

Evaluating and Optimizing Educational Content with Large Language Model Judgments

Joy He-Yueya
Stanford University
heyueya@cs.stanford.edu

Noah D. Goodman
Stanford University
ngoodman@stanford.edu

Emma Brunskill
Stanford University
ebrun@cs.stanford.edu

ABSTRACT

Creating effective educational materials generally requires expensive and time-consuming studies of student learning outcomes. To overcome this barrier, one idea is to build computational models of student learning and use them to optimize instructional materials. However, it is difficult to model the cognitive processes of learning dynamics. We propose an alternative approach that uses Language Models (LMs) as educational experts to assess the impact of various instructions on learning outcomes. Specifically, we use GPT-3.5 to evaluate the overall effect of instructional materials on different student groups and find that it can replicate well-established educational findings such as the Expertise Reversal Effect and the Variability Effect. This demonstrates the potential of LMs as reliable evaluators of educational content. Building on this insight, we introduce an instruction optimization approach in which one LM generates instructional materials using the judgments of another LM as a reward function. We apply this approach to create math word problem worksheets aimed at maximizing student learning gains. Human teachers' evaluations of these LM-generated worksheets show a significant alignment between the LM judgments and human teacher preferences. We conclude by discussing potential divergences between human and LM opinions and the resulting pitfalls of automating instructional design.¹

Keywords

Large language models, Instructional design, Educational content development, Math education

1. INTRODUCTION

Instructional design, the process of creating educational materials or experiences such as textbooks and courses, is a critical component in advancing education [41]. A significant challenge in instructional design is the need for extensive

¹Our prompts are publicly available at <https://github.com/StanfordAI4HI/ed-expert-simulator>.

J. He-Yueya, N. D. Goodman, and E. Brunskill. Evaluating and optimizing educational content with large language model judgments. In B. Paaßen and C. D. Epp, editors, *Proceedings of the 17th International Conference on Educational Data Mining*, pages 68–82, Atlanta, Georgia, USA, July 2024. International Educational Data Mining Society.

© 2024 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.
<https://doi.org/10.5281/zenodo.12729776>

studies involving real students to evaluate the effectiveness of these instructional materials [16, 26]. This traditional method is expensive, time-consuming, and fraught with logistical and ethical challenges, making it difficult to quickly innovate and implement new teaching strategies. Recent studies have explored the use of Language Models (LMs) to simulate students' interactions with educational content [21, 12], offering a less expensive alternative. However, our initial investigations reveal that the LMs (at the time) struggled to model the *dynamics* of student learning, often failing to maintain a consistent level of knowledge when simulating students' responses before and after learning interventions (see Appendix A).

In light of these challenges, our work explores the use of LMs, such as GPT-3.5 and GPT-4 [1], to evaluate and optimize educational materials. Unlike previous attempts to directly simulate student learning [21, 12, 25, 28, 24], we leverage the advanced reasoning capabilities and pedagogical knowledge of LMs and position LMs as educational experts who can assess and enhance the effectiveness of instructional content.

To validate the potential of LMs in simulating educational experts, we use GPT-3.5 to assess the overall effect of instructional materials on different student groups and find that the judgments of GPT-3.5 can replicate two well-known findings in educational psychology: the Expertise Reversal Effect [14, 43, 15, 13] and the Variability Effect [29, 42]. These results suggest that LMs have the potential to act as coherent evaluators of instructions, offering insights consistent with those obtained from human subjects research.

Building on the insight that LMs can, to some extent, mimic educational experts, along with prior work demonstrating LMs' capability for iterative improvement [20, 5, 50, 49, 36], we propose an instruction optimization approach (Figure 1). In this approach, one LM (the *optimizer*) generates new instructional materials, and another LM (the *evaluator*) evaluates these materials by predicting students' learning outcomes (e.g., post-test scores). We apply this approach to optimize math word problem worksheets, aiming to maximize students' post-test scores. External assessments like post-tests are not perfect, but they are an important tool in many educational systems to help shed insight into student progress. In each optimization step, we prompt the optimizer LM to generate new worksheets based on a list of previously generated worksheets with their corresponding post-test scores, then the new worksheets are scored by

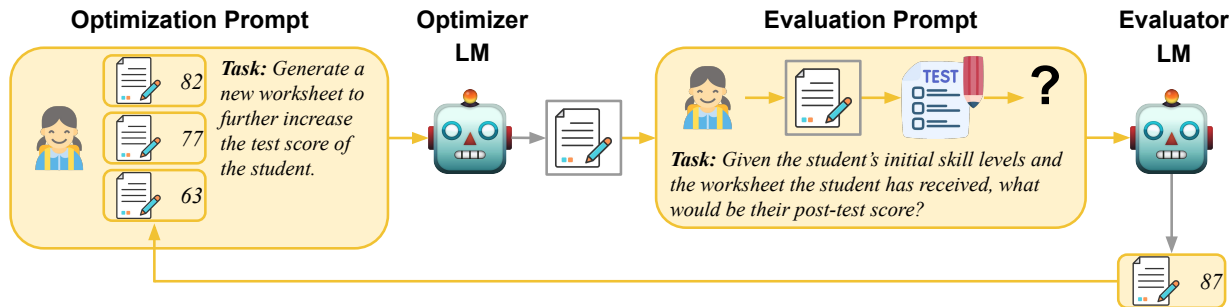


Figure 1: We use an optimizer LM to generate new educational materials (e.g., worksheets) and an evaluator LM to judge the effectiveness of these materials by predicting students’ post-test performance. In this example, the optimizer LM generates a new worksheet based on the student information and previous worksheet-score pairs, then the evaluator LM predicts the student’s post-test score given the new worksheet. The new worksheet-score pair serves as feedback for the optimizer LM to refine worksheets.

the evaluator LM. We show that the optimizer LM can refine worksheets based on the judgments of the evaluator LM. We also ask human teachers to compare worksheets generated from different stages of optimization and find a significant correlation between the evaluator LM’s judgments and human preferences. This highlights the potential of LMs in informing the design of real-life experiments and reducing the number of costly experiments in education. However, human teachers sometimes cannot distinguish worksheets that LMs perceive as different, which suggests the necessity for further investigations to ensure LMs effectively complement traditional educational research methodologies.

Our paper makes the following contributions:

- We demonstrate the potential of LMs to serve as reliable evaluators of educational content by replicating two well-established educational findings: the Expertise Reversal Effect and the Variability Effect.
- We introduce an instruction optimization approach using the LM judgments as a reward function and demonstrate the feasibility of using LMs to iteratively improve educational materials, focusing on the domain of math word problem solving.
- We recruit human teachers to evaluate LM-generated math word problem worksheets from different stages of optimization and find a significant alignment between human teacher preferences and LM judgments, highlighting a promising application of LMs in informing the design of costly human subjects experiments in education.

After describing these results, we conclude the paper with a discussion of open issues and future directions.

2. RELATED WORK

There is an extensive and growing literature on simulating human behaviors and LM-based tools for learning and instructional design.

2.1 Simulating students for education

Long before LM-based tools, researchers explored methods to build simulated students (usually machine learning systems whose behavior is consistent with data from human

students) and applications of such simulated students in education [44, 3, 8, 24, 22]. For example, [44] discusses how teachers can develop and practice their tutoring strategies by teaching simulated students, as demonstrated in recent work that develops an interactive chat-based tool that allows teachers to practice with LM-simulated students [21]. Another promising application of student simulators considered in [44] is to enable collaborative learning where students can work collaboratively with a simulated peer or by teaching a less knowledgeable simulated student [23, 19, 39, 11].

