



HTN-Based Tutors: A New Intelligent Tutoring Framework Based on Hierarchical Task Networks

Momin N. Siddiqui
msiddiqui66@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

Adit Gupta
ag3338@drexel.edu
Drexel University
Philadelphia, PA, USA

Jennifer M. Reddig
jreddig3@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

Christopher J. MacLellan
cmacellan3@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

ABSTRACT

Intelligent tutors have shown success in delivering a personalized and adaptive learning experience. However, there exist challenges regarding the granularity of knowledge in existing frameworks and the resulting instructions they can provide. To address these issues, we propose HTN-based tutors, a new intelligent tutoring framework that represents expert models using Hierarchical Task Networks (HTNs). Like other tutoring frameworks, it allows flexible encoding of different problem-solving strategies while providing the additional benefit of a hierarchical knowledge organization. We leverage the latter to create tutors that can adapt the granularity of their scaffolding. This organization also aligns well with the compositional nature of skills.

CCS CONCEPTS

• **Computing methodologies** → **Planning for deterministic actions**; • **Applied computing** → **Computer-assisted instruction**; **Interactive learning environments**.

KEYWORDS

Human-centered computing, Intelligent tutoring systems, Artificial Intelligence, Hierarchical Task Network, Scaffolding

ACM Reference Format:

Momin N. Siddiqui, Adit Gupta, Jennifer M. Reddig, and Christopher J. MacLellan. 2024. HTN-Based Tutors: A New Intelligent Tutoring Framework Based on Hierarchical Task Networks. In *Proceedings of the Eleventh ACM Conference on Learning @ Scale (L@S '24)*, July 18–20, 2024, Atlanta, GA, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3657604.3664702>

1 INTRODUCTION

Intelligent Tutoring Systems (ITSs) are computer programs that utilize AI techniques to provide personalized and adaptive learning [17]. Several randomized controlled trials have demonstrated their

effectiveness at improving student learning gains [7, 19]. Frameworks for intelligent tutoring include *constraint-based* tutoring, which uses constraint-based modeling (CBM) to specify domain principles that every correct solution must follow [18], and *example-tracing* tutoring, which uses generalized traces of problem-solving behavior [3]. One of the more well-known frameworks is *rule-based* tutoring, which uses production rules to represent the cognitive model, where every rule is tied to a student skill [2]. Some examples of rule-based tutors adopted at scale are Cognitive Tutor/MATHia [6, 21], OATutor [4], ASSISTments [13] and Apprentice Tutor [1].

One line of research concerning intelligent tutors is the granularity of instructions. Granularity, in the context of intelligent tutors, refers to the amount of reasoning that is handled by the student internally on each step [24]. Low granularity means that the system provides detailed, scaffolded steps to help guide the learner. In contrast, high granularity means that the system provides much less support—often just a single input box for the learner to enter the final answer. Research suggests that scaffolding should ideally be faded, shifting from low granularity to high granularity and reducing support as learners gain proficiency [15, 20]. However, limited research has explored ITS that dynamically changes the granularity of steps based on student skill level, functionality that we refer to as *adaptive scaffolding*. In most intelligent tutors, granularity is static. It is a feature of the user interface rather than encoded in the cognitive model, preventing adaption based on the learner’s skill level. This gap necessitates formalizing a new expert model for intelligent tutors that can represent knowledge at different levels of granularity and better support adaptive scaffolding.

Across many existing systems, each skill the user is expected to acquire is represented independently and discretely [25]. In the case of rule-based tutors, skills are represented as distinct production rules that activate when the rule’s conditions are met. We argue that the representation of skills needs to better capture their compositional and hierarchical nature [11].

To address the aforementioned challenges, we present a new framework using Hierarchical Task Networks (HTNs) that we refer to as *HTN-based* tutoring. HTNs are used in automated planning due to their human-like decomposition [16]. We leverage this property of HTNs to encode different problem-solving strategies informed by the student’s knowledge level. Each action is tied to meaningful



This work is licensed under a Creative Commons Attribution International 4.0 License.

