

Large Language Models for In-Context Student Modeling: Synthesizing Student’s Behavior in Visual Programming

Manh Hung Nguyen
MPI-SWS
manguyen@mpi-sws.org

Sebastian Tschiatschek
University of Vienna
sebastian.tschiatschek@univie.ac.at

Adish Singla
MPI-SWS
adishs@mpi-sws.org

ABSTRACT

Student modeling is central to many educational technologies as it enables predicting future learning outcomes and designing targeted instructional strategies. However, open-ended learning domains pose challenges for accurately modeling students due to the diverse behaviors and a large space of possible misconceptions. To approach these challenges, we explore the application of large language models (LLMs) for in-context student modeling in open-ended learning domains. More concretely, given a particular student’s attempt on a reference task as observation, the objective is to synthesize the student’s attempt on a target task. We introduce a novel framework, *LLM for Student Synthesis* (LLM-SS), that leverages an LLM for synthesizing a student’s behavior. Our framework can be combined with different LLMs; moreover, we fine-tune LLMs to boost their student modeling capabilities. We instantiate several methods based on LLM-SS framework and evaluate them using an existing benchmark, STUDENTSYN, for student attempt synthesis in a visual programming domain. Experimental results show that our methods perform significantly better than the baseline method NEURSS provided in the STUDENTSYN benchmark. Furthermore, our method using a fine-tuned version of the GPT-3.5 model is significantly better than using the base GPT-3.5 model and gets close to human tutors’ performance.

Keywords

Student Modeling, Generative AI, Large Language Models, Open-ended Learning Domains, Visual Programming

1. INTRODUCTION

Student modeling refers to the process of representing the current state of a learner’s knowledge, skills, preferences, and learning needs [1]. This is pivotal in developing educational systems as it allows for the personalization of learning experiences [2], catering specifically to each student’s unique abilities and growth areas, and targeted instructional strategies that can significantly enhance the learning process [3].

M. H. Nguyen, S. Tschiatschek, and A. Singla. Large language models for in-context student modeling: Synthesizing student’s behavior in visual programming. In B. Paaßen and C. D. Epp, editors, *Proceedings of the 17th International Conference on Educational Data Mining*, pages 341–348, 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.12729830>

By understanding student behavior, tutoring systems and educators can identify patterns and trends [4, 5], thereby predicting future learning outcomes [6] and providing timely support. Moreover, it allows them to detect if and when a student is losing interest or facing challenges [7], enabling them to intervene effectively [8]. In particular, student modeling is key in *open-ended learning domains* where creativity and exploratory behaviors are encouraged [9, 10]

In open-ended learning domains such as programming, students can take different learning paths and complete a task with different strategies [9]. This results in diverse behaviors and presents significant challenges to modeling a particular student’s behavior [10]. In recent years, some efforts in student modeling for open-ended learning domains have been made, such as representing knowledge and forecasting future performance using deep learning [6], investigating students’ problem-solving approaches using Natural Language Processing [12], early prediction of conceptual understanding [7], clustering-based methods for misconception discovery [4], students’ attempts synthesis in block-based visual programming [11], and predicting students’ post-test performance and interest using multimodal predictive student modeling [13]. Existing works on student modeling in open-ended learning domains often require a large behavioral dataset from students or use a complex pipeline, and sometimes, a combination of both [6, 11, 13, 14]. In this paper, we seek to leverage recent advances in generative AI and large language models (LLMs) for student modeling in open-ended learning domains and address the above-mentioned shortcomings.

In particular, LLMs have demonstrated advanced capabilities for *in-context learning* in which a model learns to solve a downstream application scenario when prompted with appropriate contextual information [15, 16]. Notably, they have been used to simulate humans for replicating human subject studies [17] and to simulate students for training teaching assistants [18]. In this work, we investigate the potential of leveraging such capabilities of LLMs for *in-context student modeling* in open-ended learning environments. In our setup, an LLM observes a student’s attempt on a reference task as the student’s behavioral context, and the objective is to synthesize the student’s attempt on a target task, reflecting the student’s problem-solving style and misconceptions observed. In essence, we seek to address the following research question: *Given a specific student’s behavioral context, are LLMs capable of effectively modeling the student and subsequently synthesizing the student’s attempt on a target task?*