Our work is closely related to prior studies on simulation-based evaluations of instructional design [27, 44, 24]. Running empirical experiments with real students to test different instructions, often called formative evaluations [26], can help instructional designers pilot-test and improve their designs. Since formative evaluations with real students are expensive, it is of interest to build simulated students who can provide cheap, fast evaluations. [27] introduces Heuristic Searcher, a system that can be used as a simulated student who learns arithmetic skills from instructions encoded as formal constraints that describe erroneous states. They show that simulations can reveal interactions of instruction with students’ prior knowledge. Similarly, [24] uses a cognitive model as a simulated student to help design instructions for training circuit board assemblers. However, these systems are limited by their inability to process instructions in natural language or adapt to broader, less structured knowledge domains. Although advancements in LMs [1, 6] offer opportunities to simulate students who can interact with educational content in natural language, recent attempts show that LM-simulated students can at times produce responses that are unrealistic or overly advanced, particularly in situations where such capability is not expected [21, 2]. Therefore, instead of focusing on simulating students who can interact with instructions, we aim to simulate educational experts capable of evaluating the quality of those instructions.

2.2 LMs as educational experts and tutors

Recent advancements in LMs [1, 6] have sparked a new wave of work that explores ways in which LMs can improve education. Some studies have considered using LMs for educational content development [38, 39, 7, 47, 51]. [38] uses LMs to produce educational resources (e.g., code ex-

amples and explanations) for an introductory programming course. They find that students perceive the quality of LM-generated resources to be similar to materials produced by their peers. [45] explores the ability of LMs to become teacher coaches who can provide teachers with pedagogical suggestions. There have been many papers that focus on using LMs to provide automated feedback and explanations to students learning math [31, 46, 17, 35], programming [30, 34], and physics [37]. None of these studies considers optimizing educational content for a given student.

2.3 LMs for simulating humans

Our work is loosely related to recent work on using LMs to simulate different aspects of human behavior. For instance, researchers have used LMs to replicate results from social science experiments and public opinion surveys [4, 2, 10]. The ability of LMs to mimic human behaviors opens up exciting possibilities in areas such as education [21], product design [33, 32], and skill training [40, 18].

3. INSTRUCTION EVALUATION

To do instructional design without running experiments with real students, we introduce a new type of evaluation for assessing the impact of instructions on student learning gains.

3.1 LM-based simulated expert evaluation

Traditionally, developing effective instructional materials involves administering pre-tests and post-tests to numerous students under different experimental conditions. However, this type of evaluation is expensive and time-consuming [44, 16]. Another option is to have pedagogical experts create content based on their (often implicit) assessment of how well it will support student learning [48]. Such human experts can also judge whether particular instructions are likely to be more or less effective.

Inspired by this, we use an LM to simulate a human pedagogical expert evaluation of a particular instructional content that can be described in text. More precisely, we task an LM to take in information about a student’s relevant background (e.g., the student’s prior knowledge), a particular set of instructional content that would be provided to the student, and predict the student’s subsequent performance on particular test questions. We outline this procedure in Alg. 1. We call this evaluation a Simulated Expert Evaluation (SEE). In general, we are interested in understanding whether a particular instructional artifact is effective for a group of students, such as eighth graders doing algebra, or fifth graders struggling with fractions. To do this evaluation, we will have a simulated expert assess the impact on many different students with different prior knowledge or background, to get an aggregate estimate of the impact of a particular instruction.

More precisely, the input of an SEE are an evaluator LM M_e (for example, ChatGPT/Claude/Gemini/etc), a text description of a student persona, a text instruction, an evaluation task description, and T test questions x_1, \dots, x_T . In an SEE, we first initialize an evaluator LM instance m_e . We create a list of prompts $\{p_1, p_2, \dots, p_T\}$ and sequentially give these prompts as input to m_e . The initial prompt p_1 (see an example in Figure 2) includes the student persona that

Algorithm 1 Simulated Expert Evaluation (SEE)

Input: Evaluator M_e , *student_persona*, *instruction*, *eval_task*, test questions x_1, \dots, x_T
Initialize m_e as an instance of M_e
Set $p_1 \leftarrow \text{student_persona} \parallel \text{instruction} \parallel x_1 \parallel \text{eval_task}$
Give p_1 as input to m_e and receive output o_1
for $i = 2, 3, \dots, T$ **do**
 Set $p_i \leftarrow x_i \parallel \text{eval_task}$
 Give p_i as input to m_e and receive output o_i
end for
Extract s_1, \dots, s_T from o_1, \dots, o_T
Set $s \leftarrow \frac{\sum_i s_i}{T}$
return s

represents the student’s proficiency in relevant skills, the instruction, the first problem on the post-test x_1 , and the task for the LM (which is to predict the probability that the student can solve the problem correctly). For example, in the figure we show a case where a student is learning to solve systems of equations. The student starts with some skill levels, and then the student is provided with a worked example, and then given a problem to work out a solution to a new problem, along with other activities (omitted for space). As shown in orange, after such “simulated” instruction, the student will be asked a test question (in orange)— here a word problem requiring the student to set up and solve two equations to work out the age of a sister. In green the LM m_e will be instructed to predict the performance on this question of the student who start from this initial background knowledge and completed the provided activities (worked example etc). More formally, the LM evaluator m_e takes in this p_1 (as shown in Figure 2) and outputs a response o_1 . We then feed the LM m_e the subsequent prompts p_2, \dots, p_T , where each prompt p_i includes an additional test problem x_i and the task description (see an example in Figure 3). Since we give all the prompts to the same LM instance m_e , m_e generates o_i in the context of $\{p_1, o_1, p_2, o_2, \dots, p_i\}$. Essentially this repeated interaction just allows us to have the LM m_e evaluate how the student persona would do on the post test, given the provided instruction.² After we collect o_1, \dots, o_T , we extract the probability of success (in percentage) for each test problem s_1, \dots, s_T and calculate the post-test score s as the average probability of success:

$$s = \frac{\sum_i s_i}{T} \tag{1}$$

We will later discuss the specific student personas (student background description) we use.

A simulated educational expert offers many benefits. For example, we can easily get a simulated expert assessment of the impact of different instructional content on the same student by altering the instructional content in the prompt. Additionally, by setting the instructional content to an empty string (see Figure 15 in Appendix), we ask the simulated expert to estimate a student’s pre-test score, allowing for the use of identical questions in both pre-tests and post-

²We also explored providing all the test questions and then asking the LM to predict the probability of the student getting each right, but we found that asking the LM to predict one by one was more effective.

Here is an 8th-grade student with the following skill levels (each skill is rated on a scale from 1 to 5):

1. Being able to set up systems of equations given a word problem: {level1}
2. Being able to solve systems of equations: {level2}

Here's the instruction that the student receives. The student is asked to study a problem and its solution. Here's the problem:

A brownie recipe is asking for 350 grams of sugar, and a pound cake recipe requires 270 more grams of sugar than a brownie recipe. How many grams of sugar are needed for the pound cake?

Here's its solution:

Step 1: Identify the amount of sugar needed for the brownie recipe, which is 350 grams.

Step 2: Understand that the pound cake recipe requires 270 more grams of sugar than the brownie recipe.

Step 3: Add the additional 270 grams of sugar to the 350 grams required for the brownie recipe.

Step 4: The total amount of sugar needed for the pound cake recipe is 350 grams + 270 grams = 620 grams.

The student is then asked to work on the following problem on their own:

The size of a compressed file is 1.74 MiB, while the size of the original uncompressed file is 5.5 times greater. What is the size of the uncompressed file, in MiB?

(... more instructions ...)

Now the student is asked to work on the following problem on a test:

Alyssa is twelve years older than her sister, Bethany. The sum of their ages is forty-four. Find Alyssa's age.

Given the student's initial skill levels and the instruction the student has received, what's the probability that the student can solve the problem correctly? Explain your reasoning and give a single number between 0 and 100 in square brackets.

Figure 2: The prompt for predicting the post-test score in a simulated expert evaluation. The blue text is the student persona. The black text is the instruction given to the student. The orange text is a problem on the post-test. The green text describes the evaluation task for the LM.

Now the student is asked to work on the following problem on a test:

{problem}

Given the student's initial skill levels and the instruction the student has received, what's the probability that the student can solve the problem correctly? Explain your reasoning and give a single number between 0 and 100 in square brackets.

Figure 3: The orange text is the test question. The green text describes the evaluation task for the LM.

tests. This approach allows us to estimate learning gains while avoiding biases associated with memorization.

