User Modeling and User-Adapted Interaction (2025) 35:7
https://doi.org/10.1007/s11257-025-09426-4

Modeling the phases of rule learning during problem solving with an interactive learning environment

Deniz Sonmez Unal[1] · Erin Solovey[2] · Catherine M. Arrington[3] · Erin Walker[1]

Received: 11 March 2024 / Accepted in revised form: 26 January 2025

Abstract
While existing student modeling methods focus on predicting students' knowledge states, they often overlook the underlying cognitive processes contributing to learning. In this work, we integrate cognitive processes, specifically phases of rule learning, into student modeling, drawing inspiration from cognitive science. Rule learning involves rule search, discovery, and following, providing a systematic framework for understanding how individuals acquire and apply knowledge. We conduct two studies to explore rule learning phases in a real-world learning context. Moreover, we present a two-step approach to first predict the phases of rule learning students experience during problem solving with an intelligent tutoring system and then estimate the time spent on each predicted phase. Furthermore, we identify the relationships between the time spent on specific phases of rule learning and student performance. Our findings underscore the importance of integrating cognitive processes into student modeling for more targeted interventions and personalized support.

Keywords Student modeling · Rule learning · Cognitive processes · Intelligent tutoring systems

✉ Deniz Sonmez Unal
des204@pitt.edu

Erin Solovey
esolovey@wpi.edu

Catherine M. Arrington
cma205@lehigh.edu

Erin Walker
eawalker@pitt.edu

[1] University of Pittsburgh, Pittsburgh, PA, USA

[2] Worcester Polytechnic Institute, Worcester, MA, USA

[3] Lehigh University, Bethlehem, PA, USA

Published online: 06 March 2025

1 Introduction

One of the primary goals of intelligent tutoring systems (ITS) is to provide students with individualized assistance to support their learning process. These systems achieve this goal by appropriate task selection and giving timely hints and feedback (VanLehn 2006). Assessing students' skills and knowledge state is the first step to design this personalized assistance and the necessary interventions. Therefore, student modeling plays a significant role in ITSs.

Various methods have been used to model student knowledge from Bayesian knowledge tracing (BKT) to logistic regression-based models such as performance factor analysis (PFA) as well as ensemble models consisting of combinations of the two, and their modern extensions using advanced machine learning techniques (d Baker et al. 2011; Liu et al. 2021). Even though these models are thoroughly studied and useful in predicting the knowledge state of a student based on their responses, they do not put as much emphasis on capturing the cognitive processes leading to learning (Piech et al. 2015; Pelánek 2017).

In this work, we incorporate these cognitive processes into student modeling, as they were conceptualized in rule learning studies of cognitive science. These studies suggest rule learning has three main phases: rule searching, rule discovery, and rule following (Crescentini et al. 2011). Rule search involves gathering information, identifying the similarities and differences within the information in order to build knowledge, rule discovery is defined as the moment when the knowledge is acquired, and rule following is when the acquired rule is applied/fired in appropriate situations. Many cognitive processes that lead to robust learning are closely aligned with the phases of rule learning, including cognitive processes such as activation and organization of prior knowledge, acquisition and integration of new knowledge, and retrieving and reinforcing the acquired knowledge. In this way, rule search involves activating and organizing prior knowledge, rule discovery involves acquisition and integration of new knowledge, and rule following involves retrieving and reinforcing the learned rule effectively. Therefore, understanding the phases of rule learning can provide insight into the cognitive processes underlying robust learning and allow us to design instructional interventions that promote effective learning.

A key advantage of identifying the phases of rule learning during learning is that they allow us to detect cognitive states not just based on correctness but on the progression of cognitive processes during learning. For example, a student might spend considerable time in the rule search phase, struggling to organize prior knowledge or identify patterns in the information presented. Traditional models might only predict an incorrect answer at the end of the problem, but our approach could intervene earlier by recognizing when a student is stuck in rule search and prompting them with hints or scaffolds to facilitate the transition to rule discovery. When the student enters the rule discovery phase, where they understand the correct rule but may still struggle to apply it consistently, the system can offer delayed feedback or reflective prompts to strengthen their comprehension. Therefore, the rule learning framework provides an appropriate level of detail not only for the purpose of explaining the underlying cognitive processes of students' behaviors, but it also allows for a continuous assessment of students' cognitive states, enabling the identification of precise moments where learning occurs. This continuous assessment can also be especially valuable during pauses in interaction, where no immediate data are generated from student responses. Detecting the phases that occur consecutively during pauses and capturing when the transition from one phase to another will happen by estimating the time spent on each consecutive phase may help better interpreting the pauses during learning, whether they indicate reflection, confusion, or preparation for rule application. This opens up the possibility for adaptive interventions at moments that would traditionally be overlooked by item response-based models.

While rule learning has been extensively studied in cognitive science, the research has largely been conducted in highly controlled laboratory settings using artificial tasks that may not fully represent the complexity of real-world learning environments. With this work, we propose an approach that allows for direct identification of these phases in a complex learning task. In the aforementioned controlled settings, the progression through rule learning phases is highly uniform, with little variation beyond the duration of specific phases based on task difficulty. In contrast, our work seeks to understand how (or if) these phases might change when applied to a more dynamic and less predictable environment. By doing so, we aim to identify potential "additional phases" or "patterns" that may emerge in such settings, which are not typically observed in controlled tasks. Our work aims to refine and possibly expand the definition of rule learning in complex, real-world learning tasks. Our findings could suggest necessary adjustments or extensions to the existing rule learning framework to better capture the nuances of learning in more realistic settings. In this paper, we answer the following research questions:

RQ 1 What do the phases of rule learning look like within a complex learning task?

RQ 2 Can we predict what phase of rule learning a student is in during a problem-solving activity?

RQ 3 Can we estimate the time spent on the different phases of rule learning during problem solving?

The paper is organized as follows: Sect. 2 gives background information about the previous efforts on student modeling. We present learning processes relevant to student modeling and how those processes were defined and supported within intelligent tutoring systems. We explain how phases of rule learning are defined in cognitive science and their relevance to skill acquisition processes. Section 3 presents our task designed to capture behaviors related to the phases of rule learning in a realistic learning environment. Section 4 presents the corpus that we collected over 2 studies. Section 5 presents our methodology to identify the phases of rule learning in a realistic learning environment. In Sect. 6, we demonstrate our efforts to predict the phases of rule learning students go through and estimate the time spent on these phases, respectively, and how our models perform on a dataset that is held out from training to show generalizability of our methods. Discussion including limitations of our study and implications for application is presented in Sect. 7. Conclusions and future work are found in Sect. 8.

2 Background

2.1 Student modeling

We review student modeling efforts in intelligent tutoring systems in two groups. The first group is the knowledge tracing methods. Traditional Bayesian knowledge tracing (BKT) is a probabilistic model that detects if a student has learned a skill/knowledge at a given time of solving a problem (Corbett and Anderson 1994). It is a special case of hidden Markov model (HMM). The model uses fixed, skill-specific parameters to estimate the probability that a student acquired the targeted skill/knowledge component. Further extensions of this model include introducing student-specific parameters (Pardos and Heffernan 2010; Lee and Brunskill 2012), item difficulty (Pardos and Heffernan 2011), and contextualized estimations of the probability of guessing and slipping behaviors (d Baker et al. 2008). More recent extensions to this model include introducing deep neural networks. Piech et al. (Piech et al. 2015) were the first to utilize deep learning in knowledge tracing. Their deep knowledge tracing model (DKT) used recurrent neural networks (RNN) to predict student performance using the history of the students' previous activity. Later work proposed incorporating attention mechanism for knowledge tracing in order to be able to model the longer-range dependencies between students' learning interactions (Wang et al. 2022; Zhang et al. 2022), and externally storing knowledge concepts in order to enhance interpretability of DKT by allowing for detecting exactly which knowledge components students mastered (Zhang et al. 2017).

The second group of student models is the logistic regression-based models. This family of models operates by calculating a skill estimate for a given knowledge component and then applying a logistic transformation on this estimate to obtain the probability of a correct answer to a problem associated with the knowledge component. Two typical instances of these models are the additive factors model (AFM) and performance factors analysis (PFA) (Cen et al. 2006; Pavlik et al. 2009). AFM suggests the students are more likely to give a correct response to a problem when they have more prior practice with the skill associated with the particular problem. This model utilizes the total number of previous practice opportunities that a student had for the given skill. PFA claims successful and unsuccessful practices may have different effects on the probability of a correct response. Therefore, it separates the successful practices from unsuccessful ones. Similar to BKT, recent extensions of PFA have utilized machine learning techniques. Asselman et al. (Asselman et al. 2021) used ensemble methods to enhance the predictive performance of PFA. They found a PFA approach based on XGBoost algorithm has outperformed the original PFA method in predictive performance.

The methods we have discussed are used in various scenarios within intelligent tutoring systems, and they have achieved state of the art performance in modeling students' knowledge state. However, these methods are not concerned with the underlying processes that contribute to learning such as encoding, interpretation, rehearsal, and automation. In contrast, modeling rule learning phases offers a more continuous assessment of students' cognitive states, capturing the states that occur not only during problem solving but also in pauses. Identifying the phases of rule learning may help categorize the cognitive states as they unfold during students' interaction with educational technology, including periods of pause where important cognitive activities like reflection, strategy formulation, or mental rehearsal might occur. This continuous monitoring could allow us to better detect moments of confusion, transition, or readiness to apply a learned rule, offering new opportunities for timely, targeted interventions that traditional models may overlook.

2.2 Learning processes relevant to student modeling

The central focus of many highly utilized student models is knowledge components (KC) (Aleven and Koedinger 2013). The knowledge-learning-instruction (KLI) framework introduced by Koedinger et al. (Koedinger et al. 2012) defines a KC as *"an acquired unit of cognitive function or structure that can be inferred from performance on a set of related tasks."* The KLI framework conceptualizes KC as a generalization for a unit of cognition/knowledge. Other popular terms in education and cognitive science literature include "skill" (Bloom et al. 1964), "production rule" (Anderson 2013), and "schema"(Cheng and Holyoak 1985; Kirschner 2002). This body of research also studies the processes of building this unit of knowledge and how these processes lead to an appropriate design of instruction and to enhancing learning. Cognitive skill acquisition has been described as happening in several phases. According to Fitts (Fitts 1964), skill acquisition happens in three stages: the cognitive stage, the associative stage, and the autonomous stage. In the cognitive stage, the knowledge needs to be encoded, interpreted and rehearsed. The task performance on the skill is slow and prone to errors. The associative stage serves as a transition between the cognitive and autonomous stages, during which knowledge is refined and errors are gradually reduced. Finally, in the autonomous stage, the improvement on the skill continues as the performance on the skill is effortless and free of errors. This three stage model was adopted in the early version of Anderson's ACT theory (Anderson 1982, 1987). While Fitts's model describes the overarching stages of skill acquisition, Anderson's ACT theory further explains these stages within its framework. ACT not only aligns with Fitts's observations but also provides a detailed explanation of the underlying cognitive processes associated with each stage. The first stage of skill acquisition was called the declarative stage. This included having learners encode gained information about the skill. This stage corresponds to the cognitive stage described in Fitts's model. The knowledge compilation mechanism in ACT, where the knowledge is converted from a declarative to a procedural form, corresponds to the associative stage. The last stage of skill acquisition was called the procedural stage. In this stage, learners achieve faster and more effortless problem solving and this corresponds to the autonomous stage in Fitts's model. Taatgen and Lee later developed a learning mechanism called production compilation (Taatgen and Lee 2003) which was based on the knowledge compilation mechanism in ACT and the chunking mechanism in Soar (Newell 1994). This improved mechanism has been integrated in the later versions of ACT-R (Anderson et al. 2004).

Within different stages of skill acquisition, various instructional techniques have been promoted to enhance learning. Since the earlier stages of skill acquisition focus on collecting and categorizing information on the skill to be learned, the emphasis is usually on constructing the schema related to the skill (Hummel and Nadolski 2002) through instructional methods such as solving analogical problems (Gick and Holyoak 1980, 1983) and worked-out examples (Cooper and Sweller 1987; Renkl 2014). In the later stages, the purpose of instructional events shifts from recognizing instances of problems to automation of the learned rules/schema for the skill (Van Merrienboer and Paas 1990), and gaining speed and accuracy through practice (Newell and Rosenbloom 1981) and problem solving (Atkinson et al. 2003).