L@S '24, July 18–20, 2024, Atlanta, GA, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0633-2/24/07.
<https://doi.org/10.1145/3657604.3664702>

abstractions representing knowledge components. HTN-based tutors capture a key characteristic of skills, they are built upon one another in compositional hierarchies.

Our contribution is twofold. First, we propose a cognitive modeling framework that can deliver adaptive scaffolding, aligning with the learner’s progress to reduce cognitive load for novices and challenge advanced learners. Second, by modeling the compositional nature of skill development, HTN-based tutors offer a more realistic and practical framework for learning complex subjects.

2 BACKGROUND

In intelligent tutors, granularity is typically a characteristic of the user interface and not the cognitive model; granularity reflects the level of reasoning participants must engage in between interaction points [24]. A larger grain size implies more reasoning with each interaction. Systems have been constructed with varying levels of granularity. Answer-based systems have the largest possible granularity, presenting learners with a problem and asking them to input the final answer without giving feedback to learners on intermediate steps. This requires the learner to do extensive mental reasoning for each response. In contrast, step-based and substep-based systems offer more frequent, detailed guidance. That is, they break problems down into multiple steps that the learner needs to perform, reducing the amount of independent reasoning required before each response. Step-based tutoring systems [23] aim to have users enter the steps they would naturally write down when solving a problem while showing their work. In contrast, substep-based tutors give scaffolding and feedback at a level of detail that is even finer than the steps students would typically exhibit when showing their work.

Hierarchical task networks (HTN) are popular for providing a robust framework for automated planning, an area of artificial intelligence research [10]. HTNs are used in robotics, video games, and military simulations, among many other applications. Seminal systems such as SHOP and SHOP2 [16] exemplify the effectiveness of HTNs in decomposing abstract tasks into concrete, manageable tasks through the use of Methods and Operators. Their hierarchical approach to task decomposition has a natural alignment with human cognitive skill organization, making it an intuitive framework for encoding task knowledge. This alignment facilitates a deeper understanding of planning as a cognitive function. While its potential as a cognitive modeling tool is evident [12], research on its use in intelligent tutors, and in the learning sciences more generally, is limited. Our proposed system builds upon the domain-independent framework of HTNs to create expert models for tutoring systems.

3 HTN-BASED TUTORS

In this section, we will outline the components of our proposed HTN-based tutoring framework. We will describe the state, the expert model, and how both support model tracing and tutoring.

3.1 State

The system utilizes a short-term working memory that contains an up-to-date description of the system’s state. The working memory represents the state using facts, which are logical predicates that describe different elements in the tutor and how they relate. An

example fact describing the value stored in the “addFraction” field might be:

```
Fact(field='addFraction', value='1/2+1/4')
```

Facts encode the details of the problem, its expected scaffolding, and the user’s expertise for each skill as calculated through knowledge tracing [8]. The knowledge of which step should be scaffolded can be dynamically inferred based on a user’s skill levels using Axioms. Axioms are described in greater detail in the following section.

3.2 Expert Model

The expert model stores problem-solving strategies, including tasks, operators, methods, and axioms.

3.2.1 Tasks. A task represents an activity to perform.¹ Tasks can be performed using an Operator or a Method. Each task has a definition associated with it, which we call the task head. The task head is essentially the name or identifier of the task. The task head corresponds to heads that appears in Operator or Method. In order to start problem-solving, a task has to be instantiated.

3.2.2 Operator. Operators represent primitive behaviors. Each operator has a head, preconditions, and effects. The preconditions are a partial description of a state that describes when the operator can be applied and The effects describe the actions to be take when applied. These usually result in some change to the state (e.g., updating a tutor field with a particular value). Operators are similar in many respects to the production rules in rule-based models. The key difference is that operators are only considered in the context of the target task that they are known to perform, as specified by their head, whereas all productions are considered in every state.