Importantly, using an LM as a simulated educational expert is different from using an LM to directly simulate a given student. Simulating a student involves modeling the student's problem-solving process, while an SEE only predicts the post-test scores without actually simulating the student's responses. In our early explorations, we found that the LMs we were using were poor simulators of student learning processes (see Appendix A). For example, while an LM might successfully offer an answer as to what a fifth grader might respond to an algebra question, if we then

provided some pedagogical material and asked the LM to directly simulate what the student would now respond to a test question, the LM frequently could "forget" its (the fifth graders') limited knowledge and answer the question perfectly. In other words, the LMs (at the time) had a poor model of the *dynamics* of student learning. To emphasize the distinction using our example, a simulated educational expert LM will predict the chance the student will correctly answer the sister age question (e.g., 70% chance the student will get it right) whereas an LM-as-student must generate the actual answer (e.g., "Alyssa is 32").

We now evaluate to what extent a given LM, such as GPT-3.5, can act as a simulated educational expert and predict the impact of instructions on student learning outcomes. In particular, such a simulated educational expert should align with prior known results about student learning from different types of instructional materials. Our primary interest is to consider if we can use LMs as simulated educational experts to evaluate and optimize new instructional content. Therefore, our focus here is not a comprehensive evaluation of the ability of LMs to match known results in the learning sciences, but to do a small set of basic tests to make sure there are reasons to believe that an LM might be a good simulated educational expert.

In particular, we run simulated expert evaluations (SEEs) over a population of student personas to see if the judg-

ments of our simulated educational expert LMs can replicate two well-known instructional effects: the Expertise Reversal Effect [14, 43, 15, 13] and the Variability Effect [29, 42]. We selected these two phenomena in part because they are compatible with text-only prompts, though in the future we hope to extend LM-based simulated educational experts to multi-modal instructional inputs.

3.2 Replicating prior studies using SEEs

Replicating previous experiments involves: 1) creating participant personas that include relevant covariates; 2) allocating participants to different experimental conditions; 3) creating instructional materials customized for each experimental group and implementing the intervention; and 4) conducting post-tests to evaluate the impact of the intervention. Optionally, conducting pre-tests to determine baseline measures could also be included.

Since prior studies did not publish their experimental materials, we create new student personas and instructional materials, focusing on the domain of math word problem solving. We create $n = 120$ student personas using a fixed template (see the blue text in Figure 2) that describes their initial proficiency in the relevant math skills: the ability to set up systems of equations and the ability to solve systems of equations. The student persona vary only in skill levels, *level1* and *level2*, for which we randomly assigned integers between 1 and 5 (inclusive) to simulate diverse student abilities. We define the experimental groups and assign student personas to these groups using the same procedures as prior studies. We simulate the pre-test stage by running one SEE (see Alg. 1) for each student persona. The instruction for all these SEEs is an empty string because the students have not received any instruction yet (see Figure 15 in Appendix). To implement the intervention and collect students’ post-test scores, we run multiple SEEs (one for each student persona) with specific instructions tailored to their assigned experimental group. After collecting post-test scores, we compare the overall effect of various instructions on different experimental groups with results from prior studies. We describe the two phenomena we replicate below.

3.3 Expertise Reversal Effect

The Expertise Reversal Effect, introduced by [14], suggests that the effectiveness of instructional strategies changes as a learner’s knowledge and skills develop. In the early stages of learning, beginners often benefit from structured guidance, which compensates for their limited background knowledge and helps manage cognitive load. However, as learners gain proficiency, the same strategies that were once helpful can become superfluous or even hinder learning by overloading the cognitive system. For learners who have achieved a higher level of expertise, minimal guidance is preferable because they have built up a sufficient framework of knowledge that allows for efficient organization and processing of new information [13].

3.3.1 Prior real-life experiments

In Experiment 3 of [15], they replicate the Expertise Reversal Effect in the domain of coordinate geometry problem solving. Their experiment consists of three stages: pre-test, instruction, and post-test.

During the first stage, they administer a pre-test to all participants (42 Year 9 students from a Sydney Catholic girls’ school). Based on the pre-test scores, they divide participants into two groups: more knowledgeable learners (upper median group) and less knowledgeable learners (lower median group). They then randomly allocate students in each of these two groups to two subgroups: one receives practice-based instruction, and the other receives worked-example-based instruction. This leads to four experimental groups: 1) Low-knowledge/practice, 2) Low-knowledge/worked example, 3) High-knowledge/practice, and 4) High-knowledge/worked example.

In the second stage, participants in the practice conditions are given 8 problems (numbered from 1 to 8) to solve on their own. Participants in the worked-example conditions are given the same set of 8 problems in the same order but with fully worked-out step-by-step solutions for the problems with odd numbers.

In the final stage, they ask participants to take a post-test. They find that for less knowledgeable (low-expertise) learners, the worked-example group performs significantly better than the practice group on the post-test. For more knowledgeable learners (high-expertise), there is no significant difference between the worked-example group and the practice group. Similarly, [43] uncovers the Expertise Reversal Effect in their experiment teaching students how to use a database program (see their Figure 1).

3.3.2 Replication

Following the procedures used in [15], we run SEEs in a different domain: math word problems involving systems of equations. We create 120 student personas that describe their initial proficiency in relevant math skills (see Section 3.2). Based on the sum of skill levels for the two skills, we divide these student personas into two groups: high-expertise learners (upper-median group) and low-expertise learners (lower-median group)³. For each group, we randomly assign half of the students to the practice condition and the other half to the worked-example condition, which leads to four experimental groups:

1. low-expertise/practice.
2. low-expertise/worked example.
3. high-expertise/practice.
4. high-expertise/worked example.

We create the two types of instructions (practice and worked example), the pre-test, and the post-test, using math word problems from the ALGEBRA dataset [9]. We randomly select 8 problems for the two instructions and a different set of 8 problems $\{x_1, \dots, x_8\}$ for the pre-test and post-test. The practice instruction contains 8 word problems with no solutions, while the worked-example instruction includes the same set of 8 problems with step-by-step solutions for the first, third, fifth, and seventh problems. Detailed instructions are available at <https://github.com/StanfordAI4HI/ed-expert-simulator>.

³In our SEEs, we do not need to divide students into low-expertise learners and high-expertise learners based on their pre-test scores since we know all students’ latent skill levels.

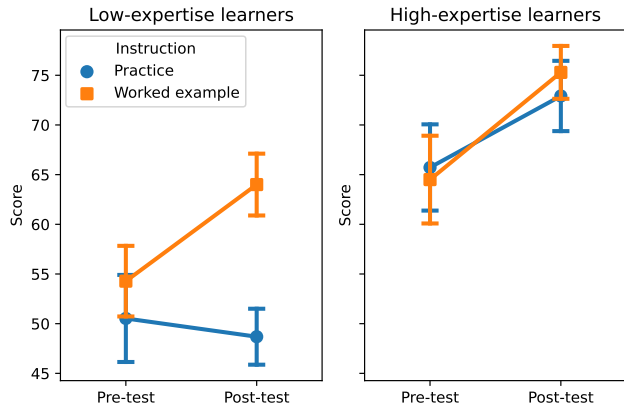


Figure 4: The LM-based simulated expert evaluations can replicate the Expertise Reversal Effect. According to the LM judgments, low-expertise learners benefit more from worked examples than practice, but there is no significant difference between practice and worked examples for high-expertise learners. Circles and squares indicate the average test score for each group. The error bars show the standard error.

We simulate the pre-tests and post-tests as described in Section 3.2. Figure 15 in the Appendix shows a prompt for the pre-test. Figure 2 shows a prompt for predicting the post-test scores for the worked-example group. We use gpt-3.5-turbo-16k as the evaluator LM for all SEEs, with the temperature parameter set to 0 for more consistent output.

3.3.3 Results

We find that LM-based SEEs can accurately replicate the Expertise Reversal Effect, as shown in our comparison of pre-test and post-test scores across all participants (see Figure 4). Among low-expertise learners, the worked-example group significantly outperforms the practice group on the post-test. This performance difference is not observed among high-expertise learners, where both practice and worked-example groups show similar outcomes.

