

Rubikon: Intelligent Tutoring for Rubik's Cube Learning Through AR-enabled Physical Task Reconfiguration

Haocheng Ren*
University of Michigan
Computer Science and Engineering
Ann Arbor, Michigan, USA
chrenx@umich.edu

Muzhe Wu*
Carnegie Mellon University
Human-Computer Interaction
Institute
Pittsburgh, Pennsylvania, USA
University of Michigan
Computer Science and Engineering
Ann Arbor, Michigan, USA
muzhew@andrew.cmu.edu

Gregory Thomas Croisdale
University of Michigan
Computer Science and Engineering
Ann Arbor, Michigan, USA
gregtc@umich.edu

Anhong Guo
University of Michigan
Computer Science and Engineering
Ann Arbor, Michigan, USA
anhong@umich.edu

Xu Wang
University of Michigan
Computer Science and Engineering
Ann Arbor, Michigan, USA
xwanghci@umich.edu

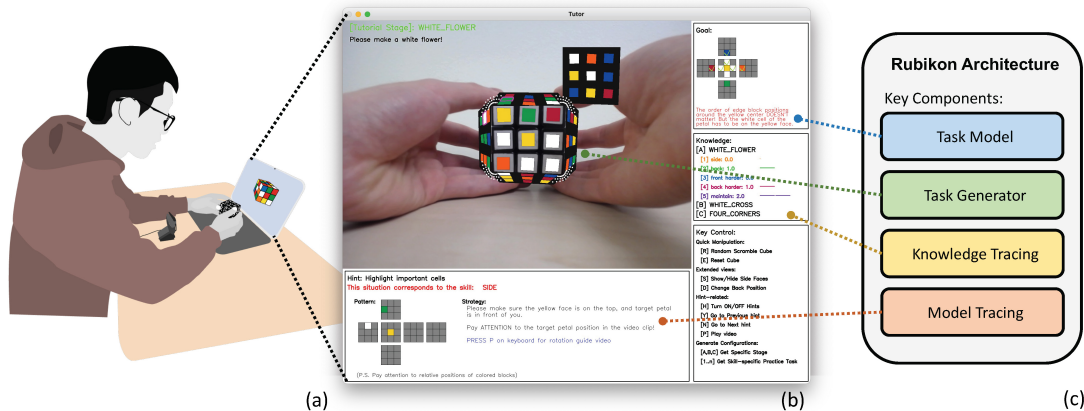


Figure 1: Rubikon is an intelligent tutoring system for Rubik's Cube learning. (a) The foundational design of Rubikon is an AR setup, where learners manipulate a physical cube with ARUco markers attached to each square, and pose a camera towards the cube to enable tracking and rendering. With this setup, learners see a rendered Rubik's Cube on a display while manipulating the physical cube in their hands. (b) Through AR rendering, Rubikon automatically generates new configurations of the Rubik's Cube for the user to practice unmastered skills. Rubikon detects the status of the cube to infer user behavior and provide immediate feedback and hints. (c) Rubikon supports the learning of a 3D physical task by integrating key design principles of cognitive tutors which have seen success in tutoring math and programming.

Abstract

Learning to solve a Rubik's Cube requires the learners to repeatedly practice a skill component, e.g., identifying a misplaced square and putting it back. However, for 3D physical tasks such as this, generating sufficient repeated practice opportunities for learners can be

challenging, in part because it is difficult for novices to reconfigure the physical object to specific states. We propose Rubikon, an intelligent tutoring system for learning to solve the Rubik's Cube. Rubikon reduces the necessity for repeated manual configurations of the Rubik's Cube without compromising the tactile experience of handling a physical cube. The foundational design of Rubikon is an AR setup, where learners manipulate a physical cube while seeing an AR-rendered cube on a display. Rubikon automatically generates configurations of the Rubik's Cube to target learners' weaknesses and help them exercise diverse knowledge components. In a between-subjects experiment, we showed that Rubikon learners scored 25% higher on a post-test compared to baselines.

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License.
DIS '25, Funchal, Portugal

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1485-6/25/07
<https://doi.org/10.1145/3715336.3735788>

CCS Concepts

- **Human-centered computing** → **Mixed / augmented reality**;
- **Applied computing** → **Interactive learning environments**.

Keywords

Augmented reality, education/learning, intelligent tutoring system, tangible, physical skill acquisition

ACM Reference Format:

Haocheng Ren, Muzhe Wu, Gregory Thomas Croisdale, Anhong Guo, and Xu Wang. 2025. Rubikon: Intelligent Tutoring for Rubik’s Cube Learning Through AR-enabled Physical Task Reconfiguration. In *Designing Interactive Systems Conference (DIS ’25)*, July 05–09, 2025, Funchal, Portugal. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3715336.3735788>

1 Introduction

The Rubik’s Cube, a 3D combination puzzle, has been widely recognized for its cognitive and educational benefits [69], including bolstering people’s spatial intelligence [28], supporting math [28], architecture [8] and engineering [59] learning, and improving people’s hand-eye coordination [64]. Conventional methods of learning to solve a Rubik’s Cube involve the use of manuals [57], video demonstrations [56], and interactive websites [58]. These methods possess notable limitations. First, it introduces significant cognitive burden on the learners to manually map the status of their cube at hand to the cube in the paper or online tutorials. Second, it lacks repeated configurations of the Rubik’s Cube for the learners to practice. Third, given the immense complexity of the Rubik’s Cube’s configuration space, it is difficult for learners to keep track of their knowledge progression and practice unmastered elements.

Prior work on cognitive tutors [7, 34] has demonstrated great success in providing personalized instruction to learners, and has been widely explored in subject domains such as math [32, 34, 43, 53] and programming [50, 55]. However, applying the design principles of cognitive tutors in building a tutoring system for the physical Rubik’s Cube poses key challenges. A key consideration is to retain the haptic experience in solving the Rubik’s Cube while providing the intelligent tutoring affordances through knowledge and model tracing [34, 65]. On the one hand, it is essential for cognitive tutors to provide repeated practice opportunities to learners on unmastered knowledge elements. However, in the context of solving a Rubik’s Cube, it requires repeated configurations of the cube in order for learners to practice on diverse problem scenarios, which is difficult to achieve, especially for novices. On the other hand, adaptive tutoring requires accurately capturing student performance on the task in order to infer their skill mastery and give feedback. However, it is challenging to acquire learner’s step-by-step performance when they are solving a Rubik’s Cube.

