

Teaching Spell Checkers to Teach: Pedagogical Program Synthesis for Interactive Learning

Momin N. Siddiqui
Georgia Institute of Technology
Atlanta, Georgia, USA
momins@gatech.edu

Abbie L. Olszewski
University of Nevada
Reno, Nevada, USA
aolszewski@unr.edu

Vincent Cavez
Stanford University
Stanford, California, USA
vcavez@stanford.edu

Srirangaraj Setlur
University at Buffalo
Buffalo, New York, USA
setlur@buffalo.edu

Sahana Rangasrinivasan
University of Buffalo
Buffalo, New York, USA
srangasr@buffalo.edu

Maneesh Agrawala
Stanford University
Stanford, California, USA
magrawala@stanford.edu

Hari Subramonyam
Stanford University
Stanford, California, USA
harihars@stanford.edu

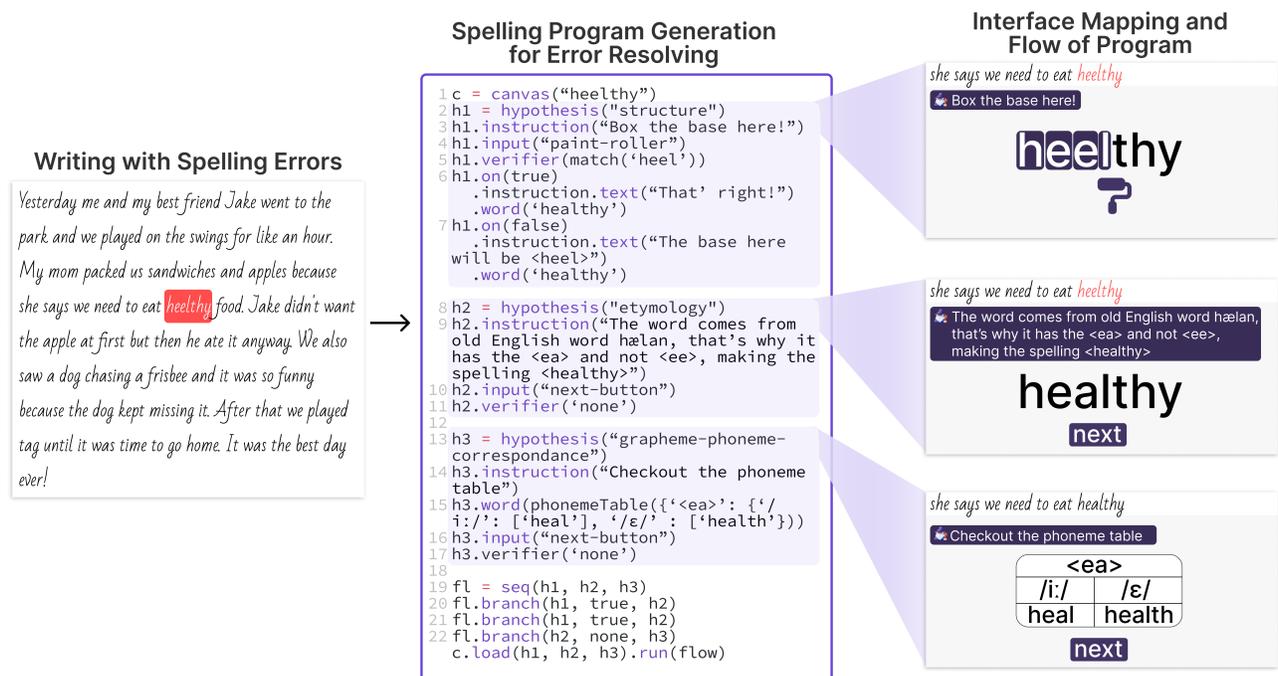


Figure 1: This figures illustrates how in SPIRE identifies misspellings (left), generates a program for intervention (center) and subsequently renders them in the interface (right) for learner's inquiry to surface word related insights

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference acronym 'XX, Woodstock, NY
© 2018 ACM.
ACM ISBN 978-1-4503-XXXX-X/18/06
<https://doi.org/XXXXXXXX.XXXXXXX>

Abstract

Spelling taught through memorization often fails many learners, particularly children with language-based learning disorders who struggle with the phonological skills necessary to spell words accurately. Educators such as speech-language pathologists (SLPs) address this instructional gap by using an inquiry-based approach to teach spelling that targets the phonology, morphology, meaning, and etymology of words. Yet, these strategies rarely appear in everyday writing tools, which simply detect and autocorrect errors. We

introduce SPIRE (Spelling Inquiry Engine), a spell check system that brings this inquiry-based pedagogy into the act of composition. SPIRE implements Pedagogical Program Synthesis, a novel approach for operationalizing the inherently dynamic pedagogy of spelling instruction. SPIRE represents SLP instructional moves in a domain-specific language, synthesizes tailored programs in real-time from learner errors, and renders them as interactive interfaces for inquiry-based interventions. With SPIRE, spelling errors become opportunities to explore word meanings, word structures, morphological families, word origins, and grapheme-phoneme correspondences, supporting metalinguistic reasoning alongside correction. Evaluation with SLPs and learners shows alignment with professional practice and potential for integration into writing workflows.

ACM Reference Format:

Momin N. Siddiqui, Vincent Cavez, Sahana Rangasrinivasan, Abbie L. Olszewski, Srirangaraj Setlur, Maneesh Agrawala, and Hari Subramonyam. 2018. Teaching Spell Checkers to Teach: Pedagogical Program Synthesis for Interactive Learning. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 18 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

School-age students often learn spelling by memorization [31]. However, the way in which a word is spelled is more than just a convention of letters. It reflects patterns of meaning, sound, structure, and spelling conventions that connect words to one another and to their histories. For instance, consider the word ‘reach’ meaning to extend or stretch out. A majority of us can spell this word with ease, and are unlikely to spell it as say <reech>, in which the long /i/ sound is mapped to the grapheme <ee>. In fact, English represents the /i/ sound in many ways, including <ee> as in see, <ea> as in reach, <ie> as in piece, and <e-e> as in these. Memorization helps us recall the correct form. However, for many learners, especially children with *language-based learning disorders*, rote memorization is insufficient because it recruits only one aspect of the lexical quality of the word (phonology or orthography/spelling) [54, 65]. What supports their learning is understanding the logic behind these alternatives; for example, noticing that ‘reach’ and ‘teach’ both use the <ea> to represent /i/, and that this pattern traces back to the Old English *reccan* [29]. Seen this way, spelling becomes less an arbitrary code to memorize and more a *system of patterns* that can be reasoned about, making the language more learnable and meaningful [57]. In fact, learners who struggle with spelling through memorization are often referred to speech and language therapy. In therapy, speech-language pathologists (SLPs) use *inquiry* as an approach to help learners make sense of how words are spelled and built. For example, if a learner writes <reech> for “reach”, the SLP may begin by asking how the learner decided on that spelling. If the learner explains that /i/ is often spelled <ee> as in “see” or “speech”, the SLP can validate this reasoning, then introduce a contrast, such as “teach” or “teacher”, to highlight the alternative <ea> spelling. With a different error, such as using <sine> for “sign”, the SLP may draw attention to the silent <g> and connect it to morphologically related words like “signal” or “signature”. This practice, known as *Structured Word Inquiry* (SWI) [8], turns spelling into a process of investigation rather than memorization, drawing on phonology (the relationship between sounds and letters), semantics (meaning of the word), morphology

(the structure and meaning of word parts such as roots, prefixes, and suffixes), and etymology (the historical origins and evolution of words) to explain why words are spelled the way they are. Through these inquiry-based interventions, SLPs help learners build *durable* knowledge that extends beyond a single correction to a broader understanding of how written language works. Consequently, they learn how to independently investigate words.

In contrast, everyday writing tools take a much narrower approach to spelling errors. A misspelling is detected, underlined, and replaced with the most probable correction. Whether a learner types “*reech*” because they overgeneralized a phonics rule, misapplied a morphological pattern, or simply hit the wrong key, the system’s response is the same. This transactional design supports error-free text, but it bypasses the chance to engage learners in noticing patterns or testing hypotheses. This is a missed opportunity because therapy time is scarce [42] and spelling occurs within these tools in the context of authentic writing tasks, such as science reports and essay writing. Without support in these everyday contexts, the carryover of SWI pedagogy from therapy to classroom writing is limited. What is needed is a way to embed the principles of SWI into the tools learners already use, turning each correction into a brief, *situated learning experience*. This motivates our central research question: *How might we bring the rich, inquiry-based pedagogy of SLPs into everyday writing tools such that spelling errors become an opportunity to learn about the writing system?*

In this work, we introduce SPIRE (Spelling Inquiry Engine), which brings structured word inquiry into everyday writing by synthesizing short, just-in-time instructional sequences from each error. Central to our approach is *pedagogical program synthesis*, a form of probabilistic program synthesis tailored to instruction: given a misspelling and writing context such as surrounding text, SPIRE infers a distribution over likely causes (e.g., grapheme-phoneme mapping, morphological structure, etc.) and then composes a bounded program consisting of three to five steps based on established SLP interventions represented as a domain-specific language (DSL). Each program is designed to guide the learner through a brief cycle of interactive inquiry and hypothesis testing before fading back into the flow of writing. Programs are optimized for just the necessary interventions and branch adaptively on the learner’s next action. This design addresses two core challenges: the system only ever sees the misspelling, not the learner’s whole reasoning, and effective instruction must change depending on how the learner responds. At runtime, SPIRE compiles the synthesized program into a compact inline panel for inquiry.

Similar to the <Healthy> inquiry in Figure 1, the program may ask the learner, through a short interactive speech prompt, to explain how they arrived at a spelling like <reech>. If the learner responds that /i/ is often spelled <ee> as in “see” or “speech,” the program can validate this reasoning and then, using simple visual widgets, introduce a contrast such as “teach” or “teacher” to highlight the <ea> pattern. It may also display the morphological family of reach (e.g., reaching, reachable, reached) or visualize its Old English root *reccan*, showing how etymology preserves the <ea> spelling. By combining spoken dialogue with lightweight visual widgets, the system mirrors how SLPs use multiple modalities in practice, including talking through different student hypotheses while also drawing or writing on paper, thereby keeping the interaction brief but engaging. We

evaluate SPIRE through expert appraisal with SLPs to examine whether synthesized programs align with professional practice using 25 authentic misspellings drawn from ecologically valid contexts. We also conducted a preliminary user study with learners to assess the usability of rich interventions within everyday writing tools.

Our key contributions include: (1) a Pedagogical Program Synthesis that represents SLP's instructional moves in a domain-specific language and assembles short, conditional programs from learner errors in real time, (2) SPIRE, a spell-checker that embeds these synthesized programs into lightweight, interactive interfaces, and (3) an empirical evaluation with experts and learners demonstrating how SPIRE aligns with professional practice and integrates into writing workflows.

2 Related Work

Our work builds on three interconnected bodies of research: research on spelling development and learning difficulties in school-age students, evidence-based interventions designed to support struggling spellers, and recent efforts leveraging AI to deliver adaptive literacy instruction.

Vellutino et al. [71] synthesized four decades of empirical work, establishing that difficulties in phonological awareness, rapid naming, and verbal working memory constitute core deficits that persist in development. Longitudinal studies of upper elementary students with reading or spelling disabilities reveal how these difficulties in phonological information processing evolve throughout formal schooling [64]. However, the presentation of symptoms varies considerably between individuals, requiring adaptive intervention approaches [38]. Corpus-based documentation allowed for an empirical understanding of characteristic dyslexic error types. DysList systematically cataloged error patterns in multiple languages [58], while subsequent analysis identified specific phoneme-grapheme correspondence failures that persist despite repeated exposure to high-frequency vocabulary [10]. These difficulties extend beyond traditional literacy contexts into digital environments, where searchers with dyslexia exhibit distinct information-seeking behaviors requiring specialized interface considerations [49].

Ehri's phase model [16, 18] describes progression through distinct developmental stages: pre-alphabetic, partial alphabetic, full alphabetic, and consolidated alphabetic, each characterized by qualitatively different approaches to word recognition. This framework establishes grapheme-phoneme knowledge as the basis for reading acquisition in alphabetic languages [17]. However, Silliman et al. [66] demonstrated that preadolescents with atypical language skills show distinct error patterns reflecting deficits in phonological, morphological, and orthographic dimensions, indicating that effective intervention must address multiple representational systems simultaneously.