While the phases of building knowledge are explored in low-level fashion in cognitive science literature as we mentioned in this section, the KLI framework describes them at a higher level. The KLI framework suggests three learning processes: induction and refinement processes, understanding and sense-making processes, and memory and fluency building processes. Induction and refinement processes involve building the knowledge for the first time and revisiting and refining the gained knowledge. These processes are facilitated by providing students with worked-out examples or giving error feedback in learning environments. Sense-making processes include students' deliberate reasoning and verbally mediated thinking efforts to understand the given learning material and build knowledge. Prompting self-explanations and accountable talk promote sense-making processes and acquisition of more complex knowledge components. Finally, memory and fluency building processes include building up the memory to allow for quick and effortless retrieval and application of the knowledge. Spacing practice and optimized scheduling are utilized in order to help memory and fluency building processes when designing instruction.

The processes described above provide guidelines for how instruction should be structured both within traditional learning environments and modern ITSs. However, the link between these processes has been rarely discussed in the context of student modeling, which is the backbone of an ITS. Pelanek (Pelánek 2017) has connected the previously developed student models with the learning processes defined in the KLI framework. They propose a mapping between the importance of different modeling aspects and the learning processes defined in KLI. For example, they suggest that a model from the logistic models family may be the better choice for modeling memory and fluency building processes because of the gradual change in the knowledge state during these processes. Similarly, they argue the BKT would be more fitting for modeling understanding and sense-making processes as BKT models a transition from unknown to known state. Pelanek (Pelánek 2017) used the KLI (Knowledge-Learning-Instruction) framework because it facilitates a direct connection between learning processes and student models. This is achieved by defining both processes and models at the same level, specifically in terms of knowledge components or skills.

Similarly, our work aims to fill the gap between student modeling and the underlying mechanisms of learning processes, in particular the mechanisms that support inductive reasoning. These mechanisms involve collecting information, generalization and categorization/chunking of that information, then recalling them in appropriate situations. Prior research provides evidence that inductive reasoning methods lead to better learning, enhanced problem-solving performance, and gaining expertise; thus, there is a need for further exploration and support of these mechanisms (Haverty et al. 2000).

Our work contributes to this purpose as we are drawing a parallel between the underlying mechanisms of knowledge/skill acquisition and the phases of inductive reasoning as defined in a rule learning paradigm within cognitive science research (Crescentini et al. 2011). This approach allows us to explore these mechanisms at a finer grain than how they are typically studied within educational contexts, where cognitive mechanisms are traditionally defined at a higher level of abstraction. This finer grain of analysis enables us to explore cognitive processes at a level that aligns with the granularity found in cognitive science. By operating at this level, we can map our observations to the level of individual knowledge components, maintaining continuity with traditional approaches, while also making a direct connection with cognitive science, leveraging insights from the field to inform our analysis. We believe detecting these mechanisms will provide a deeper insight into students' behaviors and individual needs and eventually allows us to design better tutoring systems. Moreover, this approach holds the potential to promote more interdisciplinary research by creating a bridge between cognitive science and educational data mining.

2.3 Rule learning

Inductive reasoning is an essential component of learning. It includes collecting instances of some phenomenon and creating inference rules by recognizing such instances, detecting the similarities and differences of them, understanding the associations between them, and generalizing the learned associations on new instances. The steps of this process are conceptualized as the phases of rule learning.

Within cognitive science, the way humans go through these phases has been studied using variations in a rule attainment task (Burgess and Shallice 1996; Crescentini et al. 2011). These tasks aim to detect three main phases of rule learning: rule search, rule discovery, and rule following. To do this, subjects respond to a series of stimuli making one response if they believe that the stimuli are following a rule and a different response if they are unsure of the rule or think the stimuli are changing at random. For example, in one version, subjects are given a series of numbers on a computer screen and rules are represented as a simple mathematical pattern (e.g., +2 resulting in stimuli 32, 34, 36…) (Li et al. 2012). In another version, subjects see a circular array of small circles with only one of them being blue. Similar to the first version described, subjects are asked to give a certain response if they notice the position of the blue circle changing based on a rule (e.g., counterclockwise 2) (Cao et al. 2016). Based on clearly defined behavioral indicators during the task, phases of rule learning can be identified, regardless of the stimuli presentation. This provides a standardized framework for analyzing learning processes that allows for identification of these states across different contexts. Specifically, the phases of rule learning were identified as follows:

Rule search The first response with a new rule and all responses preceding rule discovery.

Rule discovery The first response in which the participant indicates they know the rule.
Rule following The streak of correct responses after rule discovery.
Rule violation An incorrect response after rule following.

While most studies using this paradigm employ abstract tasks, our work aims to identify these phases in a realistic setting. Our work aims to understand if we can detect the mentioned phases of rule learning during mathematical problem solving. Problem-solving tasks can provide rich data on the learning process, allowing for a more detailed analysis of the different phases of rule learning. By using a problem-solving task, we keep a controlled and standardized setting for studying rule learning that also allows exploring it in a realistic learning scenario and making it easier to compare results across different studies and educational contexts. This could provide insights into the cognitive processes involved in learning, potentially leading to the development of more advanced educational technologies.

3 Problem-solving task design

Our task design aims to mirror the structure of rule learning tasks commonly used in cognitive science research, where participants respond to number sequences associated with specific rules (Crescentini et al. 2011). As our goal was to explore rule learning in a real-world learning context, we designed a probability problem-solving scenario, where each problem consists of multiple steps, with each step associated with a distinct rule. These rules correspond to the underlying skills or concepts required to solve the problem step effectively.

We designed our task within the online tutoring system, ASSISTments (Heffernan and Heffernan 2014). ASSISTments is an online platform that allows teachers to compose problems with hints, solutions, and interactive materials such as images and videos that can support problem solving while also providing timely feedback to the students and diagnostic data about students' performance to teachers.

We built a problem set consisting of 9 probability problems on ASSISTments. Each problem was divided into 3 to 4 substeps. Participants were first presented with the full problem and asked to report how confident they are in solving the given problem on a scale of 1 (not confident) to 5 (very confident). To move on and actually solve the problem, participants click the Next Problem button and solve the substeps of the full problem one by one until they reach the final step that leads them to the solution of the problem. As part of the task design, we included a prompt below the problem texts to guide participants toward focusing their attention in different ways that encourage two distinct ways of using cognitive control. This was done for the purposes of other research (Unal et al. 2020, 2023) that is exploring cognitive control in problem solving with an interactive learning environment (see Fig. 1 for a sample problem.).

Each substep of a problem was associated with one "rule." These rules represent the skills or formulas required for solving the given problems and are conceptually similar to knowledge components as both knowledge components and rules can be understood as discrete units of knowledge that need to be acquired or mastered during the learning process (Aleven and Koedinger 2013; Corbett and Anderson 1994). We

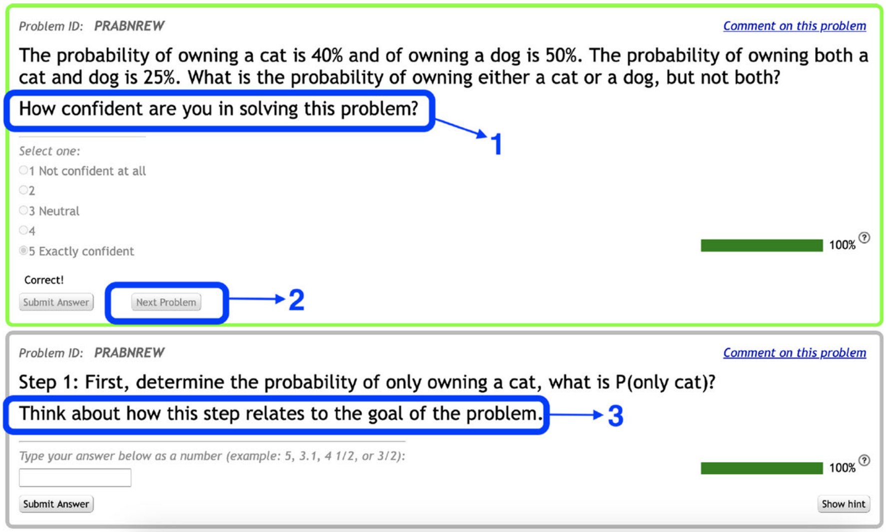

Fig. 1 Screenshot of a sample problem. 1: Participants are asked to rate their confidence level. 2: When this button is clicked, the second box with the first problem step unfolds. 3: Prompt to bias attention toward the current step of the problem

had seven different rules associated with the problem steps throughout the problem set: 1) identifying the number of favorable outcomes, 2) identifying the number of possible outcomes, 3) calculating probability of an event, 4) finding the probability of only one event when two events can occur together, 5) finding the probability of either of the two events happening but not both, 6) finding the conditional probability of an event, 7) finding the probability of dependent events. We aim to understand how the students move through the phases of rule learning for these rules as they are solving the problem steps associated with these rules. For example, for the problem given in Fig. 1, there are three substeps (one is shown for demonstration purposes). The first two substeps are associated with rule 4, finding the probability of only one event when two events can occur together ($P(\text{only } A) = P(A) - P(A \cap B)$). The third substep of that problem is associated with rule 5, finding the probability of either of the two events happening but not both ($P(A \text{ or } B \text{ not both}) = P(\text{only } A) + P(\text{only } B)$).

4 Corpus

Using the task design we described in Sect. 3, we collected student data over two studies. Study 1 aims to identify what the phases of rule learning look like in a realistic learning setting rather than a laboratory-based analog of a learning environment. Given the potential variance in behavioral indicators compared to the original rule learning paradigm, the utilization of a thinkaloud protocol serves a crucial role in validating the rule learning phases through verbalization of thought processes during task performance (Ericsson and Simon 1980). We code the thinkaloud audio data to serve as the ground truth of the rule learning phases during mathematical problem solving. Study 2 replicates the task design of Study 1, with specific adjustments to the experiment protocol including the omission of the thinkaloud method in order to collect more natural behavioral data on the phases of rule learning and minor changes to the problem set to increase the difficulty to observe a wider range of patterns/configurations of rule learning phases.

4.1 Procedures

We recruited undergraduate students from a university in the Northeastern USA. Our recruitment methods included sending emails to student mailing lists and posting flyers around the university campuses. We recruited students who had completed at most two university-level math courses to prevent the inclusion of participants with extensive mathematical proficiency.

Participants came in for an hour-long session and they completed 6 tasks after providing written consent. (1) First, participants completed a tutorial on using ASSISTments. Using a mock problem, we explained the interface and how they would submit their answers and request hints if they like. (2) After the tutorial, participants took a pre-test that consisted of six probability problems on ASSISTments. The structure of the pre-test questions mirrored that of the main problem set, as they were designed to assess students' proficiency in the same skills/rules. The pre-test included two problems on calculating basic probabilities, two on the probability of dependent events, and two on the probability of non-mutually exclusive events. Participants solved these problems as they were solving a regular test, meaning the hints and the feedback provided by ASSISTments were not available. (3) After the pre-test, participants did another tutorial following the format outlined in our task design, which consisted of problems divided into problem steps. (4) After this tutorial, participants solved 9 problems in the format they had practiced in the previous tutorial. (5) After this session, participants took a post-test that was isomorphic to the pre-test. (6) Finally, they completed a demographic questionnaire. Participants were allowed to use a pen and a scratch paper as well as the calculator on the computer. We used a screen recording tool with audio to record students' screen while using ASSIStments.

Study 1 was run in person and included the thinkaloud protocol to collect ground truth data on the phases of rule learning during problem solving. Participants were given time to practice thinking out loud during the third task, and they did the fourth task (the main problem-solving task) thinking out loud. The screen recording tool also captured the audio data of thinkalouds provided by the participants in this study. A researcher was present in the room to assist and answer questions from the participant.