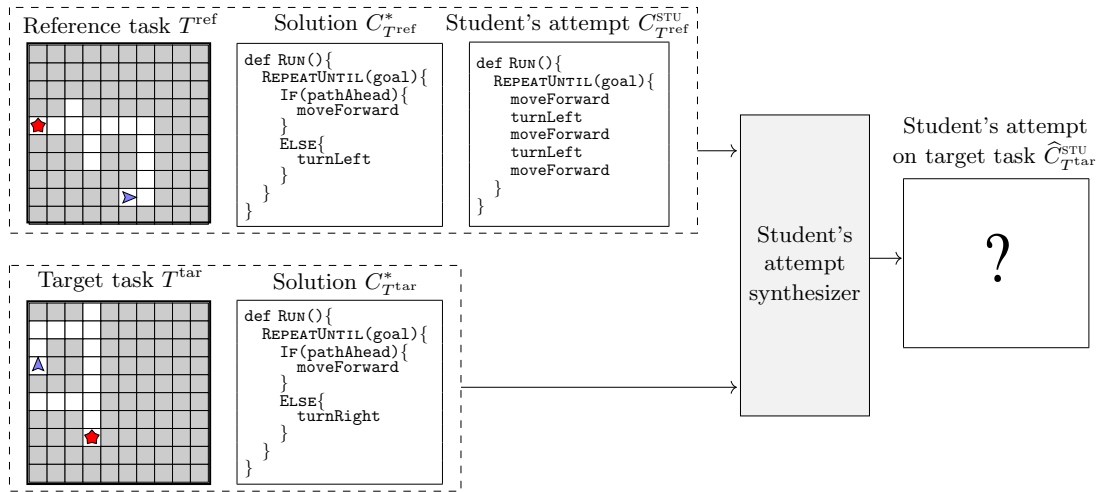


Figure 1: Illustration of our problem setup in a visual programming environment. The scenario is taken from the STUDENTSYN benchmark [11]. A synthesizer observes a tuple of $(T^{\text{ref}}, C_{T^{\text{ref}}}^*, C_{T^{\text{ref}}}^{\text{STU}})$ indicating a student STU’s behavior. Then, given a target task T^{tar} along with a solution $C_{T^{\text{tar}}}^*$, the synthesizer generates a student’s attempt $\hat{C}_{T^{\text{tar}}}^{\text{STU}}$ that imitates the student’s behavior.

To this end, we introduce a novel framework, *LLM for Student Synthesis* (LLM-SS), that leverages LLMs for modeling and synthesizing a student’s behavior. The design of our framework is inspired by *Perturbation Student Model* [19], based on the idea that a student’s knowledge can be modeled as perturbations to expert knowledge. Our framework operationalizes this idea by providing a student’s behavioral context in the prompt and improving the expert knowledge of a base LLM via fine-tuning. In summary, our main contributions are:

- I. We formalize the problem of using an LLM’s in-context learning capabilities for student modeling and behavior synthesis in open-ended learning domains.
- II. We propose a novel framework LLM-SS for synthesizing student’s behavior. Our framework can be combined with different LLMs; moreover, we fine-tune LLMs to boost their student modeling capabilities.
- III. We evaluate several methods instantiated from our framework on an existing benchmark, STUDENTSYN, for student attempt synthesis in a visual programming domain. Our results highlight that our methods perform significantly better than baselines without requiring complex pipelines or extensive datasets.
- IV. We publicly release the implementation of LLM-SS to facilitate future research.¹

2. RELATED WORK

Student modeling and synthesis in open-ended domains. As discussed in the previous section, there have been recent developments on student modeling for open-ended learning domains, with techniques ranging from misconception discovery to identification of struggling students and investigating problem-solving strategies [4–7, 11–13]. Among these recent works, our work is closer to that of [11] as we are addressing the problem of synthesizing a student’s behavior by focusing

on misconceptions in observed attempts. In fact, our evaluation is based on the STUDENTSYN benchmark from [11] that considers the problem of synthesizing a student’s attempt in visual programming domains. As part of this benchmark, [11] proposed an automated method, NEURSS, that requires extensive pre-training on expert data and continual training on real-world data from similar students. Our framework aims to avoid this complex training pipeline by leveraging the in-context learning capabilities of LLMs. Our work is also similar in spirit to contemporary works that use LLMs for simulating students to teach learners in conversational tutoring systems [20] or train human tutors [18].

LLMs in programming education. Generative AI and LLMs hold great promise in enhancing the field of education through a complementary relationship between human teachers and generative models [21, 22]. Some of the earlier works applying LLMs in educational settings focused on computing and programming education domains and studied a various of scenarios, including generating high-precision feedback [23, 24], generating programming exercises [25], repairing bugs in programming assignments [26], task synthesis for visual programming [27, 28], and benchmarking LLMs capabilities with that of human tutors [28, 29]. Our work differs from these works, given our focus on leveraging LLMs for modeling a student and synthesizing students’ attempts.

3. PROBLEM SETUP

In this section, we formalize the problem of leveraging LLMs for in-context student modeling in open-ended domains. While we focus on LLM-based methods, we provide a generic setup that encapsulates various baseline methods that do not use LLMs (e.g., baseline NEURSS used for comparison in Section 5). In particular, our problem setup is inspired by the work of [11] that will also be used later as a benchmark in our experiments.

Preliminaries and synthesis objective. Given an open-ended learning domain, there is a student, henceforth referred to as

¹Github: <https://github.com/machine-teaching-group/edm2024-llm-student-modeling>

STU, aiming to solve some tasks in the domain. We denote the space of all possible tasks by \mathbb{T} , and the space of all possible solutions and attempts by \mathbb{C} . In particular, we are given a reference task $T^{\text{ref}} \in \mathbb{T}$ of interest along with a solution $C_{T^{\text{ref}}}^* \in \mathbb{C}$.² Our main goal is to develop a synthesizer that can model the student STU by observing how STU solves T^{ref} , and subsequently synthesize an attempt on any similar target task T^{tar} , imitating STU’s behavior. More concretely, we consider the following two-step process:

- (1) First, the synthesizer observes a student’s *context* tuple $(T^{\text{ref}}, C_{T^{\text{ref}}}^*, C_{T^{\text{ref}}}^{\text{STU}})$, where $C_{T^{\text{ref}}}^{\text{STU}} \in \mathbb{C}$ is the student STU’s attempt on solving the reference task.
- (2) Next, given a target task $T^{\text{tar}} \in \mathbb{T}$ conceptually similar to T^{ref} , along with a solution $C_{T^{\text{tar}}}^* \in \mathbb{C}$, the synthesizer synthesizes a student’s attempt $\widehat{C}_{T^{\text{tar}}}^{\text{STU}}$, which should be close to how the student STU would attempt T^{tar} .³

Quality rubric for evaluation. We evaluate the performance of a synthesizer based on the quality of their synthesized student’s attempt $\widehat{C}_{T^{\text{tar}}}^{\text{STU}}$. Based on existing literature [11, 29], we quantitatively measure the generative quality using expert-based assessments w.r.t. the following quality rubric:

- Q-STU. This attribute measures whether the synthesized attempt $\widehat{C}_{T^{\text{tar}}}^{\text{STU}}$ captures the student STU’s behavior (e.g., problem-solving strategy and underlying misconceptions).
- Q-TASK. This attribute measures whether the synthesized attempt $\widehat{C}_{T^{\text{tar}}}^{\text{STU}}$ captures the characteristics of T^{tar} (e.g., partially reflecting its solution $C_{T^{\text{tar}}}^*$).
- Q-OVERALL. This attribute measures whether the synthesized attempt $\widehat{C}_{T^{\text{tar}}}^{\text{STU}}$ successfully captures both the student’s behavior and the target task’s characteristics at the same time. We will set $\text{Q-OVERALL} = \text{Q-STU} \times \text{Q-TASK}$.