Müller et al. [50] found that syllable-based reading instruction produces significant improvements in both word-reading accuracy and comprehension in second-grade students.

Meta-analytical findings demonstrate more severe reading impairments in deep orthographies such as English compared to shallow orthographies with consistent grapheme-phoneme mappings [9], suggesting that intervention approaches must account for the characteristics of the writing system. This has led to challenges to

phonics-centric instruction, with arguments that English orthography encodes morphological and etymological information often obscured by phonics alone, advocating instead for teaching the logic of the spelling system [6]. However, structured phonics approaches using word boxes and word sort procedures have been proven effective in facilitating both word recognition and spelling development [35], indicating that systematic instruction remains valuable even as the field recognizes the limitations of purely phonological approaches.

Beyond specific intervention techniques, broader pedagogical frameworks shape how spelling instruction is delivered and how learners engage with orthographic knowledge. Inquiry-based learning positions learners as active constructors of knowledge through systematic investigation rather than passive recipients of instruction. This approach follows five distinct phases that characterize the inquiry cycle: orientation, conceptualization, investigation, conclusion, and discussion [52]. Theoretical foundations align with the contemporary understanding of active knowledge construction [22], although empirical evidence reveals essential nuances. Meta-analytical findings show moderate positive effects on learning outcomes, with instructional guidance playing a critical role in determining effectiveness [39]. Unguided discovery learning proved to be less effective than structured inquiry with appropriate scaffolding, indicating that learner autonomy must be balanced with systematic support. Application to language pedagogy shows that inquiry-based teaching develops a deeper linguistic understanding compared to traditional instruction [40], but substantial barriers to implementation persist. Institutional constraints, time limitations, and teachers' resistance to constructivist approaches create friction in classroom contexts [14]. These challenges are amplified for learners with spelling difficulties who need additional scaffolding to balance open-ended exploration with systematic investigation of orthographic principles.

2.1 Intervention Methods

Intervention approaches for struggling spellers have predominantly focused on phonological processing and phonics-based instruction, building on the established relationship between phonological awareness and reading acquisition. Intensive remedial instruction targeting phonological processing can prevent reading failure in young children [69, 70], with explicit systematic phonics instruction producing better outcomes for at-risk children compared to implicit methods [21]. Early literacy support programs that combine phonological training with reading practice have shown effectiveness [30], and systematic reviews confirm that phonics training benefits English-speaking poor readers, although the effect sizes remain modest [48]. Individual response patterns vary considerably, and long-term outcomes depend heavily on the intensity and duration of the intervention. Multisensory instructional approaches have shown particular promise, with significant improvements documented when visual, auditory, and kinesthetic modalities are integrated [63]. Color-coded onset-rime decoding interventions have proven effective for first-grade students at serious risk for reading disabilities [32]. However, critical reviews note that despite advances in understanding dyslexia, the effectiveness of treatment remains variable, with many struggling readers requiring sustained and intensive intervention [1]. Evidence-based interventions must create a "virtuous circle" where gains in one area support the development of related skills [67].

Structured Word Inquiry offers a fundamentally different instructional approach that recognizes English spelling as a morphophonemic system where spellings represent sound, meaning, and history in an orderly manner. Rather than treating English orthography as a collection of exceptions and irregular rules, SWI demonstrates that the spelling system becomes highly consistent when words are contextualized within their morphology, phonology, etymology, and definition. Bowers and Kirby [8] found that morphological instruction produces significant effects on vocabulary acquisition, establishing that teaching morphemic structure supports both spelling and comprehension. This approach engages children in generating and testing hypotheses about how the spelling system works, helping them understand why words are spelled the way they are, an investigative process that aligns with inquiry-based learning principles discussed previously. Colenbrander et al. [11] conducted a randomized controlled trial assessing the effectiveness of SWI for students in grades 3 and 5 with reading and spelling difficulties, demonstrating improved performance on both trained and untrained words.

In practice, SWI instruction is typically delivered by speech-language pathologists working with small groups of students, providing the individualized attention and expert guidance necessary for effective implementation. However, the growing demand for specialized literacy support has created significant shortages of qualified SLPs [68], limiting access to these evidence-based interventions. Colenbrander et al.'s [11] findings revealed that SWI instruction delivered by teaching assistants with limited training produces significantly less effective outcomes compared to expert-delivered instruction, underscoring the expertise gap and the need for alternative delivery methods that maintain instructional quality while addressing scalability challenges.

2.2 Digital Tools for Teaching Language

Technological interventions for spelling and vocabulary instruction have evolved along several trajectories, each addressing different aspects of literacy support for children with dyslexia. One line of work has explored tangible interfaces to reduce cognitive load through concrete representations of abstract linguistic concepts. These systems demonstrated that physical manipulatives support young language learners through embodied learning experiences [19, 51]. This approach has been adapted for Chinese learners [20], addressing the distinct challenges of character-based writing systems where visual complexity and stroke order play central roles.

Adaptive spelling correction systems represent a second trajectory, accommodating characteristic dyslexic error patterns through specialized spellcheckers trained on dyslexic error corpora [59, 60]. These systems learn individual error patterns to provide personalized correction suggestions [43], representing significant advances in accommodating dyslexic spelling patterns. However, they remain fundamentally corrective rather than instructional, helping users spell correctly without developing a deeper orthographic understanding.

Gamification has emerged as a third approach, using game mechanics to maintain engagement during spelling practice [25]. Game-based systems have been developed to address spelling difficulties in different linguistic contexts [15, 55, 76]. Comprehensive reviews of these technologies note that most remain limited to phonetic-based approaches [56], failing to capture the comprehensive inquiry-based

methodology that characterizes effective SWI instruction. Context-aware writing assistance has shown promise in improving outcomes [77], yet even sophisticated adaptive systems lack frameworks to support hypothesis generation and testing that make SWI effective.

2.3 AI Pedagogical Techniques

AI-assisted literacy support has evolved from error correction systems [23] to more sophisticated instructional approaches that adapt to individual learners. The emergence of large language models has particularly enabled generative methods that create personalized learning experiences. Goodman et al. [26] designed LaMPPost, an AI-assisted writing tool for adults with dyslexia that helps compose emails while preserving user agency. Wang et al. [72] developed conceptual motivation models specifically for students with dyslexia, recognizing that practical assistive learning tools must take into account motivational factors alongside functional capabilities. Leong et al. [41] demonstrated that AI-enabled context personalization for vocabulary learning improves motivation by situating words within personally relevant scenarios. Attygalle et al. [3] explored text-to-image generation for vocabulary learning using the keyword method, finding that externalizing memorable visual links significantly increases memory retention. Peng et al. [53] developed Storyfier, exploring vocabulary support through narrative contexts generated by text models. Wang et al. [75] extended this with LearnMate, providing personalized learning plans and adaptive support. Investigation of AI-based support in speech-language pathology has documented both opportunities and concerns about bias and clinical validity [42], with particular attention to providing equitable care for culturally and linguistically diverse children. Design explorations of AI-supported home practice in speech therapy reveal how varying levels of AI oversight can provide informational, emotional, and practical support to parents implementing therapy techniques [13], although questions remain about alignment with professional therapeutic goals.

Conversational and Socratic approaches have emerged for fostering critical thinking through interactive dialogue. Ruan et al. [61] developed EnglishBot, an AI-powered conversational system for second language learning that adapts feedback based on student interaction, demonstrating greater engagement compared to traditional listen-and-repeat interfaces. Systems using AI to probe student understanding through targeted questioning enable scalable oral assessment [33], while question-driven exploration for personalized education [34] and investigations of whether AI partners can empower learners to ask critical questions [47] suggest that appropriately designed systems scaffold metacognitive skills. Metacognitive tutoring agents have proven to be effective in teaching inquiry-driven modeling in science education, demonstrating that intelligent tutoring systems can scaffold authentic scientific inquiry processes [36]. However, these conversational approaches, while effective for verbal reasoning, may not optimally support the hands-on, constructive activities central to structured word inquiry.

Program synthesis offers an approach particularly well suited to inquiry-based learning contexts requiring structured artifacts. Unlike conversational AI that responds to queries, program synthesis can generate interactive exercises and examples tailored to specific learning objectives. Theoretical foundations for program synthesis [28] have been extended by demonstrations that large language

models effectively perform synthesis tasks [4], including interactive synthesis by augmented examples [78]. Recent work has explored how hypothesis search enables inductive reasoning with language models [73], providing mechanisms for iterative refinement of generated artifacts. Synthesis and verification pipelines demonstrate how LLM-generated outputs can be iteratively corrected through automated verification and feedback loops [45], achieving high success rates through repeated refinement cycles. Kazemitabaar et al. [37] evaluated CodeAid in classroom deployment, documenting how LLM-based synthesis balances exploration with scaffolding. For inquiry-based spelling instruction, program synthesis enables dynamic generation of morphological investigations adapted to learners' current hypotheses, providing the structured yet flexible framework necessary for effective SWI implementation at scale.

3 User Experience

At a high level, SPIRE transforms traditional spell checking into an interactive process of linguistic inquiry. When a learner misspells a word, SPIRE invites them to explore its structure, meaning, and sound rather than simply replacing it with a correction. Each spelling mistake becomes an opportunity to reason about English orthography through short, guided activities embedded directly in the writing interface.

3.1 Usage Scenario

Matilda opens her word processor, which has SPIRE integrated into it. As she types, she writes: *"I like how the art of constructd."* The system flags *constructd* as a possible error by enclosing it in a red rectangle. When Matilda clicks on the word, the red highlight transitions into a red text indicator, and an inquiry canvas opens below her writing (Fig. 2a1).

The system begins by asking Matilda to use the word in a sentence—a diagnostic step to determine whether she understands its meaning. She clicks on a microphone icon (Fig. 2a1) and says aloud: *"The builder constructed the building."* The system transcribes her response and verifies that her usage reflects an appropriate understanding of *construct*. Since her sense of meaning is correct, it affirms her answer (Fig. 2a3) and proceeds to the next inquiry.

Next, SPIRE invites Matilda to *"box the base"* of the word (Fig. 2b1), a structured word inquiry practice where learners identify the morphemic base. Using a paint roller tool, Matilda drags across the portion she believes to be the base and highlights *struct* (Fig. 2b2). Although not fully correct, this action shows that she recognizes the structural pattern embedded within *construct*. The system records this as partial evidence of her morphological awareness and proceeds to test a new hypothesis related to graphemes.

At this point, the word is corrected to *constructed*, revealing how *<u>* replaces *<a>* and an *<e>* appears before *<d>* to form the correct grapheme sequence *<c><o><n><s><t><r><u><c><t><e><d>*. Rather than presenting this as a static correction, the system visually highlights these changes to draw Matilda's attention to how sound-letter patterns map across morphology and phonology.

Finally, SPIRE prompts Matilda to suggest words that share the same base *<struct>* (Fig. 2c1). She types *structure* and *instruct* (Fig. 2c2). The system affirms *structure* but flags *instruct* for having an extra *<s>*, explaining the morphological constraint (Fig. 2c3).

This final step reinforces her understanding of morphological families and related word forms before she returns to her essay.

Through these brief, multimodal interactions—speaking, dragging, typing—Matilda not only corrects her mistake but also deepens her understanding of word structure and related forms. Each step is conversational, adaptive, and pedagogically grounded, embodying how SPIRE turns everyday spelling errors into opportunities for structured linguistic exploration.

4 Pedagogical Program Language

To develop SPIRE, we need to encode the dynamic conversational practice of SLPs into something a computer can assemble in real time from a single misspelling and its context. Our pedagogical programming language (i.e., DSL) formalizes how SLPs reason, question, and guide discovery during therapy as composable program primitives that can be synthesized into concise, adaptive instructional sequences. Crucially, our DSL is grounded in practitioner concepts and built with practitioners. We first inventoried canonical SWI moves and terminology from established texts (books written by expert practitioners, instructional material used in graduate programs, and skill development workshops [5, 7, 12, 31, 44]) into a spreadsheet. Next, working with a practicing SLP, who is also a co-author, we organized the moves as program primitives, including language primitives, hypothesis templates, action primitives, and control and compositional primitives. For this collaborative work, we utilized digital whiteboards to organize the concepts, created sketches of program logic and instructional moves, and also produced storyboards of envisioned learning episodes. This was an iterative process that involved weekly virtual meetings with the research team, spanning over six months. During this process, we also consulted with two other practicing SLPs who have over 10 years of experience to receive feedback on the evolving DSL.