In this work, we are driven by two overarching research questions: (i) For tasks that require manual configuration of physical objects in order for novices to practice, how might we automatically generate targeted practice opportunities? (ii) If we are able to provide these practice opportunities, will they help novices learn? We specifically probed into these research questions in the context of learning to solve the Rubik’s Cube, a popular physical task with considerable benefits on spatial thinking skills [69], hand-eye coordination [64], math learning [28], and computational thinking [3].

Uniquely, Rubik’s Cube’s symmetrical shape makes it easier to track the positions of each square.¹ This makes the Rubik’s Cube a good candidate for developing intelligent tutoring setups, where generating new configurations of the physical object becomes possible and inferring learners’ performance becomes easier.

We present Rubikon, an intelligent tutoring system for learning to solve the Rubik’s Cube. The design of Rubikon closely follows the design principles of cognitive tutors [35], while addressing novel technical challenges when implementing them in such a 3D physical task. The foundational design of Rubikon is an AR setup, where learners manipulate a physical cube with ArUco markers attached to each square, and pose a camera towards the cube to enable tracking and rendering. With this setup, learners see a rendered Rubik’s Cube on a display while manipulating the physical cube in their hands, as shown in Figure 1. This enables Rubikon to (i) generate new configurations of the Rubik’s Cube to present targeted problem-solving opportunities to learners, which eliminates the need to manually reset or reconfigure the cube to a specific state for each new problem; and (ii) track the cube’s status change to infer users’ operations, in order to provide real-time feedback and model their learning trajectory. We implement Rubikon as a full-fledged cognitive tutor: it specifies a task model and compares user behavior against the task model to provide feedback to learners; and it uses a skillometer to track a learner’s weaknesses and prompts them to practice more when needed.

To evaluate whether Rubikon can successfully generate learning opportunities that target learners’ weaknesses and thus enhance learning, we conducted a between-subjects experiment with 36 participants who self-identified as novices in solving a Rubik’s Cube. The study also aims to investigate whether the AR setup poses extra cognitive load on the users. With these goals, we designed two baseline conditions. Baseline 1 simulates a business-as-usual learning scenario, in which learners watch online tutorial videos while manipulating a regular cube in their hands. Learners can pause at any time and re-watch the video in any way they want. In Baseline 2, learners had the same instructions as Baseline 1, but manipulated a cube augmented with ArUco markers while seeing the rendered cube on a display. The comparison between the two baselines aims to address whether the AR rendering of the Rubik’s Cube on a separate display adds to users’ cognitive load when their views and hand operations are separated. Finally, in the Rubikon condition, learners learn to solve the Rubik’s Cube with the full system functionality, including the AR setup and the intelligent tutoring capabilities, and learners receive feedback on every move and get prompted to repeatedly practice a skill until they demonstrate mastery of it.

We found that learners in the Rubikon condition scored 25% higher on the post-test compared to the traditional video tutorial baseline. We further probed into the reasons behind the higher learning gains, and found that learners in the Rubikon condition had more comprehensive and balanced coverage of the diverse skill components required to solve a Rubik’s Cube during their practice time, whereas learners in the baselines had limited exposure constrained by the physical configurations of the cube. Many participants in the baselines struggled to match the configuration of

¹A Rubik’s Cube has six faces, each face has nine squares, also referred to as “cells.”

the cube in their hands with that in the tutorial video, which increased the “preparation cost” for learners to engage in meaningful practice. Additionally, learners in the three conditions reported a similar level of cognitive load. We demonstrate the feasibility and pedagogical effectiveness of automatically generating targeted and personalized practice opportunities for Rubik's Cube learning.

With the acknowledgment that the AR-setup and user behavior tracking techniques implemented in Rubikon are specifically limited to the task of learning to solve a Rubik's Cube, the concept of generating targeted and personalized practice opportunities for physical task learning while retaining the haptic experience through AR may apply in other contexts. Through this proof-of-concept study, we encourage future Mixed Reality-based tutoring systems for physical task learning to engage with the idea of personalized task reconfiguration.

2 Related Work

2.1 Rubik's Cube Learning

In his 1981 book “Magic Cube” [60], Singmaster proposed the layer-by-layer strategy, which involves solving the Rubik's Cube starting with the first layer and then progressing to the second and third. This method has become the preferred strategy for beginners to learn. For a long time, paper manuals (e.g., [57]), typically included with a purchased Rubik's Cube, have been the primary learning resource. As technology advanced, video demonstrations (e.g., [56]) and interactive websites (e.g., [25, 58] offering virtual simulations of the cube and interactive tutorials) became increasingly popular. However, these methods require learners to constantly match their Rubik's Cube with the Rubik's Cube in the tutorials, creating extraneous cognitive load.

There are also existing commercial and research solutions for tracking the cube's state using internal [20, 47] and external sensors [9, 23, 45]. Products like GIIKER SUPER CUBE [20] and GoCube [47] enable tracking through embedded Bluetooth or IMU sensors to provide real-time guidance through connected mobile devices. Besides tracking (e.g., with RGB cameras [27, 45]), prior work in academic research has also focused on making Rubik's Cube learning more accessible and engaging. Khan et al. elicited design guidelines for enhancing learners' engagement with the Rubik's Cube through a series of workshops with designers and educators [30]. Chen and Liu proposed a mathematical method for learning path optimization that balances knowledge density with playfulness [13]. However, all the prior solutions still require the users to manually reset or reconfigure the cube to specific states in order to exercise certain knowledge components, which is particularly difficult for novices. Rubikon's AR setup using ArUco markers enables *digital reconfiguration* of the cube state without requiring any *physical manipulations*, which is complementary to the existing approaches for state tracking, and allowing for more flexible implementations of intelligent tutoring systems to target learner weaknesses by generating any targeted exercises on-the-fly.

2.2 Cognitive Tutor Design Principles

Decades of research has shown that cognitive tutors [34, 35, 44] lead to better learning in comparison with traditional methods. We

briefly introduce the design principles of cognitive tutors as summarized in prior work [35]. First, cognitive tutors represent student competence as a set of production rules. The production rules are often derived from cognitive task analysis [5] which specifies the knowledge components one needs to acquire in order to complete a task successfully. Second, cognitive tutors provide instruction in a problem-solving context [34]. Students need to apply the production rules in order to solve the problems. Students acquire the production rules, i.e., the skills, through repeatedly applying them during problem solving [33, 34]. Third, model tracing and knowledge tracing [35] are two main algorithms used in cognitive tutors. Model tracing compares the students' move with the projected move based on an expert task model to give them hints and feedback [34, 35]. Knowledge tracing evaluates students' mastery of each skill based on a sequence of student performance [35, 49]. Rubikon carefully follows the above design principles in defining a task model, generating new problem-solving scenarios through AR-enabled automatic reconfiguration, and providing knowledge and model tracing.