This finding might seem expected, considering the LM has likely been trained on data that includes descriptions of these prior studies. However, the ability of the LM to make these accurate judgments is particularly noteworthy for two reasons. First, our evaluation methods do not explicitly hint at those prior studies (for instance, we do not specify whether the instruction is practice-based or involves worked examples). Second, we develop and use new instructional materials in a domain different from those explored in prior studies. This highlights the LM’s ability to generalize its learning to novel contexts.

3.4 Variability Effect

The Variability Effect [29, 42] suggests that introducing a variety of instructional examples is not always beneficial for students. Variability in problem situations can enhance student learning only if students have sufficient working memory resources to handle the added cognitive load.

3.4.1 Prior real-life experiment

The experiment in [29] demonstrates the Variability Effect in the domain of geometrical problem-solving. The researchers divide participants (60 students from fourth-year classes of a secondary technical school) into four experimental groups:

1. low-variability/practice.
2. low-variability/worked example.
3. high-variability/practice.
4. high-variability/worked example.

Participants in the practice conditions receive a set of six problems. Participants in the worked-example conditions receive the same problems with their step-by-step solutions. The problems with odd numbers (first, third, and fifth) are identical for the low-variability and high-variability conditions. The problems with even numbers in the low-variability conditions have the same problem formats as the problems with odd numbers, but with different values. The problems with even numbers in the high-variability conditions have different values and problem formats than the problems with odd numbers. After receiving the intervention, participants take a post-test consisting of six problems.

The results show that introducing a variety of examples can significantly enhance learning and the ability to apply knowledge, especially when the cognitive load is low, such as when students learn from worked examples. However, in situations where the cognitive load is already high, like when students solve practice problems independently, there is no effect of variability because too much variability can overload a learner’s memory capacity. This finding highlights the balance needed between introducing diversity in learning materials and managing cognitive load for optimal learning outcomes.

3.4.2 Replication

We run SEEs in the domain of math word problems involving systems of equations, which is the same as Section 3.3.2. We create $n = 120$ student personas that describe their initial proficiency in relevant math skills (see Section 3.2). Following [29], we randomly divide these student personas into four experimental groups of equal size:

1. low-variability/practice.
2. low-variability/worked example.
3. high-variability/practice.
4. high-variability/worked example.

We create four types of instructions (one for each experimental group) and a post-test, using math word problems from the ALGEBRA dataset [9]. We randomly select a set of 6 problems $\{x_1, \dots, x_6\}$ for the post-test, which is identical in all conditions. In the instructions for both low-variability and high-variability conditions, the first, third, and fifth problems are identical. In the low-variability conditions, the second, fourth, and sixth problems have the same format as the odd-numbered problems but use different values. In the high-variability conditions, the even-numbered problems differ from the odd-numbered ones in both values and formats. In practice conditions, only problems are presented. In worked-example conditions, all problems and step-by-step solutions are presented. We use gpt-3.5-turbo-16k as the evaluator LM for all SEEs, with the temperature parameter set to 0 for more consistent output.

3.4.3 Results

We find that these LM-based SEEs successfully replicate the Variability Effect. We plot the post-test scores for all participants in Figure 5. High variability in problems significantly enhances performance in worked-example conditions, whereas it has no effect in practice conditions. Additionally, our data from these SEEs also replicates the Expertise Reversal Effect (see Figure 6), suggesting that the LM judgments are consistent. As discussed in Section 3.3.3, despite the LM potentially being trained on datasets that include descriptions of previous studies, our evaluations use new instructional materials from a different domain than prior studies and do not reference these studies. This demonstrates the LM’s capability to act as a reliable evaluator.




Figure 5: The LM-based simulated expert evaluations can replicate the Variability Effect. According to the LM judgments, high variability in problems significantly enhances performance in worked-example conditions, but there is no effect of variability in practice conditions. Circles and squares indicate the average test score for each group. The error bars show the standard error.




Figure 6: The simulated expert evaluations designed for replicating the Variability Effect can also replicate the Expertise Reversal Effect, demonstrating the consistency of LM judgments. Among the low-expertise learners, the worked-example group performs significantly better than the practice group on the post-test. Among the high-expertise learners, there is no significant performance difference between the practice and worked-example groups. Circles and squares indicate the average test score for each group. The error bars show the standard error.

4. INSTRUCTION OPTIMIZATION

We showed that LMs have relevant educational knowledge when used via SEEs, which suggests SEEs can be used to craft better instructional materials.

4.1 The Instruction Optimization algorithm

We introduce the Instruction Optimization algorithm (see Alg. 2), which iteratively refines instructional materials by using an evaluator LM (M_e) and an optimizer LM (M_o). In this algorithm, M_o generates new instructions based on prior instructions and resulting learning outcomes (e.g., post-test scores) of a particular student, and M_e evaluates these new instructions by predicting the student’s learning outcomes.

More precisely, the input of the algorithm are M_e , M_o , an initial text instruction c_0 , a text description of a student persona, an evaluation task description, an optimization task description, and T test questions x_1, \dots, x_T . We use an SEE (Alg. 1) to assess the quality of c_0 and receive an initial post-test score r_0 . We store the instruction-score pair (c_0, r_0) in *mem*. Following the prompt design strategy in [49], we construct the initial optimization prompt p_0 (see an example in Figure 7) that includes the student persona, the instruction-score pair(s) in *mem*, and the optimization task description for the LM (which is to generate a new instruction to further increase the post-test score of the student). In each optimization step n , we give p_{n-1} as input to M_o , which generates a new instruction. We repeat this K times and generate K new instructions for each optimization step. We evaluate each of the K instructions using Alg. 1 and store the resulting instruction-score pairs in a temporary memory *mem'*. We then add instruction-score pairs in *mem'* to *mem*. We sort instruction-score pairs in *mem* by their scores in ascending order and remove the instructions with the lowest scores from *mem* until *mem* meets the maximum length requirement (due to the LM’s prompt length limit). We update the prompt using the updated *mem* and use the new prompt for the next iteration.

Algorithm 2 Instruction Optimization

Input: Evaluator M_e , Optimizer M_o ,
Initial instruction c_0 , *student_persona*, *eval_task*,
optimization_task, test questions x_1, \dots, x_T
Evaluate c_0 using Algorithm 1 and receive score r_0
Set $\mathcal{D} \leftarrow \{(c_0, r_0)\}$
Set *mem* $\leftarrow \{(c_0, r_0)\}$
Set $p_0 \leftarrow \text{student_persona} \parallel \text{mem} \parallel \text{optimization_task}$
for $n = 1, 2, \dots, N$ **do**
 Set *mem'* $\leftarrow []$
 for $k = 1, 2, \dots, K$ **do**
 Generate $c_{(n-1)K+k}$ by feeding p_{n-1} into M_o
 Evaluate $c_{(n-1)K+k}$ using Algorithm 1 and receive
 score $r_{(n-1)K+k}$
 Append $(c_{(n-1)K+k}, r_{(n-1)K+k})$ to *mem'*
 Set $\mathcal{D} \leftarrow \mathcal{D} \cup \{(c_{(n-1)K+k}, r_{(n-1)K+k})\}$
 end for
 Set *mem* $\leftarrow \text{mem} \cup \text{mem}'$
 Sort *mem* by scores in ascending order
 Set *mem* $\leftarrow \text{mem}[: \text{max length}]$
 Set $p_n \leftarrow \text{student_persona} \parallel \text{mem} \parallel \text{optimization_task}$
end for
return \mathcal{D}

Here is an 8th-grade student with the following skill levels (each skill is rated on a scale from 1 to 5, where higher numbers indicate more proficiency):

1. Being able to set up systems of equations given a word problem: 1
2. Being able to solve systems of equations: 1

I have some worksheets along with the student’s test scores after receiving the worksheets. The worksheets are arranged in ascending order based on their scores, where higher scores indicate better quality.

Worksheet:

You need to study a problem and its solution. Here’s the problem: A brownie recipe is asking for 350 grams of sugar, and a pound cake recipe requires 270 more grams of sugar than a brownie recipe. How many grams of sugar are needed for the pound cake? Here’s its solution: Step 1: Identify the amount of sugar needed for the brownie recipe, which is 350 grams. Step 2: Understand that the pound cake recipe requires 270 more grams of sugar than the brownie recipe. Step 3: Add the additional 270 grams of sugar to the 350 grams required for the brownie recipe. Step 4: The total amount of sugar needed for the pound cake recipe is 350 grams + 270 grams = 620 grams.