In what follows, we describe each category of primitives that form the basis of our DSL. Collectively, these primitives keep the logic and vocabulary of practice intact while making them computationally operable for just-in-time, in-context human learning.

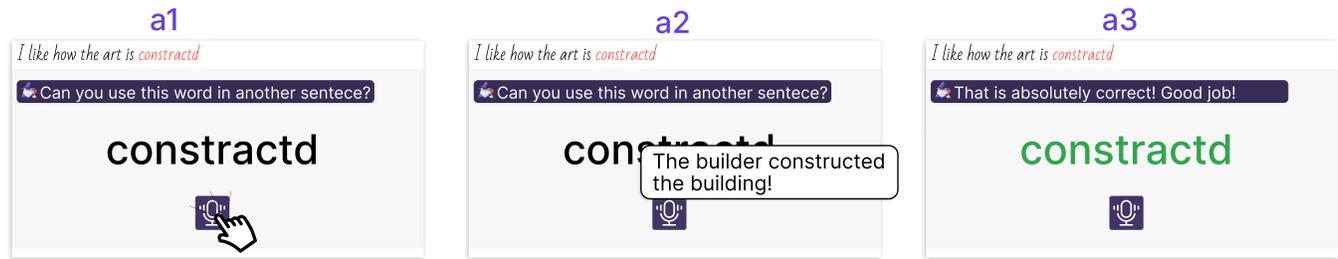
4.1 Language Knowledge Primitives

This set of primitives defines the linguistic substrate that each pedagogical program operates on. Concretely, they represent the structured linguistic information an SLP implicitly draws upon when reasoning about a misspelling, including knowledge of a word's morphology, graphemes, phonemes, word families, and etymology. In our DSL, this knowledge is encoded as a structured record `word_properties` whose named attributes (shown in Table 1) correspond directly to the analytic categories used in SWI. Just as with SLPs, these attributes enable the pedagogical programs to reason about the logic behind a spelling and pose the same kinds of questions a practitioner might ask, such as "What is the base of the word?", "What sound do these letters represent?", and "Which other words share this base?"

4.2 Hypotheses Template Primitives

Our DSL formalizes the inquiry logic of structured word inquiry as a set of reusable **hypothesis templates**. Each template represents a recurring pattern of *reasoning* that enables the learner to test and refine a hypothesis about how a word works. It expresses a

a Is sense/meaning of the word known?



b Is the base of the word known?



c are morphological relatives known

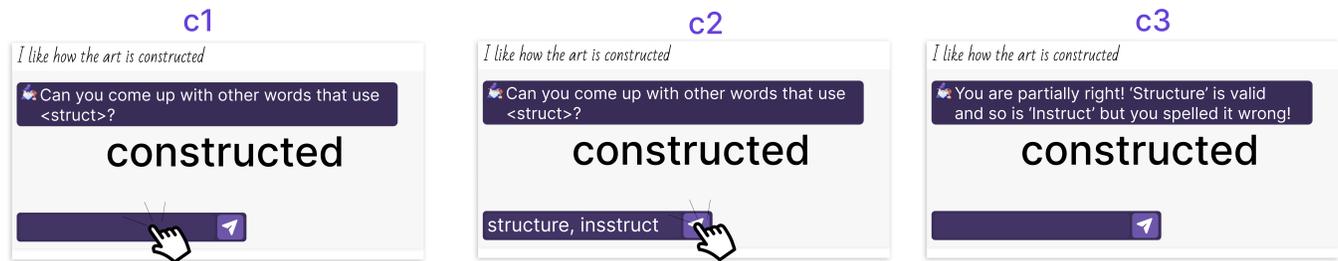


Figure 2: This figure demonstrates the usage scenario. We showcase the learner going through three interventions for the following inquiries: (a) “is the sense/meaning of the word known?”, (b) “is the base of the word known?” and (c) “are the morphological relatives known?”

single pedagogical move typical in SWI as a structured *quintuple* consisting of preconditions, evidential features, action base, warrant, and learning effect. *Preconditions* state when a line of inquiry is warranted. They are conditional predicates over an observable situation, including the type of deviation or error in a learner’s spelling production, the linguistic characteristics of the intended word, and basic usage context, such as semantic appropriateness. Concretely, they are represented as a combination of interpretable Boolean predicates computed over the linguistic primitives as well as diagnostic descriptors in natural language. The predicate provides a formal check for whether the template applies, while the descriptor captures practitioner judgment that is hard to encode as a formal rule.

For instance, in the grapheme identification hypothesis template (H8)(Table. 2), the symbolic guard checks that the that the deviation is *grapheme-local* rather than structural: morphology and boundary errors are not implicated ($\neg(\text{prefix_error} \vee \text{suffix_error}$

$\vee \text{segmentation_error})$), no suffixing convention applies ($\neg \text{suffixing_change_applies}$), the phoneme sequence of the attempt essentially matches the target ($\text{phoneme_match} \leq \epsilon$), the mismatch is small and local ($\text{grapheme_mismatch_count} \in \{1, 2\}$), and morpheme boundaries are preserved ($\text{morpheme_boundaries_preserved}$). The paired diagnostic descriptor states this in practitioner terms e.g., “sounds correct but uses a nonstandard letter team.” Thus *reech*→*reach* satisfies the guard (phonology intact, one letter–team substitution), whereas *running*→*running* is excluded by the suffixing-change check and *alot*→*a lot* by the segmentation predicate. At runtime, an LLM can assign soft confidence to descriptors, allowing preconditions to serve as both interpretable constraints and probabilistic cues for selecting the most pedagogically relevant hypothesis.

Evidential features specify what linguistic features must be inspected in order to carry out the inquiry. In our DSL, evidential features are typed references to the linguistic primitives listed in

DSL Attribute	Pedagogical Role in SLP Instruction	Example (“reaching”)
morphemes	Meaning-bearing parts composing the word; focus of “box the base” and word-sum construction.	["reach", "-ing"]
bases	Core lexical root carrying meaning across a family; anchors family-based pattern recognition.	["reach"]
prefixes	Bound morphemes before the base; invite prefix-meaning inquiries and family extension.	none
suffixes	Bound morphemes following the base; used in suffixing-convention and inflectional change prompts.	["-ing"]
graphemes	Letter or letter-team units representing phonemes; supports grapheme identification and comparison.	["r", "ea", "ch", "i", "ng"]
phonemes	Sound units of the spoken form; validate or contrast the learner’s sound–spelling reasoning.	['/r/', '/i:/', '/tʃ/', '/r/', '/ŋ/']
related_words	Morphological or orthographic relatives for matrices and “in/out” sorting; surface consistent base spelling.	“reach”, “reaching”, “overreach”; contrasts: “teach”, “preach”
etymology	Historical origin explaining modern spelling stability; supports etymological reasoning.	Old English <i>reccan</i> (“to stretch out”)
homophones	Same sound, different spelling or meaning; prompts meaning-based distinction.	none
semantic_appropriateness	Ensures the written word matches intended meaning; drives meaning-first inquiry when mismatched.	Intended “extending or stretching out” (“She is reaching for the shelf.”)
syntactic_correctness	Checks grammatical fit; distinguishes morphological from typographic errors.	Progressive verb fits: “She is reaching. . .”
meaning_understood	If unclear, prioritizes definition/denotation prompts before structure/GPC work.	Learner typed “reaching” intending “reaching”.
context_sentence	Provides authentic sentence evidence; used in meaning or boundary reasoning.	“She was reaching for the book.”

Table 1: List of Language knowledge primitives encoded in word_properties.

Table 1. For instance, for grapheme inquiry, the evidence is the pair of grapheme inventories and their alignment to the phoneme sequence; Crucially, evidential features do *not* enumerate examples themselves; they expose the handles that the *program synthesis* stage will later use to instantiate a small contrast set. For grapheme inquiry, for instance, the evidence is just the learner and target grapheme sequences aligned to the same phoneme sequence; during synthesis, these handles permit selecting minimal contrasts for the contested mapping (e.g., attested spellings of /i/ such as <ee>, <ea>). Analogously, for suffixing, evidence is the base and the suffix with relevant orthographic context, etc. In short, evidential features say “what to look at” and program synthesis later decides which concrete examples to show.

The *Action Base* describes the type of learner moves, i.e., established spelling interventions to carry out the inquiry. Typical action types include identifying graphemes in one’s own spelling, mapping phonemes to letter teams, boxing the base, analytically removing a suffix, constructing a forward word-sum, assembling a compact family matrix, and sorting IN/OUT relatives, or orally segmenting a run-together token. Different interfaces might realize the action types differently (e.g., highlight or tap to identify graphemes, drag-to-align or speak-to-align for phoneme→grapheme mapping) while preserving identical pedagogical semantics. The *Learning Effect* records the conceptual outcome that the move is designed to produce, i.e., an update to the learner’s knowledge state. Typical effects include meaning clarified, base verified, word-sum assembled or deconstructed, convention induced, boundaries corrected, grapheme–phoneme mapping aligned, family confirmed, etymological rationale established, homophone distinguished, or visual discrimination achieved. These intended learning effects not only describe the outcome of the current inquiry but also serve as preconditions for subsequent hypotheses

and as features for downstream selection without committing to any particular control policy.

Finally, *Warrants* name the explanatory link that connects the action’s observables to the learning effects. Formally, it is an explanatory operator with a type (e.g., grapheme comparison) and parameters drawn from the evidential features (which items to contrast, which rule to name, which relatives to surface). Crucially, the warrant is what generates the *just-in-time conversation* for guided inquiry toward the learning effect. Because the warrant specifies the explanatory move and the slots it fills, the system can produce the next utterance that validates the learner’s work, introduces the right contrast, and generalizes the insight. For *grapheme identification* (H8), once the learner has identified letter teams, the warrant instantiates a comparison: “Now line up your teams with the target; notice that <ea> is the stable team in this family—so here we write *reach*.” When the observed evidence does not support the intended explanation, the warrant also determines the conversational pivot (e.g., from comparison to meaning-first or to relatives), keeping the guidance evidential rather than corrective.

4.3 Compositional Primitives

This set of primitives describes how discrete acts of inquiry unfold coherently over time. In structured word inquiry, learning progresses flexibly, where the learner advances a hypothesis, inspects evidence, articulates what that evidence reveals, and then extends or closes the inquiry. In SPIRE, our goal is that given the current misspelling (i.e., attempt state) and the intended canonical form (i.e., target state), a plan is generated in the form of a sequence of inquiry steps, in which each step yields a learning effect that reduces some part of the observed discrepancy or produce the specific evidence needed for

the next justified step. In effect, the plan serves as a small *proof-by-inquiry* that bridges the gap from attempt to target. *Compositional* primitives specify the legality of moves as dependencies among the hypothesis template fields. In a simple *sequence*, meaning clarification can precede base identification; grapheme identification naturally precedes a grapheme comparison; verification of a base licenses testing a suffixing condition. *Nesting* reflects the way SLPs zoom in: a suffixing test sits inside a word-sum construction; verifying likely relatives sits inside a family-pattern generalization. *Branching* captures the exploration of alternatives SLPs routinely entertain: try two plausible vowel teams or contrast two related families, then merge back once one branch yields a learning effect or both are ruled out. *Closure* occurs when the discrepancies relevant to the learner’s attempt have been resolved or when the inquiry has produced the named understanding that the SLP was aiming for.

Collectively, these primitives let SPIRE treat instruction as constrained search under uncertainty. Given an attempt state and a target state, the synthesizer explores the legal space defined by the control and compositional primitives. It proposes short sequences of hypothesis templates whose preconditions are satisfied, whose actions bind the needed evidence, whose warrants generate just-in-time explanations, and whose learning effects shrink the discrepancy. The following section presents our program-synthesis architecture that operationalizes this process.

5 System Architecture

The pedagogical program language presented in Section 4 defines what constitutes valid SWI instruction—the hypotheses, evidence, and warrants that structure expert practice. This section describes how these abstract specifications are operationalized at runtime. The system architecture translates the DSL’s declarative primitives into executable instructional sequences, automatically synthesizing programs from observed spelling errors and rendering them as interactive interventions within the writing interface.