3.2.3 Method. Methods represent non-primitive behaviors. Like operators, each method has a head that defines the task it performs and preconditions that describe when it is applicable. However, unlike operators, methods have subtask decompositions instead of effects. There can exist multiple methods with the same head to describe different ways of decomposing a given task under different conditions. This allows the encoding of multiple strategies for performing a task. The method is essentially a higher-level abstraction that break down a task into simpler subtasks.

3.2.4 Axiom. Axioms are used to infer facts that are not explicitly asserted in the current state. Each axiom has a head and preconditions. However, unlike methods and operators, the head refers to a new fact that can be inferred when its preconditions are satisfied.

3.3 Model Tracing with HTNs

Model tracing is a technique for inferring a student’s mental operations given their observed behavior on a problem [5], enabling real-time and contextual feedback on the student’s inferred state.

To support model tracing with HTN-based cognitive models, all problems start with a single higher-level task that is linked to the problem (e.g., reduce the logarithmic expression). The HTN model tracing system searches for a way to recursively decompose this task until it grounds out into operators that describe actions to be taken in the tutor. For each task, the system retrieves all the

¹Here we use the HTN definition of a *task* [16], which is different from typical tutoring system definition [23].

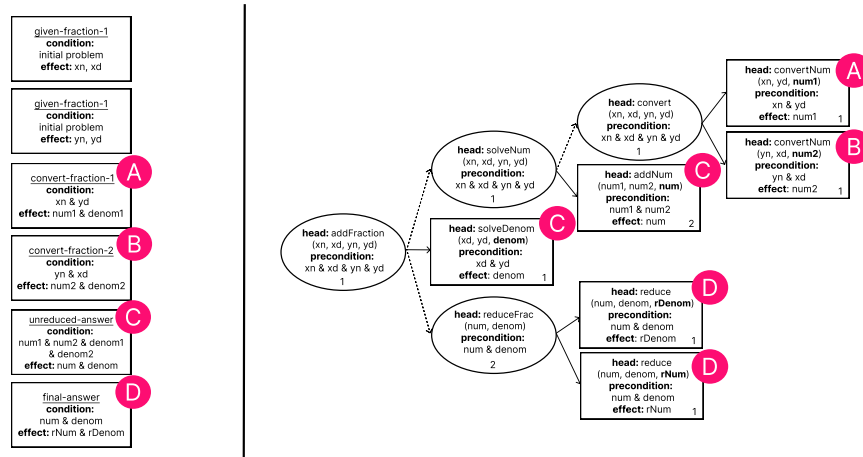


Figure 1: A representation of fraction addition problem-solving knowledge in rule (left) and HTN formats (right), showing methods (ellipse) and operators (rectangle) for head tasks, with lettered callouts indicating equivalent steps in both frameworks.

methods or operators that can achieve it and checks to see if their preconditions are satisfied. If any of the methods and/or operators are applicable, then it selects one (this choice is represented as an OR decomposition in Figure 2). If a method is selected, then it decomposes the current task into a conjunction of subtasks, each of which the system tries to satisfy via recursive decomposition (the conjunction of subtasks to be satisfied is represented as an AND decomposition in Figure 2). If an operator is selected, then it waits for the student to take action corresponding to the effects described by the operator. If the expected effects are not observed, then the system backtracks to find an alternative decomposition that matches the observed student action. When no decomposition matches, the system identifies the student action as incorrect. However, if the student takes an expected action, then model tracing continues and the system moves on to the next subtask that needs to be decomposed.

In each state, the HTN-based system only considers methods that satisfy the current task/subtask. As a result, it evaluates fewer Methods/Operators than rule-based tutors, which evaluate all production rules in every state. Figure 1 compares a simplified rule-based expert model, comparable to what is used in a Cognitive Tutor [6] (on the left) with a simplified HTN-based expert model (on the right).

4 EXAMPLE

Having described the framework of HTN-based tutors, it is worthwhile to showcase an example in the context of a tutor. Figure 2 shows a tutor for reducing logarithmic expression. The tutor (right) displays different granularity levels based on the user’s estimated knowledge of each component (e.g., as estimated via knowledge tracing [8]). A student who is highly proficient (marked by the green progress bar) sees an interface similar to an answer-based tutor. A student with intermediate skills (yellow progress bar) sees an interface similar to step-based, while for a novice, the interface resembles a sub-step-based tutor, outlining even the steps that are more traditionally done mentally.

