this information, the algorithm can generate ever-more efficient plans to achieve its goals.

With *Velocity*, Enlearn uses reinforcement learning to understand what problems within a concept to give next to students.

With *Velocity*, Enlearn uses reinforcement learning to understand what problems within a concept to give next to students, with the goal of eventually answering the most difficult or complex problem correctly. Let's look at a type of comprehension problem from *Velocity*, reading short passages and finding the moral of the stories. This type of problem can be easier or harder along several independent dimensions. For instance, one passage might contain more advanced vocabulary, another might use more complicated sentence structure, another might require more background knowledge about American culture, and yet another might use satire or other rhetorical devices that obscure the main idea.

Dimensions span a wide range of difficulty or complexity, and the goal of reinforcement learning is to give problems of different difficulty or complexity along specific dimensions to students until they are able to do the problems with the greatest level of difficulty or complexity across all the dimensions (reward). To accomplish this, the algorithm can give a new problem that is more or less difficult or complex along one of the dimensions than the previous problem along one of the dimensions (action). The system measures whether students were able to successfully complete the steps of the problem, in this case, selecting the sentences pertaining to the moral from the passage, and then choosing the right moral from a multiple-choice list (observation). The system remembers the difficulty or complexity of the previous several questions and whether the student was able to answer them correctly (state) and uses this information in subsequent actions.

A complete picture shows the reinforcement learning algorithm is constantly analyzing the sequences of detailed responses it has seen. With this information, it offers new patterns of problems to students in the hopes of discovering better learning pathways. For example, it might learn that students who correctly answered two problems with high text difficulty but low required background knowledge are now well-equipped to answer a question with high text difficulty and medium required background knowledge. With this type of information, the system can plan a series of problems to teach a particular concept, such as the moral of a story, to students so they will be able to answer the hardest problem in minimum time, adjusting its plan depending on student responses.

This system is allowed to give new problems to students that are similar in difficulty to the ones before, adjusting difficulty up or down across one or more dimensions. It can thus learn which dimensions of difficulty should be made easier or harder, depending on what students get right or wrong. This very fine-grain form of additivity constrains the system from making wild leaps with problems suddenly getting much easier or harder.

Since the system learns which dimensions should be increased or decreased, given the dimension difficulty of the past few problems the student did and whether or not they were solved correctly, it can discover interesting things about the dimensions. For instance, the algorithm might learn that some dimensions should be kept at the lowest level until all the others are mastered, or when a student is struggling with problems that are difficult among all dimensions, the best thing to do is to make the text difficulty dimension easier. The algorithm can only do this because it assumes most students respond in similar ways, or in other words, it can generalize across students. It currently cannot learn that Elena is more adventurous than Tommy and so she should be given harder problems even when she answers one incorrectly. Whether or not adding these types of student parameters would be worthwhile and something we should attempt to model, is something we would like to determine in the future.

Enlearn is a research-driven organization and while reinforcement learning is not new and has been used in many applications, the devil is in the details. Enlearn created this specific version of algorithm based on a study (Mandel, Liu, Brunskill, & Popović, 2016) in which a reinforcement learning algorithm was invented to serve math problems adaptively in an online fraction game. The reinforcement learning algorithm was superior in the experiments to several other options, including a standard learning progression. The algorithm used by *Velocity*, is a variation on a class of algorithms best-suited for problems where it is unknown at the start what will happen when we try different actions (Osband, Russo, & Van Roy, 2013). This class of algorithms is probably optimal at still doing a good job even while learning how the world works. Finally, Enlearn scientists have invented ways to use past data to assess how new algorithms would behave had they been used instead (Liu, Mandel, Brunskill, & Popović, 2014). This ability allows new algorithms and approaches to be tested and validated (or refined or discarded) prior to directly involving students.

English Language Arts Content

The ELA content in *Velocity* is based on a pedagogically sound, well-designed, standards-based, research-driven curriculum.

In a typical ELA program for struggling or non-struggling students, one or two lessons about a particular skill are available, which students might work on until mastery is achieved or there might be some repetition of the skill in later lessons. Either way, there is a limited or finite number of lessons about that particular skill. The sequence of lessons is fixed, with little room for modification to meet a particular student's need. Even in most digital programs, there is a predetermined path or paths that students are expected to take, again with little room for modification when the paths do not meet the needs of all students. This is not the case for *Velocity*, where the lessons are abundant and sequence is not predetermined.