Study 2 was run online without the thinkaloud protocol in order to collect more natural behavioral data on the phases of rule learning. We used Zoom to communicate with the participants. Participant shared their screen as they were performing the tasks quietly on ASSISTments to allow us to observe and record their actions.

4.2 Data

In Study 1, we collected data from 20 undergraduate students (6 male, 14 female) between 18 and 23 years old ($M = 19.45, SD = 1.27$). Each student solved 28 problem steps. One student's data were removed from the dataset as they performed 100% on the pre-test. We had 532 data points from the remaining 19 students and 28 problem steps at the end. Each data point represents a student's actions and response times on a problem step.

In Study 2, we collected data from 56 students (14 male, 40 female, 2 non-binary) without the thinkaloud protocol. Similar to Study 1, students solved 28 problem steps, which largely mirrored those in the previous study, with minor substitutions of certain steps as previously noted. In addition to the problem steps, they responded to self-report prompts after each problem step to indicate their knowledge on the corresponding problem steps. The responses to these prompts served as the ground truth for the phases of rule learning in this study, in the absence of the thinkaloud data. Four students were removed from the dataset due to performing 100% correctly in the pre-test. In total, this dataset had 2912 data points.

5 RQ 1: identifying the phases of rule learning

In this section, we outline our methodology for addressing "RQ 1: What do the phases of rule learning look like within a complex learning task?" To answer this question, we conducted two studies. In Study 1, we used a thinkaloud protocol to capture rich, qualitative data on students' thought processes as they engaged in problem solving, allowing us to directly observe and code instances of rule learning phases. This approach let us investigate whether the rule learning phases defined in cognitive science adequately describe what happens during a real-world task and assess if these phases are experienced in the same way. As part of understanding rule learning phases in a realistic learning setting, we also explored whether there is a direct relationship between any of these phases and learning gains. Since the thinkaloud protocol could interfere with the natural progression of participants' problem solving, in Study 2, we removed the thinkaloud protocol to collect more naturalistic data. We aimed to observe how participants progressed through the phases without external prompting, allowing us to assess if the rule learning phases still emerged similarly and to confirm whether we could identify the same states in a more authentic problem-solving context.

5.1 Study 1: eliciting behaviors related to the phases of rule learning during problem solving with thinkaloud data

This study elicited in-depth information by asking participants to verbalize their thoughts while performing a problem-solving task in a realistic learning environment. We identified the phases of rule learning by coding the thinkaloud data we received. In the coding scheme, we had four different labels (*rule search*, *rule discovery*, *rule following* and *rule violation*), corresponding to the phases of rule learning as defined

Table 1 Coding scheme labels and descriptions

Label	Description	Example
Rule search	Substeps where participants clearly indicate they are not sure about the solution or they do not know how to solve the problem. The participant is simply guessing the answer or trying to figure out the right way to solve it	"I don't quite remember how to solve this but I'm going to try multiplying them before I take the hint."
Rule discovery	Substeps where the participant has just discovered a rule	"Oh! So, we add the probability of just A and just B."
Rule following	Substeps where the participant explains how they got to an answer and gives their response without hesitation. "Rule following" instances are representing participants who have correct prior knowledge on the rule that is associated with the problem, in other words, participants who followed the correct rule were assigned this label	"We multiply the two probabilities together and that is 5/44."
Rule violation	Substeps where the participant gave an incorrect response to a previously acquired rule due to math errors	"Oh I read the wrong thing."
Follow wrong	A special case of "rule following." Participants who have misconceptions about the rule would still verbalize their answer without hesitation. In this case, they would follow a wrong rule they think is correct	"We multiply the two probabilities together and that is 0.12." when the correct strategy is adding the two probabilities and the correct answer is 0.7

in formal rule learning tasks (Crescentini et al. 2011) that we described in Sect. 2.3. In addition to those, we added a special case of the rule following phase that would represent prior incorrect knowledge on a rule (*follow wrong*). Unlike formal rule learning tasks, where rules are presented randomly and learned solely during the task, participants may have pre-existing misconceptions about the mathematical problem solving rules presented to them. The labels in our coding scheme and their descriptions with examples are found in Table 1.

We annotated the thinkaloud data at the problem step level (i.e. each substep of the problems in the task), where each student utterance was assigned one or more of the labels corresponding to the different phases of rule learning that we had defined.

In both the formal rule learning task from cognitive science and our real-world problem-solving task, participants respond to stimuli displayed on the screen, each associated with a rule. While participants encounter repeated instances of these stimuli, the manner of experiencing the phases differs between tasks. In the formal rule learning task, each response to a stimulus corresponds to a single rule learning phase, meaning each stimuli can be labeled as one phase. Conversely, in our problem-solving task, which involves more complex stimuli (i.e., problem steps vs. a single letter or spatial location), transitions between rule learning phases can occur within the response time to a single stimulus. In the problem-solving task, participants may start from a rule search phase for a given problem step but can transition to discovering the rule by carefully analyzing the information provided in the step or seeking a hint, without needing to see the next problem step associated with the same rule. This contrasts with the formal rule learning task, where participants must encounter multiple instances of stimuli associated with the same rule to reach rule discovery. Therefore, we observe patterns of rule learning phases, in addition to single phases per problem step as participants may transition between phases within a single problem step.

In Study 1 data, we observed 9 unique patterns or configurations of phases at the problem step level. These could either be sequences of phases or a single phase experienced during solving the corresponding problem step. Four of these configurations occurred more frequently than the others and were: *rule following*, *rule search–rule discovery* (searching and discovering a rule), *follow wrong–rule search–rule discovery* (following a wrong rule, searching for the correct rule then discovering), and *rule violation–rule following* (violating a previously learned rule then following). Less frequent combinations shared similarities with these predominant patterns. We revisited the thinkaloud data to investigate whether the frequent and scarce patterns represented distinct phases. We found that students were progressing through the same phases as the major patterns; however, their articulation of certain phases was less explicit. Thus, we viewed the minor patterns as truncated expressions of the main ones, where students exhibited the same cognitive processes but in a more limited or less verbalized way. For example, one of the main patterns includes "follow wrong, rule search, rule discovery." However, we observed a minor pattern that only showed "follow wrong." Given that students needed to eventually discover the correct rule in order to move on to the next problem in our problem-solving task, this minor pattern is essentially an incomplete articulation of the full process. Behavioral data confirmed that students did go through "rule discovery," even if they didn't explicitly articulate it during the thinkalouds. This was evidenced by the fact that students not only found the correct answer but also consistently applied the same rule in subsequent occurrences of similar problems. The four predominant groups were used as overarching labels for the whole data set. Accordingly, data points with less frequent combinations were labeled based on their resemblance to the predominant patterns. Figure 2 illustrates the frequencies of each pattern. The results of our data coding efforts showed that in addition to the "rule search–rule discovery–rule following" sequence that is typical in the standardized task (Crescentini et al. 2011), participants exhibited different combinations of the rule learning phases. Participants who already possess familiarity with a particular rule associated with a problem step in our study may consistently remain in the rule following phase throughout all problem-solving steps linked to that specific rule. Furthermore, we have observed instances where participants initially hold incorrect prior knowledge of certain rules, subsequently transitioning from following an incorrect rule to searching for the correct one, and ultimately experiencing the phases of rule discovery and rule following during later encounters with problems linked to that particular rule.

Moving forward, we explore the relationship between the time spent on the phases of rule learning and learning gains. To measure learning gains, we took the difference between the proportion of correct answers in the post-test and the proportion of correct answers in the pre-test. We extracted the time spent on corresponding utterances in

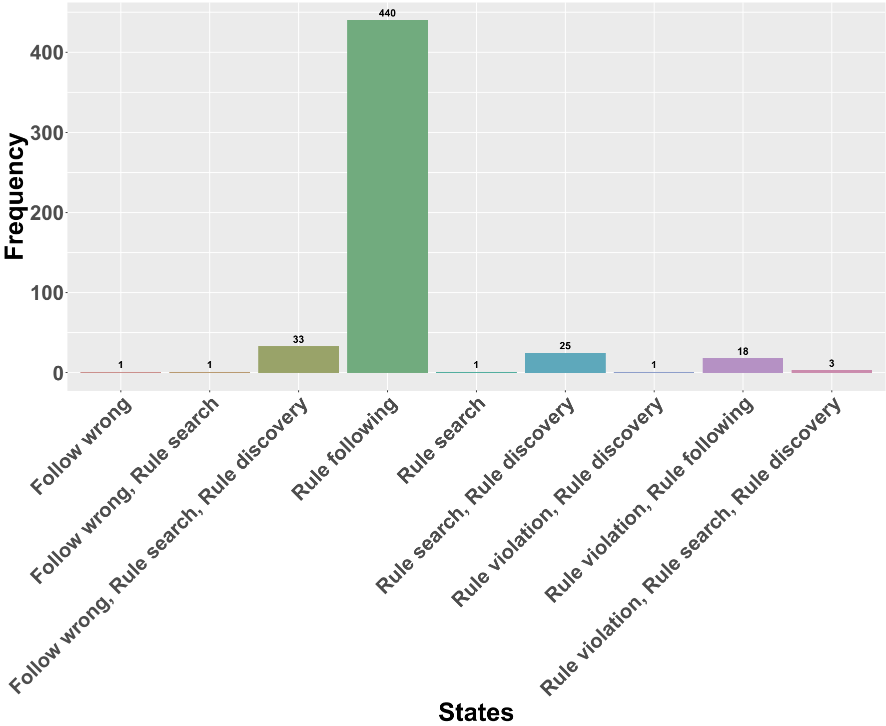

Fig. 2 Unique configurations of rule learning phases observed at the problem step level

the thinkalouds for the time spent on individual rule learning phases. As time spent on the rule learning phases may change based on external factors such as individual differences or the problem step on hand, we converted the data into proportions of time on different rule learning phases for each participant to have a more uniform measure. Figure 3 illustrates the relationship between the proportional time spent on the rule learning phases with learning gain. Zero proportions were excluded as they indicate that the student did not experience that cognitive state, and including them would skew the analysis by introducing irrelevant data points. This is particularly relevant for sparsely occurring phases like rule violation. If a student never encountered a rule violation, including them in the correlation analysis could skew the results by introducing an artificial baseline, making it difficult to compare with phases that are more frequently experienced. Results indicate the proportion of time spent on rule discovery is significantly correlated with learning gain $R = 0.53$, $p = 0.021$. This result highlights the role of the rule discovery phase in students' learning progress. Although we did not observe significant correlations between learning gains and time spent in other phases, distinguishing between these phases is still valuable. Accurately detecting how long students spend in earlier, consecutive phases such as "follow wrong" or "rule search" can enhance the prediction of when the rule discovery phase will occur.