Figure 2 also shows the expected model tracing path corresponding to each user (left). The numbering at the bottom of nodes represents the ordering sequence of their respective AND branches. Each path is marked with colors matching their progression. The green yields no scaffolding, while yellow and red have provide progressively more scaffolding in each case. The plus button on the right of the input field expands the scaffolding manually if the student would like more support. Red implies scaffolding is available, while grey implies scaffolding has already been expanded for that step.

5 DISCUSSION

5.1 Adaptive Scaffolding

In intelligent tutors, the focus is typically on providing contextualized hints and personalized practice sequences rather than adaptive scaffolding [4, 6]. These systems utilize two primary mechanisms: the outer loop and the inner loop [23]. The outer loop selects problems that learners have not yet mastered [22], focusing their practice where it is needed rather than on skills they already know. The inner loop offers hints and immediate feedback directly through the tutor interface, tailoring to the specific context. Sometimes students with low skills do not have the metacognition to implement the hints and explanations they are given [9]. They require explicit problem-solving instruction and modeling of their problem-solving strategies to make progress. Our cognitive models enable alternative methods to administer adaptive scaffolding.

5.1.1 Granularity. Granularity is explicitly represented in HTN’s task hierarchy, where higher-level tasks are less granular and lower-level tasks represent more granularity. This allows multiple subtask lists to be authored, catering to learners of varying skill levels, as seen in figure 2. This approach not only addresses cognitive load but also enhances learning by dynamically adjusting the scaffolding based on the HTN model to meet the learner’s evolving needs. Adaptive granularity enables the tutor to reduce the degrees of freedom [26] for problem-solving while retaining higher-level task context by means of tailored instructions. Effective scaffolding

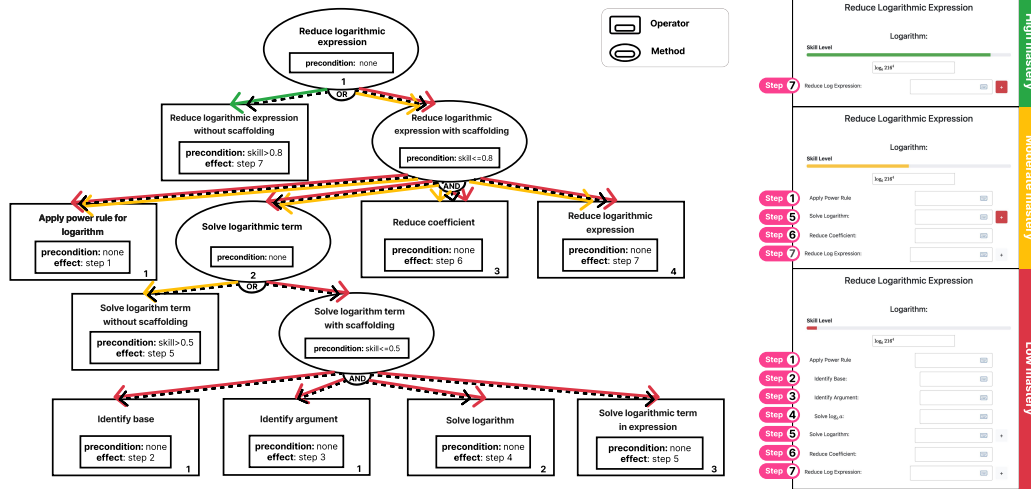


Figure 2: An HTN-based tutor implementation for solving logarithmic expressions is shown. On the left, various paths for model tracing correspond to different skill levels: green for high-skill, yellow for moderate, and red for low-skill students.

should gradually fade as learners gain proficiency [15], allowing them to become more independent in their problem-solving. Our system progressively reduces the level of scaffolding provided. As learners advance, the tutor transitions to an answer-based format, fostering greater learner autonomy and reinforcing mastery. By giving students practice in a format that more closely resembles the format of problems on exams and tests, we hypothesize our tutors will yield better performance gains under testing conditions than typically seen with ITS [14].