This architecture addresses a fundamental limitation in existing spelling tools. Conventional spell checkers detect errors and suggest corrections based on string similarity, treating spelling as a mechanical transcription problem. SPIRE must instead infer why a misspelling occurred, what it reveals about the learner’s linguistic reasoning, and which line of inquiry will advance their orthographic understanding. The spelling inquiry engine, implemented in Python, orchestrates this process through five modules: a spell checker identifies target words from context, a word property synthesizer extracts morphological and phonological features, an error analyzer diagnoses likely causes, a spelling inquiry engine composes hypothesis templates into executable programs following the DSL notation, and an interface generator renders these programs as interactive components. Together, these modules transform the DSL’s declarative specifications into adaptive, just-in-time instruction tailored to individual learner errors.

5.1 Spell Checker

The spell checker is the entry point of SPIRE, responsible for detecting misspellings and identifying their most likely target words. It operates continuously in the background and is activated after every second long pause in writing. Upon activation, it analyzes the

surrounding sentence to locate potential spelling errors and predict their intended forms. For example, when a learner writes “Reechable,” the spell checker identifies the error and generates “Reachable” as the target word. We designed this module to address the limitations of conventional spell checkers, which rely primarily on string similarity and often fail when a learner’s attempt deviates substantially from the correct form (e.g., *yous* → *use*). Even commercial tools like Grammarly [27] tend to substitute alternative word choices rather than identifying the learner’s intended word. Because SPIRE aims to reason about the learner’s linguistic understanding, it requires target predictions that are both contextually grounded and semantically faithful. To achieve this, the spell checker uses a large language model (LLM) to perform contextual error detection and correction. The LLM considers both the local sentence and the broader discourse context to infer the intended word rather than relying solely on edit distance. Once the target word is identified, both the attempted and target forms—along with their sentence context—are passed to the next module, the word property synthesizer, which begins the process of linguistic analysis.

5.2 Word Property Synthesizer

The word property synthesizer extracts the linguistic features of the target word that are relevant for analysis and instruction. Specifically, it identifies morphological, phonological, and orthographic properties such as morphemes (meaning-bearing units), graphemes (letters or letter teams representing sounds), and phonemes (sound units). These features serve as the linguistic foundation on which the system builds hypotheses about the learner’s spelling behavior. We designed this module to capture the kinds of linguistic reasoning speech-language pathologists (SLPs) use when teaching spelling. Standard spell checkers only produce surface-level corrections, but effective instruction requires understanding why a word is spelled the way it is—its structure, sound patterns, and etymology. To operationalize this expert reasoning, the synthesizer must produce interpretable, structured representations of words that downstream modules can reason over. To achieve this, the word property synthesizer uses a large language model (LLM) prompted with a structured specification and few-shot examples. It receives the target word and its writing context, then outputs a set of linguistic predicates in the form of (*field-name*, *value*) tuples (e.g., (*'morpheme'*, *'reach'*), (*'suffix'*, *'able'*))—instantiating the language knowledge primitives defined in Table 1. A corpus of English phoneme–grapheme mappings is provided as grounding context to ensure validity. The resulting property set—representing the linguistic knowledge an SLP would reference—is then passed to the error analyzer for diagnosing the underlying causes of the misspelling.

5.3 Error Analyzer

The error analyzer identifies the most plausible causes of a learner’s misspelling. It takes as input the linguistic properties generated by the word property synthesizer and compares them against an annotated list of common error types, such as phoneme–grapheme mismatches, morphological confusions, or suffixing errors. This module bridges linguistic analysis and instructional reasoning by explaining why the error occurred rather than merely labeling it. Understanding the underlying cause is essential for selecting pedagogically meaningful

ID	Template	Key Preconditions	Action Base	Learning Effect
H1	Meaning Clarification	Meaning unclear or semantically inappropriate.	Define intended meaning; contextual prompt.	Meaning aligned; enables next inquiry.
H2	Base Denotation	Base + affix present; meaning known.	Identify main meaning part (base).	Semantic core anchored in base.
H3	Morphological Structure	Affix ≥ 1 ; single base.	“Box the base” or decompose parts.	Base verified; structure understood.
H4	Word-Sum Construction	Decomposable into morphemes.	Build word step-by-step (synthetic sum).	Morphemes combined \rightarrow correct form.
H5	Spelling Convention	Suffixing rule applies (e.g., doubling, e-drop).	Inspect suffix rule via word-sum.	Orthographic rule induced.
H6	Word Matrix	Base stable across ≥ 2 relatives.	Assemble word matrix grid.	Family pattern generalized.
H7	Word Boundaries	Segmentation or spacing error.	Segment aloud, retype spaced form.	Correct word boundaries restored.
H8	Grapheme Identification	Small grapheme mismatch (≤ 2).	Identify graphemes, compare to target.	Grapheme–phoneme alignment achieved.
H9	Etymology Inquiry	Historical root available.	Trace origin; discuss spelling stability.	Origin-based spelling rationale.
H10	Morphological Relatives	Base shared by other words.	Sort IN/OUT relatives; construct family.	Base spelling reinforced by family.
H11	Morpheme Verification	Base/affix uncertainty.	Verify morphemes across known words.	Morpheme consistency confirmed.
H12	Etymological Relatives	Common ancestor offers insight.	Compare historical cousins.	Etymological family explained.
H13	False Relatives	Optional; visual similarity only.	Contrast look-alike vs. true relatives.	False relatives excluded.
H14	Phoneme–Grapheme Mapping	Phonology \approx target.	Map sounds to graphemes (phoneme table).	Sound–spelling map aligned.
H15	Pronunciation Variation	Family shows sound shifts.	Compare relatives with shared grapheme.	Stable spelling despite sound change.
H16	GPC Pattern Consistency	Consistent spelling, varied pronunciation.	Compare family spellings.	Pattern consistency generalized.
H17	Homophone Distinction	≥ 1 homophone exists.	Sort by meaning; contrast forms.	Correct form tied to meaning.
H18	Similar Word Comparison	Similar orthography detected.	Visual contrast task.	Subtle form difference noticed.

Table 2: This table consists of the abbreviated summary of SPIRE Hypothesis Templates.

interventions. To achieve this, the error analyzer uses a large language model (LLM) prompted with the annotated error taxonomy and few-shot examples. The LLM evaluates how the learner’s attempted and target words diverge across linguistic dimensions and returns a ranked list of likely error categories. This structured output is then passed to the spelling inquiry engine for generating the corresponding guided intervention program.

5.4 Spelling Inquiry Engine

The spelling inquiry engine instantiates the pedagogical program language described in Section 4, generating concrete intervention

programs by composing hypothesis templates into executable sequences that SPIRE delivers to the learner. It determines what instructional steps to present and in what order, transforming linguistic analyses and diagnosed error types into a coherent sequence of inquiry activities. Each program represents a short cycle of hypothesis testing designed to help the learner reason about the structure, sound, or meaning of the target word. We designed this module to model how speech-language pathologists (SLPs) guide learners through discovery-based spelling instruction. Rather than offering fixed feedback, the engine composes adaptive instructional sequences that reflect pedagogical intent—selecting, sequencing, and framing

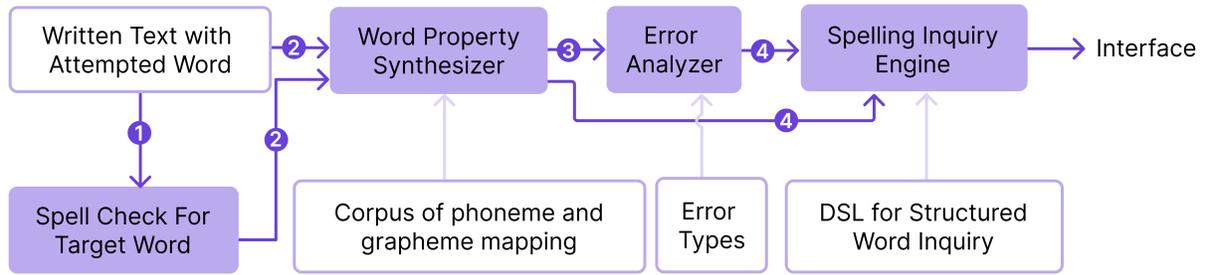


Figure 3: Overview of the System Architecture of SPIRE. The number labels on the edges mark the timestep at which each module is activated.

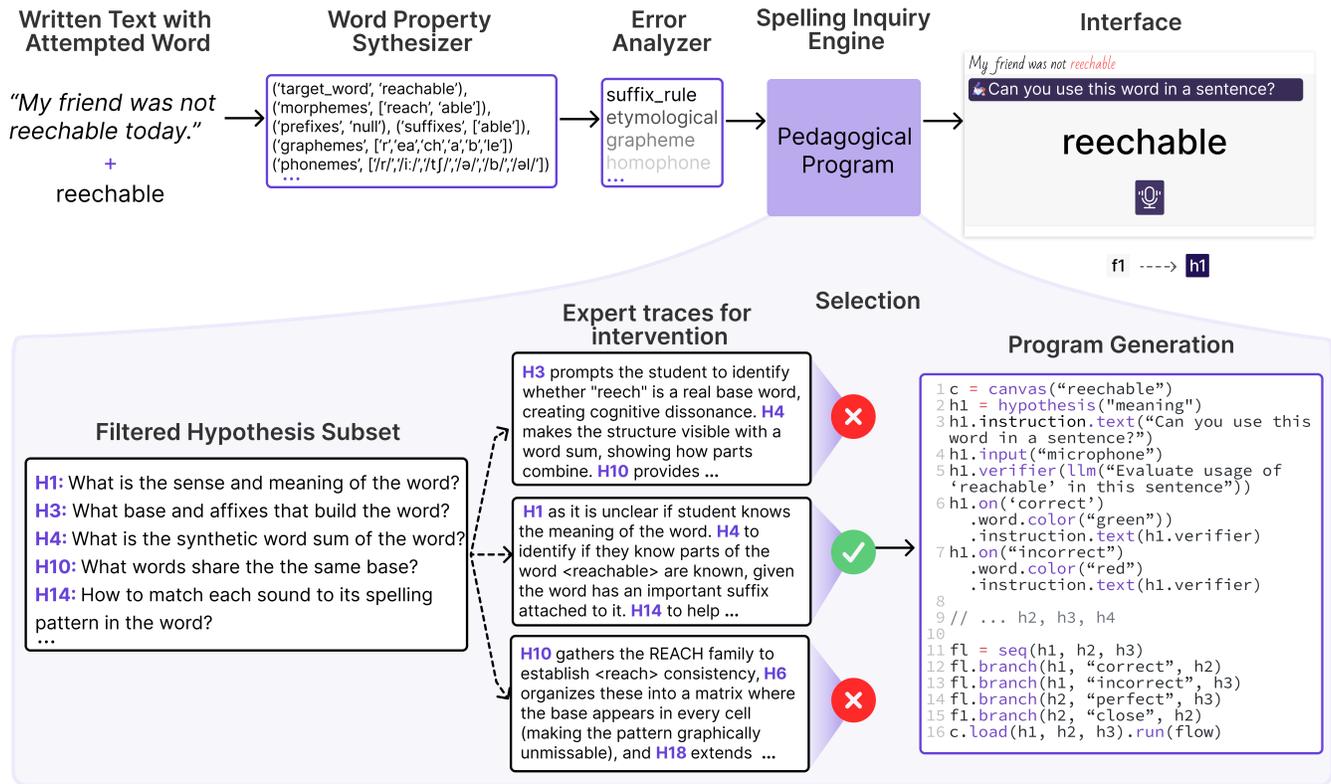


Figure 4: This figure showcases an example for the pipeline, providing each step’s output.

interventions based on the learner’s observed error and linguistic context. To achieve this, the engine integrates symbolic reasoning with language-model inference. Implemented in Python, it converts the filtered set of hypotheses into a structured program using our domain-specific language (DSL), which defines the syntax and logic of instructional moves. The resulting program is serialized and rendered in React as an interactive interface. As shown in Figure 4, the engine consists of four submodules: (A) Hypothesis Filter, (B) Trace Generator, (C) Trace Selector, and (D) Program Generator, which together simulate the expert reasoning process of an SLP.