Nevertheless, the implementation of cognitive tutor design principles on a 3D physical task (such as solving Rubik's Cube) poses several challenges. First, providing learners with problem-solving tasks in an authentic context can be challenging when the physical objects require repeated configuration. Second, both model tracing and knowledge tracing require capturing student moves, i.e., their performance on a specific task. In prior tutors, e.g., algebra [44], programming [50], and chemistry [31] tutors, learners can demonstrate their performance through web-based [53, 55] or conversational interfaces [21]. However, for 3D physical tasks, how to capture learners' performance accurately becomes a challenge.

2.3 Mixed Reality-based Physical Task Guidance and Tutoring Systems

People traditionally turn to text- and video-based tutorials for guidance on physical tasks [14, 67]. As an emergent technology, Mixed Reality (MR) shows prospects of revolutionizing the physical task learning experience, allowing students to perceive virtual instructions while interacting with the tangible world, leading to increased engagement [46] and improved learning outcomes [19, 51, 61]. Additionally, prior work has shown the importance of immersion for physical task learning in which learners are required to physically coordinate their body and hands [48]. Here, we review prior MR-based physical task guidance and tutoring systems on three aspects: (i) how physical tasks are reconfigured, (ii) how learners' performance is captured, and (iii) what feedback is provided to learners.

Many prior physical task performance support systems are built around linear and single-path tasks, e.g., machine operation [10, 15, 26, 37, 38] and assembly [66, 67]. These tutorial systems do not have the affordance for reconfiguration. Other systems focus on non-linear tasks with divergent procedures, including learning musical instruments [63], physical exercises [6, 18, 41], and combination games [22]. However, these systems often do not generate problem-solving scenarios that pertain to the learners' prior knowledge. They either require the learners to select tasks themselves or demand human tutors' real-time, one-on-one instruction which is not

always accessible [18, 63]. Thus, automatically generating configurations of physical objects to enable deliberate practice is relatively under-explored in MR-based guidance and tutoring systems.

Most existing MR physical task tutors employ sensor-based methods [38], computer vision models [26, 41], and other advanced motion-trackers [18, 63] to monitor users' behavior, which often demands expensive setup and calibration efforts, and is prone to errors. Rubikon explores a lightweight and accurate tracking method inferring users' behavior through ArUco markers attached to each square of a Rubik's Cube, which also enables task reconfiguration on-the-fly.

Prior work has extensively examined the effectiveness of different types of feedback in MR-based guidance and tutoring systems, including visual cues (e.g., words [10, 52, 66], 2D figures and videos [6, 10, 15, 62], 3D arrows [11, 15], and avatars [11, 17]) for motion guidance. Recent work Paper Trail [52] allows teachers to author explanations in an AR environment. However, most of the referenced systems offer such feedback as static information prior to students' actions and are not responsive. A small fraction of prior work provides interactive feedback [22, 38, 41, 54]. The most common format is corrective feedback (e.g., warning signs [38]) and on-demand hints (e.g., animations [26]). Through model tracing, Rubikon provides interactive feedback and multi-level hints to users in support of their learning.

3 Rubikon: An Intelligent Tutoring System for Rubik's Cube Learning

We introduce Rubikon, an intelligent tutoring system for learning to solve the Rubik's Cube, that is built using Python and OpenCV. The design of Rubikon closely follows the design principles of cognitive tutors [35], while addressing novel technical challenges when implementing such principles in a 3D physical task to learn to solve a Rubik's Cube. Note, here are the often-used terminologies of a Rubik's Cube: (i) square: a Rubik's Cube has six faces, and each face has nine squares; (ii) block: a sub-cube of the Rubik's Cube. A Rubik's Cube has 27 blocks. A block can have three squares when it is at the corner, two squares when it is in the middle, or one square when it is in the center of a face; and (iii) petal: it refers to the blocks with two squares that connect two faces of the cube.

3.1 AR Rendering in Support of Task Generation

It is uniquely challenging for tutoring systems of 3D physical tasks to provide deliberate practice opportunities as it requires substantial efforts to reconfigure the physical objects to a certain state. In the case of solving a Rubik's Cube, it would be overwhelmingly difficult for novices to reset the cube to some stage due to their lack of skills.

Rubikon enables automatic task generation through AR rendering. ArUco markers are attached to each square of a Rubik's Cube. Users can pose a camera towards the cube, and the camera detects the markers and renders the color of each square in real time. The learner sees a fully-rendered Rubik's Cube through a display, while manipulating the physical ArUco marker cube in their hands. This setup enables Rubikon to automatically configure the status of a Rubik's Cube and generate tasks for users to solve. A fully rendered Rubik's Cube is shown in Figure 6.

With this setup, Rubikon generates tasks that target the skills the user has not demonstrably mastered. For example, if the user has not mastered solving petals in the BACK HARDER position (see Figure 3), Rubikon generates Rubik's Cube configurations that put a petal in that position for the learner to practice.

3.2 Capturing Learner Performance to Enable Intelligent Tutoring

It is critical for intelligent tutoring systems to track and analyze students' problem-solving processes to offer personalized feedback, and dynamically present new problems to students based on their prior knowledge. Implementing model tracing and knowledge tracing on 3D physical task tutoring is challenging since there are no given ways to capture learner performance on a task. Rubikon uses the status of the Rubik's Cube to infer the learner's behavior, which is convenient and accurate. With the AR setup, the status of the physical cube is accurately tracked with the ArUco markers, making it possible to capture learner behavior and task performance.

3.3 Key Components in Rubikon

With the AR setup as the foundation, we develop Rubikon following the design principles of cognitive tutors [35]. Rubikon has the following key components, as shown in Figure 2:

- (1) **Task Model:** It specifies the solution paths to reach a desired state.
- (2) **Model Tracing:** Rubikon captures users' rotations through detecting the changes of the cube's status. This allows Rubikon to provide real-time feedback on a user's moves by comparing them against the pre-defined task model.
- (3) **Knowledge Tracing:** Rubikon uses a skillometer to track a user's weaknesses and prompts them to practice more when needed.
- (4) **Task Generation:** Rubikon provides learners with repeated and personalized practice opportunities by generating new configurations of the Rubik's Cube through AR rendering.