5.1.2 Strategy Recognition. It is valuable for students to receive feedback on the strategies they use, as this enables them to meta-cognitively assess their problem-solving. Due to strategy being encoded directly in HTN methods, our model tracing system can identify the approach taken by students and provide feedback on their strategic choices. This should help improve student’s current understanding and enhance their outcomes [26].

5.2 Representation Power

Skills are compositional in nature [11], serving as building blocks for higher-level skills. In contrast to production rule engines, which process rules individually without a hierarchical knowledge structure, our HTN-based approach organizes knowledge hierarchically, reflecting the natural organization of skills. The flexible nature of the HTN-based tutor allows for the reusability of methods and operators while not compromising the hierarchy of skills. Moreover, the tasks in the HTN do not have to be strictly sequential and can have interleaved dependency. This makes it easier to encode complex problems where tasks do not follow a linear sequence of steps (e.g., where the subtasks can be performed in any order).

6 LIMITATION & FUTURE WORK

While HTN-based tutors show promise, they have yet to be tested and integrated within a large-scale deployment. Integrating our cognitive modeling approach within a platform to support standard

inner and outer loop features would allow for a more holistic evaluation. We will investigate the following questions: 1) Does adaptive scaffolding, through adjusting granularity, improve student learning? 2) What is the optimal strategy for adaptive granularity? To address the first question, we will compare the performance of intelligent tutors with static granularity to HTN-based tutors with adaptive granularity through controlled experiments. Students will be randomly assigned to one of three groups: an HTN-based tutor with static scaffolding, one with adaptive scaffolding, or a control group without tutoring. We will evaluate effectiveness using learning gains measured by pre- and post-tests. For the second question, we will use a similar design with two treatment groups: one where granularity follows a U-shaped curve (starting large, decreasing, then increasing) and another where granularity follows a sigmoid curve (starting small and increasing). Outcomes will be assessed using pre- and post-tests.

7 CONCLUSION

To address the limitations of existing tutoring systems, we present an alternative framework for tutor cognitive models that uses HTNs. HTN-based tutors provide better personalization by enabling adaptive scaffolding functionalities that have thus far been under explored. The hierarchical nature of the system enables reusability of skills across tutors while maintaining their hierarchical association. It allows the composition of skills to formulate new higher-level skills, enabling better knowledge compilation [11]. The human-like encoding of skills opens the opportunity to incorporate several features that will help us deliver a better experience for learners.