The ELA content in *Velocity* is based on a pedagogically sound, well-designed, standards-based, research-driven curriculum. Additionally, *Velocity* is organized in a logical, predictable way using strands that contain similar skills. The strands are Comprehension, Foundational Skills, Language, and Vocabulary and Word Study. This is where the similarities between traditional programs and *Velocity* diverge.

For each skill in *Velocity*, the following process is used to prepare content for students:

» Based on the expertise of seasoned curriculum developers, skills are identified for inclusion and initial design ideas are brainstormed and collected.

» A structure for teaching that skill, called a Problem Type (PT), is designed by curriculum specialists, user-experience designers, and software engineers. A PT is used to create as many lessons, called problems, as needed to meet the needs of all students. Additionally, multiple PTs could be built for a single skill when there are multiple pedagogically-sound ways to approach the teaching of the skill.

» Within the PT, there are dimensions defined, unique to the nature of that specific skill, that serve to make each of the problems more or less cognitively difficult or complex. As each problem is built, the appropriate dimensions are assigned by the problem's author and are related to the particular content within that problem.

» Each PT will contain some number of steps. Each step is related to the metacognition and/or mental model of the skill being encoded in the PT. Steps also provide a source of data for the system to interpret student performance and identify the best next step. This use of steps is also is known as formative assessment or embedded assessment and it is nearly continuous as the student moves through *Velocity*, eliminating the need to stop and administer a separate assessment.

» Scaffolds, specific to each step in a problem, are incorporated to add support and/or additional instruction. Examples of scaffolds include: providing a mental model; turning on or off audio support; controlling the amount of text shown to the student; removing or revealing an answer to a particular question; providing audio or video hints or metacognitive support; showing the paragraph where an answer can be found; or completing part or all of a graphic organizer.

» When the PT and the problems have been built and tested, they are added to the ecosystem containing the existing PTs and problems. Then, these become available to students through the system.

Problems that share a PT and the same dimensions are called an equivalence set because these problems are all at an equivalent level of cognitive difficulty. Equivalence sets allow students to receive more than one problem at a particular level of cognitive difficulty, without repeating the exact same content over and over. For example, a particular student is having difficulty with a certain skill because she has a weakness in a prerequisite skill. The engine will recognize this weakness, based on student performance (embedded assessment), and serve up the prerequisite skill next. After successfully remediating the prerequisite skill, the engine will take the student back to the original problem, but instead of serving up the same content, will pick an equivalent problem for the second time through. Through the use of PTs, problems, and equivalence sets, almost unlimited content could be available for students and it is unlikely students with similar needs will see the exact same content.

Velocity's real-time data will be available to teachers so they can see not only who is online at any given moment, but also who is struggling and in need of teacher attention.

Teachers

It might seem like there is no need for the teacher with *Velocity*, but that could not be farther from the truth. Enlearn and Voyager Sopris Learning are firm believers in the fact that teachers make a difference in the success of students. Our collective goal is to help teachers so more students can be helped by teachers. During the spring 2014 classroom trials with Enlearn's generative adaptive platform, one of the central findings about the platform's real-time data was teachers were able to assist individual students three times more frequently than occurred in the traditional paper-based classrooms, leading to better student outcomes (Catalano, 2015; Gottlieb, 2016).

Velocity's real-time data will be available to teachers so they can see not only who is online at any given moment, through the use of dashboards, but also who is struggling and in need of teacher attention. In fact, the amount of data collected and used to make decisions about the students using *Velocity*, is almost unbelievable and certainly more than could be managed by any teacher, especially with a class full of students. The real-time data will also be serving up notifications to the teacher about student events that deserve recognition and action, such as a student showing persistence with a particularly difficult problem or a small group of students who are ready for a teacher-led lesson.

Teacher-led lessons are designed to mirror the metacognition and mental models used in the online problems, but are written for teacher delivery to a student or a small group of students. Teacher-led lessons serve three purposes: (1) introduction of a skill that either needs additional explanation to be understandable by students or one that is not appropriate for an online environment, such as a speaking and listening skill where a collaborative conversation is required; (2) remediation of a skill when a student or students are not quite grasping a concept; and (3) demonstration of mastery and/or application of a skill when students have reached the greatest difficulty on all dimensions for a particular PT.