Illustrative example for visual programming domain. In our experimental evaluation (Section 5), we will consider an existing benchmark, STUDENTSYN [11], for student attempt synthesis in a visual programming domain of *Hour of Code: Maze Challenge by Code.org* (HoCMaze) [30]. As an illustrative example, Figure 1 shows a concrete scenario for our problem setup.

4. OUR LLM-SS FRAMEWORK

In this section, we propose a novel framework, namely LLM-SS, for in-context student modeling and synthesizing students’ attempts. It is inspired by the *Perturbation Student Model* as discussed below (Section 4.1). Afterward, we delve into two components of LLM-SS: providing student’s context (Section 4.2) and providing domain expertise (Section 4.3).

²A task T can have multiple solutions, and C_T^* refers to any solution codes written by experts being provided as input.

³There are different granularity levels at which we can synthesize the student STU’s behavior, including: (a) a coarse-level binary prediction of success/failure, (b) a medium-level prediction w.r.t. predefined misconceptions; (c) a fine-level synthesis of student’s attempt. Here, we focus on this fine-level objective of synthesizing a student’s attempt.

4.1 Perturbation Student Model

Perturbation Student Model is based on the idea that a student’s knowledge can be modeled as perturbations to expert knowledge [19]. This model was introduced as an extension of the *Overlay Model* [31, 32] – it allows modeling a student’s misconceptions and “buggy” knowledge that deviates from expert knowledge. It assumes that incorrect behaviors of a student can be caused by systematically applying a set of perturbations to domain expertise.

In our LLM-SS framework, we use an LLM to model a student in an open-ended learning domain following the same idea of *Perturbation Student Model*. More concretely, we provide a student’s knowledge by a behavioral context in a prompt to LLM (Section 4.2), and provide domain-specific expertise through fine-tuning the LLM on expert data (Section 4.3).

4.2 Providing Student’s Context

Next, we discuss how to provide a student’s context to an LLM and leverage the LLM’s capabilities of in-context learning. Again, the goal of a student’s context is to give an LLM information about the student, which may include the student’s background, preferences, learning history, and problem-solving trajectories on multiple tasks. This information can be provided to a given LLM as a context in a prompt – existing works have shown that LLMs can effectively learn from such contextual information without explicit training or further parameter updates [15, 16].

In our framework, the prompt includes a student’s context in the form of a problem-solving attempt on a reference task, which is represented by an information tuple $(T^{\text{ref}}, C_{T^{\text{ref}}}^*, C_{T^{\text{ref}}}^{\text{STU}})$; see Section 3. We expect the LLM to infer the student’s misconceptions from the observed attempt along with the necessary perturbations to obtain $C_{T^{\text{ref}}}^{\text{STU}}$ from $C_{T^{\text{ref}}}^*$. Subsequently, the LLM is asked to play the role of this student and synthesize an attempt for a target task T^{tar} , which should reflect the student’s behavior. This is when the LLM should apply the same perturbations to obtain $\widehat{C}_{T^{\text{tar}}}^{\text{STU}}$ from $C_{T^{\text{tar}}}^*$.

Figure 2 shows an example of our main prompt template for providing the student STU’s context and synthesizing the student’s attempts. We note that our LLM-SS framework can accommodate multiple solutions for a task and richer representations of the student’s context as input by appropriately adapting the prompt. In this template example, we have shown a single solution for a task and one student’s attempt, as considered in our experimental evaluation; see Section 5.

4.3 Providing Domain Expertise

Next, we discuss how to provide domain-specific expertise to an LLM for student modeling. In general, datasets used for pre-training LLMs may not contain data coming from specialized open-ended learning domains such as interactive educational games [10], physics simulations [33], or visual programming [30]. Consequently, LLMs could be far from experts in these domains; for instance, even state-of-the-art models like GPT-4 perform poorly in synthesizing solutions for visual programming tasks [28]. In such settings, we need to enhance an LLM domain-specific knowledge to effectively model a student as per the *Perturbation Student Model*. In particular, we will enhance an LLM’s domain expertise via fine-tuning

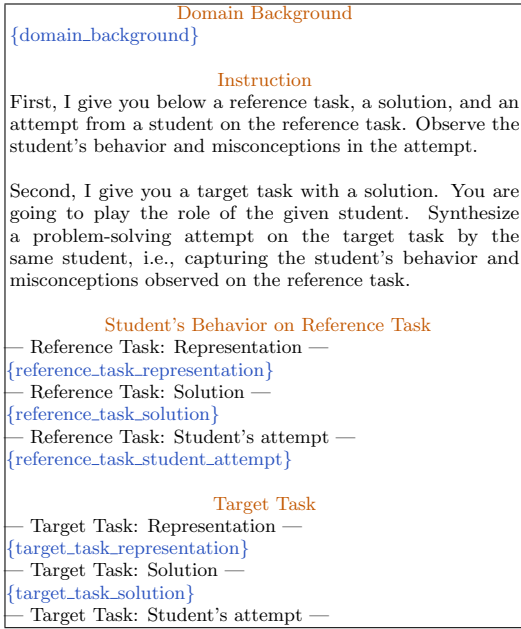


Figure 2: Prompt template used in LLM-SS framework. {placeholders} are used to include details for each scenario.