Test score: 20

Worksheet:

{worksheet2}

Test score: 34

Worksheet:

{worksheet3}

Test score: 66

(...more worksheets and scores...)

Generate a new worksheet to further increase the test score of the student. You will be evaluated based on this score function:

```
“python
{utility_string}
”
```

The new worksheet should begin with <WORKSHEET> and end with </WORKSHEET>.

Figure 7: The optimization prompt for generating a new instruction (e.g., a worksheet) based on prior instructions and post-test scores for a given student. The blue text is the student persona. The orange text contains the previous worksheet-score pairs, sorted in ascending order. The green text describes the optimization task for the LM. The score function is a Python program that illustrates how post-test scores are computed without revealing the content of the post-test.

4.2 Optimizing math worksheets

We demonstrate the effectiveness of our Instruction Optimization algorithm (Alg. 2) on optimizing math word problem worksheets. Starting with an initial worksheet that yields a low post-test score (see the first worksheet-score pair in Figure 7), we demonstrate that the optimizer LM (GPT-4) can improve the post-test score of the generated worksheets until convergence (see Figure 8). Figure 9 shows two examples of LM-generated worksheets.

We describe the implementation details below. We use gpt-3.5-turbo-16k as M_e and gpt-4 as M_o . We use temperature = 0 for gpt-3.5-turbo-16k, and we use temperature = 1 for gpt-4 to encourage more diverse generations of new worksheets. We generate $K = 3$ worksheets per optimization step. We run a total of $N = 19$ steps. When we evaluate each worksheet using Alg. 1, we run three independent evaluations and take the average post-test score across the three evaluations as the final post-test score for each worksheet, which allows M_o to get more stable reward signals. We randomly select $T = 6$ problems from the ALGEBRA




Figure 8: LMs can iteratively improve math word problem worksheets using their own judgments as the reward function. Each dot represents the average post-test score of 3 worksheets generated at each optimization step, with error bars indicating the standard deviation.

Post-test score: 34

Title: Mastering Systems of Equations

Let's delve into more elaborate examples to enhance your skills in setting up and solving systems of equations.

Problem 1:

A movie theater sells two types of tickets: adults and children. The price of an adult ticket is \$15 and the price of a child ticket is \$8. If the total income from the sale of tickets on a certain day was \$1200 and a total of 100 tickets were sold, how many adult and children tickets were sold that day?

Remember the steps:

- Identify your variables.
- Write the system of equations based from the problem.
- Solve the system of equations.

Problem 2:

In a class, there are twice as many girls as boys. If the total number of students in the class is 45, how many boys and girls are there in the class? Set up and solve a system of equations to answer the question.

Problem 3:

A birdwatcher saw a group of birds in a tree. She counted 16 bird heads and 40 bird legs. If there were only parrots and crows on the tree, with parrots having 2 legs and crows having 4, how many parrots and crows did the birdwatcher see?

Problem 4:

Tim and Jim are saving money. Tim saves \$5 per day while Jim saves \$7 per day. After some days, Tim noticed that the total amount of money saved by them is \$240. If they saved the same number of days, how many days did it take them to save that amount and how much did each of them save?

Remember the method:

- Step 1: Understand what's being asked in the problem.
- Step 2: Identify the unknowns and assign variables to them.
- Step 3: Write down a system of equations that represent the problem.
- Step 4: Solve the system of equations to find the solution.

Practicing these problems will strengthen your skills in both setting up and solving systems of equations. Don't get discouraged if you have difficulties at first. Keep practicing, and you'll improve!

Post-test score: 87

Title: Equation Excellence - Mastering Systems of Equations

Objective: Your journey to master the art of translating word problems into systems of equations and solve them confidently is about to level up. Dive into these hand-picked problems to hone your skills.

Problem 1:

In a zoo, there are giraffes and zebras. The number of giraffes is two more than three times the number of zebras. If there are 29 animals in total, how many giraffes and zebras are there in the zoo?

Hints for Success:

1. Assign 'G' to represent Giraffes and 'Z' for Zebras.
2. Derive the equations from the problem: ' $G + Z = 29$ ' and ' $G = 3Z + 2$ '.
3. Solve these equations to find the number of giraffes and zebras.

Problem 2:

A magic show sells two types of tickets - regular seats for \$2 and VIP seats for \$5. If the total sales from 30 tickets were \$80, can you figure out how many regular and VIP seats were sold?

Hints for Success:

1. Let's denote the number of regular seats as 'R' and VIP seats as 'V'.
2. From the conditions given in the problem, form these equations: ' $R + V = 30$ ' and ' $2R + 5V = 80$ '.
3. Solve these equations to find the number of regular and VIP seats sold.

Problem 3:

In a garden, the number of daffodils is 6 less than twice the number of tulips. Altogether, there are 26 flowers in the garden. How many daffodils and tulips are in the garden?

Hints for Success:

1. Let 'D' represent Daffodils and 'T' represent Tulips.
2. According to the problem, we can formulate these equations: ' $D + T = 26$ ' and ' $D = 2T - 6$ '.
3. Solve these equations to uncover the number of daffodils and tulips.

Remember that every problem is a stepping stone to do better. Embrace challenges and use the provided hints to reinforce the process of setting up and solving systems of equations. Enjoy your journey to becoming an expert!

Figure 9: Two examples of worksheets generated by the optimizer LM. As predicted by the evaluator LM, the left one leads to a post-test score of 34, and the right one gives a post-test score of 87.

dataset [9] as the post-test. The post-test is identical in all evaluations. The max length limit for *mem* is set to be 8 worksheet-score pairs. We use a fixed student persona, where the student's initial skill levels are low (see the blue text in Figure 7), to simulate a personalized instructional design process. In the optimization task description, we use a Python program to show how post-test scores are computed without showing specific problems on the post-test (see Figure 16 in Appendix).

4.3 Human evaluation of worksheets

We conducted a human teacher evaluation on worksheets generated by the LM, involving 95 participants from Prolific. All participants met the following criteria:

1. The participant has a background in teaching.
2. The participant is based in the U.S.
3. The participant is fluent in English.
4. The participant has an approval rate over 95% on Prolific.

All participants provided informed consent before participation following an approved institutional review board (IRB) protocol⁴. We paid participants at a rate of 12 USD per hour.

⁴The risks associated with this study are minimal. Participants were told that the study data would be stored se-

We asked participants to do pairwise comparisons where they need to indicate their preferences between two worksheets (worksheet A and B) using a continuous slider scale, with a rating of -1 assigned for a strong preference towards the first worksheet, a rating of 1 for a strong preference towards the second worksheet, and a rating of 0 for no preference (see the task description in Figure 10).

We evaluated a total of 10 worksheets, of which one was the initial worksheet we provided to the optimizer LM, and the rest were randomly selected from LM-generated worksheets. There are 45 unique pairs of worksheets, and we randomly assigned 4 or 5 pairs of worksheets to each participant. We also assigned an additional catch-trial question after the first three pairs of worksheets to filter participants who did not pay attention. We collected a total of 465 ratings from $n = 95$ participants, and each pair of worksheets has at least 7 ratings. After filtering out participants who failed the catch trial, we have a total of 270 ratings from 56 participants. We calculate the human preference score for each worksheet c , $hps(c)$, by accumulating the negative ratings when a worksheet was less preferred and the positive ratings when it was more preferred, normalized by the total number of times each worksheet was compared:

curely, minimizing the risk of confidentiality breach. Their individual privacy is maintained during the research and in all published and written data resulting from the study.