Students

In *Velocity*, the content is not the only aspect being analyzed and adjusted based on student performance. Student engagement is equally important in *Velocity*. *Velocity* pushes students to become self-regulated learners who monitor the quality of their work, while at the same time supporting, encouraging, and rewarding students for their effort and productive struggle while working. This engagement strategy is based on growth mindset (Dweck, 2006) and incentive structures work (O'Rourke, Haimovitz, Ballwebber, Dweck, & Popović, 2014). Those rewards are called *Velocity* Points and can be used within the system for an ever-growing list of items and activities. While in *Velocity*, whimsical characters guide students through their online experience, sometimes providing help and hints, at other times providing moments of delight and intrigue, especially during visits to other worlds that exist in the Nexus.

The Curriculum Content Research

Due to a lack of consensus of the most appropriate methods of teaching reading, a mandate was issued by the U.S. Congress for a synthesis of the research about reading instruction, with the hope of influencing practice (Wanzek & Vaughn, 2008). This synthesis of research was the National Reading Panel (NRP; 2000) report. This report identified five critical components of reading instruction necessary to gain reading skill and reach the goals of reading fluently and comprehending text. These five critical components are: (a) phonological awareness, (b) phonics, (c) fluency, (d) vocabulary, and (e) comprehension. Wanzek and Vaughn present other summaries of research about effective reading instruction and effective intervention of reading and they report these summaries concur with the NRP, "explicit instruction in components of reading involving decoding words effectively, fluency and comprehension" (p. 126) are necessary components for students to achieve success with reading.

Based on research, the NRP components, influenced by national and state standards, are included in the *Velocity* curriculum.

PHONOLOGICAL AWARENESS

Phonemes are the basic building blocks of spoken language and words, or the smallest units that make one word different from another (Moats, 2000). The awareness of the smaller segments of individual speech sounds is referred to as phonemic awareness, and is part of the more encompassing term, phonological awareness. Phonological awareness refers to "an awareness of, and ability to manipulate, the sound structure of oral language" (Ouellette & Haley, 2013, p. 30).

The ability to hear and manipulate the individual sounds in words accounts for significant differences between good and poor readers (Melby-Lervåg, Lyster, & Hulme, 2012). Those students who do not attend to individual sounds and do not recognize the relationship of sounds to letters as they begin to read will be at risk of reading difficulties. In other words, for students to become readers, they must discover the words they hear come apart into smaller pieces of sound (Shaywitz, 2003). Shaywitz and Shaywitz (2004) confirm, based on highly reliable scientific evidence, reading disability correlates specifically to a weakness in accessing the sounds of spoken language.

The critical skills in phonological awareness are blending, combining sounds together to make a larger segment such as a syllable or word, and segmenting, the ability to break a word down into individual sounds (Ouellette & Haley, 2013). Studies (Anthony, Lonigan, Driscoll, Phillips, & Burgess, 2003; Lonigan et al., 2009) noted that blending and segmenting develop differently, with blending skills developing before segmenting. Even though students may begin to blend and segment phonemes without connecting the sounds to letters, phonemic awareness training works best and is most beneficial when it is combined with practice in connecting sounds to letters (Linan-Thompson & Vaughn, 2007).

Phonological and phonemic awareness are found in the Foundational Skills strand in *Velocity*.

PHONICS AND WORD STUDY

Phonics instruction is the systematic use of sound-symbol relationships to teach the reading and writing of words, and builds on phonemic and phonological awareness (Mesmer & Griffith, 2005; Moats, 2000). The goal of phonics instruction is to teach students the relationship between spoken sounds and printed letters for use in decoding and spelling words (Chall, 1967). Phonics is an important component of literacy instruction because English is fundamentally a code, an alphabetic code, where spoken language is represented using letters, also known as graphemes, which correspond with sounds, or phonemes (Mesmer & Griffith, 2005; Moats, 2000).

As young students become aware of the alphabetic principle, the relationship of sounds to letters, they need to learn the single-letter consonants as well as consonant blends and consonant digraphs. Beginning readers need to learn short and long vowels, the -r controlled vowels, and vowel digraphs. Knowing these common patterns automatically enhances both decoding and spelling (Henry, 2003; Moats, 2000). An additional useful strategy in phonics instruction is the onset-rime (Bryant, 2002). Bryant suggests onset-rime has two effects: (1) helping children, who hardly know phonemes, to be ready to divide up units into smaller units of phonemes, and (2) helping children to learn about letter sequences, such as "-ight," when it is time to do so.