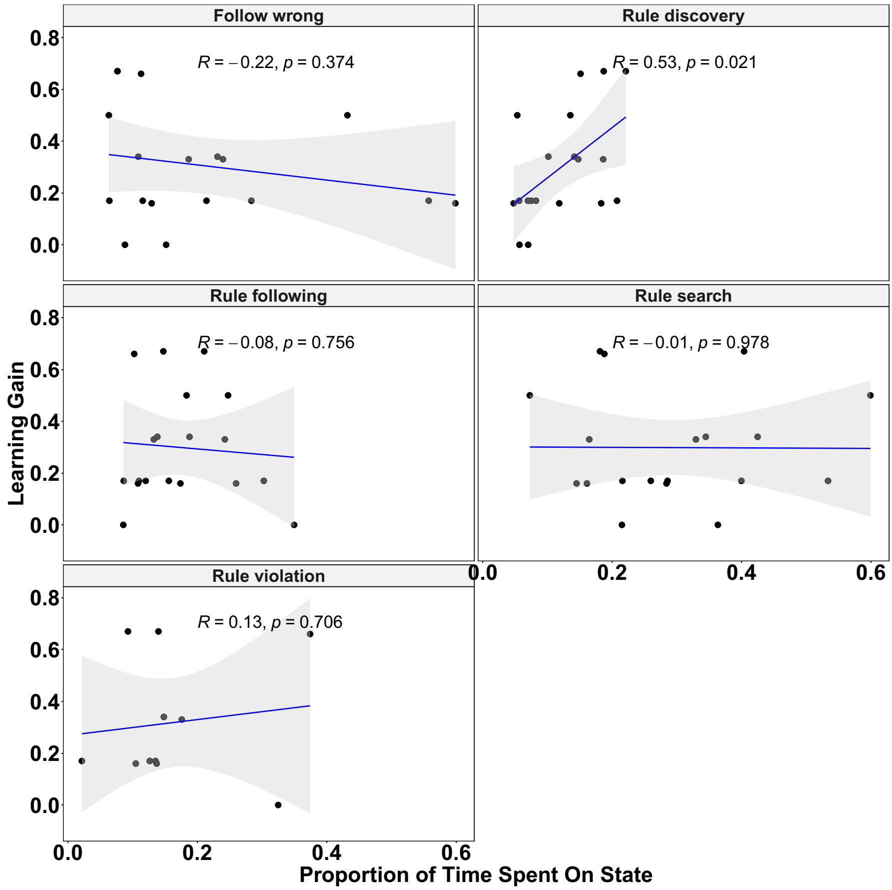

Fig. 3 Correlation between the proportion of time spent on each rule learning phase and learning gains in Study 1. Each point represents an individual student's data

5.2 Study 2: collecting naturalistic behavioral data on rule learning phases

Study 1 revealed insights into rule learning phases in a real-world learning task, using thinkaloud data as a ground truth. A downside of thinkaloud protocols is that they can slow the participants down or disrupt their natural flow (Ericsson and Simon 1980). This can affect task completion times and potentially influence their behavior as they may become more cautious. In Study 2, we replaced the thinkaloud protocol with lightweight self-report before and after the task, in order to collect more natural in-task behavioral data on the phases of rule learning. In order to obtain a more varied distribution of the rule learning phases and to gather ground truth data for the rule learning phases in the absence of thinkaloud data, we implemented two modifications to the primary task. First, we exchanged the first two problems in the problem set for two harder problems. We noticed rule following is the most common phase throughout Study 1. We expected harder problems would bring more instances of rule search and discovery behaviors. Our second change was to move from thinkaloud to self-report prompts after each problem step for participants to indicate their knowledge on the problem steps they passed (see Fig. 4). Participants were asked to type all the options that applied to them.

Each option was designed to represent a rule learning phase. We came up with the options by extracting the most common statements from the participant thinkalouds for each rule learning phase in Study 1. Options 1 to 5 indicate rule following, rule search, rule discovery, following the wrong rule, and rule violation phases, respectively. We integrated the prompt menu directly into the ASSISTments environment (i.e., by including them as "problems" in the problem set), so it appeared automatically after each problem was solved, which helped streamline the data collection process. We did not observe any clear evidence that the prompts interfered with students' progression through the experiment. Since the prompts were introduced during the practice sessions, students became accustomed to them early on and generally used the prompt menu as intended.

We used participant responses to these questions as the ground truth of the rule learning phases they went through on the problem step. Figure 5 illustrates the distribution of the most frequent participant responses to this prompt. We first checked whether the modification of problems changed the rule learning phases observed. The inclusion of two harder problems at the start seems successful as we see more responses that correspond to the rule learning phases that were observed less frequently in Study 1. However, rule following was still the most frequent phase similar to what we found in our first study. Furthermore, the combination of options that were chosen by the participants aligned with the sequences of rule learning phases for the most part.

One distinction that we noticed is that participants did not always choose all the options reflecting the phases that occurred together in our first study. For example, a participant who entered a rule search phase would also experience rule discovery at the same problem step as they need to find the correct answer to the problem step before moving on with the task according to our task design, but in the second study we see some participants chose options 2 (rule search) and 3 (rule discovery) exclusively. This could be due to how participants perceived the options differently. Some participants might be more inclined to break down their thought processes into distinct phases, while others may be more inclined to perceive rule search and rule discovery as intertwined. Moreover, some combinations of options show participants may have had different interpretations of the options than what we intended. For example, including option 5 (rule violation) in one's selection probably indicates different cognitive states when we compare two combinations such as 1 (rule following) and 5 (rule violation), and 4 (following the wrong rule) and 5 (rule violation). Choosing options 1 (rule following) and 5 (rule violation) would be identical to choosing option 5 (rule violation) exclusively, indicating an occasional slip of a rule. The nuance is that the participant wanted to emphasize that they actually knew the rule when they chose option 1 (rule following), too. On the other hand, if a participant chose option 5 (rule violation) with option 4 (following the wrong rule) they would acknowledge that they made a mistake without indicating it was a slip as option 4 (following the wrong rule) means they did not expect their response to be incorrect.

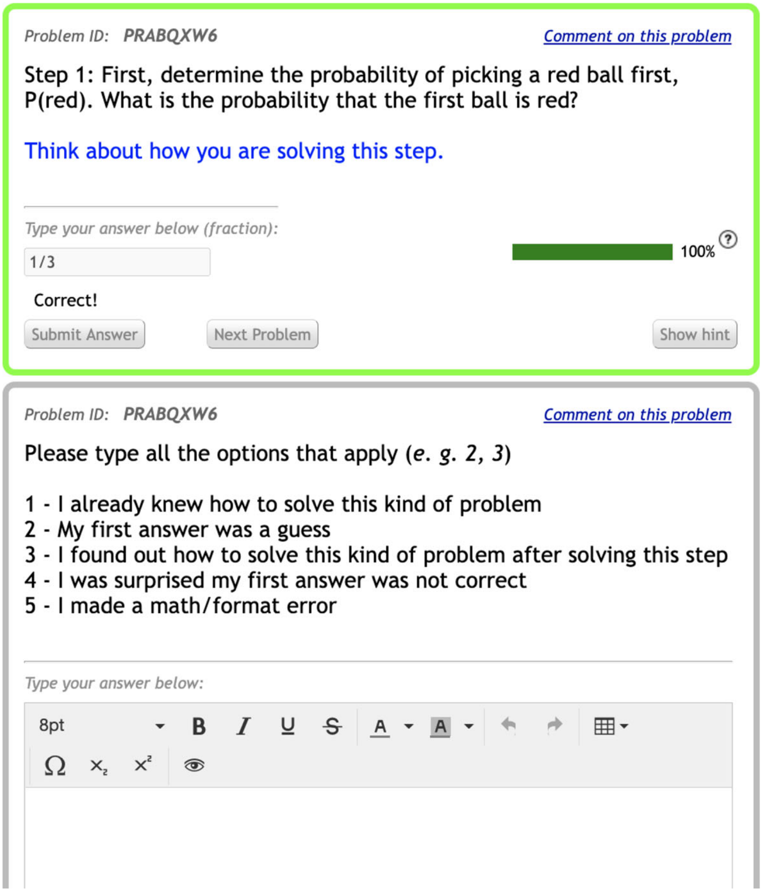

Fig. 4 Students respond to self-report knowledge indication prompts after each substep of a problem

Even though the participant interpretations of the options may slightly vary, the distribution of phases mostly aligns with our findings from our first study, so we use the same 4 overarching labels, *rule following*, *rule search–rule discovery* (searching and discovering a rule), *follow wrong–rule search–rule discovery* (following a wrong rule, searching for the correct rule then discovering), and *rule violation–rule following*, to label the self-report responses for the rule learning phases.

In response to RQ 1, our findings indicate that the traditional phases of rule learning as defined in cognitive science are not sufficient to capture the complexity of students' cognitive processes in a realistic learning task. Unlike controlled lab settings, students in our study exhibited rule learning phases in different sequences, with some transitions occurring in unexpected orders that are not described in the cognitive science literature. For instance, we observed patterns where students moved from following an incorrect rule to rule search and then to rule discovery, as well as cases where they remained in the rule following phase throughout the task. These findings highlight the need for an expanded framework to accurately represent rule learning phases in real-world tasks.

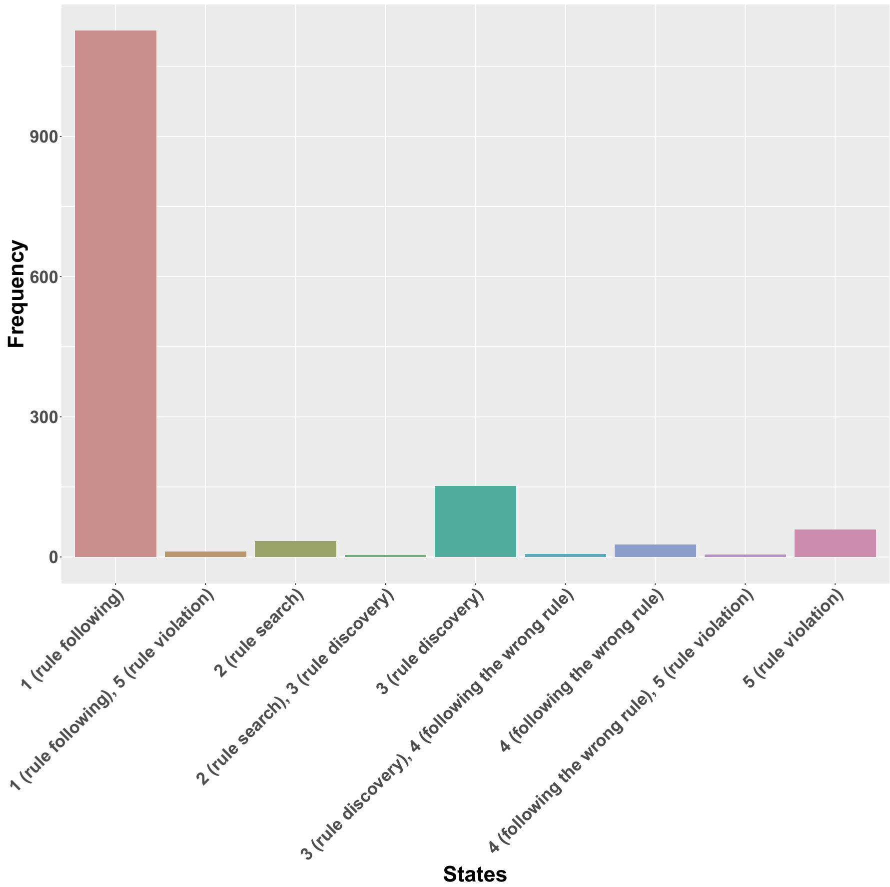

Fig. 5 Distribution of participant responses that reflect the rule learning phase they experience in Study 2. Numbers in the x-axis represent the combinations of options selected by the participants in the self-report prompts (see Fig. 4). The combination of options does not indicate a sequence but rather a co-occurrence. Some combination of responses that occurred less than 4 times were excluded from this plot

The thinkaloud protocol used in Study 1 provided essential insights, helping to establish a coding scheme that identifies verbal cues associated with each phase. However, it also introduced limitations by potentially altering task timing and encouraging students to verbalize even when they may have otherwise paused. The predominance of rule following in our data further suggests that the problems may have been somewhat straightforward, potentially limiting the presence of more exploratory phases. The verbalizations captured in Study 1 can inform a standardized approach for future research, such as using multiple-choice options for self-reporting cognitive states. This informed the design of Study 2, where we replaced the thinkaloud component to capture more naturalistic data and assess the applicability of rule learning phases in an environment closer to real-world learning.

6 RQ 2: predicting phases of rule learning on assistments data

To answer "RQ 2: Can we predict what phase of rule learning a student is in during a problem-solving activity?" we demonstrate our methods to predict the phases of rule learning for a given student and problem step. We present a two-step approach to achieve this goal. The first step is to assign one of the 4 labels to each data point that consists of a participants' actions and response times. The labels are the different combinations of phases that we identified at the problem step level in the previous section: *rule following* (RF), *rule search–rule discovery* (RS-RD), *follow wrong–rule search–rule discovery* (FW-RS-RD), and *rule violation–rule following* (RV-RF). We show the results of two different methods to solve this classification problem. As our second step, to identify when each individual phase occurs within the predicted labels, we build a model to estimate the time spent on each individual rule learning phase. We train and evaluate all models on Study 1 data and we hold out Study 2 data from all training and use it as a test set as we do not have the ground truth for the second step of our modeling pipeline (i.e., the exact time spent on individual rule learning phases) for this dataset. This absence of ground truth in Study 2 stems from the nature of the data collection process. Unlike Study 1, Study 2 did not incorporate thinkaloud protocols. In Study 1, we obtained ground truth data for the timing of individual phases from the thinkaloud data, where the duration of students' utterances corresponding to specific phases served as a proxy for the time spent on each phase. However, since Study 2 did not include thinkaloud protocols, we do not have comparable ground truth data for this aspect of our modeling pipeline in that dataset. This limits our ability to calculate prediction error for the second step of the model. However, we instead present correlations between model predictions and learning gains in Study 2 and compare these to the corresponding correlations from Study 1, where ground truth data are available.