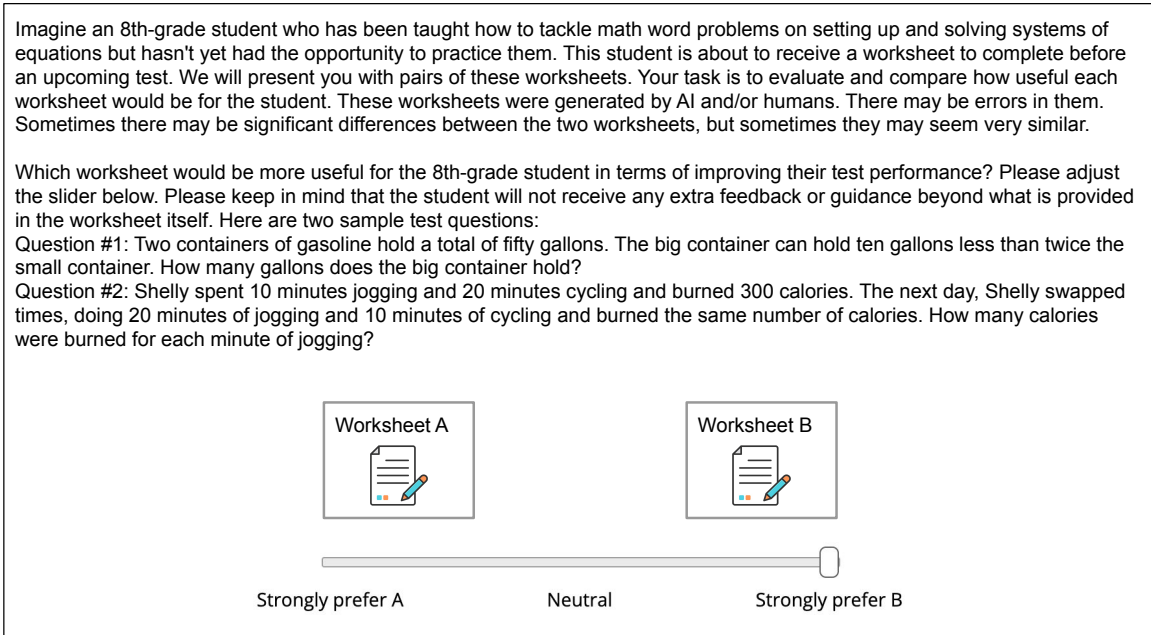


Figure 10: The evaluation task description for teachers. In each evaluation, we ask the teacher to compare two worksheets in terms of their effectiveness in improving the student’s test performance, and we show the teacher two sample test questions. Teachers can indicate their preference by adjusting the slider.

$$hps(c) = \frac{\sum_{(a,b) \in \mathcal{S}_A^c} -R(a,b) + \sum_{(a,b) \in \mathcal{S}_B^c} R(a,b)}{|\mathcal{S}_A^c| + |\mathcal{S}_B^c|}, \quad (2)$$

where \mathcal{S}_A^c denotes all the pairs of worksheets where c appears as the first worksheet (which means we need to negate the rating to indicate how much c is preferred over the other worksheet), and \mathcal{S}_B^c denotes all the pairs of worksheets where c appears as the second worksheet. $|\mathcal{S}_A^c|$ and $|\mathcal{S}_B^c|$ denote the size of the each set.

We calculate the Pearson correlation between post-test scores predicted by the LM and the human preference scores for all worksheets and find a significant correlation ($r = 0.661, p < 0.05$) between the LM judgments and human teacher preferences (see Figure 11). However, human teachers sometimes cannot distinguish worksheets that LMs identify as distinct, particularly those with predicted post-test scores ranging from 60 to 90.

5. LIMITATIONS

There are several limitations to our work. First, the LMs we used were likely trained on data containing descriptions of experiments similar to those we replicated. Therefore, we developed instructional materials in a domain different from previous studies.

Second, one potential concern is that LMs do not consistently replicate established educational results. Our experiments focused solely on mathematics and only tested the performance of GPT-3.5 and GPT-4 with a limited number of straightforward prompts. Exploring all available LMs would have been prohibitively costly, so we chose GPT-3.5 and GPT-4 due to their popularity and robust performance.

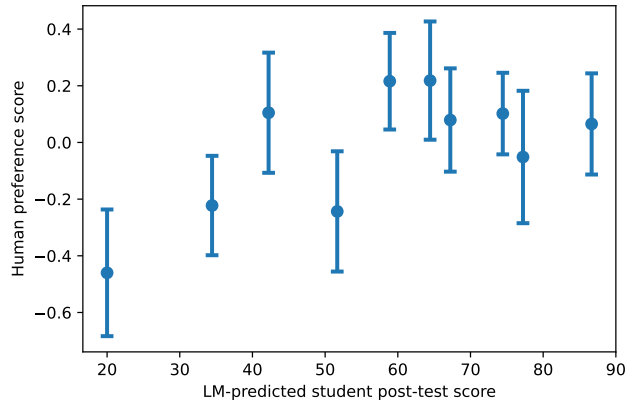


Figure 11: Evaluations by human teachers ($n = 56$) show a significant correlation between the LM judgments and human teacher preferences (Pearson $r = 0.661, p < 0.05$). The error bars are bootstrapped 95% confidence intervals.

We used GPT-3.5 as the evaluator LM because it is more cost-effective for handling long prompts. GPT-4 served as our optimizer LM, chosen for its capability in critical assessment, a feature typically reliable only in the most advanced models. Our goal was to demonstrate the potential of LMs in these tasks, not to claim these LM choices as the definitive best.

Third, there are discrepancies between the evaluations made by LMs and those by human teachers, who sometimes fail to perceive differences that LMs identify. These discrepancies

raise concerns about the models' decision-making processes and the potential impact of biases in their training data. Furthermore, our evaluations are limited to the perspective of teachers. However, ultimately, the effectiveness of the educational materials depends on how students engage with them. While LMs could aid in accelerating the design process by providing preliminary evaluations, they should not be seen as replacements for direct evaluations involving human students.

6. CONCLUSION

We show that LMs can act as evaluators on a couple of well-known findings for pedagogical effectiveness, including results that require understanding that different students can be impacted differently by different content. Even if the LM succeeds at these judgments due to reading relevant studies, it is still remarkable for two reasons: 1) Nothing in the evaluations directly cued those studies, and 2) We developed new instructional materials in a domain different from the ones explored in prior studies. This shows that the LM can read and synthesize educational science reliably into useful evaluators. In addition, LM evaluators correlated positively with expert humans in their evaluations of the effectiveness of new content.

However, there may be other known educational findings that LM evaluators do not replicate, and human teachers sometimes cannot distinguish instructional materials that LMs perceive as different and may disagree with LMs on what is optimal. Therefore, the LM may be useful as a coarse evaluator to help speed up content design as a human augmentor but not to replace human evaluators or empirical studies. An interesting open question is whether we can extend LM-based simulated educational experts to multi-modal instructional input.

7. ACKNOWLEDGEMENTS

We would like to thank Dongwei Jiang, Rose E. Wang, Allen Nie, and Shubhra Mishra for their feedback on our work. This work was supported by the Jungle Corporation Stanford Graduate Fellowship and a Stanford Hoffman-Yee grant.