5.4.1 Hypothesis Filter. The hypothesis filter performs the first selection step in the spelling inquiry engine. It narrows the full set of

pedagogical hypotheses to those that are relevant for the learner’s specific word and error context. Each hypothesis encodes a possible instructional focus (Section 4.2)—such as analyzing a base word, exploring a grapheme–phoneme mapping, or inspecting a suffixing convention. This filtering process ensures that the system considers only meaningful and non-trivial inquiries. Without it, the program could surface redundant or pedagogically irrelevant interventions (e.g., asking about suffixing when none applies). The goal is to approximate how an SLP intuitively rules out lines of inquiry that do not advance the learner’s understanding. To achieve this, the filter uses a unification-based pattern matcher to align the predicate tuples from the word property synthesizer with the preconditions defined

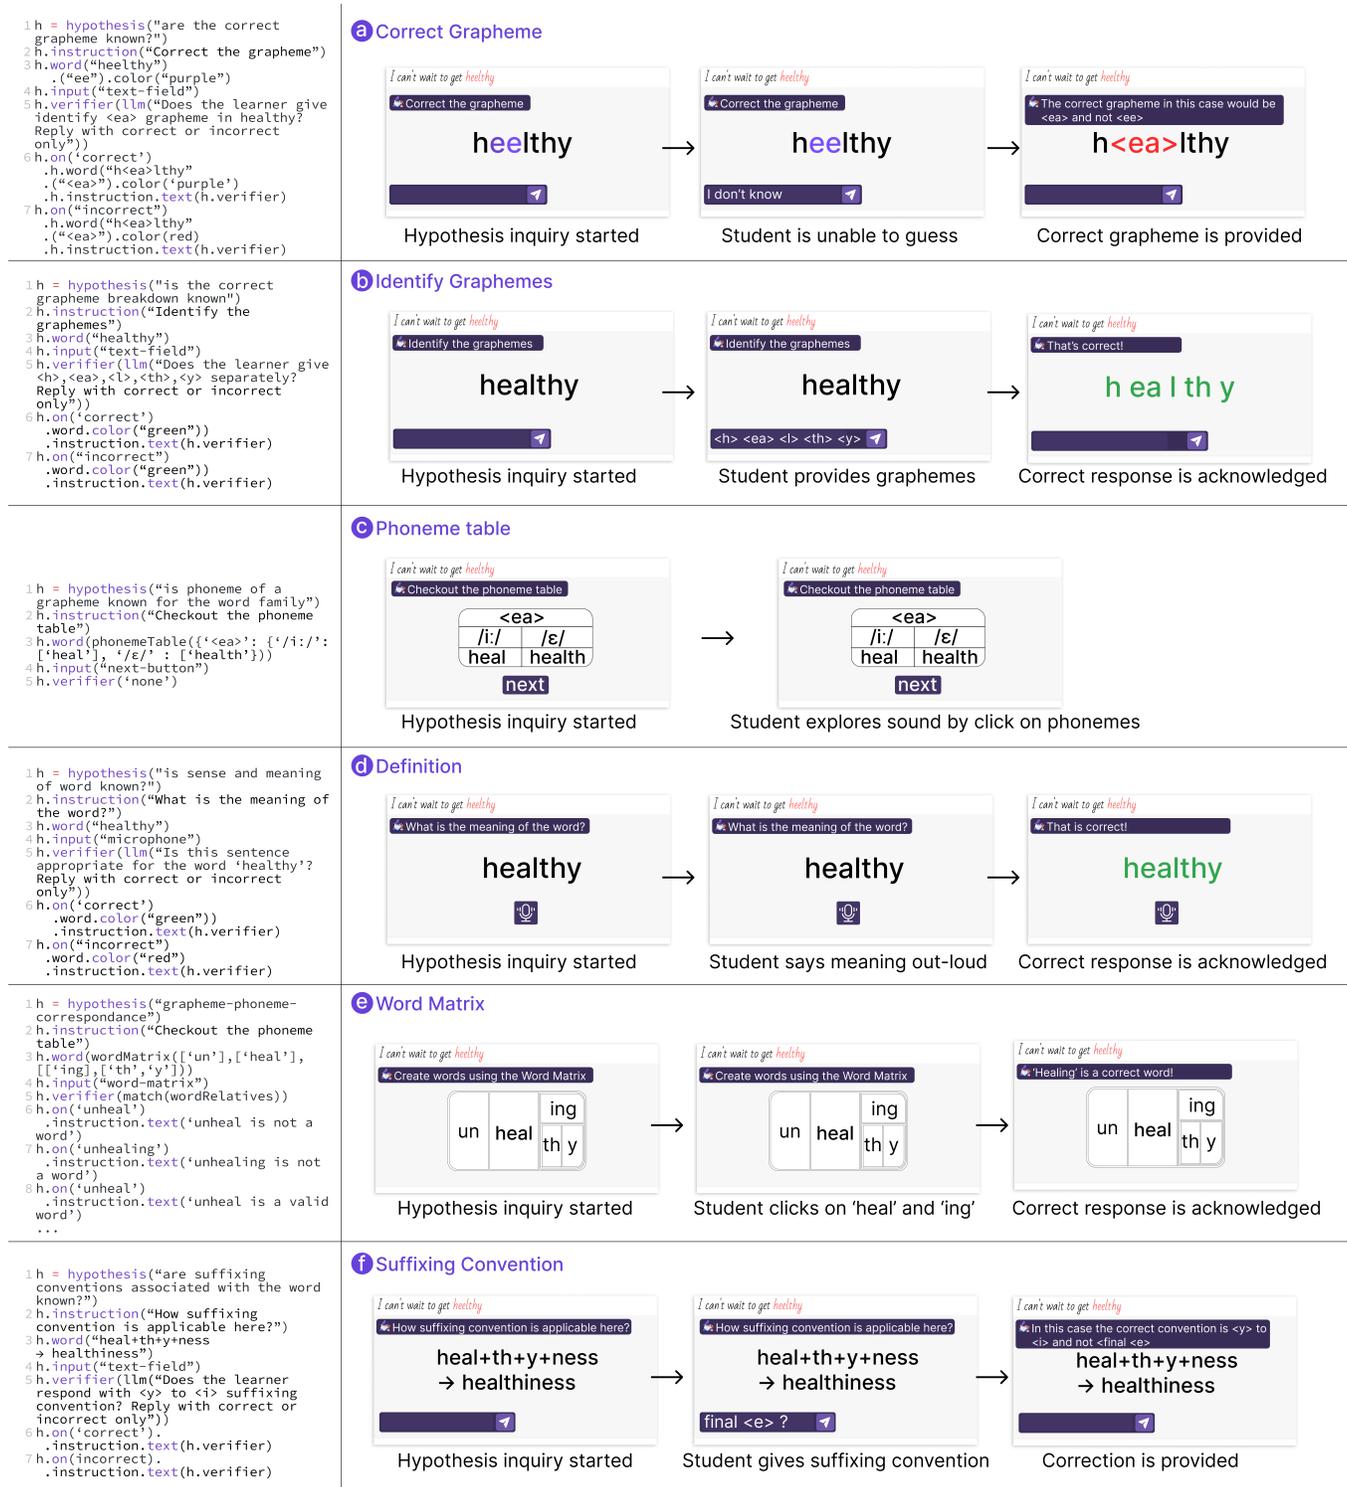


Figure 5: This figure showcases some example hypothesis testing in the interface (right) and the program that guides that interface (left).

in each hypothesis. Guards and diagnostic descriptors encoded in the hypotheses are evaluated symbolically or semantically using a language model. The output is a refined set of hypotheses that satisfy all constraints, forming the candidate pool for trace generation in the next stage.

5.4.2 Trace Generator & Trace Selector. The trace generator and trace selector together form the core reasoning mechanism of the spelling inquiry engine. Their role is to construct and refine instructional sequences—called traces—that guide the learner through 2–5 hypothesis-driven steps of exploration. The generator proposes multiple alternative instructional paths, while the selector evaluates and ranks these paths to identify the one most pedagogically sound. This two-phase design is modeled after human inductive reasoning and draws inspiration from the hypothesis search framework [74], which separates divergent generation from convergent evaluation. Effective instruction requires both creativity and precision: educators first imagine several possible ways to guide a learner, then choose the one that best advances understanding. By mirroring this cognitive process, SPIRE balances exploration (diversity of ideas) with instructional rigor (quality of reasoning).

To achieve this, we implement a multi-instance LLM architecture. The Trace Generator launches three parallel LLM instances, each independently exploring distinct instructional traces grounded in the learner’s linguistic and error profiles. This parallelism promotes divergent reasoning, surfacing multiple pedagogically plausible sequences that link the attempted and target words through morphological or phonological insight. The resulting traces are then passed to the Trace Selector, a single LLM instance that performs convergent reasoning. It reviews and compares the candidate traces using criteria such as pedagogical validity, conceptual coherence, and instructional clarity. The top-ranked trace—the one most likely to foster meaningful understanding—is selected and forwarded to the Program Synthesizer for conversion into executable DSL code. This architecture ensures both breadth and depth: breadth through generative exploration of multiple possible instructional paths, and depth through rigorous evaluation and selection. Together, these modules enable SPIRE to emulate the reflective judgment of an expert SLP—creatively hypothesizing instructional strategies and critically choosing the most effective one to guide the learner’s inquiry.

5.4.3 Program Synthesizer. The program synthesizer converts the selected instructional trace into an executable pedagogical program following the DSL notation defined in Section 4. It operationalizes the reasoning captured in hypothesis templates—their preconditions, evidential features, actions, warrants, and learning effects—transforming abstract hypotheses into concrete, interactive steps that guide the learner through linguistic inquiry. Each synthesized program specifies instructional actions, learner inputs, and conditional feedback, following the notation defined in our domain-specific language (DSL) (see “Generated Program” in Figure 4).

We designed this module to bridge pedagogical intent and interface behavior. While the preceding modules determine what to teach and why, the synthesizer determines how that reasoning manifests interactively. It ensures that every instructional step—such as asking the learner to identify a base or compare graphemes—is both executable and adaptable to learner responses. Technically, the synthesizer accepts the best trace, word properties, and DSL schema as

input context. Using structured prompting, it generates a declarative program with branching logic that adapts dynamically to user input. Each branch encodes conditional feedback and next-step transitions, enabling the interface to respond to learner actions in real time. To ensure reliability, the synthesizer validates the generated program for syntax and runtime consistency; if compilation fails, the system automatically regenerates it. The resulting program is serialized and passed to the interface generator for rendering.

5.5 Interface Generator

The interface generator renders the synthesized pedagogical program as an interactive learning experience. It translates the declarative program produced by the synthesizer into dynamic interface components that guide the learner through inquiry-based spelling correction. Each hypothesis node in the program corresponds to a visual interaction module that binds instructional text, learner input, and feedback logic, allowing the system’s reasoning to unfold directly within the writing environment.

We designed this module to preserve an isomorphic relationship between pedagogical logic and interface behavior. Rather than displaying static feedback, the interface visually manifests the structure of the instructional program—showing learners not just the correction, but the reasoning process behind it. This design enables the same program to operate seamlessly in both headless (text-based) and interactive (graphical) modes, making it adaptable across different learning contexts. Technically, the interface generator parses the JSON execution plan produced by the program synthesizer and compiles it into React components. Each component corresponds to a hypothesis in the instructional trace (e.g., h1, h2, h3), specifying its instructional text, input affordances (such as text fields or visual tools like paint-rollers), and verification conditions. Conditional branches (`on(true)`, `on(false)`) define how the interface adapts to learner responses in real time, ensuring pedagogical continuity across different interaction outcomes. This direct coupling between program logic and user interaction enables SPIRE to render the reasoning process of an SLP as an interactive, scaffolded learning experience.

5.6 Implementation Details

SPIRE integrates a Python backend with a React frontend. The spelling inquiry engine is implemented in Python using a unification-based pattern matcher [46] and Anthropic’s Claude Sonnet [2] for hypothesis filtering, generation, and selection. The synthesized instructional programs are serialized as JSON execution plans, which the React interface renders as interactive components mirroring the system’s pedagogical logic. To support multimodality, SPIRE employs Google Gemini’s speech-to-text for capturing spoken responses and Gemini’s text-to-speech for verbalizing instructions and feedback, enabling voice-based interaction alongside text input. This modular setup ensures consistent execution across conversational and interactive modes while remaining adaptable for integration with future writing or tutoring systems.

6 Technical Evaluation

To validate our technical pipeline, we conducted an expert evaluation study with specialists in structured word inquiry and spelling instruction for children.