– existing works have shown that pre-trained LLMs can be tailored to specific domains via fine-tuning [34, 35]

In our framework, we aim to improve a given LLM’s capability of generating solutions C_T^* for any task T similar to the reference task T^{ref} . Once the LLM acquires a better understanding of how to solve tasks in the domain, it is expected to better infer the student’s behavior from the context provided in Section 4.2. More concretely, we use pairs of (task T , solution C_T^*) in the domain to create a fine-tuning dataset $\mathbb{D}_{\text{ft}} = \{\mathbf{x}^{(k)}, \mathbf{y}^{(k)}\}$, where $\mathbf{x}^{(k)}$ is an input prompt containing a task to be solved and $\mathbf{y}^{(k)}$ is the desired solution generated by the LLM. We consider an LLM parameterized by θ , with p_θ denoting conditional probability distribution of sampling responses. We perform supervised fine-tuning to adjust θ through gradient descent, with the objective of minimizing the negative log-likelihood loss given by $L_{\text{ft}}(\theta) := -\mathbb{E}_{(\mathbf{x}^{(k)}, \mathbf{y}^{(k)}) \sim \mathbb{D}_{\text{ft}}} [\log p_\theta(\mathbf{y}^{(k)} | \mathbf{x}^{(k)})]$ [35].

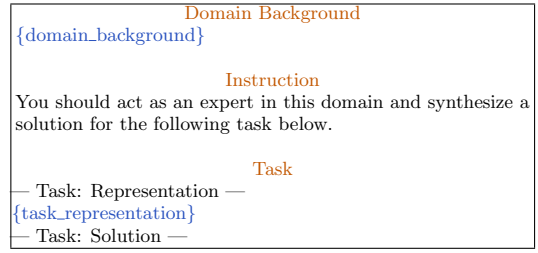
Figure 3 shows the pipeline overview of fine-tuning an LLM in our framework along with an example of fine-tuning prompt template. In each prompt $\mathbf{x}^{(k)}$, we first start by describing the domain background (same as in Figure 2). Then, we use an instruction to steer the LLM’s behavior to act as a domain expert and solve a task. The last part of the prompt is a representation of the task to be solved.

5. EXPERIMENTAL EVALUATION

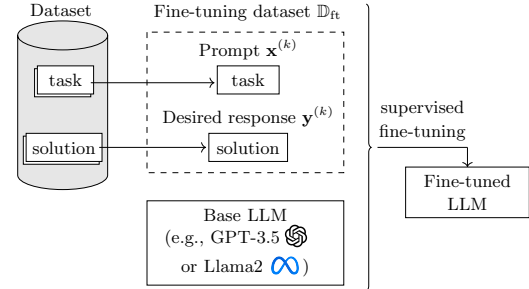
This section presents our experimental evaluation, including description of STUDENTSYN benchmark with baseline methods from [11] (Section 5.1), evaluated methods (Section 5.2), evaluation procedure (Section 5.3), and results (Section 5.4).

5.1 STUDENTSYN Benchmark and Baselines

We use the STUDENTSYN benchmark from [11], designed to evaluate student’s attempt synthesis methods in the visual



(a) Prompt for fine-tuning.



(b) Fine-tuning an LLM using synthetic data.

Figure 3: Fine-tuning an LLM using expert knowledge in LLM-SS framework.


block-based programming domain of *Hour of Code: Maze Challenge* by *Code.org* (HoCMaze) [30]. This programming domain has been popularly used in several existing works [11, 36–38]. Figure 1 shows an example of task T^{ref} along with a solution $C_{T^{\text{ref}}}^*$ – a task in HoCMaze is specified by a visual grid containing an avatar (blue arrow), a goal (red star), and some walls (gray cells); a solution code brings the avatar to the goal’s location while avoiding hitting the walls. STUDENTSYN is a challenging benchmark for our problem setup, as evidenced by the huge performance gap between human tutors and automated methods proposed in [11].

Benchmark scenarios. This benchmark comprises two reference tasks T^{ref} , namely HoCMaze-4 and HoCMaze-18 [30], and three target tasks T^{tar} associated with each reference task. In our illustration of problem setup in Figure 1, we use HoCMaze-18 as T^{ref} . The benchmark considers six types of misconceptions, such as confusion between left/right directions when turning, writing repetitive turn commands, and ignoring the If-Else/While structure. The benchmark provides a set of scenarios comprising a student STU with a specific misconception, one coding attempt $C_{T^{\text{ref}}}^{\text{STU}}$ on each reference task, and the STU’s attempt on each target task $C_{T^{\text{tar}}}^{\text{STU}}$ serving as a ground-truth. In total, we evaluate on these 36 scenarios ($2 T^{\text{ref}} \times 3 T^{\text{tar}} \times 6 \text{STU}$).

Dataset for fine-tuning. Along with benchmark scenarios, [11] also provides a synthetic dataset consisting of (task, solution) pairs, where tasks are similar to either HoCMaze-4 or HoCMaze-18. Here, task similarity is measured conceptually by edit distance in solution codes. This synthetic dataset was created and used for pre-training models introduced in [11]. In our framework, we will use it to fine-tune a base LLM to boost its domain expertise. In total, there are 10,000 training tasks and 500 validation tasks for HoCMaze-4, and 40,000 training tasks and 500 validation tasks for HoCMaze-18.