Students in the upper grades also benefit by learning syllable patterns as well as the common morpheme (or meaning) units such as prefixes, suffixes, and some Latin and Greek root words (Ehri, 2005; Mesmer & Griffith, 2005). Learning these patterns assists students in not only decoding and spelling words, but in understanding the meaning of those words. This link to vocabulary is especially useful and valuable. Words with Latin roots are found extensively in literature and social studies text, while the Greek roots aid comprehension of math and science terms.

Phonics and word study are found primarily in the Foundational Skills strand, but also in the Vocabulary and Word Study strands of *Velocity*.

FLUENCY

According to Rasinski (2006), "fluency is, in a sense, a bridge between phonics and word decoding on the one hand, and vocabulary (word meaning) and comprehension (passage meaning) on the other" (p. 62). Rasinski (2012) elaborates on fluency by stating "fluency has two essential components: automaticity and prosody" (p. 517). Rasinski (2012) defines automaticity as the ability to recognize words automatically, without effort, and prosody as the connection to comprehension. Rasinski and colleagues (2016) further clarify the definition of fluency to include three components: "(1) accuracy in word decoding, often referred to as competency in phonics, (2) automaticity in word recognition, and (3) appropriate use of prosodic features (oral expression) such as stress, pitch and suitable phrasing" (p. 165).

Wolf (2007) describes fluency as "not a matter of speed; it is a matter of being able to utilize all the special knowledge a child has about a word—its letters, letter patterns, meanings, grammatical functions, roots, and endings—fast enough to have time to think and comprehend" (pp. 130-131). Fluency and comprehension, although separate processes (Wolfe & Nevills, 2004), are mutually intertwined. That is, fluency is both a cause and a consequence of comprehension.

Fluency is incorporated into the Comprehension, Language, and Vocabulary and Word Study strands in *Velocity*.

VOCABULARY

Vocabulary refers to the words a person understands and uses in listening, speaking, reading, and writing. Vocabulary is directly related to reading comprehension as students try to make meaning of the words in text. Students learn word meanings through direct and indirect, also known as incidental, experiences with oral and printed language (Beck, McKeown, & Kucan, 2002; National Reading Panel, 2000; Phythian-Sence & Wagner, 2007). Carlisle (2007) states "learning an unfamiliar word begins when it is encountered in an oral or written language context and when understanding of that word matters to the listener or reader" (p. 82).

Phythian-Sence and Wagner (2007) suggest direct instruction of words needs to include repeated exposures to words, their definitions, and contextual information, allowing students to explore the meaning of new words instead of just trying to memorize them. Spencer, Goldstein, and Kaminski (2012) concur that effective vocabulary instruction is explicit, intentionally designed, with carefully selected vocabulary targets (most important words to teach), and linked to instruction.

Phythian-Sense and Wagner (2007) also state "comparing estimates of the number of words that children acquire through their school years with estimates of the number of words they acquire through direct instruction suggests an important role for incidental learning" (p. 7). Repeated read-alouds, especially with some direct explanation of word meanings in context with young students (5 to 8 years old) showed the opportunity for incidental learning. Children's acquisition of word meaning in oral contexts was found to extend to written contexts as they learned to read, finding that students could acquire knowledge of unknown words in context-rich paragraphs without explicit instruction.

Vocabulary instruction is mainly found in the Vocabulary and Word Study strand, but is also an essential part of the Comprehension strand in *Velocity*.

COMPREHENSION

Comprehension is the action or capability of understanding something. Listening comprehension involves multiple processes in the understanding and making sense of spoken language. Reading comprehension is defined as the level of understanding of a text that comes from the interaction between the words that are written and the way the words trigger knowledge outside the text. Reading comprehension is a direct result of active reading in which readers think about their reading, making connections and inferences to understanding text. Duke and Pearson (2008) summarize the characteristics of good readers to include: good readers are active readers; have clear goals and evaluate whether they are meeting their goals; preview and make predictions; try to determine the meaning of unfamiliar words in the text; integrate their prior knowledge with material in the text; identify main ideas and summarize; engage in self-questioning and visualization; and more.