3.3.1 Task Model. The task model is core to any given cognitive tutor, which specifies the skill components and solution paths needed to successfully perform the task. In designing cognitive tutors, people perform a cognitive task analysis of a task domain in order to specify the key components one needs to master, which are referred to as knowledge components [36]. In designing Rubikon, we performed a theoretical cognitive task analysis [5, 39] based on the 8-stage conventional solution [25] of a Rubik's Cube. We outline 11 knowledge components for solving the first layer of a Rubik's Cube (as shown in Figure 3) and split them into three sequential stages, White Flower, White Cross, and Four Corners. Each knowledge component defines a sequence of moves a user needs to apply and master in order to solve a unique pattern of the cube. It is worth noting that not all knowledge components listed in Figure 3 may be encountered by a learner during practice. Depending on the initial configuration of the Rubik's Cube, learners may only encounter the easiest knowledge component, such as SIDE in the first stage. This introduces limitations that even when learners can successfully complete a White Flower, they may not have fully exercised all the knowledge components in this stage.

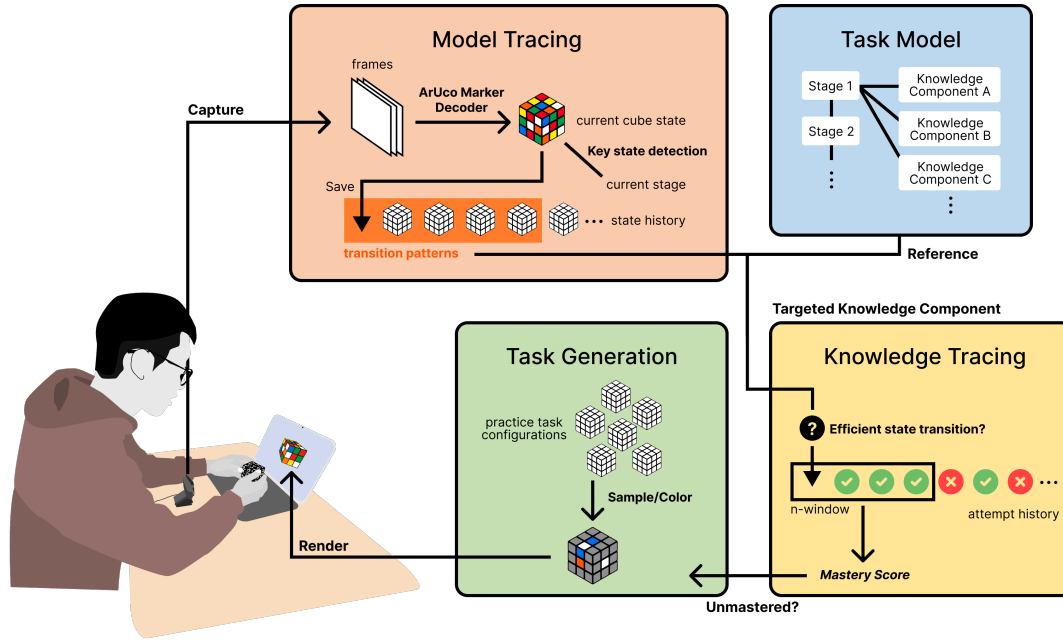


Figure 2: Rubikon enables personalized learning through a closed-loop design. While the learner manipulates the cube, the Model Tracing module parses frames streamed from a webcam into cube states. Informed by a Task Model that defines the movement sequences learners need to perform, the Knowledge Tracing module analyzes the transitions between cube states to assess mastery. For knowledge components not yet mastered, the Task Generation module correspondingly samples cube configurations and renders them on the screen for deliberate practice.

Knowledge components. Each knowledge component corresponds to a unique pattern of configuration on the Rubik's Cube, and learners need to know what sequence of moves to apply in each situation. There are 11 knowledge components representing over 700K possible configurations for the first layer of the cube since many configurations share the same pattern [29]. For instance, as shown in Figure 4, for the knowledge component SIDE, although cubes in (a) and (b) are two different configurations, they follow the same pattern. There are 8 possible configurations (indicated by white stars and the white square position in (b)'s pattern) that would allow a learner to exercise the knowledge component SIDE. They share the pattern that the target white square is located in

Task Model (for first layer of the cube)	
1. White Flower Stage	
skill_1: SIDE	*
skill_2: BACK	*
skill_3: FRONT HARDER	**
skill_4: BACK HARDER	**
skill_5: MAINTAIN	**
2. White Cross Stage	
skill_1: MATCH	**
3. Four Corners Stage	
skill_1: LEFT CORNER	*****
skill_2: RIGHT CORNER	*****
skill_3: TOP LAYER	*****
skill_4: UNDERNEATH	*****
skill_5: MISMATCH	*****

Figure 3: The Rubikon task model outlines 11 essential components for solving the first layer of the Rubik's Cube. Each component corresponds to a specific move sequence necessary to master distinct patterns. Difficulty levels are indicated by the number of *.

the center on adjacent sides of the yellow-center face. Similarly, (c)'s pattern shows 4 configurations that would allow a learner to exercise the knowledge component BACK HARDER. They share the pattern that the target white square is located at the same horizontal or vertical level as the yellow square, on adjacent sides of the yellow-center face, and has 3 other-colored squares between the white and yellow squares. These patterns essentially define the knowledge components and serve as the basis for Rubikon to generate new configurations that can help learners practice.

Difficulty level of knowledge components. The difficulty levels of the knowledge components are determined by the number of times a block needs to be switched from one layer to a different layer during rotations. For example, in Figure 4, (a) represents the knowledge component SIDE, which has a difficulty level of 1 star (the easiest) because learners only need to rotate the top layer counterclockwise to align the target squares. On the other hand, (c) represents the knowledge component BACK HARDER, which requires users to rotate the rightmost layer clockwise and then the top layer clockwise, making it harder.

3.3.2 Model Tracing. When a learner uses Rubikon, the model tracing algorithm detects and tracks the status of the cube through the ArUco markers. The algorithm then infers the user movement (e.g., left-face 90-degree clockwise rotation) based on the cube status change and the rotation templates. The system stores the states of the cube after every user move, which we refer to as cube state history. When a block is detected to be put in place, Rubikon gives positive audio feedback to users. When users are struggling, they can request hints. Rubikon provides multi-level hints to the user as in traditional cognitive tutors.