	HoCMaze-4			HoCMaze-18		
	Q-OVERALL	Q-STU	Q-TASK	Q-OVERALL	Q-STU	Q-TASK
GPT-3.5-SS	0.28	0.56	0.50	0.14	0.61	0.25
GPT-4-SS	0.61	0.86	0.72	0.51	0.81	0.58
GPT-3.5ft-SS	0.64	0.69	0.75	0.82	0.92	0.86
Llama2-7B-SS	0.08	0.14	0.44	0.08	0.25	0.39
Llama2-70B-SS	0.36	0.58	0.50	0.26	0.56	0.50
Llama2-7Bft-SS	0.52 (0.05)	0.55 (0.07)	0.90 (0.05)	0.30 (0.08)	0.66 (0.11)	0.39 (0.09)
Llama2-70Bft-SS	0.65 (0.08)	0.87 (0.05)	0.73 (0.05)	0.53 (0.03)	0.83 (0.02)	0.63 (0.03)
NEURSS	0.43	0.56	0.67	0.25	0.78	0.36
TUTORSS	0.84	0.92	0.89	0.85	0.89	0.95

(a) Detailed results w.r.t. to each attribute in the quality rubric: Q-OVERALL, Q-STU, and Q-TASK. Fine-tuned models are highlighted in green. Fine-tuning improves LLMs’ capabilities of capturing both student’s behavior and target task’s characteristics.



(b) Overall performance (Q-OVERALL). Green areas correspond to fine-tuning improvements. TUTORSS (red lines) serves as a performance upper bound.

Figure 4: (a) shows performances of methods w.r.t. individual attributes in our quality rubric. (b) shows the overall performance of capturing both student’s behavior and target task’s characteristics. Human tutors (TUTORSS) serve as an oracle. For methods using a fine-tuned LLM, we report numbers averaged over three fine-tuning runs with standard errors (except GPT-3.5ft-SS with only one run, due to the high costs of using fine-tuning APIs from OpenAI).

Baseline methods. We compare our framework with baseline method NEURSS [11], an LSTM-based neural network pre-trained on expert knowledge and continually trained on real students’ attempts. We also compare our framework with human tutors in the visual programming domain, referred to as TUTORSS in [11]. Here, TUTORSS can be considered an oracle that provides performance upper bounds. We re-use the students’ attempts synthesized by NEURSS and TUTORSS from [11], and re-assess them w.r.t. our rubric in Section 5.3.

5.2 Methods Based on LLM-SS Framework

Methods using a base LLM without fine-tuning. Based on our LLM-SS framework, we develop the following concrete methods using base models without fine-tuning step described in Section 4.3: GPT-3.5-SS using GPT-3.5 [39], GPT-4-SS using GPT-4 [40], Llama2-7B-SS using Llama2-7B-Chat [35], and Llama2-70B-SS using Llama2-70B-Chat [35].

Methods using a fine-tuned LLM. We further develop the following three concrete methods by fine-tuning three models: GPT-3.5ft-SS using fine-tuned GPT-3.5 [39], Llama2-7Bft-SS using fine-tuned Llama2-7B-Chat [35], and Llama2-70Bft-SS using fine-tuned Llama2-70B-Chat [35]. We did not fine-tune the GPT-4 model as APIs are not publicly available. Details of our fine-tuning procedure are explained in Section 4.3.⁴

5.3 Evaluation Procedure

For each scenario from the STUDENTSYN benchmark (see Section 5.1), we create a prompt following the template in Figure 2 to use it as input to an LLM. We use the domain background representation for HoCMaze based on prompts in recent works [27]. Subsequently, all scenarios together with student attempts synthesized by the LLM are presented to two independent experts for assessment – these two experts

⁴For fine-tuning Llama2-70B models, we used a cluster of 2×36 cores, 2.40 GHz Intel Xeon Platinum Processor 8360Y, and $8 \times$ Nvidia A100 80GB, with parallelization under a 64-bit Debian. We fine-tuned a model for each reference task separately, and one run on a reference task took up to 35 hours. For GPT-3.5, we fine-tuned the GPT-3.5-turbo-0613 model for each reference task separately, and one run on a reference task took up to 7 hours. We paid about 1000\$ in total for using fine-tuning APIs provided by OpenAI.

have extensive knowledge in computer science and visual programming domains, and follow the evaluation rubric from Section 3. The annotation process is done in a blind condition, in which experts do not know from which method a coding attempt is synthesized. In total, there are about 500 codes to be annotated by each expert corresponding to different scenarios and methods.

These two experts annotated the synthesized codes by using binary values $\{0, 1\}$ for annotation, i.e., each quality attribute could take a value of 0 (bad) or 1 (good). Concretely, $Q\text{-STU} = 1$ means that $\hat{C}_{T^{\text{STU}}}^{\text{STU}}$ captures the student STU’s behavior in terms of the problem-solving strategy and underlying misconceptions, and otherwise $Q\text{-STU} = 0$; similarly, $Q\text{-TASK} = 1$ means that $\hat{C}_{T^{\text{tar}}}^{\text{STU}}$ captures the characteristics of target task T^{tar} , and otherwise $Q\text{-TASK} = 0$. $Q\text{-OVERALL}$, defined as $Q\text{-STU} \times Q\text{-TASK}$, takes values of $\{0, 1\}$. We validate the expert annotations w.r.t. $Q\text{-OVERALL}$ using Cohen’s kappa inter-agreement reliability [41], obtaining a value of 0.71, indicating *substantial agreement* between two experts.