6.1 Rule-based classification

We start by examining whether the established identifiers of rule learning phases from cognitive science-developed in controlled, artificial tasks-can accurately determine rule learning phases in a realistic problem-solving environment. The conventional definitions for rule learning phases are well established and linked to response correctness and rule order (Crescentini et al. 2011). We test the applicability of this model to real-world learning contexts by implementing a simple rule-based classifier aligned with these standard definitions. This approach allows us to assess if these predefined cognitive phases are sufficient for predicting rule learning in a complex setting, especially given our findings that the sequence and nature of rule learning phases can differ in real-world tasks. The rule-based classifier we implemented operates by the following rules: If the rule associated with the problem step was seen for the first time and a correct response was given, then the student is assigned a RF (rule following) label for that problem step. For subsequent occurrences of the same rule, if a correct response is given, the student is assigned a RF label again. In the case of an incorrect response, we first check the correctness of the response given to the previous problem step asso-

Table 2 Confusion matrix of the rule-based classifier

	Predicted Label				
		RF	RV-RF	RS-RD	FW-RS-RD
Actual Label	RF	440	3	0	0
	RV-RF	0	17	5	0
	RS-RD	3	0	23	0
	FW-RS-RD	0	4	31	0

ciated with the same rule. If the previous response was correct, we assume that the student was previously in a rule following state and that the current state represents a violation of the rule (RV-RF). Finally, if two consecutive responses are incorrect at first try for problem steps associated with the same rule, the student is assigned an RS-RD (rule search then rule discovery) label (rule discovery was included as the participant needs to find the correct answer eventually in order to move to the next problem step). The algorithm for this method is as follows:

```
if correct response then
    label problem step as RF
end if
if incorrect response then
    if previous response to same rule was correct then
        label problem step as RV-RF
    else
        label problem step as RS-RD
    end if
end if
```

The confusion matrix of predicted labels from this model and actual labels from the thinkaloud data coding results is given in Table 2.

This model is easy to implement and it predicts three labels associated with the rule learning phases in the literature well without needing an additional data source. However, since the model is based on the correctness of the responses, it fails to capture the phase when the correctness of the response was not indicative of the actual rule learning phase. Specifically, when we answered our first research question, we discovered a pattern that is not defined in the rule learning paradigm. We found that participants who give incorrect responses to problem steps are not always in rule search phase. They may be in a rule following phase where they follow a wrong rule that they believe to be correct. We named this specific phase "follow wrong" indicating students who have incorrect prior knowledge on a rule associated with a problem step. In the rule-based model we present in this section, we do not have a rule for labeling the follow wrong phase as the correctness of the student's response is not sufficient and human insight into student misconceptions or a more sophisticated data-driven approach is necessary to detect these subtler, less-defined cognitive states.

6.2 Sequence modeling with hidden Markov models

In this section, we explore a more advanced approach to address the limitations of the previous rule-based model and to improve our ability to predict subtler cognitive states. We introduce a hidden Markov model (HMM) to improve our prediction of rule learning phases, specifically the phases that are not easily identifiable based solely on response correctness.

An HMM is a statistical model that represents systems with hidden (unobserved) states, where the system transitions between these states based on certain probabilities. HMMs are commonly used in modeling sequences of observations, making it a natural choice for our problem of modeling sequences of problem steps. Moreover, the hidden states in an HMM can be interpreted as latent cognitive states, which we hypothesize will map to the phases of rule learning.

We divide the dataset into two groups "follow" (sequences with rule following behaviors with occasional rule violation) and "acquisition" (sequences where rule search has been experienced at least once) based on observed patterns in participants' behaviors. This grouping approach emerged from an exploratory analysis, where we experimented with different data divisions to identify the best fit for our modeling task. Initial trials included training a single model on the full dataset and dividing the data into three groups to capture diverse behavior patterns, but cross-validation results indicated that dividing the data into two groups provided the most robust model performance.

We then trained separate models to predict the four labels, we associated with the phases of rule learning. For the "follow" group, we trained a model tailored to sequences where students mainly engaged in rule-following behaviors, with occasional rule violations, as rule search phases were infrequent and thus less relevant for this subset. Similarly, the "acquisition" group consists of the student rule pairs where the student went through a rule acquisition process (i.e., experiencing rule search and rule discovery phases) in the sequence of problem steps associated with the given rule. Within this setting, rule violations followed by rule following are not prevalent as much as the acquisition phases.

In these models, the problem steps were represented with a set of features such as behavioral indicators like whether the student answered correctly, time spent on the step, first response time, attempt count, first action, as well as contextual features like problem and step number in order to provide information about the difficulty of the problem or how far along the student is. While the features we used are naturally specific to the ASSISTments platform, these features have direct analogs in other learning environments and subject domains. For example, other ITSs such as Cognitive Tutor leverage correctness of responses to adjust task difficulty, while platforms such as ALEKS and Carnegie Learning's MATHia use time spent on a step and attempt counts to personalize student learning paths (Anderson et al. 1995; Fang et al. 2019; Ritter and Fancsali 2016). We used generic features rather than problem set-specific ones to make our approach more generalizable. While problem set-specific information such as categorizing/annotating common errors students make on the problem set could offer additional insight, this would make the model too specific to the current problem

Table 3 Features that are used in the hidden Markov models to predict rule learning phases during problem solving

Feature	Description
Correct	Whether the participant got the problem step correct at first try
Time spent on step	Normalized time spent on the problem step
First response time	Normalized time spent until the first action of the participant
Attempt count	The number of times the participant attempted the problem step
First action	Whether the first action of the participant was an attempt or hint request
Problem	Problem number
Problem step	Problem step number

set. Moreover, categorizing this kind of information would require substantial human effort. Therefore, we used HMMs with more generic features and we hypothesize their natural ability to capture the sequential dynamics of observations will allow us to track how students transition between different cognitive states, including more difficult-to-detect ones like misconceptions (e.g., "follow wrong"), based on their actions over time. The full list of features and their descriptions are given in Table 3.

For the hidden states in our models, we proposed two states for the "follow" group model, as students in this group typically alternated between "rule following" (RF) and occasional "rule violation–rule following" (RV-RF) states. For the "acquisition" group, we proposed three hidden states, as students could progress through the sequence of "follow wrong–rule search–rule discovery" (FW-RS-RD), "rule search–rule discovery" (RS-RD), or remain in a "rule following" (RF) state. After cross-validating configurations with two to five hidden states, this structure provided the best fit for our data.

To evaluate the models' performance on a test participant and rule pair, the first step is to decide which of the two HMMs to use. We apply both models on the test instance and the model that gives the highest log likelihood for the given test instance determines whether the test instance belongs in the "follow" or "acquisition" groups and the corresponding model is chosen. For example, for a given participant and rule pair, if the HMM that was trained on the "follow" group has a higher log likelihood compared to the other model, the given participant was predicted to be in the "follow" group for the given rule. We evaluated this classification scheme under tenfold cross-validation with random over-sampling for the minority class. We achieved 84% and 89% classification accuracy for the "follow" and "acquisition" groups, respectively. The next step is to interpret the hidden states of both models. We print the frequency counts of hidden states and the rule learning sequences coded in the thinkalouds in contingency tables to understand how well hidden states estimated from HMMs could determine the actual rule learning phases (see Tables 4, 5).

Table 4 shows S_1 and S_2 predominantly map to RV-RF and RF states. Similarly, Table 5 illustrates the rule learning state and hidden state mapping for the participants classified as belonging in acquisition group. We see a clear mapping for RF (S_2)

Table 4 Mapping of observed rule learning states to hidden states for the follow group

Observed Rule Learning States	S_1	S_2
Rule following (RF)	2	290
Rule violation, rule following (RV-RF)	14	0
Follow wrong, rule search, rule discovery (FW-RS-RD)	2	0

This table illustrates how the hidden states S_1 and S_2 capture the behavior of students primarily in the rule following (RF) phase, with occasional rule violation–rule following (RV-RF) instances. The two-state model is appropriate here as this group rarely transitions through other rule learning phases. The distribution of observations across S_1 and S_2 reflects the predominance of rule following behavior in this group

Table 5 Mapping of observed rule learning states to hidden states for the acquisition group

Observed Rule Learning States	S_1	S_2	S_3
Rule following (RF)	1	150	0
Rule violation, rule following (RV-RF)	8	0	0
Follow wrong, rule search, rule discovery (FW-RS-RD)	33	0	0
Rule search, rule discovery (RS-RD)	13	3	10

This table shows how the hidden states S_1, S_2, and S_3 capture the sequences such as follow wrong–rule search–rule discovery (FW-RS-RD), rule search–rule discovery (RS-RD), and other transitions not present in the follow group

and FW-RS-RD (S_1) states. However, even though S_3 seems to map to RS-RD state perfectly, we still have some RS-RD instances in S_1 as well.

When we investigated these instances specifically, we found two different behaviors in RS-RD state. The first group of students indicate that they do not remember how to solve the current problem step but they will still make a guess. Notice this is different than our FW-RS-RD (follow wrong–rule search–rule discovery) state as students in this state respond to the problem step confidently then realize their answer was wrong. In the second group, the participants indicate they do not remember the solution similar to the first group, but instead of making an attempt, they make a hint request. We found that the first group of students was predicted to be in the FW-RS-RD state, although their true state was RS-RD. These states can only be differentiated through the thinkaloud data, posing a challenge for our model in distinguishing between them.

6.3 RQ3: estimating the time spent on the phases of rule learning

Since rule learning phases frequently manifest in sequences/combinations during problem steps in our task, determining the precise timing of each phase becomes essential to understand how each phase individually influences learning. To address "RQ 3: Can we estimate the time spent on the different phases of rule learning during problem solving?" we next estimate the time spent on each individual phase within these combinations to detect when transitions between phases occur."

We build a simple model that learns the proportions of time spent on each individual rule learning phase based on some grouping factors. We use the rule type,

Table 6 Root mean squared errors for time spent on rule learning phases

State	RMSE	Mean	Stdev	Min	Max
Follow wrong	26.31	21.13	28.30	3.36	163.51
Rule discovery	12.05	10.24	9.39	1.17	55.77
Rule following	7.37	13.40	10.22	1.66	106.23
Rule search	22.67	28.24	28.08	1.06	143.81
Rule violation	13.11	13.54	14.46	2.07	57.38

the number of times the corresponding rule was seen, and the cognitive control condition that the participant was assigned. As part of a separate set of analyses, this study included a cognitive control manipulation based on the dual mechanisms of cognitive control framework (Braver 2012). Participants were assigned to one of two conditions: proactive or reactive cognitive control. In the proactive condition, participants were instructed to continuously think about how each step of the problem related to the overall goal they had been presented at the start, promoting active goal maintenance. In contrast, participants in the reactive condition were asked to focus solely on the current problem step, without considering the overall goal, promoting a more just-in-time manner of goal maintenance. Although we hypothesized that the proactive control condition might enhance learning, our analysis found no significant differences in learning gains between the two conditions. However, participants in the proactive condition spent more time on the problem steps on average, likely due to the added cognitive demand of linking each step back to the main goal. We included the condition as part of the feature set for our analysis, as it influenced time spent on problems and the overall student experience. Further details on this manipulation can be found in our previous work (Unal et al. 2020). We included the condition in the grouping factors as the time spent on problem steps differ between the conditions. In the training phase, within each group, we compute the proportional time spent on the seen rule learning phases. To give a prediction on the time spent on the seen phases, we multiply the total time spent on a particular problem step by the learned proportions. The training is performed on Study 1 data. We first evaluated the model under a tenfold cross-validation scheme. Cross-validation was done at the participant level, meaning two participants' data were left out of training in each fold. The results of this model are presented in Table 6. The RMSE of the predictions were very close to the standard deviation of the training data for each rule learning phase. To further evaluate the predictions of the model, we picked two samples that yielded the best and the worst prediction error during cross-validation, and we plotted the actual and predicted values against each other on these cases (Fig. 6). We divided the instances where the predicted time falls between certain percentiles, such as 10-20%, into the same decile. Subsequently, we create a calibration plot that plots the average predicted time against the actual time spent for each decile. We achieved a well-calibrated prediction, reflecting an accurate alignment between predicted and actual outcomes, would result in all deciles closely aligning with a 45-degree line on the calibration plot even for the worst case of a train/test split in our cross-validation experiments.