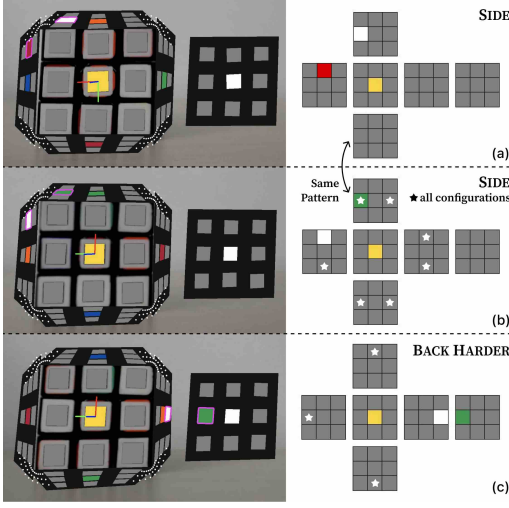


Figure 4: Each knowledge component in Rubikon's task model has a unique pattern of configuration on the Cube. (a) show a specific configuration of the Rubik's Cube corresponding to the knowledge component SIDE. (b) shows that there are 8 configurations that follow this pattern, which all enable the learners to exercise the knowledge component SIDE (indicated by the white stars). (c) shows 4 configurations of the knowledge component BACK HARDER (indicated by the white stars).

3.3.3 Knowledge Tracing. Rubikon detects the knowledge component the user just exercised for the purpose of knowledge tracing. The knowledge tracing algorithm analyzes the cube state history and identifies sequences of states where the initial and ending states match those of a specific knowledge component, as specified in the task model. The identified sequence of states signifies an attempt A at a knowledge component. To “grade” whether this attempt is correct, the knowledge tracing algorithm examines the sequence of states the Rubik's Cube goes through during this attempt: it first computes the minimum number of steps $\text{MINSTEPS}(\cdot)$ it takes to transition each state S_i to the final state S_k ; it then calculates the ratio of adjacent states (S_i, S_{i+1}) where the pre-calculated minimum number of steps is decreasing. If and only if the ratio is higher than a threshold $t_1 = 0.8$, the algorithm considers this attempt at the target knowledge component to be successful. With this, the algorithm can effectively identify situations where the learner, lacking a mental model of which moves to apply, randomly manipulates the cube and accidentally puts a block back in place.

$$\text{Ratio} = \frac{\sum_{i=1}^{k-1} \mathbb{I}[\text{MINSTEPS}(S_i, S_k) < \text{MINSTEPS}(S_{i+1}, S_k)]}{k}, \quad (1)$$

$$A = \mathbb{I}[\text{Ratio} > t_1]. \quad (2)$$

To model the learner's mastery of each knowledge component, the knowledge tracing algorithm aggregates the most recent n (set at 3) attempts $A\{n\}$, each multiplied by a weight $w(H_i)$ determined by the hint requested H_i , if any, for that attempt. Different hint levels correspond to different weights ($w(\text{None}) = 1.0$; $w(1) = 0.8$; $w(2) = 0.5$; $w(3) = 0$). The learner is considered to have mastered the

knowledge component if and only if the score for that knowledge component exceeds the threshold t_2 (set at 2.4).

$$\text{Mastery} = \mathbb{I}[\text{Score} > t_2] = \mathbb{I}[\sum_{i=1}^n w(H_i) \cdot A_i > t_2]. \quad (3)$$

The user's progress is visualized through a skillometer as shown in Figure 5, and all parameters (n , $w(\cdot)$, t_1 , t_2) are empirically determined through pilot studies. It is worth noting that knowledge tracing refers to a family of methods on observing, representing and quantifying a student's knowledge states [1, 16], i.e., the student's mastery levels on the knowledge and skills to be taught. We selected a straightforward method to infer the student's mastery level on each knowledge component, i.e., the student needs to demonstrate several “good” sequences of moves. Through our pilot studies, we found it worked well and supported our goals. Future work could use more sophisticated methods as summarized in this recent survey [1] on knowledge tracing to more accurately model a student's knowledge states.

3.3.4 Task Generation. With knowledge tracing, the system keeps track of the skills the learner has not mastered and generates new configurations of the Rubik's Cube to target the learner's weaknesses. For example, if the learner has not demonstrably mastered the skill of BACK HARDER, Rubikon will generate configurations of the Rubik's Cube that match the pattern of BACK HARDER for the learner to practice. These configurations are randomly sampled from a set of cube states associated with each knowledge component, with the specific color arrangements also randomized at runtime, ensuring task variety.

3.4 Rubikon User Interface

The Rubikon user interface contains 5 sections, as shown in Figure 5:

- (1) **The Playground.** The top-left section displays the rendered Rubik's Cube with extended views.
- (2) **Hint Panel.** The bottom-left section describes the hints, i.e., what the user should do to achieve the goal.
- (3) **Goal Image.** The top-right section provides a 2D view of the desired cube state to achieve in the current stage. Only the relevant squares are shown, while the rest of the squares are grayed out.
- (4) **Skillometer.** The middle-right section visualizes the skillometer based on the knowledge tracing algorithm.
- (5) **Controls.** The bottom-right section lists the key-press events users can trigger. This includes resetting and scrambling the cube, turning on and off extended views, requesting hints, and generating tasks.

3.4.1 Extended Views. Rubikon provides extended views of the cube to reduce the efforts on frequent rotations. These views, along with AR rendering, provide a comprehensive multi-angular perspective, revealing hidden sides. Through iterative design, we crafted the extended views: 4 trapezoidal shapes for the left, right, up, and down sides, and a shrunk mirrored image for the back side (Figure 6c). While the default back view location is top-right, users can relocate it for convenience. Spatial connectivity between squares of the same block in the extended views is denoted by arc-shaped