Nevertheless, further investigation into the annotations revealed that the majority of disagreements between two experts were borderline cases where the quality attribute value was unclear. This motivated us to refine the scale of assessment where $Q\text{-STU}$ and $Q\text{-TASK}$ would take values of $\{0, 0.5, 1\}$, with 0.5 now indicating partially capturing the student’s behavior or the target task’s characteristics. We note that $Q\text{-OVERALL}$, defined as $Q\text{-STU} \times Q\text{-TASK}$, now takes values of $\{0, 0.25, 0.5, 1\}$. With this refined scale, one expert did the entire annotations again and the final results reported in Section 5.4 are based on these new annotations. We report averaged results in the range $[0.0, 1.0]$ by aggregating across all scenarios for a given reference task.

5.4 Results

Without fine-tuning: GPT-4-SS outperforms NEURSS. Among our methods that use a base LLM model without fine-tuning, GPT-4-SS achieved the highest scores in all quality attributes across both reference tasks, followed by Llama2-70B-SS (see Figure 4a). Additionally, GPT-4-SS performs significantly better than the NEURSS baseline w.r.t. $Q\text{-OVERALL}$ in both




Figure 5: Losses and evaluations during fine-tuning our two best-performing methods GPT-3.5ft-SS and Llama2-70Bft-SS. We plot data per 0.1 epoch. Losses are plotted on log scale for better visibility of dynamics. Validation BLEU/accuracy metrics are decided by the fine-tuning library/platform and shown as a sanity check and are not used for optimization. For GPT-3.5ft-SS, the number of epochs depends on the budget spent for OpenAI APIs; we spent roughly half of the total budget for each task. For Llama2-70Bft-SS, the number of epochs is determined by generative performance on a small validation set of examples.

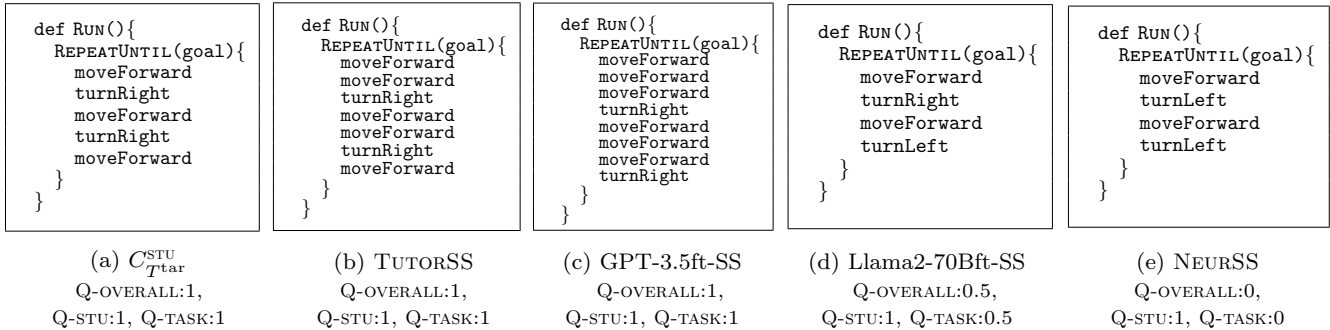


Figure 6: Student STU’s attempts for the scenario shown in Figure 1. (a) shows ground-truth student STU’s attempt $C_{T^{\text{tar}}}^{\text{STU}}$ provided in the STUDENTSYN benchmark. (b-e) show synthesized student STU’s attempts $\hat{C}_{T^{\text{tar}}}^{\text{STU}}$ provided by different methods.

reference tasks ($p \leq 0.05$), based on the χ^2 test [42].⁵ Q-OVERALL scores of GPT-3.5-SS and Llama2-7B-SS are lower than that of the baseline NEURSS (which also motivates why we need to do fine-tuning discussed in Section 4.3).

Fine-tuning shows significant improvements. Our methods using fine-tuning, namely GPT-3.5ft-SS, Llama2-7Bft-SS, and Llama2-70Bft-SS, demonstrate significant improvements compared to using their base versions without fine-tuning ($p \leq 0.05$), as shown in Figure 4b. Remarkably, for HoCMaze-18, there is no significant difference between the performances of GPT-3.5ft-SS and human tutors in TUTORSS ($p > 0.05$). We observe that fine-tuning enhances a base LLM’s ability to capture the target task’s structure (Q-TASK), as shown in Figure 4a – this improvement is expected given they are fine-tuned to generate solutions for tasks. More importantly, their ability to capture the student’s behavior (Q-STU) also increases across all reference tasks and fine-tuned models. Figure 5 provides insights into the fine-tuning process.

Example of synthesized student’s attempt. In Figure 6, we investigate the scenario for HoCMaze-18 from Figure 1. In this scenario, the student STU’s misconception is ignoring conditionals when attempting to solve the given task. Figure 6(a) shows student code $C_{T^{\text{tar}}}^{\text{STU}}$ for the target task provided in the

benchmark. Figures 6(b-e) show student codes $\hat{C}_{T^{\text{tar}}}^{\text{STU}}$ synthesized by different methods. Student code synthesized by GPT-3.5ft-SS has the same misconception observed in $C_{T^{\text{ref}}}^{\text{STU}}$, while adapting to T^{tar} (Q-OVERALL = 1). Notably, it is very close to the code written by human tutors in TUTORSS (Q-OVERALL = 1). Llama2-70Bft-SS synthesized a code that captures the student’s misconception, but only partially reflects the target task’s characteristics (Q-OVERALL = 0.5). In contrast, the NEURSS baseline synthesized a code that overfits $C_{T^{\text{ref}}}^{\text{STU}}$, failing to reflect the layout of T^{tar} as it uses `turnLeft` blocks instead of `turnRight` blocks (Q-OVERALL = 0).