8. REFERENCES

- [1] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [2] G. V. Aher, R. I. Arriaga, and A. T. Kalai. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, pages 337–371. PMLR, 2023.
- [3] J. R. Anderson, A. T. Corbett, K. R. Koedinger, and R. Pelletier. Cognitive tutors: Lessons learned. *The journal of the learning sciences*, 4(2):167–207, 1995.
- [4] L. P. Argyle, E. C. Busby, N. Fulda, J. R. Gubler, C. Rytting, and D. Wingate. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3):337–351, 2023.
- [5] Y. Bai, S. Kadavath, S. Kundu, A. Askell, J. Kernion, A. Jones, A. Chen, A. Goldie, A. Mirhoseini, C. McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- [6] S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- [7] P. Denny, H. Khosravi, A. Hellas, J. Leinonen, and S. Sarsa. Can we trust ai-generated educational content? comparative analysis of human and ai-generated learning resources. *arXiv preprint arXiv:2306.10509*, 2023.
- [8] J. E. Greer and G. I. McCalla. *Student modelling: the key to individualized knowledge-based instruction*, volume 125. Springer Science & Business Media, 2013.
- [9] J. He-Yueya, G. Poesia, R. E. Wang, and N. D. Goodman. Solving math word problems by combining language models with symbolic solvers. *arXiv preprint arXiv:2304.09102*, 2023.
- [10] J. J. Horton. Large language models as simulated economic agents: What can we learn from homo silicus? Technical report, National Bureau of Economic Research, 2023.
- [11] H. Jin, S. Lee, H. Shin, and J. Kim. "teach ai how to code": Using large language models as teachable agents for programming education. *arXiv preprint arXiv:2309.14534*, 2023.
- [12] S. Jinxin, Z. Jiabao, W. Yilei, W. Xingjiao, L. Jiawen, and H. Liang. Cgmi: Configurable general multi-agent interaction framework. *arXiv preprint arXiv:2308.12503*, 2023.
- [13] S. Kalyuga. Expertise reversal effect and its implications for learner-tailored instruction. *Educational psychology review*, 19:509–539, 2007.
- [14] S. Kalyuga, P. Ayres, P. Chandler, and J. Sweller. The expertise reversal effect. *Educational Psychologist*, 38(1):23–31, 2003.
- [15] S. Kalyuga and J. Sweller. Measuring knowledge to optimize cognitive load factors during instruction. *Journal of educational psychology*, 96(3):558, 2004.
- [16] P. K. Komoski. Eric/avcr annual review paper: An imbalance of product quantity and instructional quality: The imperative of empiricism. *AV Communication Review*, 22(4):357–386, 1974.
- [17] H. Kumar, D. M. Rothschild, D. G. Goldstein, and J. M. Hofman. Math education with large language models: Peril or promise? Available at SSRN 4641653, 2023.
- [18] R. Liu, H. Yen, R. Marjeh, T. L. Griffiths, and R. Krishna. Improving interpersonal communication by simulating audiences with language models. *arXiv preprint arXiv:2311.00687*, 2023.
- [19] Q. Ma, H. Shen, K. Koedinger, and T. Wu. Hypocompass: Large-language-model-based tutor for hypothesis construction in debugging for novices. *arXiv preprint arXiv:2310.05292*, 2023.
- [20] A. Madaan, N. Tandon, P. Gupta, S. Hallinan, L. Gao, S. Wiegrefe, U. Alon, N. Dziri, S. Prabhunoye, Y. Yang, et al. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*, 2023.

- [21] J. M. Markel, S. G. Opferman, J. A. Landay, and C. Piech. Gpteach: Interactive ta training with gpt-based students. In *Proceedings of the tenth acm conference on learning@ scale*, pages 226–236, 2023.
- [22] N. Matsuda, W. W. Cohen, J. Sewall, G. Lacerda, and K. R. Koedinger. Evaluating a simulated student using real students data for training and testing. In *International Conference on User Modeling*, pages 107–116. Springer, 2007.
- [23] N. Matsuda, V. Keiser, R. Raizada, A. Tu, G. Stylianides, W. W. Cohen, and K. R. Koedinger. Learning by teaching simstudent: Technical accomplishments and an initial use with students. In *Intelligent Tutoring Systems: 10th International Conference, ITS 2010, Pittsburgh, PA, USA, June 14–18, 2010, Proceedings, Part I 10*, pages 317–326. Springer, 2010.
- [24] J. S. Mertz. Using a simulated student for instructional design. *International Journal of Artificial Intelligence in Education (IJAIED)*, 8:116–141, 1997.
- [25] M. H. Nguyen, S. Tschischek, and A. Singla. Large language models for in-context student modeling: Synthesizing student’s behavior in visual programming from one-shot observation. *arXiv preprint arXiv:2310.10690*, 2023.
- [26] N. Nieveen and E. Folmer. Formative evaluation in educational design research. *Design Research*, 153(1):152–169, 2013.
- [27] S. Ohlsson, A. M. Ernst, and E. Rees. The cognitive complexity of learning and doing arithmetic. *Journal for Research in Mathematics Education*, 23(5):441–467, 1992.
- [28] A. Opedal, A. Stolfo, H. Shirakami, Y. Jiao, R. Cotterell, B. Schölkopf, A. Saparov, and M. Sachan. Do language models exhibit the same cognitive biases in problem solving as human learners? *arXiv preprint arXiv:2401.18070*, 2024.
- [29] F. G. Paas and J. J. Van Merriënboer. Variability of worked examples and transfer of geometrical problem-solving skills: A cognitive-load approach. *Journal of educational psychology*, 86(1):122, 1994.
- [30] M. Pankiewicz and R. S. Baker. Large language models (gpt) for automating feedback on programming assignments. *arXiv preprint arXiv:2307.00150*, 2023.
- [31] Z. A. Pardos and S. Bhandari. Learning gain differences between chatgpt and human tutor generated algebra hints. *arXiv preprint arXiv:2302.06871*, 2023.
- [32] J. S. Park, J. O’Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–22, 2023.
- [33] J. S. Park, L. Popowski, C. Cai, M. R. Morris, P. Liang, and M. S. Bernstein. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pages 1–18, 2022.
- [34] T. Phung, V.-A. Pădurean, A. Singh, C. Brooks, J. Cambronero, S. Gulwani, A. Singla, and G. Soares. Automating human tutor-style programming feedback: Leveraging gpt-4 tutor model for hint generation and gpt-3.5 student model for hint validation. *arXiv preprint arXiv:2310.03780*, 2023.
- [35] E. Prihar, M. Lee, M. Hopman, A. T. Kalai, S. Vempala, A. Wang, G. Wickline, A. Murray, and N. Heffernan. Comparing different approaches to generating mathematics explanations using large language models. In *International Conference on Artificial Intelligence in Education*, pages 290–295. Springer, 2023.
- [36] R. Pryzant, D. Iter, J. Li, Y. T. Lee, C. Zhu, and M. Zeng. Automatic prompt optimization with” gradient descent” and beam search. *arXiv preprint arXiv:2305.03495*, 2023.
- [37] A. V. Riabko and T. A. Vakaliuk. Physics on autopilot: exploring the use of an ai assistant for independent problem-solving practice. *Educational Technology Quarterly*, 2024.
- [38] S. Sarsa, P. Denny, A. Hellas, and J. Leinonen. Automatic generation of programming exercises and code explanations using large language models. In *Proceedings of the 2022 ACM Conference on International Computing Education Research-Volume 1*, pages 27–43, 2022.
- [39] R. Schmucker, M. Xia, A. Azaria, and T. Mitchell. Ruffle&riley: Towards the automated induction of conversational tutoring systems. *arXiv preprint arXiv:2310.01420*, 2023.
- [40] O. Shaikh, V. Chai, M. J. Gelfand, D. Yang, and M. S. Bernstein. Rehearsal: Simulating conflict to teach conflict resolution. *arXiv preprint arXiv:2309.12309*, 2023.
- [41] P. L. Smith and T. J. Ragan. *Instructional design*. John Wiley & Sons, 2004.
- [42] J. Sweller, J. J. van Merriënboer, and F. Paas. Cognitive architecture and instructional design: 20 years later. *Educational psychology review*, 31:261–292, 2019.
- [43] J. E. Tuovinen and J. Sweller. A comparison of cognitive load associated with discovery learning and worked examples. *Journal of educational psychology*, 91(2):334, 1999.
- [44] K. VanLehn, S. Ohlsson, and R. Nason. Applications of simulated students: An exploration. *Journal of artificial intelligence in education*, 5:135–135, 1994.
- [45] R. E. Wang and D. Demszky. Is chatgpt a good teacher coach? measuring zero-shot performance for scoring and providing actionable insights on classroom instruction. *arXiv preprint arXiv:2306.03090*, 2023.
- [46] R. E. Wang, Q. Zhang, C. Robinson, S. Loeb, and D. Demszky. Step-by-step remediation of students’ mathematical mistakes. *arXiv preprint arXiv:2310.10648*, 2023.
- [47] Z. Wang, J. Valdez, D. Basu Mallick, and R. G. Baraniuk. Towards human-like educational question generation with large language models. In *International conference on artificial intelligence in education*, pages 153–166. Springer, 2022.
- [48] D. D. Williams, J. B. South, S. C. Yanchar, B. G. Wilson, and S. Allen. How do instructional designers evaluate? a qualitative study of evaluation in practice.