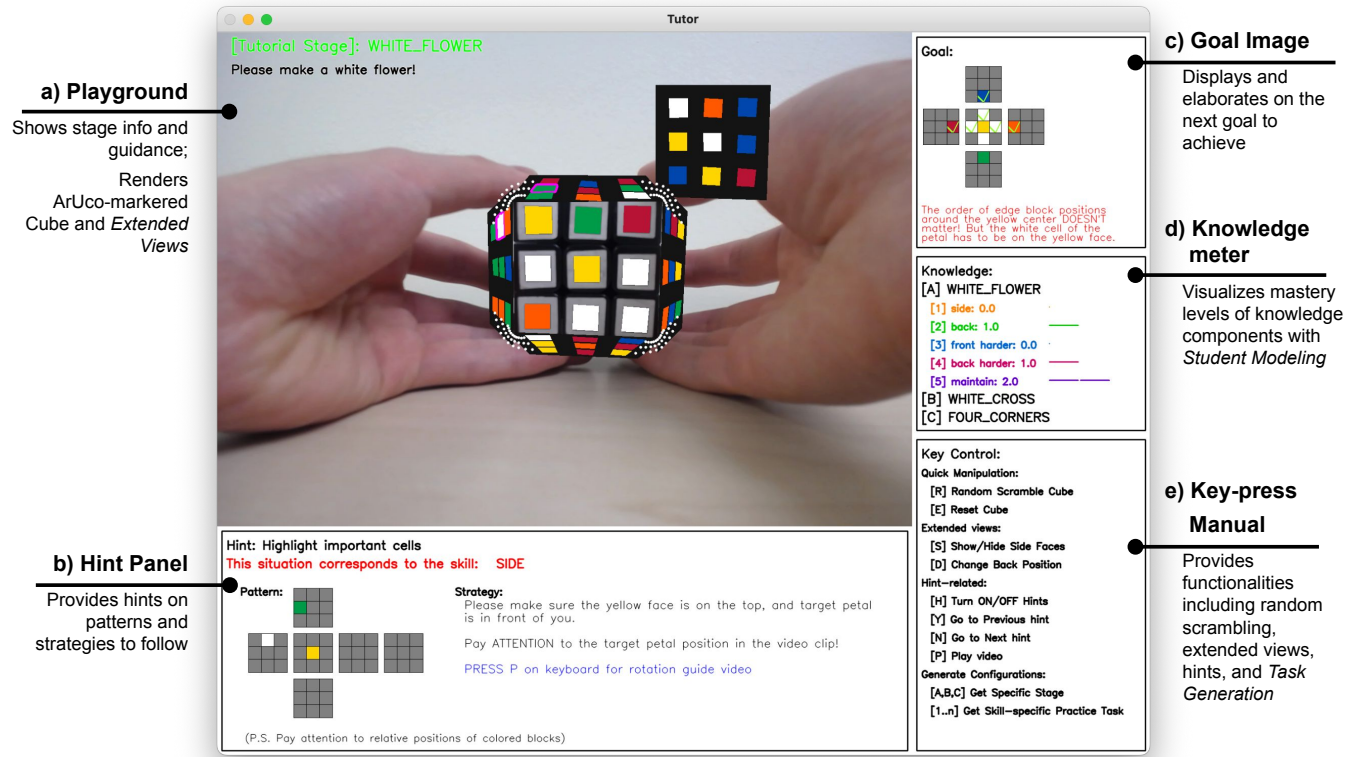


Figure 5: The Rubikon user interface has 5 key sections.

dotted lines in Figure 7. These views are dynamically computed and adapted to various cube orientations.

3.4.2 Multi-level Hints. Rubikon provides three levels of hint messages (Figure 7):

- (1) **Level 1:** Highlight the target block and destination.
- (2) **Level 2:** Gray out unimportant blocks.
- (3) **Level 3:** Explicitly show step-by-step arrow guidance (bottom-up hints).

The information contained in each level of hints increases based on the design principles of cognitive tutors [35]. Users have full control over when to request a hint, and at which level. Requesting




Figure 6: (a) A physical Rubik's Cube with ArUco markers attached to all squares. (b) A fully-rendered Rubik's Cube on the display with AR rendering. (c) A fully-rendered Rubik's Cube with extended views.




Figure 7: Three levels of hint messages: (a) No hint. (b) Hint Level 1: target block is highlighted. (c) Hint Level 2: irrelevant blocks are grayed out. (d) Hint Level 3, step-by-step arrow guidance.

a hint will lower the user's score on the corresponding knowledge component (refer to Section 3.3.3 for details) indicating the need for further practice.

3.4.3 Exploration and Practice Modes. Users will by default start using Rubikon in the exploration mode, guided by a goal image

of the current stage. They can then freely rotate the Rubik's Cube to achieve that goal. Throughout the process, users may request hints and get feedback. After a full exploration, users are directed to the practice mode, during which they are given additional practice opportunities on the knowledge components that they have not demonstrably mastered. Users can also manually request to start the practice mode at any time.

4 Evaluation Study

The evaluation study aims to investigate whether the AR setup in Rubikon AR can successfully generate learning opportunities that target learners' weaknesses and enhance learning. We compare Rubikon with two baselines and address the following research questions:

- RQ1: When users learn to solve a Rubik's Cube with the three methods, business-as-usual video tutorials with a regular Rubik's Cube, video tutorials with a rendered Rubik's Cube, and the Rubikon system, which leads to higher learning gains and why?
- RQ2: Whether the AR setup poses extra cognitive load on the users?
- RQ3: What are users' experiences when using Rubikon? What challenges do they encounter and what are the design implications for developing intelligent tutoring systems for learning 3D physical tasks?

4.1 Study Design

To evaluate Rubikon, we designed a between-subjects experiment in which participants learned to solve a Rubik's Cube in one of three conditions, as shown in Figure 8. We adopted a business-as-usual baseline (Baseline 1) in which participants learned to solve the Rubik's Cube by watching online video tutorials, while manipulating a regular cube. In Baseline 2, participants manipulated an AR-rendered cube and watched the same online video tutorials as in Baseline 1. This allows us to investigate whether AR renderings would add excessive cognitive load to the learning process since the user's views and hand operations are separated.

4.2 Participants

We recruited 36 participants through mailing lists at a large public university in the United States. We included a screening survey in the original recruitment email to select participants who were novices on solving the Rubik's Cube. The survey contains several




Figure 8: Three conditions in the evaluation study. (a) Condition 1: Baseline 1 where users watch video tutorials and use a regular Rubik's Cube. (b) Condition 2: Baseline 2 where users watch video tutorials and use an AR-rendered cube. (c) Condition 3: Rubikon.

questions on solving a 3-by-3 Rubik's Cube. From a pool of 184 survey respondents who met the inclusion criteria of being a novice on the task, we sent email invitations following their response time to the screening survey. In the end, 36 people participated in the study, with 15 female, 20 male, 1 undisclosed, and an average age of 22, ranging from 18 to 33.