6. CONCLUDING DISCUSSIONS

We proposed a novel LLM-based framework, LLM-SS, for in-context student modeling in open-ended learning domains. The results showcase that methods instantiated from LLM-SS are capable of modeling a student’s observed behavior and synthesizing the student’s attempt on a target task. We also highlight that fine-tuning a base LLM using expert knowledge in a given open-ended learning domain significantly improves its effectiveness in student modeling. More importantly, our framework does not require building a complex training pipeline as existing works, making it broadly applicable to new domains. In summary, our work demonstrates the potential of using LLMs for in-context student modeling, especially in challenging open-ended learning domains.

Next, we discuss some limitations of our current work and

⁵ χ^2 tests reported here are computed on aggregated data across both the reference tasks.

ideas to tackle them in the future. First, our framework was evaluated on one visual programming domain, and the scenarios we considered do not fully capture the wide spectrum of open-ended learning domains; it would be interesting to evaluate our framework in other open-ended learning domains (e.g., algebra or text-based programming). Moreover, it would also be useful to do a more systematic analysis to see which misconceptions or students' behaviors are not well captured by our framework. Second, we provided a student's context through only one example of a problem-solving attempt; it would be interesting to evaluate the effectiveness of our framework when the student's context contains richer information, including the student's background and attempts on different tasks. Third, we evaluated our framework on student modeling metrics but have not evaluated how this modeling helps improve the performance of downstream applications; as future work, it would also be important to investigate the usefulness of our modeling framework directly in downstream applications, such as performance prediction, task recommendation, or synthetic behavioral dataset generation for training data-intensive models. In particular, since our framework allows fine-grained synthesis of a student's attempts beyond binary performance prediction, it would be interesting to see how our framework can potentially be applied for providing finer-grained feedback to the student about possible misconceptions.

Finally, we note that there are several ethical implications regarding the use of LLMs for student modeling. For instance, the attempts synthesized by LLMs may not accurately reflect a student's understanding or ability. Moreover, LLMs are prone to hallucination and might generate inaccurate information. Therefore, it is crucial to implement appropriate validation mechanisms and safeguards when deploying LLM-based student modeling techniques in classrooms.

Acknowledgement

Funded/Co-funded by the European Union (ERC, TOPS, 101039090). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

- [1] Kurt VanLehn. Student Modeling. *Foundations of Intelligent Tutoring Systems*, pages 55–78, 2013.
- [2] Konstantina Chrysafiadi and Maria Virvou. Student Modeling for Personalized Education: A Review of the Literature. *Advances in Personalized Web-Based Education*, 78:1–24, 2015.
- [3] Anna N. Rafferty, Rachel Jansen, and Thomas L. Griffiths. Using Inverse Planning for Personalized Feedback. In *Proceedings of the International Conference on Educational Data Mining (EDM)*, 2016.
- [4] Andrew Emerson, Andy Smith, Fernando J. Rodríguez, Eric N. Wiebe, Bradford W. Mott, Kristy Elizabeth Boyer, and James C. Lester. Cluster-Based Analysis of Novice Coding Misconceptions in Block-Based Programming. In *Proceedings of the Technical Symposium on Computer Science Education (SIGCSE)*, 2020.
- [5] Yang Shi, Krupal Shah, Wenran Wang, Samiha Marwan, Poorvaja Penmetsa, and Thomas W. Price. Toward Semi-Automatic Misconception Discovery Using Code Embeddings. In *Proceedings of the International Learning Analytics and Knowledge Conference (LAK)*, 2021.
- [6] Lisa Wang, Angela Sy, Larry Liu, and Chris Piech. Learning to Represent Student Knowledge on Programming Exercises Using Deep Learning. In *Proceedings of the International Conference on Educational Data Mining (EDM)*, 2017.
- [7] Jade Cock, Mirko Marras, Christian Giang, and Tanja Käser. Early Prediction of Conceptual Understanding in Interactive Simulations. In *Proceedings of the International Conference on Educational Data Mining (EDM)*, 2021.
- [8] Ahana Ghosh, Sebastian Tschischek, Sam Devlin, and Adish Singla. Adaptive Scaffolding in Block-Based Programming via Synthesizing New Tasks as Pop Quizzes. In *Proceedings of the International Conference on Artificial Intelligence in Education (AIED)*, 2022.
- [9] Michael J. Hannafin, Craig Hall, Susan Land, and Janette Hill. Learning in Open-Ended Environments: Assumptions, Methods, and Implications. *Educational Technology*, 34(8):48–55, 1994.
- [10] Tanja Käser and Daniel L. Schwartz. Modeling and Analyzing Inquiry Strategies in Open-Ended Learning Environments. *International Journal of Artificial Intelligence in Education (IJAIED)*, 30(3):504–535, 2020.
- [11] Adish Singla and Nikitas Theodoropoulos. From {Solution Synthesis} to {Student Attempt Synthesis} for Block-Based Visual Programming Tasks. In *Proceedings of the International Conference on Educational Data Mining (EDM)*, 2022.
- [12] ByeongJo Kong, Erik Hemberg, Ana Bell, and Una-May O'Reilly. Investigating Student's Problem-solving Approaches in MOOCs using Natural Language Processing. In *Proceedings of the International Learning Analytics and Knowledge Conference (LAK)*, 2023.
- [13] Andrew Emerson, Wookhee Min, Jonathan P. Rowe, Roger Azevedo, and James C. Lester. Multimodal Predictive Student Modeling with Multi-Task Transfer Learning. In *Proceedings of the International Learning Analytics and Knowledge Conference (LAK)*, 2023.
- [14] Lauren Fratamico, Cristina Conati, Samad Kardan, and Ido Roll. Applying a Framework for Student Modeling in Exploratory Learning Environments: Comparing Data Representation Granularity to Handle Environment Complexity. *International Journal Artificial Intelligence in Education (IJAIED)*, 27(2):320–352, 2017.
- [15] Tom B. Brown et al. Language Models are Few-Shot Learners. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [16] Sébastien Bubeck et al. Sparks of Artificial General Intelligence: Early Experiments with GPT-4. *CoRR*, abs/2303.12712, 2023.