Educational Technology Research and Development,
59:885–907, 2011.

- [49] C. Yang, X. Wang, Y. Lu, H. Liu, Q. V. Le, D. Zhou, and X. Chen. Large language models as optimizers. *arXiv preprint arXiv:2309.03409*, 2023.
- [50] E. Zelikman, E. Lorch, L. Mackey, and A. T. Kalai. Self-taught optimizer (stop): Recursively self-improving code generation. *arXiv preprint arXiv:2310.02304*, 2023.
- [51] E. Zelikman, W. A. Ma, J. E. Tran, D. Yang, J. D. Yeatman, and N. Haber. Generating and evaluating tests for k-12 students with language model simulations: A case study on sentence reading efficiency. *arXiv preprint arXiv:2310.06837*, 2023.

APPENDIX

A. MODELING COGNITIVE PROCESSES OF LEARNING DYNAMICS

We simulate a scenario where LMs play the role of students learning from scratch, engaging with a Khan Academy video transcript to understand how to set up and solve systems of equations (see our prompt in Figure 12). To test the models' understanding, we give them a quiz with 14 math word problems related to the topic, evaluating their performance after they have read every few sentences of the transcript (see the first 30 sentences of the transcript in Figure 13).

We find that LMs would often transition from a lack of algebraic knowledge to proficiency after a minimal interaction (e.g., like reading just a few sentences from a lecture), which is not a realistic representation of how humans learn. We plot LMs' performance accuracy on the quiz over the number of sentences LMs have read in Figure 14. GPT-4, despite being asked to act as a novice, solves about 40% of the problems on the quiz after encountering the first 6 sentences of the video transcript.

You are a student who is initially ignorant. You are learning by watching a video. Here's the video:

In this video, we're gonna get some more practice setting up systems of equations.
So we're told Sanjay's dog weighs five times as much as his cat.
His dog is also 20 kilograms heavier than his cat.

Let's stop the video. Remember, you've only been taught what was shown in the video. It typically takes you 5-10 practices to learn a new skill, and you often need to rewatch a lecture and do practice problems. Remember you have done 0 examples, and that is less than 5! Can you apply what you've just learned to solve the following problem? If you do not know how to solve it, don't try to guess and ask to keep watching the video. If you do know how to solve it, then place the final answer (a number) in square brackets. Here's the problem:
Alyssa is twelve years older than her sister, Bethany. The sum of their ages is forty-four. Find Alyssa's age.

Figure 12: The black text is the video transcript. In this example, we show the LM the first three sentences of the transcript.

```
["In this video, we're gonna get some more practice setting up systems of equations.",
 "So we're told Sanjay's dog weighs five times as much as his cat.",
 'His dog is also 20 kilograms heavier than his cat.',
 "Let c be the cat's weight and let d be the dog's weight.",
 "So pause this video and see if you can set up a system of equation, two linear equations with two unknowns that w
e could use to solve for c and d. All right so let's do it together.",
 "So, what I like to do is usually there's a sentence or two that describes each of the equations we wanna set u
p.",
 "So this first one tells us Sanjay's dog weighs five times as much as his cat.",
 'So how much does his dog weigh?',
 'So his dog weighs d, so we know d is going to be equal to five times as much as his cat weighs.',
 'So his cat weighs c, so d is going to be equal to five times as much as his cat weighs.',
 "So we have  $d = 5c$ . So that's one linear equation using d and c. And so what's another one?",
 'Well, then we are told his dog is also 20 kilograms heavier than his cat.',
 "So we could say that the dog's weight is going to be equal to the cat's weight plus what?",
 'Plus 20 kilograms.',
 'So we have  $d = c + 20$ .',
 "We're assuming everything's in kilograms, so I don't have to write the units.",
 "But there you have it, I have just set up two equations in two unknowns, two linear equations, based on the infor
mation given in this word problem, which we could then solve, and I encourage you to do so if you're curious.",
 "But sometimes, the difficult part is just to find, is to re-express the information that you're given in a mathem
atical form. So that it's less likely that we get shown up by talking birds in the future, we've set a little bit of
exercise for solving systems of equations with substitution.",
 'And so this is the first exercise or the first problem that they give us.',
 '-3x-4y=-2 and y=2x-5 So let me get out my little scratch pad and let me rewrite the problem.',
 'So this is  $-3x-4y=-2$  and then they tell us  $y=2x-5$ .',
 "So, what's neat about this is that they've already solved the second equation.",
 "They've already made it explicitly solved for y which makes it very easy to substitute for.",
 "We can take this constraint, the constraint on y in terms of x and substitute it for y in this first equation and
then solve for x.",
 "So let's try it out.",
 "So this first equation would then become  $-3x-4$  but instead of putting a y there the second constraint tells us th
at y needs to be equal to  $2x-5$ .",
 "So it's  $4(2x-5)$  and all of that is going to be equal to  $-2$ .",
 "So now we get just one equation with one unknown.",
 'and now we just have to solve for x.',
 "So, let's see if we can do that."]
```

Figure 13: The first 30 sentences in the video transcript.

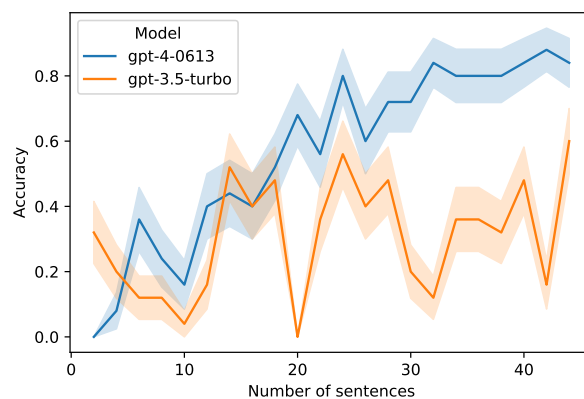


Figure 14: Simulating the dynamics of learning with language models (LMs) poses a challenge. Unlike human learners, LMs can shift from having no knowledge to being proficient after only a brief interaction, such as reading just a few sentences from a lecture.

B. PROMPTS

Prompts used in this work are publicly available at <https://github.com/StanfordAI4HI/ed-expert-simulator>.

Here is an 8th-grade student with the following skill levels (each skill is rated on a scale from 1 to 5):

1. Being able to set up systems of equations given a word problem: {level1}
2. Being able to solve systems of equations: {level2}

Now the student is asked to work on the following problem on a test:

{problem}

Given the student's initial skill levels, what's the probability that the student can solve the problem correctly? Explain your reasoning and give a single number between 0 and 100 in square brackets.

Figure 15: The prompt for predicting the pre-test score in a simulated expert evaluation. The blue text is the student persona. The orange text is the test. The green text describes the evaluation task for the LM.

```
def utility(worksheet: str):
    """
    Evaluates the worksheet in terms of test performance. Returns final test score of the student.
    """
    messages = []
    for j in range(len(posttest_problems)):
        test_problem = posttest_problems.iloc[j]['question']
        test_instruction = f'Now the student is asked to work on the following problem on a test: {test_problem}
Given the student's initial skill levels and the worksheet the student has received, what's the probability that th
e student can solve the problem correctly? Explain your reasoning and give a single number between 0 and 100 in squ
are brackets.'
        if j == 0:
            user_message = f'""{persona}
Here's a worksheet that the student receives:
{worksheet}

{test_instruction}
""
        else:
            user_message = test_instruction
            messages.append({"role": "user", "content": user_message})
            chat_completion = openai.ChatCompletion.create(model=model_id, messages=messages, temperature=temperature)
            output = chat_completion.choices[0].message.content
            messages.append({"role": "assistant", "content": output})
    scores = re.findall(r'\.[*?]\d+', messages_string)
    sum_score = 0
    for s in scores:
        sum_score += try_convert_to_float(s)
    avg_score = sum_score / len(scores)

    return avg_score
```

Figure 16: The utility string is a Python program that shows how the post-test score of a given worksheet is computed.