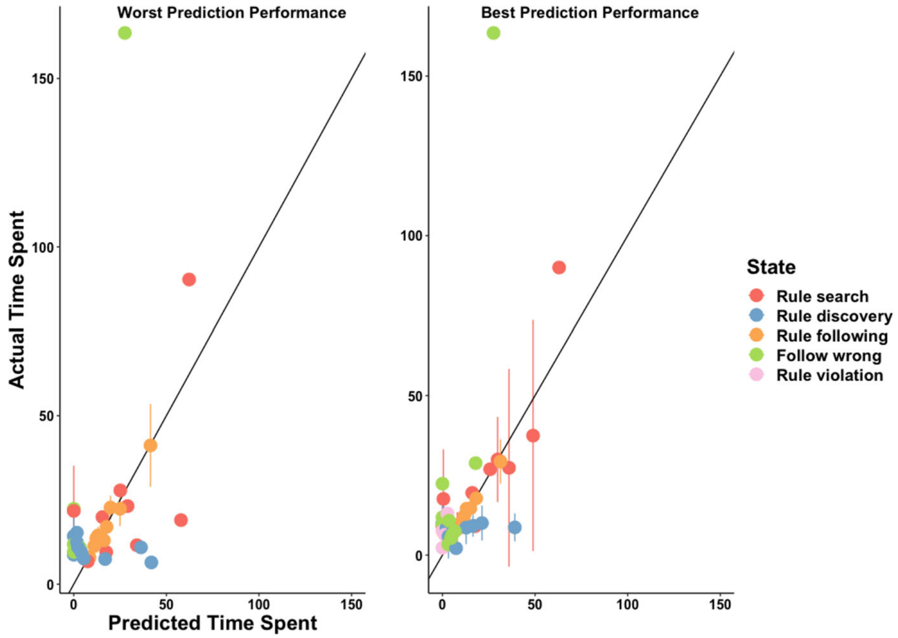

Fig. 6 Predicted and actual values of time spent on the rule learning phases demonstrated using samples that provided the worst and the best prediction performance in cross validation

6.4 Model evaluation on held-out study 2 data

We have built and evaluated our models to predict the phases of rule learning on the data that we collected from our first study. The next step is to test this model on our second study's dataset which we completely held out from all training activities. Firstly, we run the HMMs on this dataset to determine the rule learning phases. Based on the rule learning phases found from the first step, we print the hidden states of the corresponding models and we evaluate the test performance of the models by exploring the alignment between the predicted hidden states and the actual rule learning phases in the test data (i. e., responses to self-report prompts in between problem steps).

Table 7 shows the mapping between the hidden states predicted by the model and the actual rule learning phases for the participants in the "follow" group. Hidden states S_1 and S_2 determine RV-RF, and RF states, respectively. With this model, we achieved a clear mapping between the hidden states and the rule learning phases as we had in the training data.

Our results revealed a similar pattern within the "acquisition" group to what we observed in training as well. Table 8 indicates a clear mapping between S_2 and RF, and S_3 and RS-RD phases. However, considering the prediction of S_1 among the different rule learning phases, we understand that the model has a harder time distinguishing between the participants in the RS-RD phase and the participants in FW-RS-RD phase.

We initiated our two-step pipeline by leveraging the HMMs that we trained on the data from Study 1 to predict the sequences of rule learning phases for each problem

Table 7 Rule learning state and hidden state mapping for follow group on test data

Indicated Rule Learning States	S_1	S_2
1	7	604
2, 3	2	30
4	2	1
5	19	2

1: Rule following, 2,3: rule search–rule discovery, 4: follow wrong–rule search–rule discovery, 5: rule violation–rule following

Table 8 Rule learning state and hidden state mapping for acquisition group on test data

Rule Learning States	S_1	S_2	S_3
1	20	492	3
2, 3	83	51	31
4	31	0	0
5	53	0	2

1: Rule following, 2,3: rule search–rule discovery, 4: follow wrong–rule search–rule discovery, 5: rule violation–rule following

step in the Study 2 dataset. As the next step, in order to explore when individual rule learning phases take place, we estimate the time spent on the phases within the sequences that we predicted in our first step using the model we described in Sect. 6.3. We learn the proportions of time spent from Study 1 data and use them to estimate the individual time spent on the predicted phases of rule learning in Study 2 data. Then, similar to Study 1 data, we convert the predictions of time spent into proportions to have a uniform measure across participants and we explore the relationship between the predicted proportional time spent on the individual rule learning phases and learning gain (Fig. 7). Results were in line with the analyses from Study 1 data, indicating a significant relationship between time spent on rule discovery phase and learning gain. In addition to rule discovery, time spent on rule search also showed a significant positive correlation with learning gain. Moreover, we noticed participants in Study 2 spent much less time on all the phases other than rule following. These phases involve scenarios where students are challenged by prior misconceptions (“follow wrong”), gaps in their knowledge (“rule search”) or slipping an acquired rule (“rule violation”). The decrease in time spent on these phases may be attributed to the changes we introduced in the problem set in Study 2. The inclusion of more challenging problems likely afforded students increased opportunities for practice, and enabled them to achieve rule following phase more quickly, spending less amount of time on rule acquisition-related phases compared to the participants in Study 1.

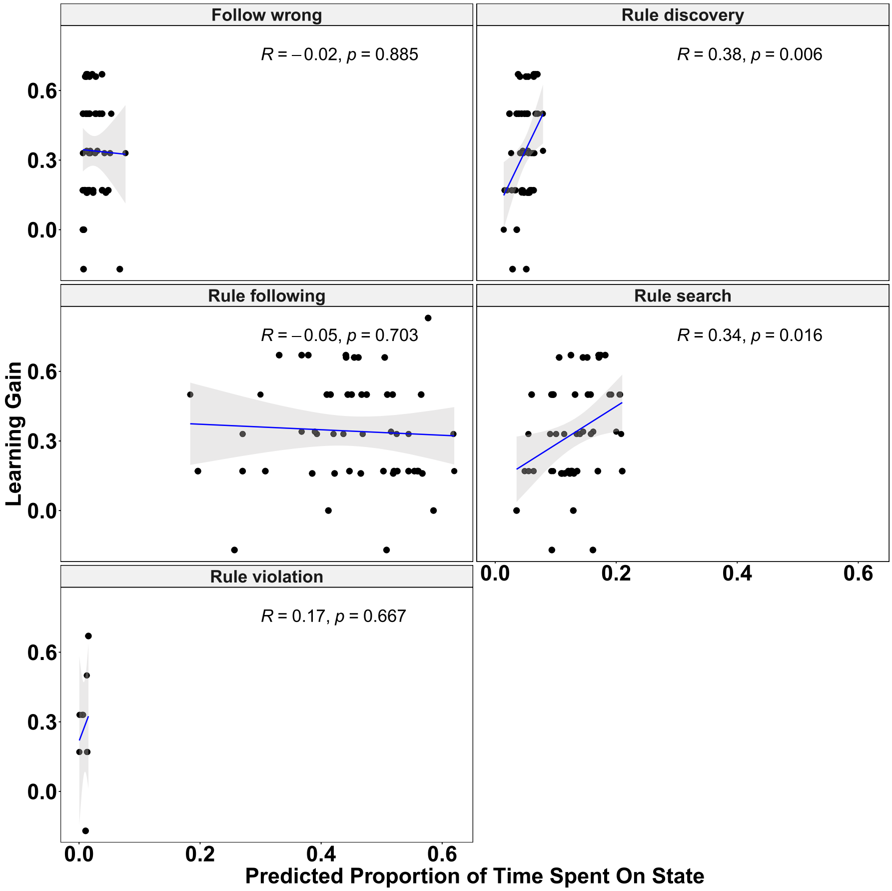

Fig. 7 Correlation between the predicted proportion of time spent on each rule learning phase and learning gains in Study 2. Each point represents an individual student's data. Similar to Study 1, zero predictions were removed to only include data from students who experienced the cognitive phase at least once

7 Discussion

7.1 Summary of findings

While instructional interventions within ITSs are grounded in the cognitive processes that underlie learning, the main focus in student modeling has often been on assessing whether students will answer the next problem correctly, rather than predicting the underlying cognitive mechanisms (Pelánek 2017). In this work, our overarching goal was to demonstrate methods to detect these kind of mechanisms in a realistic learning environment. We used phases of rule learning as defined in a rule learning paradigm (Crescentini et al. 2011), as a proxy for these mechanisms. We made this choice because the phases of rule learning provide an appropriate granularity, specifically at the level of knowledge components and skills, which aligns with the levels of detail typically examined within learning science and intelligent tutoring systems where cognitive states are traditionally defined at a higher level of abstraction (Koedinger et al. 2012).

We presented two studies we conducted to explore the phases of rule learning in a real-world learning context. Study 1 involved participants thinking aloud while solving problems, which were then coded to identify rule learning phases. A key finding that emerged from this study was that students demonstrated diverse combinations of rule learning phases not typically observed in formal rule learning tasks. This included scenarios where participants, already familiar with a specific rule, remained in the rule following phase throughout problem solving, as well as instances where they initially held incorrect prior knowledge, transitioning from following an incorrect rule to seeking the correct one, then experiencing the phases of rule discovery and rule following. These findings enhanced our understanding of how rule learning manifests in real-life scenarios and played a crucial role in structuring our predictive models. While our contributions may not redefine the phases of rule learning, they offer a formalization of additional behaviors or phases that can be expected in real-world learning contexts. In our case these were the "follow wrong, rule search, rule discovery" pattern and remaining in "rule following" phase throughout a problem. This formalization, grounded in experimental observation, provides a roadmap for researchers and designers who aim to incorporate rule learning into their models of student behavior.

With Study 2, we collected more natural behavioral data by modifying the task design by removing the thinkaloud protocol and introducing harder problems. We used participant responses to self-report prompts as ground truth for the phases of rule learning.The prompts were always answered after a problem was solved, so there was no interruption to the problem solving process or to the observation of the phases of rule learning. While it may seem unnatural to have these prompts between problems, students were introduced to this protocol during the practice phase, so it became part of their expected routine. These prompts could potentially replace thinkaloud protocols to some extent. They are certainly less demanding for participants and require less effort to analyze. However, there is a risk of students misinterpreting the options provided. This highlights the need for careful consideration of the clarity of the prompts. To mitigate this, we referred back to the thinkaloud data from Study 1 and used the most frequently mentioned descriptions of students' cognitive processing during different phases of rule learning. Even so, we still encountered instances where students interpreted the options differently than intended. So, while prompts could replace thinkalouds in some contexts, we should still ensure that misinterpretation is minimized.

Participants' responses to the self-report prompts aligned with the rule learning phases observed in Study 1, although some variations in the interpretation of phase options were noted. Overall, the distribution of phases in Study 2 largely resembled the findings from Study 1, leading to the adoption of the same labels for rule learning phases that we later use for prediction. One key difference that we observed was that students in Study 2 tended to spend less time on the tasks. Study 2 was conducted online due to COVID-19, and one might assume that the observed reduction in time spent on tasks was caused by frustration from extensive online learning. However, it is more likely attributable to the removal of the thinkaloud protocol, as prior research indicates that engaging in thinking out loud while performing a complex task can influence task completion time (Ericsson and Simon 1980; Gill and Nonnecke 2012). In contrast, the informal feedback from participants was generally positive, with many indicating that the problem-solving activity helped them learn or recall the topics.