4.3 Study Procedure

The study sessions lasted up to 2 hours in person. We administered a pre-test at the start of the session. Participants were presented with a goal image of a red cross and were asked to achieve this goal using a regular Rubik's Cube within 2.5 minutes. The time limit was set based on our pilot studies. Participants were scored based on their performance. Specifically, a correct square in place counts as one point, making the total score of the task 8 points. The pre-test served as a proxy of the participants' prior knowledge of solving the Rubik's Cube, which was also used to filter out those who were already experts.

Participants were randomly assigned to one of the three conditions and were given 45 minutes to learn to solve the first layer of a Rubik's Cube. In the two baseline conditions, users could watch the tutorial video freely. We selected a popular online video tutorial [56] with over 570K views that clearly breaks down the task into three stages. We further added captions to this online video tutorial to label the three stages and the 11 knowledge components as specified in the task model of Rubikon. Following the 45-min practice, participants answered the NASA TLX questionnaire [24] to assess the cognitive load required by the task. We chose NASA-TLX because it is a widely used, validated tool for assessing cognitive workload in a variety of task contexts [42, 68]. Participants then completed a post-test. We performed a short interview at the end to probe into the participants' learning experiences. For participants in the Rubikon condition, we asked their thoughts on the AR setup, the adaptive task generation functionality, and the perceived accuracy of the model tracing and knowledge tracing algorithms.

With the random assignment to conditions, there were 12 participants in each of the three conditions. The study was approved by our institution's IRB. All participants were compensated with a \$30 gift card.

4.4 Learning Outcome Measure

We administered a post-test after the practice asking participants to solve the first layer of a Rubik's Cube. We consider this to be an authentic post-test that reflects the tutoring objectives. We broke down the post-test into three tasks, aligned with the three learning objectives, namely solving White Flower, White Cross, and Four Corners. For each task, we configured the Rubik's Cube so that all participants went through the exact same post-test. Participants had 2.5 minutes to work on each task independently. We used a similar scoring mechanism as in the pre-test, i.e., when one square is back in place, it counts as one point, with a total score of 16 points. We also recorded the task completion time on the post-test.

4.5 Learning Process Measure

To understand how learners learned during the 45-min practice time, we defined a process measure "exercise opportunity" (see




Figure 9: The learning processes of participants are divided into segments, each representing an exercise opportunity targeting a knowledge component defined by the starting and ending cube states.

Figure 9). Exercise opportunities are associated with each knowledge component. To label an exercise opportunity for a knowledge component, we need to identify the start and ending time of the exercise period. As introduced earlier, each knowledge component has a unique initial configuration and a target state. During practice, when the learner's cube matches the initial configuration of a knowledge component and the learner starts to change the position of the target block, it marks the start time for this exercise opportunity. When the target block is back in place, it marks the ending time. During this time interval, no other knowledge components would be exercised. This allows us to derive (i) the number of knowledge components exercised during practice, and (ii) the time duration for each exercise opportunity.

During the practice time, a camera was positioned behind the users to record the study session, as shown in Figure 8. In the two baseline conditions, two researchers watched the video and manually labeled the exercise opportunities. In Rubikon, the system detects and logs the start and end of each exercise opportunity automatically.

4.6 Interview Analysis

The interview recordings were transcribed and analyzed using affinity diagramming [40]. Two researchers interpreted the transcripts, iteratively grouped the interpretation notes, and identified emerging themes from the data.

5 Results

We report our results in response to each research question. Before we ran any analysis, we did a randomization check to make sure that there was no systematic difference in the participants' prior knowledge across the three conditions. A one-way ANOVA test on the pre-test scores suggested that there was no significant difference across the three conditions ($F(2, 33) = 0.11, p = 0.90$).

5.1 The AR Setup Did Not Add To Learners' Cognitive Load

In order to understand whether the AR-rendered cube posed extra cognitive load on users, we ran an ANOVA test on the NASA TLX scores across the three conditions. Results showed that there was no significant difference on the cognitive load users reported across conditions ($F(2, 33) = 0.79, p = 0.46$), as shown in Figure 10a.

5.2 Rubikon Led to Higher Learning Gains

In order to investigate the learning benefit of Rubikon, we built a linear regression model, with the post-test score as the dependent

variable, and the condition as the fixed effect. We also included the learners' pre-test score, and the time they took on completing the post-test as covariates. We found that users in Rubikon had significantly higher learning gains compared to the two baselines, as shown in Figure 10b. With the pre-test score and post-test task completion time controlled, the regression model suggested that, when a user was assigned to the Rendered Cube condition compared to the Rubikon condition, they on average achieved 2.9 points lower out of 16 points in the post-test ($t = 2.43, p = 0.02$); when a user was assigned to the Regular Cube condition, they on average achieved 3.8 points lower in the post-test ($t = 3.25, p = 0.002$). This was a considerable effect size as learners achieved 3-4 points higher out of 16 points (25% of the whole task) through Rubikon in 45 minutes of the learning time. To explain what 25% of the entire task might look like, it could mean being able to solve all four corners instead of just two, or being able to complete the full white cross compared to being unable to move any pieces of the cross back into place.

5.3 Learners in Rubikon Had a More Comprehensive and Balanced Coverage of Knowledge Components During Practice

We further analyzed learners' learning processes to understand the reasons why users in the Rubikon condition learned more than the two baselines.

5.3.1 Learners exercised more knowledge components in Rubikon.

We examined the average number of knowledge components exercised by users in each condition. Our hypothesis was that since Rubikon generated new problem-solving activities based on learners' mastery, learners might get to exercise more knowledge components during the practice session. We ran a one-way ANOVA test and found there was a significant difference across the three conditions on the number of knowledge components exercised ($F = 29.52, p < 0.0001$). A pairwise post-hoc Tukey's HSD test indicated that Rubikon learners exercised significantly more knowledge components compared to the two baselines ($p < 0.0001$ for both), as shown in Figure 10c. On average, learners in Rubikon exercised 3 more knowledge components in comparison to the baselines. This suggested that learners in Rubikon practiced in more diverse situations than what they would have seen from a video tutorial.

5.3.2 Learners in Rubikon distributed their time more evenly among knowledge components.

Figure 11 shows the average frequency participants exercised each knowledge component across the three conditions. The number of asterisks indicated the difficulty level of each knowledge component (see definition in Section 3.3.1). A Shapiro-Wilk test showed that the distribution of users' exercise opportunities in the Rubikon condition followed a normal distribution ($p = 0.35$) whereas the two baseline conditions were skewed towards the easier knowledge components ($p < 0.001$ for both tests).