Biemiller (2003) provides a summary of the development of comprehension, starting with listening comprehension that begins at about 12 months of age, followed by the beginning of reading comprehension in kindergarten and first grade. At this early point in the development of reading comprehension, listening comprehension is far superior to reading comprehension for the majority of students, with reading comprehension lagging behind listening comprehension past the third grade. It is through the interactions with people and texts that experiences with new vocabulary, concepts, and language structures occurs.

It is necessary and important, according to Biemiller, for students to have growth opportunities in comprehension, to be composed of both reading experiences and non-print sources (oral and audio) through the elementary years.

McLaughlin (2012) identifies the role of the teacher as creating "experiences and environments that introduce, nurture, or extend students' abilities to engage with text" (p. 434) and that this must be done using explicit instruction, which includes modeling, scaffolding, facilitating, and participating. Moats (2005) states students need to learn the elements of both narrative and expository texts. Additionally, explicit comprehension instruction should include the use of modeling, thinking aloud, questioning, summarizing, and other techniques that promote active construction of meaning.

Comprehension is explicitly taught in the Comprehension strand in *Velocity*, but is supported by the development of vocabulary in the Vocabulary and Word Study strand, the explicit teaching of grammar and language structures in the Language strand, and the development of the ability to read fluently developed by instruction in the Foundational Skills strand.

The Classroom Experience

The latest statistics reported by the National Center for Education Statistics (NCES; www.nces.ed.gov) for the 2011-2012 school year, reported the average class size for teachers in elementary school in self-contained classes was 20.3 students. Since this was an average, we know half the classes were below 20.3 students and half had more than 20.3 students. As noted by Dixon, Yssel, McConnell, and Hardin (2014), "because any classroom with more than one student presents a range of diverse learning needs, teachers often struggle to provide all their students with focused learning activities specifically designed to what works best for them" (p. 112). Moats (2009), on a survey of 52 graduate students who were practicing teachers, found "significant difficulty on items that asked them to identify words with consonant blends, consonant digraphs, inflectional and derivational morphemes, and position-based spelling patterns such as the use of the spelling"—ck (p. 387). Moats (2009) goes on to explain "levels of knowledge were not related to whether the teachers were in special or regular education roles, or how many years they had taught" (p. 387).

Fortunately, prevention and intervention requires the same focus, with the understanding of the prerequisite skills needed to become a good reader (Kilpatrick, 2015). However, even a well-trained, knowledgeable teacher with 20 students in her class will find it difficult to deal with the disparate groups in the class, the grade-level, above grade-level, and below grade-level students, remembering that what works for one student in one of these groups will not necessarily work for the others.

Tomlinson and Jarvis (2009) explained "differentiation is an approach to curriculum and instruction that systematically takes student differences into account in designing opportunities for each student to engage with information and ideas and to develop essential skills" (p. 599). Dixon et al. (2014) states "teachers who do not recognize ways to differentiate or who do not feel capable of instructing different groups at the same time struggle with differentiating instruction" (p.113).

Since 1984, when Bloom's students originally found that the average student who received one-on-one tutoring from an adult was two standard deviations above the average of the control classes, with about 30 students per teacher who periodically received tests for scoring purposes, one-to-one tutoring has been a comparison to other types of interventions with students having difficulty learning. Elbaum, Vaughn, Hughes, and Moody (2000) indicate when one-to-one instruction is provided in a regular classroom, it usually amounts to about a one-minute interaction that often clarifies information, answers student questions, or checks for understanding, instead of providing systematic, remedial instruction. According to Fuchs, Fuchs, and Vaughn (2014), based on their experience in schools, few educators know how to develop and deliver intensive interventions. Further, they state "many schools do not have the know-how to provide specialized intensive intervention and, therefore, cannot offer "full-spectrum" instruction to all its students" (Fuchs et al., 2014, p. 15).

THE PROBLEM:

- Classes in most elementary schools contain about 20 students.
- Many, if not most, of the students struggle and need something other than one-size-fits-all instruction.
- Even if teachers are adept at identifying what each student needs to be successful, the preparation and delivery of differentiated instruction to all students in need in the class is a struggle, if not impossible.
- One-to-one instruction works, but when used only lasts about a minute and usually does not provide the systematic, remedial instruction required by students.

With this problem in mind, we asked ourselves, how can we help students who struggle with reading and English Language Arts and at the same time help teachers?