6.1 Material

We collected 10 writing samples from children in grades 2-5, each containing one or more spelling errors. Using the SPIRE pipeline, we generated 25 conversations between a simulated expert tutor and simulated student, with each conversation addressing one misspelled word. Samples with multiple errors were corrected across successive conversations. Each conversation included 2-5 just-in-time interventions. For each intervention, the system provided: (1) the targeted SWI question type (meaning, structure, relatives, or grapheme-phoneme correspondence), (2) the selected hypothesis for testing, (3) the rationale for this selection, and (4) the instructional dialogue. The system's responses dynamically adjusted to simulated learner responses following the branching logic expressed in our domain-specific language.

6.2 Participants

We recruited 5 experts through word-of-mouth and mailing lists. The demographic information pertinent to each expert are present in the table 3. The number of participants is in line with prior work that has required domain-experts/teachers [24, 62].

6.3 Study Protocol

Experts first attended a Zoom orientation session explaining the evaluation task. They then rated conversation transcripts through a custom web interface using 5-point Likert scales for each intervention: (Q1) "The action selected is appropriate"; (Q2) "The reason for action selection is valid"; (Q3) "I feel confident in my assessment." Experts could also add qualitative notes. After completing one practice transcript with guided walkthrough, participants independently rated all 25 transcripts over multiple days (2-6 hours total). Participants received \$50 compensation.

6.4 Results

6.4.1 Pedagogical Reasoning Validity. Experts demonstrated exceptionally high consensus on the soundness of the system's pedagogical reasoning (Q1), with 78.9% perfect agreement across all five raters ($\mu = 4.93$, $\sigma = 0.41$, average item $\sigma = 0.14$). The near-ceiling ratings indicate strong expert agreement that the system consistently provided appropriate pedagogical explanations for student errors.

6.4.2 Instructional Action Appropriateness (Q1). Actions received consistently positive ratings ($\mu = 4.38$, $\sigma = 1.15$) with moderate interrater agreement (23.9% perfect agreement). The higher variance (average item $\sigma = 0.78$) reflects the pedagogically complex nature of tutoring, where multiple valid instructional approaches often exist for addressing student errors. Qualitative notes revealed that disagreement primarily centered on: (1) pedagogical sequencing (e.g., whether to address morphology before grapheme-phoneme correspondence), (2) depth of explanation appropriate for different age groups, and (3) terminology preferences (e.g., "base" vs. "base word," "vowel digraph" vs. "vowel team").

6.4.3 Expert Confidence. Raters reported high confidence in their evaluations ($\mu = 4.93$, $\sigma = 0.41$, 78.9% perfect agreement), validating the reliability of the expert assessment process.

6.4.4 Interpretation. The contrasting patterns of expert agreement reveal two important characteristics of the system. The high agreement on reasoning validity (78.9% perfect agreement in Table 4) demonstrates that the system exhibits strong pedagogical competence—it consistently generates sound rationales that experts recognize as appropriate. In contrast, the moderate agreement on instructional actions (23.9% perfect agreement in Table 4) reflects genuine pedagogical flexibility rather than a system limitation. Effective tutoring inherently admits multiple valid approaches, and expert teachers frequently disagree on optimal instructional sequences even while concurring on underlying pedagogical principles.

7 Preliminary Feedback from Children

We present findings from our preliminary usability evaluation with 7 children, focusing on their experience with the spelling tool and overall usability perceptions.

7.1 Participants

We recruited 7 children (ages 7–11, grade 2–5) through word-of-mouth and mailing lists. Participant demographics included 6 native English speakers and 1 native Russian speaker, all currently learning to spell in English. When asked about their familiarity with structured word inquiry, no participant reported prior exposure to this approach. Parents reported varied spelling difficulty levels for their children (range 1–5 on a 5-point scale, median = 3), indicating a diverse sample in terms of spelling proficiency.

7.2 Study Procedure

We conducted 45-minute remote sessions via Zoom with children and their parents present. After obtaining parental consent and child assent, we began with a brief demonstration of the tool's features. Children then selected a word to investigate from a set of five words (selected from the from High Frequency Words Project [44])—*friend*, *definitely*, *favorite*, *because*, *beautiful*—to explore using our spelling tool. During the intervention, children were provided with a pre-written sample where their selected word was intentionally misspelled for their morphological investigation.

Following the intervention, we conducted a retrospective think-aloud protocol, asking children to reflect on the decisions they made while using the tool. Children then completed a 10-item system usability scale questionnaire adapted for children using the Flesch-Kincaid readability scale, using a 5-point Likert scale. Parents observed throughout and provided informal feedback. The children received a small toy as compensation.

7.3 Results

Our preliminary evaluation revealed several key findings about the tool's usability and pedagogical potential. Despite the small sample size ($n = 7$), which limits statistical generalization, consistent patterns emerged that provide actionable insights for design iteration.

7.3.1 Engagement and Perceived Value. Children expressed enthusiasm for continued use of the tool (median = 4, range = 3–5). More

ID	Age	Spelling Instruction Exp.	Structured Word Inquiry Exp.
E1	71	20+ years	8 years
E2	34	9 years	2 years
E3	55	15 years	4 years
E4	47	15 years	8 years
E5	47	20+ years	7 years

Table 3: Participant Demographics and Experience

Dimension	Mean (SD)	% ≥4	% =5	t-test vs. neutral (μ=3)	Perfect Agreement
Reasoning	4.93 (0.41)	98.7%	95.4%	t=109.93*	78.9%
Action	4.38 (1.15)	83.1%	70.5%	t=28.14*	23.9%
Expert	4.93 (0.41)	98.7%	95.4%	t=109.93*	78.9%

Table 4: Inter-rater agreement and rating distributions across evaluation dimensions. N = 545 ratings (109 responses × 5 experts per dimension). “% ≥ 4” shows the proportion of ratings at 4 (agree) or 5 (strongly agree); “% = 5” shows the proportion at 5 only. t-tests assessed whether mean ratings exceeded the neutral midpoint (3 = “neither agree nor disagree”). *p < .001.

significantly, during the retrospective think-aloud, all seven participants spontaneously discussed how the tool could help them learn words beyond just fixing spelling errors. A fourth grade participant shared, “I can learn new words with this!”, referring to the word-matrix mini-game he had to complete while learning about “because” and how its base “cause” can be used to make many new words like “causal” and “causes”. Parents echoed this enthusiasm, with one parent observing, “At school I don’t think he gets taught this way.”

7.3.2 *Usability and Learning.* Children responded positively to the tool’s usability across multiple dimensions (Figure 6). Most children found the interface easy to use (median = 5, range = 3–5) and felt confident while using it (median = 5, range = 4–5). One fourth-grade participant stated, “[It was easy because] everything was right there and I didn’t have to click too much.” Children generally felt the tool was not overly complicated (median = 2, range = 1–3), and that other children could learn to use it quickly (median = 3, range = 1–5), suggesting the tool has a small learning curve and can work across age and grade levels. A second-grade participant noted, “First it was confusing but then I could do it myself.”

7.3.3 *Discovering Orthographic Structure.* A striking finding emerged regarding participants’ prior spelling instruction: none of the children reported having been explicitly taught English orthography through morphological analysis. When they made spelling errors in school, they typically received either direct corrections or phonetic explanations. One participant who has been struggling with spelling recollected, “My teacher usually tells me to sound the word out.” All participants unanimously stated that this morphological approach made learning words more enjoyable and meaningful.

7.3.4 *Interface Design.* Children appreciated how the tool guided them through instructions autonomously, making the experience feel more dynamic than filling out a worksheet. Observing the incorrect spelling instance transform into the correct spelling helped them in most instances understand why their hypothesis regarding the spelling was wrong. They liked the different modalities available for input—draggable icons, audio input, and text fields. However, a parent suggested having more animations in the interface to make it more playful and fun for children.

These preliminary findings suggest that the tool addresses a genuine gap in current spelling instruction and demonstrates baseline usability for the target age group. While our sample size prevents statistical claims, the consistent positive responses and identification

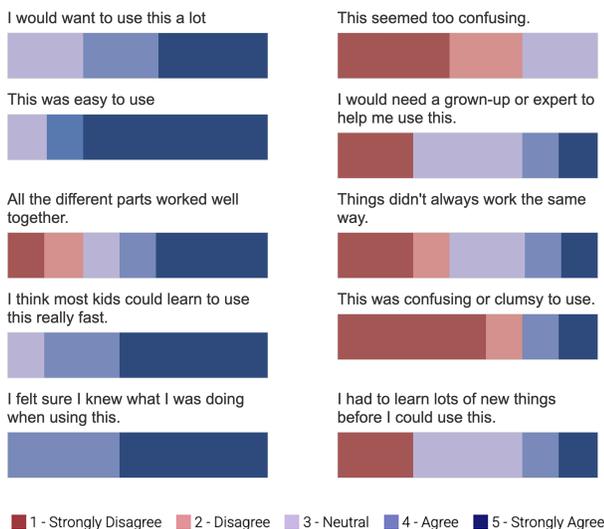


Figure 6: System Usability Survey responses from 7 children. On the left column, higher is better, while on the right column, lower is better.

of morphological learning as novel indicate the tool’s potential for broader impact.

8 Discussion

8.1 Pedagogical Program Synthesis for Personalized and Individualized Learning

Traditional rule-based systems for instructional support perform well in domains with well-defined constraints and limited solution spaces, such as mathematical problem solving, where errors follow predictable patterns. However, literacy instruction, particularly spelling instruction grounded in morphological, phonological, and etymological analysis, presents fundamentally different challenges. The cognitive processes expert educators employ when diagnosing spelling errors and constructing instructional responses cannot be adequately captured through explicit rules alone. A single misspelling can involve multiple overlapping dimensions: grapheme-phoneme mappings, morphological boundaries, suffixing conventions, semantic appropriateness, or etymological patterns. The space of pedagogically valid responses is not enumerable in advance, and the optimal instructional sequence depends on nuanced inference about the learner’s underlying reasoning. This inference requires flexible interpretation rather than pattern matching against predefined templates.

Pedagogical Program Synthesis addresses this limitation by combining the structured reasoning of domain-specific languages with the interpretive flexibility of large language models. The DSL provides instructional scaffolding, encoding hypothesis templates, evidential features, and compositional primitives that preserve pedagogical validity, while LLMs enable dynamic adaptation to individual learners’ contexts. This architecture opens possibilities that purely corrective systems cannot achieve. A learner initially seeking to correct a single grapheme error may discover, through synthesized inquiry sequences, the morphological family structure explaining that grapheme choice, the etymological origins preserving it across related words, and the systematic patterns governing similar cases. The learning effect extends far beyond the immediate correction, building metalinguistic awareness that may be transferred to future encounters.

Current classroom instruction rarely provides this form of structured, individualized morphological investigation. SPIRE demonstrates that Pedagogical Program Synthesis can serve as a vessel for structured literacy experiences that do not currently exist at scale, capturing the dynamic and responsive processes that expert educators employ while making them accessible within everyday writing contexts. The approach suggests a broader applicability to instructional domains that require flexible and context-sensitive reasoning about learner understanding.

8.2 Clinical Implications

SPIRE demonstrates broad applicability across instructional settings aligned with Multi-Tiered Systems of Supports frameworks commonly implemented in schools. The system architecture enables deployment at multiple intervention levels: classroom teachers can integrate it during universal instruction (Tier 1), specialists can use it for targeted small group support (Tier 2), and speech-language pathologists can incorporate it within intensive individualized intervention (Tier 3). Beyond educator-mediated instruction, the system

supports independent student practice, facilitating distributed practice across various learning contexts, a condition known to enhance skill generalization beyond clinical settings.

A critical limitation in current writing instruction concerns the temporal misalignment between error production and corrective feedback. Traditional classroom workflows require educators to collect student writing, perform a delayed review, and provide asynchronous feedback. This sequence separates the correction from the motivational and cognitive context in which the errors occurred. SPIRE addresses this temporal gap by providing instruction precisely when learners make spelling errors, capitalizing on the moment of authentic confusion to introduce systematic orthographic investigation. This immediacy transforms errors from failures that require correction into diagnostic opportunities for structured inquiry, aligning the pedagogical intervention with the readiness and intrinsic motivation of the learner.

The system also addresses a persistent resource constraint in specialized literacy support. Speech-language pathologists face significant service demands that limit their capacity to deliver intensive intervention to all students who need support. By automating the synthesis of instructional sequences grounded in pedagogy, SPIRE does not replace clinical judgment, but rather extends the SLP capacity by operationalizing established instructional principles within everyday writing tools where students spend substantial time.