We introduced a two-step approach in which we leverage HMMs to predict sequences of rule learning phases that students experienced while tackling complex problems. Even though we could predict the usual phases of rule learning by a simple rule-based classifier based on the definitions of these states in the literature (Crescentini et al. 2011), using HMMs allowed us to give predictions for the unusual states that we discovered during a realistic learning task. The hidden states, as revealed by the HMMs, aligned well with the observed sequences of rule learning phases. Specifically using this approach we provided reasonably accurate predictions for participants who initially followed incorrect rules, shedding light on our capacity to detect student misconceptions. Identifying and resolving these misconceptions that are caused by faulty rules or their misapplications is crucial for effectively supporting learning (Zeller and Schmid 2016). While our simple rule-based model could potentially be extended to address student misconceptions through methods like incorporating an error library with a set of hard-coded faulty rules that reflect common misconceptions or errors that students may make on problems that are associated with particular skills (Burton 1982) or comparing student answers to expert knowledge (Zinn 2014), such extensions pose challenges. The existing error libraries may be insufficient in addressing the specific misconceptions observed in our task. Alternatively, relying on expert knowledge would require substantial human input, presenting practical limitations. In contrast, using HMMs introduced a data-driven approach to solve this problem. This approach addresses some limitations associated with traditional ways of identifying student misconceptions while also providing an extension to predicting the phases of rule learning.

The predictions from our model naturally contained some error. When we investigated the errors, we found out the model distinguished between different forms of rule search instances, namely help-seeking and guessing behaviors. We encountered some instances where the model mistakenly identified “follow wrong” occurrences as the guessing form of rule search. Addressing this challenge may require further refinement, potentially involving additional feature engineering or the inclusion of supplementary data sources to more accurately distinguish between guessing behaviors and instances of following the wrong rule. This issue was more pronounced in the test error. However that could also be due to self-report nature of the ground truth data that we had in Study 2. Even though self-report measures are very common for investigating behaviors and mental states in human-focused research, the subjectivity and response biases inherent in self-report data introduce risk and inaccuracies, challenging its suitability as ground truth (Gao et al. 2021).

The second step of our approach was to pinpoint the timing of each phase within the predicted sequences. This decision was motivated by insights from Study 1, where we identified that the time spent on rule discovery significantly predicted learning gain. We built a model that predicts the time spent on the individual rule learning phases for a given participant and problem step instance based on the proportions of time spent on the rule learning phases on a rule associated with the corresponding problem from the training data. The model demonstrated overall good prediction accuracy, although slightly less for rarer phases. The errors observed were within the range of the standard deviation of the data. Moving on, we applied the model trained on Study 1 data on Study 2 data to investigate the relationship between the predicted time spent on rule learning phases and learning gains in Study 2. Our findings revealed a consistent pattern, with predictions for rule discovery significantly correlating with learning gain in Study 2 data as well. One difference we observed was that rule search also showed a significant positive correlation with learning gain. This difference prompted us to consider the possibility that the distinction between rule search and rule discovery might have been blurred. This could be attributed to coding errors in the thinkaloud data or potential errors carried by the first step of our analysis, predicting rule learning sequences with the HMM, for Study 2. The results suggest that what predicts learning gain may be the time spent on an integrated cognitive state that is the combination of rule search and discovery phases, referred to as the rule acquisition stage in literature (Crescentini et al. 2011).

7.2 Limitations

It is important to note that our results serve as a proof of concept, and there is room for refinement and enhancement of the model to predict the time spent on individual phases and its associated features to further improve its accuracy and reliability. Despite these promising results, our work is not without limitations. First, our sample size was relatively small, which is common in studies of this nature (Du et al. 2020). However, the number of problems each participant solved provided a sufficient amount of data and this made the dataset reasonably robust in size. Additionally, the thinkaloud data from Study 1 provided rich, qualitative information, which helped us gain deeper insights into students' cognitive processes despite the small number of participants. Moreover, the findings from Study 1 were confirmed by the behavioral data in the larger Study 2, which reinforces the validity of the conclusions drawn. To mitigate the risk of overfitting due to the small sample size, we implemented cross-validation at each step of the modeling pipeline. We performed cross-validation at the participant level to ensure that data from the same participant were never split between training and test sets. This prevented data leakage and guaranteed that model performance was assessed on an unseen set of participants' data in each fold. Additionally, all model training was conducted on the dataset from Study 1, and we validated the model on the entirely independent dataset from Study 2. This further ensured that our findings were generalizable across different dataset.

Second, we were limited in the control we had over the ASSISTments interface, which restricted our ability to manipulate what was shown on the screen. However, this did not appear to affect students' experience, as they did not interact with components that were not introduced during the practice runs of the experiment.

Another limitation of our study is the simplicity of the tasks, which was intentional as we aimed to translate rule learning phases from controlled cognitive science experiments into more realistic problem-solving contexts. In cognitive science, rule learning phases are typically examined through highly simplified tasks, such as identifying patterns in sequences of events like numbers or shapes, with the goal of isolating specific cognitive processes. These tasks provide clear insights into participants' phases of rule learning but do not reflect the complexity of real-world problem solving (Crescentini et al. 2011). When adapting this framework to a problem-solving environment, we faced the challenge of maintaining the simplicity necessary to detect rule learning phases while increasing the task's realism. We opted for a middle ground, designing tasks that were more complex than traditional cognitive science experiments but still structured enough to enable detection of these phases. Our contribution lies in successfully adapting the rule attainment structure to a problem solving setting, while balancing the need for detectability of lower-level cognitive phases with the complexity of authentic tasks.

Finally, we focused on a single mathematical learning domain probability problem solving—which limits the generalizability of our findings. However, we hypothesize that rule learning offers a unified framework for modeling skill acquisition across domains. Future work should explore how this framework can be extended to other subject areas in order to provide further evidence of its domain agnostic potential.

7.3 Implications for intervention and adaptation

One of the most significant advantages of modeling rule learning phases is its ability to track the learner's cognitive state continuously throughout the problem-solving process. Unlike traditional models that rely on performance over multiple problems or interactions, our approach captures shifts in cognitive processes within a single task. This granular understanding opens up possibilities for more responsive, real-time interventions without the need for extensive error libraries or predefined misconceptions.

If we take the misconception case, for example, a student might believe that probability of two independent events is calculated by adding the two probabilities. This misconception leads the student to consistently apply the wrong rule across problems. In a traditional model, these incorrect answers would be recorded, and perhaps, after enough errors, the system would adapt by offering more practice or general hints. Indeed, hard coding of the common errors, defining the buggy procedures as triggers, and their interventions within the tutoring system may overcome this issue (VanLehn 1990); however, these kinds of solutions require substantial human effort and the defined error libraries might not always be sufficient to capture a broad range of misconceptions that students might have. In contrast, our model detects when the student is in a misconception state—what we name the "following a wrong rule" phase without relying on external information about the buggy rules. This phase captures the critical cognitive moment where the student is applying an incorrect strategy or heuristic. Upon detecting this pattern, the system can intervene with a targeted counterexample that directly challenges the misconception. For instance, the system might show the student that probability of independent events is actually calculated by multiplication. To deepen this intervention, the system could prompt the student to reflect: "Why do you think you have to multiply the two probabilities? How does this differ from your initial expectation?" Previous research (Chi 2009) highlights the importance of confronting misconceptions directly to promote learning. This cognitive conflict through a counterexample approach could create the conditions for the student to revise their faulty understanding.

When the student is detected to be in the rule search phase, the system might provide scaffolding through hints, guiding the student to discover the correct rule on their own. For example, it could break down the problem into simpler steps, helping the student reapply the correct rule to more straightforward cases before gradually increasing the complexity, in alignment with overlapping waves theory (Siegler 1995), which suggests that learners often switch between incorrect and correct strategies during problem solving. This type of intervention upon detecting rule search could support productive exploration rather than waiting for the student to stumble upon the right rule after repeated errors.

Once the student moves into the rule discovery phase, where they have found the correct rule but may not fully grasp how to apply it consistently, the system can further enhance learning through delayed feedback or metacognitive prompts. Here, the goal is to help the student reflect on their newfound understanding and consider how the rule applies across different contexts. For example, after discovering that multiplying the two probabilities of the independent events, the system could prompt the student: "Can you explain why this rule works in this situation? How might it apply to other multiplication problems?" Encouraging reflection in this way draws from earlier research on the self-explanation effect (Chi et al. 1994), which has been shown to help students consolidate their knowledge and apply it more effectively in future tasks.

Finally, as the student progresses through the rule following phase, where they begin applying the correct rule but may still encounter difficulties in more complex scenarios, the system can offer additional support through worked examples or analogical reasoning tasks as Gentner et al. (Gentner et al. 2003) have shown that analogical reasoning enhances rule-based learning, enabling students to apply their knowledge in new, more challenging contexts. This intervention would ensure that learning is robust and transferable.

Our current detection model has yielded promising results, demonstrating that the system can reliably identify distinct cognitive phases such as misconception states and rule discovery. This provides a critical foundation for implementing the interventions described above. While these interventions are not yet fully deployed in practice, the model's accuracy in phase detection indicates that real-time interventions are within reach.

Our model's ability to detect and respond to distinct rule learning phase enables a level of precision and opens the way for a level of responsiveness that traditional models cannot achieve. This includes the identification of 'additional phases' or patterns that may not appear in more controlled settings, such as the "following the wrong rule, rule search, and rule discovery" pattern, or instances where students remain in the "rule following" phase throughout a problem cannot achieve. Additionally, because these phases can be detected continuously, even during pauses or moments of reflection, the system offers a seamless and ongoing estimation of cognitive states without needing extensive external information or error annotation. This continuous tracking provides a more holistic and adaptive learning experience, allowing for interventions that are tailored to the student's moment-to-moment cognitive process.

8 Conclusion and future work

In conclusion, our work advances the understanding of cognitive processes that contribute to robust learning by drawing parallels between rule learning mechanisms and student modeling. By incorporating these cognitive processes into student modeling, we open the way for more targeted interventions and enhanced support for learners at different stages of skill acquisition. In initial stages, the focus is on tasks that are similar to a “rule search” phase, such as gathering and organizing information related to the target skill, complemented by strategies like offering students worked-out examples within ITSs (Renkl 2014). As learners progress, more suitable approaches involve providing opportunities for practice and real problem solving. This shift aligns with the goal of automating acquired skills and refining both speed and accuracy (Atkinson et al. 2003).

Our work has the potential to provide this adaptation in a more continuous manner, not just across multiple tasks, but even within individual pauses, as the phases of rule learning occur within the duration of a single problem and they provide an opportunity to categorize different cognitive activities occur during the pauses.

Moreover, our work introduces an approach for exploring and supporting the underlying mechanisms of learning, ultimately fostering the development of more effective and personalized educational experiences as well as interdisciplinary research for development of these innovations. One promising avenue for future research is to extend this work through the integration of multimodal analysis, particularly incorporating brain-sensing data. Studies in cognitive neuroscience have already identified distinct neural patterns corresponding to different phases of rule learning (Cao et al. 2016; Li et al. 2012). Future research can explore how cognitive processes during learning manifest both behaviorally and neurologically by aligning these neural patterns with the behavioral data we captured. This interdisciplinary approach opens up exciting new opportunities for collaboration between AI in education and cognitive neuroscience. For example, intelligent tutoring systems could dynamically adjust their interventions in real time based on both behavioral and neural data by detecting neural correlates of rule learning phases. Such advancements could revolutionize adaptive learning strategies by offering more continuous support that responds not only to what students are doing but also to what they are thinking. Our work thus lays a foundation for future studies that aim to combine behavioral and neural data to refine student modeling especially for multimodal applications.

Acknowledgements This work was supported by the National Science Foundation under awards #1835307 and #1912474.

Author Contributions DU wrote the main manuscript text. All authors reviewed and edited the manuscript.

Declarations

Conflict of interest The authors declare no Conflict of interest.