THE ANSWER:

- Take the best elements from differentiated instruction and one-to-one tutoring, and
- Combine that with the most innovative machine learning algorithms, embedded assessment, innovative curriculum design, and teacher-led lessons, then
- Provide teachers with real-time data about students on easy-to-use dashboards and reports, and
- Motivate and engage students to keep working at the right level of instruction on the skills they need to learn to become stronger students and better readers.

In other words, the answer is *Velocity*.

References

Anthony, J. L., Lonigan, C. J., Driscoll, K., Phillips, B. M., & Burgess, S. R. (2003). Phonological sensitivity: A quasi-parallel progression of word structure units and cognitive operations. *Reading Research Quarterly, 38*(4), 470-487.

Beck, I. L., McKeown, M. G., & Kucan, L. (2002). *Bringing words to life: Robust vocabulary instruction*. New York, NY: Guilford Press.

Biemiller, A. (2003). Oral comprehension sets the ceiling on reading comprehension. *American Educator, 27*(1), 23-23.

Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13*(6), 4-16.

Bryant, P. (2002). It doesn't matter whether onset and rime predicts reading better than phoneme awareness does or vice versa. *Journal of Experimental Child Psychology, 82*(1), 41-46.

Carlisle, J. F. (2007). Fostering morphological processing, vocabulary development, and reading comprehension. In R. K. Wagner, A. E. Muse, and K. R. Tannenbaum, (Eds.), *Vocabulary acquisition: Implications for reading comprehension* (pp. 78-103). New York, NY: Guilford Press.

Catalano, F. (2015, March 5). *Seattle nonprofit Enlearn tackles thorny task: Personalize school with technology.* Retrieved from http://www.geekwire.com/2015/seattle-nonprofit-enlearn-tackles-thorny-task-personalize-school-with-technology.

Chall, J. S. (1967). *Learning to read: The great debate.* New York, NY: McGraw-Hill.

Dixon, F. A., Yssel, N., McConnell, J. M., & Hardin, T. (2014). Differentiated instruction, professional development, and teacher efficacy. *Journal for the Education of the Gifted, 37*(2), 111-127.

Duke, N. K., & Pearson, P. D. (2008). Effective practices for developing reading comprehension. *The Journal of Education, 189*(1/2), 107-122.

Dweck, C. (2006). *Mindset: The new psychology of success.* New: Random House.

Ehri, L. C. (2005). Learning to read words: Theory, findings, and issues. *Scientific Studies of Reading, 9*(2), 167-188.

Elbaum, B., Vauchn, S., Hughes, M. T., & Moody, S. W. (2000). How effective are one-to-one tutoring programs in reading for elementary students at risk for reading failure? A meta-analysis of the intervention research. *Journal of Educational Psychology, 92*(4), 605-619.

Fuchs, D., Fuchs, L. S., & Vaughn, S. (2014). What is intensive instruction and why is it important? *Teaching Exceptional Children, 46*(4), 13-18.

Gottlieb, S. (2016, August 18). *Enlearn's spring 2014 classroom trials demonstrate more effective teaching plus increased student problem-solving and mastery.* Retrieved from http://www.prweb.com/releases/2014/07/prweb12014421.htm.

Henry, M. K. (2003). *Unlocking literacy: Effective decoding and spelling instruction*. Baltimore, MD: Paul H. Brooks Publishing.

Karpicke, J., & Bauernschmidt, A. (2011). Spaced retrieval: Absolute spacing enhances learning regardless of relative spacing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 37*(5), 1250-1257.

Karpicke, J., & Roediger, H. (2010). Is expanding retrieval a superior method for learning text materials? *Memory & Cognition, 38*(1), 116-124.

Kilpatrick, D. A. (2015). *Essentials of assessing, preventing, and overcoming reading difficulties*. Hoboken, NJ: John Wiley & Sons.

Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research, 32*(11), 1238-1274.

Linan-Thompson, S., & Vaughn, S. (2007). *Research-based methods of reading instruction for English language learners, Grades K-4*. Alexandria, VA: ASCD.

Liu, Y. E., Mandel, T., Brunskill, E., & Popović, Z. (2014). Toward automatic experimentation of educational knowledge. In P*roceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 3349-3358). New York, NY: AMC.

Lonigan, C. J., Anthony, J. L., Phillips, B. M., Purpura, D. J., Wilson, S. B., & McQueen, J. D. (2009). The nature of preschool phonological processing abilities and their relations to vocabulary, general cognitive abilities, and print knowledge. *Journal of Educational Psychology, 101*, 345-358.