A unique and significant contribution of SPIRE lies in its dual function as both a student intervention and a professional learning resource. It provides a practical, contextualized model for understanding the English orthographic system, emphasizing the critical interrelationship between morphology, etymology, and phonology. Given that many educator training programs fail to provide explicit instruction on the logical and structural principles of English spelling, often leaving the impression that the system is illogical, SPIRE addresses directly this pervasive training gap. By allowing educators to learn and apply these principles alongside their students, it supports continuous professional development and enhances the fidelity of literacy instruction in the school environment.

8.3 Limitations and Future Work

SPIRE’s current architecture relies extensively on API calls to commercial large language models: Anthropic’s Claude Sonnet for hypothesis filtering, trace generation, and program synthesis, and Google’s Gemini for speech-to-text and text-to-speech processing. This dependency introduces two practical constraints. First, the cumulative cost of API calls scales linearly with usage, creating potential barriers to widespread deployment in resource-constrained educational settings. Second, these general-purpose models contain vastly more parameters and capabilities than required for the specific tasks SPIRE performs, representing computational overhead that lighter, task-specialized models could avoid. Local models fine-tuned for particular subtasks, such as extracting morphological properties from target words or filtering hypothesis templates based on error patterns, could substantially reduce both latency and cost while maintaining pedagogical quality. However, the primary focus of this work was establishing the viability of Pedagogical Program Synthesis as an instructional paradigm and demonstrating alignment with expert practice. System optimization through custom model

training represents an important engineering direction for future work, but was outside the scope of this initial investigation.

Our empirical evaluation, while providing valuable insight into the pedagogical validity and usability of SPIRE, involved modest sample sizes: five speech-language pathology experts and seven children. The expert evaluation yielded strong evidence of pedagogical soundness, with near-ceiling ratings and high inter-rater agreement on reasoning validity. The preliminary user study with children revealed consistent patterns that suggested good usability and engagement. However, these sample sizes exclude robust quantitative claims about effectiveness or generalizability in diverse populations of learners, grade levels, and ranges of spelling proficiency. A larger-scale study would strengthen statistical power and enable a more nuanced analysis of how SPIRE's effectiveness varies between learner characteristics.

Most critically, our evaluation captured only immediate usability and pedagogical appropriateness, not learning outcomes over time. A longitudinal deployment study tracking children's spelling development, transfer of morphological reasoning to untrained words, and sustained engagement across authentic writing tasks would provide compelling evidence of SPIRE's instructional impact. Such work would also illuminate how repeated exposure to structured word inquiry through SPIRE influences learners' metalinguistic awareness and independent word investigation strategies—outcomes central to the pedagogical philosophy motivating this research.

9 Conclusion

This work addressed a fundamental gap in literacy instruction: while speech-language pathologists employ inquiry-based approaches to teach spelling through structured word investigation, everyday writing tools remain limited to detection and autocorrection. We introduced SPIRE (Spelling Inquiry Engine), which brings structured word inquiry into the act of composition through Pedagogical Program Synthesis, a novel approach that operationalizes the inherently dynamic pedagogy of spelling instruction. Central to our contribution is a domain-specific language that represents SLP instructional moves as composable primitives, enabling the system to synthesize tailored programs in real-time from learner errors and render them as interactive interfaces. Each synthesized program guides learners through brief cycles of hypothesis testing on word meanings, structures, morphological families, origins, and grapheme-phoneme correspondences, transforming spelling errors into opportunities for metalinguistic reasoning alongside correction. Evaluation with five speech-language pathology experts demonstrated strong pedagogical validity, with near-ceiling ratings on reasoning appropriateness and high interrater agreement. A preliminary usability assessment with seven children revealed positive engagement and identification of morphological investigation as a novel and meaningful learning experience absent from current classroom instruction. SPIRE demonstrates that program synthesis can capture expert pedagogical reasoning that purely rule-based systems cannot, offering a vessel for structured literacy experiences at scale. This work suggests broader applicability of Pedagogical Program Synthesis to instructional domains requiring flexible, context-sensitive reasoning about learner understanding, extending beyond spelling to other areas where expert practice must adapt dynamically to individual learner needs.

10 GenAI Usage Disclosure

In compliance with the ACM policy on the usage of Generative AI, we disclose that ChatGPT (OpenAI GPT-4/5) and Claude (Sonnet 4-5) was used for text refinement and code scaffolding. The authors maintained full intellectual control over the research design, analysis, interpretation, and writing. All generated content was reviewed and revised to ensure factual accuracy and originality.

References

- [1] A. W. Alexander and Anne-Marie Slinger-Constant. 2004. Current Status of Treatments for Dyslexia: Critical Review. *Journal of child neurology* (2004). <https://doi.org/10.1177/08830738040190100401>
- [2] Anthropic. 2025. Claude Sonnet 4.5. <https://www.anthropic.com>. Model: claude-sonnet-4-5-20250929. Accessed: 2025-10-10.
- [3] Nuwan T. Attygalle, Matjaž Kljun, Aaron Quigley, Klen Čopič Pucihar, Jens Grubert, Verena Biener, Luis A. Leiva, Juri Yoneyama, Alice Toniolo, Angela Miguel, Hirokazu Kato, and Maheshya Weerasinghe. 2025. Text-to-Image Generation for Vocabulary Learning Using the Keyword Method. In *Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI '25)*. Association for Computing Machinery, New York, NY, USA, 1381–1397. <https://doi.org/10.1145/3708359.3712073>
- [4] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and Charles Sutton. 2021. Program Synthesis with Large Language Models. *arXiv: Programming Languages* (2021).
- [5] Lisa Yambrick Barnett and Katie E. Squires. 2024. *Vocabulary & Morphology using Structured Word Inquiry: Teacher Guide, Lesson Planning Template, & Student Activity Booklet of Tier 2 Vocabulary from Core Subject Areas, Elementary - Middle School+*. Independently published.
- [6] J. Bowers and Peter N. Bowers. 2017. Beyond Phonics: The Case for Teaching Children the Logic of the English Spelling System. (2017). <https://doi.org/10.1080/00461520.2017.1288571>
- [7] Peter N. Bowers. 2006. *Teaching How the Written Word Works: Word Study Lessons Using a Model of Structured Word Inquiry*. WordWorks Literacy Centre.
- [8] Peter N. Bowers and John R. Kirby. 2010. Effects of morphological instruction on vocabulary acquisition. *Reading and Writing* (2010). <https://doi.org/10.1007/s11445-009-9172-z>
- [9] Desiré Carioti, Marta Franca Masia, Simona Travellini, and Manuela Berlinger. 2021. Orthographic depth and developmental dyslexia: a meta-analytic study. *Annals of Dyslexia* (2021). <https://doi.org/10.1007/s11881-021-00226-0>
- [10] Luciana Cidrim, Andrea Oliveira Batista, Francisco Madeiro, and S. Capellini. 2024. Analysis of spelling errors from the 'dyslexic sight words' list. *Frontiers in Psychology* (2024). <https://doi.org/10.3389/fpsyg.2024.1160247>
- [11] Danielle Colenbrander, L. Parsons, J. Bowers, and C. Davis. 2021. Assessing the Effectiveness of Structured Word Inquiry for Students in Grades 3 and 5 With Reading and Spelling Difficulties: A Randomized Controlled Trial. *Reading Research Quarterly* (2021). <https://doi.org/10.1002/rrq.399>
- [12] Jennifer A. Constantine and Kara B. Lee. 2024. *Structured Word Inquiry: An Implementation Guide for Teachers*. Independently published.
- [13] Aayushi Dangol, Aaleyah Lewis, Hyewon Suh, Xuesi Hong, Hedda Meadan, James Fogarty, and Julie A. Kientz. 2025. "I Want to Think Like an SLP": A Design Exploration of AI-Supported Home Practice in Speech Therapy. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 380, 22 pages. <https://doi.org/10.1145/3706598.3713986>
- [14] Jean-Luc Dorier and Francisco García García. 2013. Challenges and opportunities for the implementation of inquiry-based learning in day-to-day teaching. *Zdm* (2013). <https://doi.org/10.1007/s11858-013-0512-8>
- [15] P. Dymora and Karol Niemiec. 2019. Gamification as a Supportive Tool for School Children with Dyslexia. *Informatics* (2019). <https://doi.org/10.3390/informatics6040048>
- [16] L. Ehri. 1995. Phases of development in learning to read words by sight. (1995). <https://doi.org/10.1111/j.1467-9817.1995.tb00077.x>
- [17] L. Ehri. 1998. Grapheme-phoneme knowledge is essential for learning to read words in English. (1998). <https://doi.org/10.4324/9781410602718-6>
- [18] L. Ehri. 2005. Learning to Read Words: Theory, Findings, and Issues. (2005). https://doi.org/10.1207/s1532799xssr0902_4
- [19] Min Fan, Alissa N. Antle, Maureen Hoskyn, and Carman Neustaedter. 2018. A design case study of a tangible system supporting young English language learners. *International Journal of Child-Computer Interaction* (2018). <https://doi.org/10.1016/j.ijcci.2018.08.001>
- [20] Min Fan, Jianyu Fan, A. Antle, Sheng Jin, Dongxu Yin, and P. Pasquier. 2019. Character Alive: A Tangible Reading and Writing System for Chinese Children At-risk for Dyslexia. *CHI Extended Abstracts* (2019). <https://doi.org/10.1145/3290607.3312756>