- [17] Gati V Aher, Rosa I. Arriaga, and Adam Tauman Kalai. Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2023.
- [18] Julia M. Markel, Steven G. Opferman, James A. Landay, and Chris Piech. GPTEach: Interactive TA Training with GPT-based Students. In *Proceedings of the Conference on Learning @ Scale (L@S)*, pages 226–236, 2023.
- [19] Robert Kass. Student Modeling in Intelligent Tutoring Systems — Implications for User Modeling. In *User Models in Dialog Systems*, pages 386–410, 1989.
- [20] Robin Schmucker, Meng Xia, Amos Azaria, and Tom Mitchell. Ruffle&Riley: Towards the Automated Induction of Conversational Tutoring Systems. *NeurIPS’23 Workshop on Generative AI for Education (GAIED)*, 2023.
- [21] Paul Denny, Sumit Gulwani, Neil T. Heffernan, Tanja Käser, Steven Moore, Anna N. Rafferty, and Adish Singla. Generative AI for Education (GAIED): Advances, Opportunities, and Challenges. *CoRR*, abs/2402.01580, 2024.
- [22] Jaeho Jeon and Seongyong Lee. Large Language Models in Education: A Focus on the Complementary Relationship Between Human Teachers and ChatGPT. *Education and Information Technologies*, 28(12):15873–15892, 2023.
- [23] Tung Phung, José Cambronero, Sumit Gulwani, Tobias Kohn, Rupak Majumdar, Adish Singla, and Gustavo Soares. Generating High-Precision Feedback for Programming Syntax Errors using Large Language Models. In *Proceedings of the International Conference on Educational Data Mining (EDM)*, 2023.
- [24] Tung Phung, Victor-Alexandru Pădurean, Anjali Singh, Christopher Brooks, José Cambronero, Sumit Gulwani, Adish Singla, and Gustavo Soares. Automating Human Tutor-Style Programming Feedback: Leveraging GPT-4 Tutor Model for Hint Generation and GPT-3.5 Student Model for Hint Validation. In *Proceedings of the International Learning Analytics and Knowledge Conference (LAK)*, 2024.
- [25] Sami Sarsa, Paul Denny, Arto Hellas, and Juho Leinonen. Automatic Generation of Programming Exercises and Code Explanations Using Large Language Models. In *Proceedings of the Conference on International Computing Education Research (ICER)*, 2022.
- [26] Jialu Zhang, José Cambronero, Sumit Gulwani, Vu Le, Ruzica Piskac, Gustavo Soares, and Gust Verbruggen. Repairing Bugs in Python Assignments Using Large Language Models. *CoRR*, abs/2209.14876, 2022.
- [27] Victor-Alexandru Pădurean, Georgios Tzannetos, and Adish Singla. Neural Task Synthesis for Visual Programming. *Transactions on Machine Learning Research (TMLR)*, 2024.
- [28] Adish Singla. Evaluating ChatGPT and GPT-4 for Visual Programming. In *Proceedings of the Conference on International Computing Education Research (ICER) - Volume 2*, 2023.
- [29] Tung Phung, Victor-Alexandru Pădurean, José Cambronero, Sumit Gulwani, Tobias Kohn, Rupak Majumdar, Adish Singla, and Gustavo Soares. Generative AI for Programming Education: Benchmarking Chatgpt, GPT-4, and Human Tutors. In *Proceedings of the Conference on International Computing Education Research (ICER) - Volume 2*, 2023.
- [30] Code.org. Hour of Code: Classic Maze Challenge. <https://studio.code.org/s/hourofcode>, 2012.
- [31] Brian Carr and Ira P. Goldstein. Overlays: A Theory of Modelling for Computer Aided Instruction. <https://hal.science/hal-00702959>, 1977. AI Memo 406.
- [32] Gordon McCalla Jim E. Greer. *Student Modelling: The Key to Individualized Knowledge-Based Instruction*. Springer-Verlag, 1994.
- [33] Carl E. Wieman, Wendy K. Adams, and Katherine K. Perkins. PhET: Simulations That Enhance Learning. *Science*, 322(5902):682–683, 2008.
- [34] Long Ouyang et al. Training Language Models to Follow Instructions with Human Feedback. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2022.
- [35] Hugo Touvron et al. Llama 2: Open Foundation and Fine-Tuned Chat Models. *CoRR*, abs/2307.09288, 2023.
- [36] Chris Piech, Mehran Sahami, Jonathan Huang, and Leonidas J. Guibas. Autonomously Generating Hints by Inferring Problem Solving Policies. In *Proceedings of the Conference on Learning @ Scale (L@S)*, 2015.
- [37] Aleksandr Efremov, Ahana Ghosh, and Adish Singla. Zero-shot Learning of Hint Policy via Reinforcement Learning and Program Synthesis. In *Proceedings of the International Conference on Educational Data Mining (EDM)*, 2020.
- [38] Umair Z. Ahmed, Maria Christakis, Aleksandr Efremov, Nigel Fernandez, Ahana Ghosh, Abhik Roychoudhury, and Adish Singla. Synthesizing Tasks for Block-based Programming. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [39] OpenAI. OpenAI GPT-3.5. <https://platform.openai.com/docs/models/gpt-3-5-turbo>, 2023.
- [40] OpenAI. GPT-4 Technical Report. *CoRR*, abs/2303.08774, 2023.
- [41] Jacob Cohen. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20:37 – 46, 1960.
- [42] William G. Cochran. The χ^2 Test of Goodness of Fit. *The Annals of Mathematical Statistics*, 23(3):315–345, 1952.