8 ACKNOWLEDGEMENT

This project is supported by National Science Foundation under Grant No. 2247790 and Grant No. 2112532. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] AI-ALOE. 2024. Technologies. <https://aialoe.org/technologies>. Accessed: April 15, 2024.
- [2] Vincent Alevan. 2010. Rule-based cognitive modeling for intelligent tutoring systems. In *Advances in intelligent tutoring systems*. Springer, 33–62.
- [3] Vincent Alevan, Bruce M McLaren, Jonathan Sewall, Martin Van Velsen, Octav Popescu, Sandra Demi, Michael Ringenberg, and Kenneth R Koedinger. 2016. Example-tracing tutors: Intelligent tutor development for non-programmers. *International Journal of Artificial Intelligence in Education* 26 (2016), 224–269.
- [4] Ioannis Anastasopoulos, Shreya Sheel, Zachary Pardos, and Shreya Bhandari. 2023. Introducing an open-source adaptive tutoring system to accelerate learning sciences experimentation. In *Proceedings of the Tenth ACM Conference on Learning@ Scale*. 251–253.
- [5] John R Anderson, C Franklin Boyle, Albert T Corbett, and Matthew W Lewis. 1990. Cognitive modeling and intelligent tutoring. (1990).
- [6] John R Anderson, Albert T Corbett, Kenneth R Koedinger, and Ray Pelletier. 1995. Cognitive tutors: Lessons learned. *The journal of the learning sciences* 4, 2 (1995), 167–207.
- [7] Carole R Beal, Rena Walles, Ivon Arroyo, and Beverly P Woolf. 2007. On-line tutoring for math achievement testing: A controlled evaluation. *Journal of Interactive Online Learning* 6, 1 (2007), 43–55.
- [8] Albert T Corbett and John R Anderson. 1994. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction* 4 (1994), 253–278.
- [9] Annalisa Cusi, Agnese Ilaria Telloni, et al. 2020. Re-design of digital tasks: the role of automatic and expert scaffolding at university level. In *Proceedings of the 10th ERME Topic Conference' Mathematics Education in the Digital Age'(MEDA)*. Johannes Kepler University, 159–166.
- [10] Kutluhan Erol, James A Hendler, and Dana S Nau. 1994. *Semantics for hierarchical task-network planning*. University of Maryland College Park.
- [11] Kurt W Fischer. 1980. A theory of cognitive development: The control and construction of hierarchies of skills. *Psychological review* 87, 6 (1980), 477.
- [12] Bradley Hayes and Brian Scassellati. 2016. Autonomously constructing hierarchical task networks for planning and human-robot collaboration. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 5469–5476.
- [13] Neil T Heffernan and Cristina Lindquist Heffernan. 2014. The ASSISTments ecosystem: Building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching. *International Journal of Artificial Intelligence in Education* 24 (2014), 470–497.
- [14] James A Kulik and John D Fletcher. 2016. Effectiveness of intelligent tutoring systems: a meta-analytic review. *Review of educational research* 86, 1 (2016), 42–78.
- [15] Katherine L McNeill, David J Lizotte, Joseph Krajcik, and Ronald W Marx. 2006. Supporting students' construction of scientific explanations by fading scaffolds in instructional materials. *The journal of the Learning Sciences* 15, 2 (2006), 153–191.
- [16] Dana S Nau, Tsz-Chiu Au, Okhtay Ilghami, Ugur Kuter, J William Murdock, Dan Wu, and Fusun Yaman. 2003. SHOP2: An HTN planning system. *Journal of artificial intelligence research* 20 (2003), 379–404.
- [17] Hyacinth S Nwana. 1990. Intelligent tutoring systems: an overview. *Artificial Intelligence Review* 4, 4 (1990), 251–277.
- [18] Stellan Ohlsson. 2016. Constraint-based modeling: from cognitive theory to computer tutoring—and back again. *International Journal of Artificial Intelligence in Education* 26 (2016), 457–473.
- [19] John F Pane, Beth Ann Griffin, Daniel F McCaffrey, and Rita Karam. 2014. Effectiveness of cognitive tutor algebra I at scale. *Educational Evaluation and Policy Analysis* 36, 2 (2014), 127–144.
- [20] Sadhana Puntambekar and Roland Hubscher. 2005. Tools for scaffolding students in a complex learning environment: What have we gained and what have we missed? *Educational psychologist* 40, 1 (2005), 1–12.
- [21] Steven Ritter and Stephen Fancsali. 2016. MATHia X: The Next Generation Cognitive Tutor. In *EDM*. ERIC, 624–625.
- [22] Steve Ritter, Michael Yudelson, Stephen E Fancsali, and Susan R Berman. 2016. How mastery learning works at scale. In *Proceedings of the Third (2016) ACM Conference on Learning@ Scale*. 71–79.
- [23] Kurt VanLehn. 2006. The behavior of tutoring systems. *International journal of artificial intelligence in education* 16, 3 (2006), 227–265.
- [24] Kurt VanLehn. 2011. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational psychologist* 46, 4 (2011), 197–221.
- [25] Kurt VanLehn and Zhendong Niu. 2001. Bayesian student modeling, user interfaces and feedback: A sensitivity analysis. *International Journal of Artificial Intelligence in Education* 12, 2 (2001), 154–184.
- [26] David Wood, Jerome S Bruner, and Gail Ross. 1976. The role of tutoring in problem solving. *Journal of child psychology and psychiatry* 17, 2 (1976), 89–100.