- [21] B. Foorman, D. Francis, J. Fletcher, C. Schatschneider, and P. Mehta. 1998. The Role of Instruction in Learning To Read: Preventing Reading Failure in At-Risk Children. (1998). <https://doi.org/10.1037/0022-0663.90.1.37>
- [22] Sharon Friesen and David Scott. 2013. Inquiry-Based Learning: A Review of the Research Literature. (2013).
- [23] Michael Gamon, Jianfeng Gao, Chris Brockett, Alexandre Klementiev, William B. Dolan, Dmitriy Belenko, and Lucy Vanderwende. 2008. Using Contextual Speller Techniques and Language Modeling for ESL Error Correction. *International Joint Conference on Natural Language Processing* (2008).
- [24] Xinya Gong, Wenhui Tao, and Yuxin Ma. 2025. CalliSense: An Interactive Educational Tool for Process-based Learning in Chinese Calligraphy. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [25] D. Gooch, A. Vasalou, L. Benton, and Rilla Khaled. 2016. Using Gamification to Motivate Students with Dyslexia. *International Conference on Human Factors in Computing Systems* (2016). <https://doi.org/10.1145/2858036.2858231>
- [26] Steven M. Goodman, Erin Buehler, Patrick Clary, Andy Coenen, Aaron Donsbach, Tiffanie N. Horne, Michal Lahav, Robert Macdonald, R. Michaels, Ajit Narayanan, Mahima Pushkarna, Joel Riley, A. Santana, Lei Shi, Rachel Sweeney, Phil Weaver, Ann Yuan, and M. Morris. 2022. LaMPost: Design and Evaluation of an AI-assisted Email Writing Prototype for Adults with Dyslexia. *International ACM SIGACCESS Conference on Computers and Accessibility* (2022). <https://doi.org/10.1145/3517428.3544819>
- [27] Grammarly Inc. 2025. Grammarly. <https://app.grammarly.com/> Accessed: 2025-10-10.
- [28] Sumit Gulwani, Oleksandr Polozov, and Rishabh Singh. 2017. Program Synthesis. *Software Systems Safety* (2017). <https://doi.org/10.3233/978-1-61499-385-8-43>
- [29] Douglas Harper. 2001–2025. *Reach*. Online Etymology Dictionary. <https://www.etymonline.com/search?q=reach> Accessed: 2025-10-11.
- [30] P. Hatcher, Kristin Goetz, M. Snowling, C. Hulme, S. Gibbs, and Glynnis Smith. 2006. Evidence for the effectiveness of the Early Literacy Support programme. *The British journal of educational psychology* (2006). <https://doi.org/10.1348/000709905x39170>
- [31] Sue Scibetta Hegland. 2021. *Beneath the Surface of Words: What English Spelling Reveals and Why It Matters*. Independently published.
- [32] S. Hines. 2009. The Effectiveness of a Color-Coded, Onset-Rime Decoding Intervention with First-Grade Students at Serious Risk for Reading Disabilities. (2009). <https://doi.org/10.1111/j.1540-5826.2008.01274.x>
- [33] John Hung, Christopher Cui, Diana M. Popescu, Saurabh Chatterjee, and Thad Starner. 2024. Socratic Mind: Scalable Oral Assessment Powered By AI. *ACM Conference on Learning @ Scale* (2024). <https://doi.org/10.1145/3657604.3664661>
- [34] Jason Jabbour, Kai Kleinbard, Olivia Miller, Robert C Haussman, and V. Reddi. 2025. SocratiQ: A Generative AI-Powered Learning Companion for Personalized Education and Broader Accessibility. (2025).
- [35] Laurice M. Joseph. 2002. Facilitating Word Recognition and Spelling Using Word Boxes and Word Sort Phonic Procedures. (2002). <https://doi.org/10.1080/02796015.2002.12086146>
- [36] David A. Joyner and Ashok K. Goel. 2015. Improving Inquiry-Driven Modeling in Science Education through Interaction with Intelligent Tutoring Agents. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (Atlanta, Georgia, USA) (IUI '15). Association for Computing Machinery, New York, NY, USA, 5–16. <https://doi.org/10.1145/2678025.2701398>
- [37] Majeed Kazemitabaar, Runlong Ye, Xiaoning Wang, Austin Z. Henley, Paul Denny, Michelle Craig, and Tovi Grossman. 2024. CodeAid: Evaluating a Classroom Deployment of an LLM-based Programming Assistant that Balances Student and Educator Needs. *International Conference on Human Factors in Computing Systems* (2024). <https://doi.org/10.1145/3613904.3642773>
- [38] A. Kuerten, M. Mota, and K. Segaert. 2019. Developmental dyslexia: a condensed review of literature. *Ilha do Desterro A Journal of English Language, Literatures in English and Cultural Studies* (2019). <https://doi.org/10.5007/2175-8026.2019v72n3p249>
- [39] A. Lazonder and Ruth Harmsen. 2016. Meta-Analysis of Inquiry-Based Learning. (2016). <https://doi.org/10.3102/0034654315627366>
- [40] Horng-Yi Lee. 2014. Inquiry-based Teaching in Second and Foreign Language Pedagogy. (2014). <https://doi.org/10.4304/jltr.5.6.1236-1244>
- [41] Joanne Leong, Pat Pataranutaporn, Valdemar Danry, Florian Perteneder, Yaoli Mao, and Pattie Maes. 2024. Putting Things into Context: Generative AI-Enabled Context Personalization for Vocabulary Learning Improves Learning Motivation. *International Conference on Human Factors in Computing Systems* (2024). <https://doi.org/10.1145/3613904.3642393>
- [42] Aaleyah Lewis, Aayushi Dangol, Hyewon Suh, Abbie Olszewski, James Fogarty, and Julie A. Kientz. 2025. Exploring AI-Based Support in Speech-Language Pathology for Culturally and Linguistically Diverse Children. *International Conference on Human Factors in Computing Systems* (2025). <https://doi.org/10.1145/3706598.3714131>
- [43] Alberto Quattrini Li, Licia Sbattella, and Roberto Tedesco. 2013. PoliSpell: An Adaptive Spellchecker and Predictor for People with Dyslexia. *Lecture Notes in Computer Science* (2013). https://doi.org/10.1007/978-3-642-38844-6_27
- [44] Rebecca Loveless and Fiona Hamilton. 2024. *The High Frequency Word Project*. Independently published.
- [45] Jiaju Ma and Maneesh Agrawala. 2025. MoVer: Motion Verification for Motion Graphics Animations. *ACM Trans. Graph.* 44, 4, Article 33 (July 2025), 17 pages. <https://doi.org/10.1145/3731209>
- [46] Christopher J. MacLellan. 2024. py_plan: A Python Planning Library. https://github.com/cmaclell/py_plan. Accessed: 2025-10-10.
- [47] Pratyusha Maiti and Ashok K. Goel. 2025. Can an AI Partner Empower Learners to Ask Critical Questions? *Proceedings of the 30th International Conference on Intelligent User Interfaces* (2025). <https://doi.org/10.1145/3708359.3712134>
- [48] G. McArthur, Yumi Sheehan, N. Badcock, Deanna A. Francis, Hua-Chen Wang, Saskia Kohnen, Erin Banales, T. Anandakumar, E. Marinus, and A. Castles. 2018. Phonics training for English-speaking poor readers. *Cochrane Database of Systematic Reviews* (2018). <https://doi.org/10.1002/14651858.cd009115.pub3>
- [49] M. Morris, Adam Fournay, Abdullah Ali, and Laura Vonessen. 2018. Understanding the Needs of Searchers with Dyslexia. *International Conference on Human Factors in Computing Systems* (2018). <https://doi.org/10.1145/3173574.3173609>
- [50] Bettina Müller, Tobias Richter, and Panagiotis Karageorgos. 2020. Syllable-based reading improvement: Effects on word reading and reading comprehension in Grade 2. *Learning and Instruction* (2020). <https://doi.org/10.1016/j.learninstruc.2020.101304>
- [51] Sumit Pandey and S. Srivastava. 2011. SpellBound: a tangible spelling aid for the dyslexic child. *IndiaHCI* (2011). <https://doi.org/10.1145/2407796.2407813>
- [52] Margus Pedaste, Mario Mäeots, Leo A. Siiman, Ton de Jong, Siswa A. N. van Riesen, Ellen T. Kamp, Constantinos C. Manoli, Zacharias C. Zacharia, and Eleftheria Tsourlidaki. 2015. Phases of inquiry-based learning: Definitions and the inquiry cycle. *Educational Research Review* (2015). <https://doi.org/10.1016/j.edurev.2015.02.003>
- [53] Zhenhui Peng, Xingbo Wang, Qiushi Han, Junkai Zhu, Xiaojuan Ma, and Huamin Qu. 2023. Storyfier: Exploring Vocabulary Learning Support with Text Generation Models. *arXiv.org* (2023). <https://doi.org/10.48550/arxiv.2308.03864>
- [54] Charles Perfetti. 2007. Reading ability: Lexical quality to comprehension. *Scientific studies of reading* 11, 4 (2007), 357–383.
- [55] M. Rauschenberger, Silke Füchsel, Luz Rello, Clara Bayarri, and A. Gorriz. 2015. A Game to Target the Spelling of German Children with Dyslexia. *ASSETS* (2015). <https://doi.org/10.1145/2700648.2811345>
- [56] Maria Rauschenberger, Ricardo Baeza Yates, and Luz Rello. 2019. Technologies for Dyslexia. *Human-computer interaction series* (2019). https://doi.org/10.1007/978-1-4471-7440-0_31
- [57] Real Spelling. 2025. Real Spelling Toolbox 2. <https://www.tbox2.com/> Accessed: 2025-10-10.
- [58] Luz Rello, Ricardo Baeza-Yates, and Joaquim Llisterri. 2014. DysList: An Annotated Resource of Dyslexic Errors. *Language Resources and Evaluation* (2014).
- [59] Luz Rello, Miguel Ballesteros, and Jeffrey P. Bigham. 2015. A Spellchecker for Dyslexia. *International ACM SIGACCESS Conference on Computers and Accessibility* (2015). <https://doi.org/10.1145/2700648.2809850>
- [60] Luz Rello, Clara Bayarri, Yolanda Otal, and Martin Pietlot. 2015. A Computer-Based Method to Improve the Spelling of Children with Dyslexia. *arXiv: Human-Computer Interaction* (2015). <https://doi.org/10.1145/2661334.2661373>
- [61] Sherry Ruan, Liwei Jiang, Qianyao Xu, Zhiyuan Liu, Glenn M Davis, Emma Brunskill, and James A. Landay. 2021. EnglishBot: An AI-Powered Conversational System for Second Language Learning. In *Proceedings of the 26th International Conference on Intelligent User Interfaces* (College Station, TX, USA) (IUI '21). Association for Computing Machinery, New York, NY, USA, 434–444. <https://doi.org/10.1145/3397481.3450648>
- [62] Joni Salminen, Chang Liu, Wenjing Pian, Jianxing Chi, Essi Häyhänen, and Bernard Jansen. 2024. Deus ex machina and personas from large language models: Investigating the composition of ai-generated persona descriptions. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [63] Nora Schlesinger and Shelley Gray. 2017. The impact of multisensory instruction on learning letter names and sounds, word reading, and spelling. *Annals of Dyslexia* (2017). <https://doi.org/10.1007/s11881-017-0140-z>
- [64] Claudia Schmidt, Janin Brandenburg, Jenny Busch, Gerhard Büttner, Dietmar Grube, Claudia Mähler, and Marcus Hasselhorn. 2021. Developmental Trajectories of Phonological Information Processing in Upper Elementary Students with Reading or Spelling Disabilities. *Reading Research Quarterly* (2021). <https://doi.org/10.1002/rrq.299>
- [65] Mark S Seidenberg and James L McClelland. 1989. A distributed, developmental model of word recognition and naming. *Psychological review* 96, 4 (1989), 523.
- [66] E. Silliman, R. Bahr, and M. Peters. 2006. Spelling Patterns in Preadolescents With Atypical Language Skills: Phonological, Morphological, and Orthographic Factors. *Developmental neuropsychology* (2006). https://doi.org/10.1207/s15326942dn2901_6
- [67] M. Snowling and C. Hulme. 2011. Evidence-based interventions for reading and language difficulties: creating a virtuous circle. *The British journal of educational psychology* (2011). <https://doi.org/10.1111/j.2044-8279.2010.02014.x>

- [68] Katie Squires. 2013. Addressing the Shortage of Speech-Language Pathologists in School Settings. *Journal of the American Academy of Special Education Professionals* (2013).
- [69] J. Torgesen, R. Wagner, C. A. Rashotte, E. Rose, P. Lindamood, T. Conway, and C. Garvan. 1999. Preventing reading failure in young children with phonological processing disabilities: Group and individual responses to instruction. (1999). <https://doi.org/10.1037/0022-0663.91.4.579>
- [70] Joseph K. Torgesen, Ann W. Alexander, Richard K. Wagner, Carol A. Rashotte, Kytja K. S. Voeller, and Tim Conway. 2001. Intensive Remedial Instruction for Children with Severe Reading Disabilities: Immediate and Long-term Outcomes From Two Instructional Approaches. *Journal of Learning Disabilities* (2001). <https://doi.org/10.1177/002221940103400104>
- [71] F. Vellutino, J. Fletcher, M. Snowling, and D. Scanlon. 2004. Specific reading disability (dyslexia): what have we learned in the past four decades? *Journal of child psychology and psychiatry, and allied disciplines* (2004). <https://doi.org/10.1046/j.0021-9630.2003.00305.x>
- [72] Ruijie Wang, Liming Chen, Ivar Solheim, Trenton Schulz, and Aladdin Ayesh. 2017. Conceptual Motivation Modeling for Students with Dyslexia for Enhanced Assistive Learning. In *Proceedings of the 2017 ACM Workshop on Intelligent Interfaces for Ubiquitous and Smart Learning (Limassol, Cyprus) (SmartLearn '17)*. Association for Computing Machinery, New York, NY, USA, 11–18. <https://doi.org/10.1145/3038535.3038542>
- [73] Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah Goodman. 2024. Hypothesis Search: Inductive Reasoning with Language Models. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=G7UtlGQmjm>
- [74] Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah D Goodman. 2023. Hypothesis search: Inductive reasoning with language models. *arXiv preprint arXiv:2309.05660* (2023).
- [75] Xinyu Jessica Wang, Christine Lee, and Bilge Mutlu. 2025. LearnMate: Enhancing Online Education with LLM-Powered Personalized Learning Plans and Support. *arXiv.org* (2025). <https://doi.org/10.1145/3706599.3719857>
- [76] Yadeesha Weerasinghe, Kahandakanaththage Maduni Pramuditha Perera, Sanka Chandika Gallage, Soundarya Ramesh, Kasun Karunanayaka, and Suranga Nanayakkara. 2025. LexiLearn: Gamifying Learning for Dyslexia in a Novel Linguistic Context. *CHI Extended Abstracts* (2025). <https://doi.org/10.1145/3706599.3720169>
- [77] Shaomei Wu, Lindsay Reynolds, Xian Li, and Francisco Guzmán. 2019. Design and Evaluation of a Social Media Writing Support Tool for People with Dyslexia. *International Conference on Human Factors in Computing Systems* (2019). <https://doi.org/10.1145/3290605.3300746>
- [78] Tianyi Zhang, London Lowmanstone, Xinyu Wang, and Elena L. Glassman. 2020. Interactive Program Synthesis by Augmented Examples. *UIST* (2020). <https://doi.org/10.1145/3379337.3415900>