References

Aleven, V., Koedinger, K.R.: Knowledge component (KC) approaches to learner modeling. Des. Recommend. Intell. Tutor. Syst. **1**, 165–182 (2013)

Anderson, J.R.: Acquisition of cognitive skill. Psychol. Rev. **89**(4), 369 (1982)

Anderson, J.R.: Skill acquisition: compilation of weak-method problem situations. Psychol. Rev. **94**(2), 192 (1987)

Anderson, J.R.: The Architecture of Cognition. Psychology Press, East Sussex (2013)

Anderson, J.R., Corbett, A.T., Koedinger, K.R., et al.: Cognitive tutors: lessons learned. J. Learn. Sci. **4**(2), 167–207 (1995)

Anderson, J.R., Bothell, D., Byrne, M.D., et al.: An integrated theory of the mind. Psychol. Rev. **111**(4), 1036 (2004)

Asselman, A., Khaldi, M., Aammou, S.: Enhancing the prediction of student performance based on the machine learning XGBoost algorithm. interactive learning environments, pp 1–20 (2021)

Atkinson, R.K., Renkl, A., Merrill, M.M.: Transitioning from studying examples to solving problems: effects of self-explanation prompts and fading worked-out steps. J. Educ. Psychol. **95**(4), 774 (2003)

d Baker, R.S., Corbett, A.T., Aleven, V.: More accurate student modeling through contextual estimation of slip and guess probabilities in Bayesian knowledge tracing. In: International Conference on Intelligent Tutoring Systems, pp. 406–415 (2008)

d Baker, R.S., Pardos, Z.A., Gowda, S.M., et al.: Ensembling predictions of student knowledge within intelligent tutoring systems. In: International Conference on User Modeling, Adaptation, and Personalization, pp. 13–24 (2011)

Bloom, B.S., Engelhart, M.D., Furst, E.J., et al.: Taxonomy of Educational Objectives, vol. 2. Longmans Green, New York (1964)

Braver, T.S.: The variable nature of cognitive control: a dual mechanisms framework. Trends Cognit. Sci. **16**(2), 106–113 (2012)

Burgess, P.W., Shallice, T.: Bizarre responses, rule detection and frontal lobe lesions. Cortex **32**(2), 241–259 (1996)

Burton, R.R.: Diagnosing bugs in a simple procedural skill. In: Intelligent Tutoring Systems, pp 157–183 (1982)

Cao, B., Li, W., Li, F., et al.: Dissociable Roles of Medial and Lateral PFC in Rule Learning. Brain Behav. **6**(11), e00551 (2016)

Cen, H., Koedinger, K., Junker, B.: Learning factors analysis—a general method for cognitive model evaluation and improvement. In: International Conference on Intelligent Tutoring Systems, pp. 164–175 (2006)

Cheng, P.W., Holyoak, K.J.: Pragmatic reasoning schemas. Cognit. Psychol. **17**(4), 391–416 (1985)

Chi, M.T.: Three types of conceptual change: belief revision, mental model transformation, and categorical shift. In: International Handbook of Research on Conceptual Change. Routledge, pp. 89–110 (2009)

Chi, M.T., De Leeuw, N., Chiu, M.H., et al.: Eliciting self-explanations improves understanding. Cognit. Sci. **18**(3), 439–477 (1994)

Cooper, G., Sweller, J.: Effects of schema acquisition and rule automation on mathematical problem-solving transfer. J. Educ. Psychol. **79**(4), 347–412 (1987)

Corbett, A.T., Anderson, J.R.: Knowledge tracing: modeling the acquisition of procedural knowledge. User Model. User Adapt. Interaction **4**(4), 253–278 (1994)

Crescentini, C., Seyed-Allaei, S., De Pisapia, N., et al.: Mechanisms of rule acquisition and rule following in inductive reasoning. J. Neurosci. **31**(21), 7763–7774 (2011)

Du, X., Yang, J., Hung, J.L., et al.: Educational data mining: a systematic review of research and emerging trends. Inf. Discov. Delivery **48**(4), 225–236 (2020)

Ericsson, K.A., Simon, H.A.: Verbal reports as data. Psychol. Rev. **87**(3), 215 (1980)

Fang, Y., Ren, Z., Hu, X., et al.: A meta-analysis of the effectiveness of ALEKS on learning. Educ. Psychol. **39**(10), 1278–1292 (2019)

Fitts, P.M.: Perceptual-motor skill learning. In: Categories of Human Learning. Academic Press, pp. 243–285 (1964)

Gao, N., Saiedur Rahaman, M., Shao, W., et al.: Investigating the reliability of self-report data in the wild: the quest for ground truth. In: Adjunct Proceedings of the 2021 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2021 ACM International Symposium on Wearable Computers, pp 237–242 (2021)

Gentner, D., Loewenstein, J., Thompson, L.: Learning and transfer: a general role for analogical encoding. J. Educ. Psychol. **95**(2), 393 (2003)

Gick, M.L., Holyoak, K.J.: Analogical problem solving. Cognit. Psychol. **12**(3), 306–355 (1980)

Gick, M.L., Holyoak, K.J.: Schema induction and analogical transfer. Cognit. Psychol. **15**(1), 1–38 (1983)

Gill, A.M., Nonnecke, B.: Think aloud: effects and validity. In: Proceedings of the 30th ACM International Conference on Design of Communication, pp. 31–36 (2012)

Haverty, L.A., Koedinger, K.R., Klahr, D., et al.: Solving inductive reasoning problems in mathematics: not-so-trivial pursuit. Cognit. Sci. **24**(2), 249–298 (2000)

Heffernan, N.T., Heffernan, C.L.: The assistments ecosystem: building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching. Int. J. Artif. Intell. Educ. **24**(4), 470–497 (2014)

Hummel, H.G., Nadolski, R.J.: Cueing for schema construction: designing problem-solving multimedia practicals. Contemp. Educ. Psychol. **27**(2), 229–249 (2002)

Kirschner, P.A.: Cognitive load theory: implications of cognitive load theory on the design of learning. Learn. Instruct. **12**(1), 1–10 (2002)

Koedinger, Kenneth R., Corbett, Albert T., Perfetti, Charles: The knowledge-learning-instruction framework: bridging the science-practice chasm to enhance robust student learning. Cognit. Sci. **36**(5), 757–798 (2012)

Lee, J.I., Brunskill, E.: The impact on individualizing student models on necessary practice opportunities. In: International Educational Data Mining Society (2012)

Li, F., Cao, B., Gao, H., et al.: Different brain potentials evoked at distinct phases of rule learning. Psychophysiology **49**(9), 1266–1276 (2012)

Liu, Q., Shen, S., Huang, Z., et al.: A Survey of Knowledge Tracing. arXiv preprint arXiv:2105.15106 (2021)

Newell, A.: Unified Theories of Cognition. Harvard University Press, Harvard (1994)

Newell, A., Rosenbloom, P.: Cognitive Skills and Their Acquisition, Chapter Mechanism of Skill Acquisition and the Law of Practice (1981)

Pardos, Z.A., Heffernan, N.T.: Modeling Individualization in a bayesian networks implementation of knowledge tracing. In: International Conference on User Modeling, Adaptation, and Personalization, pp. 255–266 (2010)

Pardos, Z.A., Heffernan, N.T.: KT-IDEM: Introducing item difficulty to the knowledge tracing model. In: International Conference on User Modeling, Adaptation, and Personalization, pp. 243–254 (2011)

Pavlik, P.I., Cen, H., Koedinger, K.R.: Performance factors analysis—a new alternative to knowledge tracing. In: Artificial Intelligence in Education, pp 531–538 (2009)

Pelánek, R.: Bayesian knowledge tracing, logistic models, and beyond: an overview of learner modeling techniques. User Model. User Adapt. Interact. **27**(3), 313–350 (2017)

Piech, C., Bassen, J., Huang, J., et al.: Deep knowledge tracing. Adv. Neural Inf. Process. Syst. **28** (2015)

Renkl, A.: Toward an instructionally oriented theory of example-based learning. Cogni. Sci. **38**(1), 1–37 (2014)

Ritter, S., Fancsali, S.: MATHia X: The next generation cognitive tutor. In: Educational Data Mining, ERIC, pp. 624–625 (2016)

Siegler, R.S.: How does change occur: a microgenetic study of number conservation. Cognit. Psychol. **28**(3), 225–273 (1995)

Taatgen, N.A., Lee, F.J.: Production compilation: a simple mechanism to model complex skill acquisition. Human Factors **45**(1), 61–76 (2003)

Unal DS, Arrington, C.M., Solovey, E., et al.: Using Thinkalouds to understand rule learning and cognitive control mechanisms within an intelligent tutoring system. In: International Conference on Artificial Intelligence in Education, pp. 500–511 (2020)

Unal DS, Arrington, C.M., Solovey, E., et al.: Eliciting proactive and reactive control during use of an interactive learning environment. In: International Conference on Artificial Intelligence in Education, pp. 753–759 (2023)

Van Merrienboer, J.J., Paas, F.G.: Automation and schema acquisition in learning elementary computer programming: implications for the design of practice. Comput. Human Behav. **6**(3), 273–289 (1990)

VanLehn, K.: Mind Bugs: The Origins of Procedural Misconceptions. MIT Press, Cambridge (1990)

VanLehn, K.: The behavior of tutoring systems. Int. J. Artif. Intell. Educ. **16**(3), 227–265 (2006)

Wang, X., Zheng, Z., Zhu, J., et al.: What is wrong with deep knowledge tracing? Attention-based knowledge tracing. Appl. Intell., pp. 1–12 (2022)

Zeller, C., Schmid, U.: Automatic generation of analogous problems to help resolving misconceptions in an intelligent tutor system for written subtraction. In: Workshops Proceedings for the 24th International Conference on Case-Based Reasoning (2016)

Zhang, J., Shi, X., King, I., et al.: Dynamic key-value memory networks for knowledge tracing. In: Proceedings of the 26th International Conference on World Wide Web, pp. 765–774 (2017)

Zhang, M., Chang, L., Liu, T., et al.: Augmenting convolution neural networks by utilizing attention mechanism for knowledge tracing. In: International Conference on Intelligent Information Processing, pp. 80–86 (2022)

Zinn, C.: Algorithmic debugging and literate programming to generate feedback in intelligent tutoring systems. In: Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz), pp. 37–48 (2014)

Deniz Sonmez Unal is a Ph.D. candidate in Intelligent Systems at the University of Pittsburgh. Her research focuses on modeling student cognitive states using multimodal data, including interaction logs, verbal protocol data, and neural signals to enhance the diagnostic capabilities of intelligent tutoring systems.

Dr. Erin Treacy Solovey is an associate professor of computer science at Worcester Polytechnic Institute and a 2024-25 Harvard-Radcliffe Institute Fellow. Her research is at the intersection of human-computer interaction, artificial intelligence (AI), and neuroscience, often applying machine learning methods to develop intelligent interfaces that respond to a person's needs. She also investigates improving information accessibility for deaf individuals as well as effective human interaction with complex and autonomous systems and vehicles.

Dr. Catherine Arrington is Professor of Psychology and Cognitive Science and Director of the Institute for Data, Intelligent Systems, and Computation at Lehigh University. Her research interests include cognitive processes engaged during human multitasking in both laboratory and real-world settings.

Dr. Erin Walker is the Associate Dean for Research in the School of Computing and Information at the University of Pittsburgh, and is an Associate Professor with joint appointments in Computer Science and the Learning Research and Development Center. She completed her PhD in 2010 in Human-Computer Interaction from Carnegie Mellon University, and was subsequently awarded a Computing Innovations Postdoctoral Fellowship. In 2013, she became faculty in the School of Computing, Informatics, and Decision Systems Engineering at Arizona State University, and moved to the University of Pittsburgh in 2019. Her research spans human-computer interaction, artificial intelligence, and educational technology. More specifically, she uses interdisciplinary methods to improve the design and implementation of educational technology, and then to understand when and why it is effective. Her current focus is two-fold: Examine how artificial intelligence techniques can be applied to support social human-human and human agent learning interactions, and use educational data mining approaches to develop deeper understanding of learners' cognitive, metacognitive, and motivational states in the use of learning technologies. Her work has resulted in over fifteen journal articles and forty peer-reviewed full conference papers, including a best paper award at Creativity and Cognition, a best poster award at AIED, best young researcher's track paper award at AIED, and additional award nominations at ITS, AIED, and CSCL.