Mandel, T., Liu, Y. E., Brunskill, E., & Popović, Z. (2016). Efficient Bayesian clustering for reinforcement learning. *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, 25, 1830-1838.

McLaughlin, M. (2012). Reading comprehension: What every teacher needs to know. *The Reading Teacher, 65*(7), 432-440.

Melby-Lervåg, M., Lyster, S. A. H., & Hulme, C. (2012). Phonological skills and their role in learning to read: A meta-analytic review. *Psychological Bulletin, 138*(2), 322.

Mesmer, H. A. E., & Griffith, P. L. (2005). Everybody's selling it—But just what is explicit, systematic phonics instruction? *The Reading Teacher, 59*(4), 366-376.

Moats, L. C. (2000). *Speech to print: Language essentials for teachers*. Baltimore, MD: Paul H. Brookes Publishing.

Moats, L. C. (2005). *Language essentials for teachers of reading and spelling: Module 6 Digging for meaning: Teaching text comprehension*. Boston, MA: Sopris West Educational Services.

Moats, L. C. (2009). Knowledge foundations for teaching reading and spelling. *Reading and Writing, 22*(4), 379-399.

National Reading Panel. (2000). *Report of the National Reading Panel: Teaching children to read: An evidence-based assessment of the scientific research literature on reading and its implications for reading instruction: Reports of the subgroups.* Washington, D.C.: National Institute of Child Health and Human Development, National Institutes of Health.

O'Rourke, E., Haimovitz, K., Ballweber, C., Dweck, C., & Popović, Z. (2014, April). Brain points: A growth mindset incentive structure boosts persistence in an educational game. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, 32, 3339-3348.

Osband, I., Russo, D., & Van Roy, B. (2013). (More) efficient reinforcement learning via posterior sampling. *Advances in Neural Information Processing Systems,* 26, 3003-3011.

Ouellette, G. P., & Haley, A. (2013). One complicated extended family: The influence of alphabetic knowledge and vocabulary on phonemic awareness. *Journal of Research in Reading, 36*(1), 29-41.

Phythian-Sence, C., & Wagner, R. K. (2007). Vocabulary acquisition: A primer. In R. K. Wagner, A. E. Muse, and K. R. Tannenbaum, (Eds.), *Vocabulary acquisition: Implications for reading comprehension* (pp. 1-14). New York, NY: Guilford Press.

Rasinski, T. V. (2006). Fluency: An oft-neglected goal of the reading program. In C. Cummins (Ed.), *Understanding and implementing Reading First initiatives: The role of administrators* (pp. 60-71). Newark, DE: International Reading Association.

Rasinski, T. V. (2012). Why Reading Fluency Should Be Hot! *The Reading Teacher, 65*(8), 516-522.

Rasinski, T. V., Rupley, W. H., Pagie, D. D., & Nichols, W. D. (2016). Alternative text types to improve reading fluency for competent to struggling readers. *International Journal of Instruction, 9*(1), 163-178.

Shaywitz, S. (2003). *Overcoming dyslexia: A new and complete science-based program for reading problems at any level.* New York, NY: Knopf.

Shaywitz, S. E., & Shaywitz, B. A. (2004). Reading disability and the brain. *Educational Leadership, 61*(6), 6-11.

Spencer, E. J., Goldstein, H., & Kaminski, R. (2012). Teaching vocabulary in storybooks: Embedded explicit vocabulary instruction for young children. *Young Exceptional Children, 15*(1), 18-32.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1, No. 1). Cambridge, MA: MIT Press.

Tomlinson, C. A., & Jarvis, J. (2009). Differentiation: Making curriculum work for all students through responsive planning and instruction. In J. S. Renzulli, E. J. Gubbins, K. S. McMillen, R. D. Eckert, & C. A. Little (Eds.), S*ystems and models for developing programs for the gifted and talented* (pp. 559-628). Storrs, CT: Creative Learning Press.

Wanzek, J., & Vaughn, S. (2008). Response to varying amounts of time in reading intervention for students with low response to intervention. *Journal of Learning Disabilities, 41*(2), 126-142.

Wolf. M. (2007). *Proust and the squid: The story and science of the reading brain.* New York, NY: Harper.

Wolfe, P., & Nevills, P. (2004). *Building the reading brain, preK-3.* Thousand Oaks, CA: Corwin Press.