Notably, learners in Rubikon exercised the difficult knowledge components more than the two baselines. We observed that learners in the two baseline conditions did not want to reconfigure the Rubik's Cube manually to exercise a specific knowledge component because reconfiguration might mess up a step they had completed. Moreover, learners also did not know how to reconfigure the Rubik's Cube manually to a specific state. In these cases, learners could

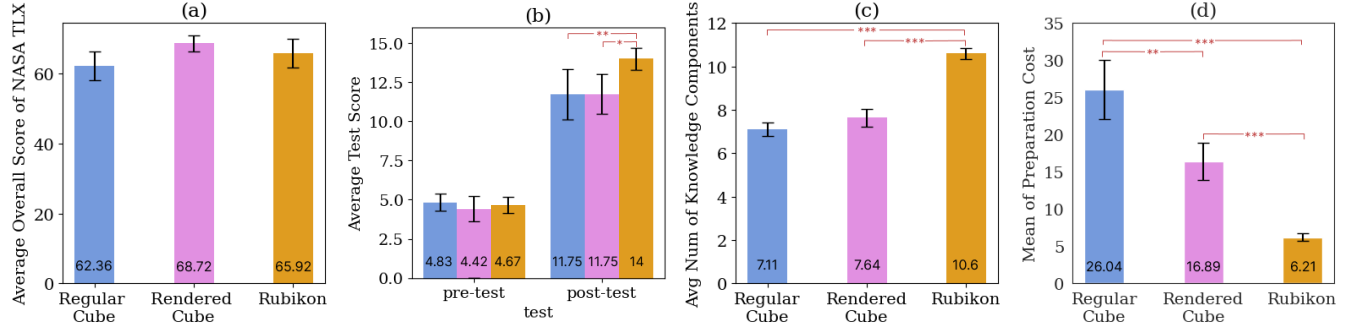


Figure 10: Evaluation Study Results. (a) Learners in the three conditions reported similar levels of NASA TLX overall score after the 45-min practice period. (b) Learners in Rubikon had significantly higher learning gains compared to learners in the two baselines. (c) Learners in Rubikon exercised significantly more knowledge components during the 45-min practice time. (d) Learners in Rubikon had significantly lower “preparation cost” — the time required for cube reconfiguration to exercise a knowledge component. *** = $p < 0.001$, ** = $p < 0.01$, and * = $p < 0.05$.




Figure 11: The frequency users exercised each knowledge component across the three conditions. * in the y-axis indicates the difficulty level of each knowledge component (see definition in Section 3.3.1). Rubikon learners had a more balanced coverage in their practice and a bigger edge in comparison with the baselines on the difficult knowledge components.

only watch the video tutorial without hands-on practice. In contrast, in Rubikon, users got prompted to exercise every knowledge component and were given repeated practice opportunities when they had not demonstrated independent success on them.

5.3.3 Learners in Rubikon spent less time on preparation and configuration before each exercise opportunity. When learning to solve the Rubik’s Cube with video tutorials, learners needed to align their cube with the status of the cube in the video. The mapping of the two representations took skill, effort and time. We further labeled “exercise time” and “preparation time” (see Figure 9). “Exercise time” is the duration of an exercise opportunity, and “preparation time” is the interval between the previous and current exercise opportunities. Since the time spent configuring the Rubik’s Cube to a specific state constitutes “preparation” rather than “learning”, we introduce the metric “preparation cost”, defined as the ratio of “preparation time” to “exercise time.” This metric allows us to evaluate whether Rubikon improves learning efficiency by reducing the time needed to configure the cube. When the “preparation cost” is lower, it means the learners need lower preparation effort to engage in meaningful

practice. We ran a one-way ANOVA test and found there was a significant difference across the three conditions on the “preparation cost” ($F = 28.87$, $p < 0.001$). A pairwise post-hoc Tukey’s HSD test suggested that Rubikon users had significantly lower “preparation cost” in comparison with the two baselines ($p < 0.0001$ for Regular Cube condition, $p = 0.0003$ for Rendered Cube condition), as shown in Figure 10d. The AR setup enabled the system to automatically reconfigure the rendered Rubik’s Cube so that the learners did not need to waste time reconfiguring the cube to a specific state.

5.4 Subjective User Experiences

We summarized the themes from the affinity diagram of the interview transcripts.

5.4.1 The AR setup incurred higher cognitive load. Even though we did not see a difference in the cognitive load users reported through the NASA TLX questionnaire, in the interviews, participants did mention that the AR setup introduced higher cognitive demand. Some participants shared that the AR setup constrained the way they manipulated the cube, which resulted in higher physical and mental efforts. On the other hand, multiple participants shared that the cognitive demand decreased as they became familiar with the system. For example, P10 commented, “Rubikon was new so it took some time before I could get used to it and use it as well as the normal Rubik’s Cube.”

5.4.2 Extended views provided useful information. Multiple participants shared that they were able to use the extended views to reconstruct the same 3D space as when using a regular cube, specifically with the visual aids Rubikon provided. “I think I got the point where the virtual cube has the same spatial awareness as the physical cube. (P14)” Many participants also consistently mentioned that the extended views provided just enough information that was easily understandable, “Helpful to present more information, not too overwhelming, straightforward. (P7)” With the extended views, participants could quickly locate the squares they were looking for without the need to physically manipulate the cube to another side. On the other hand, some participants shared concerns about the usability of the extended views. P22 and P35 mentioned that the squares in the extended views were too small, and P29 said that “it was hard for me to visualize the blocks on the back side.”

5.4.3 Learners appreciated the automatic configuration of the Rubik's Cube in Rubikon. In the baseline conditions, users had to match the configuration of the Rubik's Cube in their hands with that in the tutorial video, which required time, effort, and skill. Many participants found this process frustrating due to mismatches between the cube in their hand and the one in video. *"When mismatches occur, I don't know how I can get there. (P27)"* They struggled to progress smoothly through the tutorial, as it presented the solving process linearly and chronologically, resulting in either wasted time or incomplete understanding.

In contrast, Rubikon users did not face such issues. They appreciated Rubikon's automatic generation of problems targeting each knowledge component (P1, P25) and valued the additional practice opportunities it provided, leading to a more efficient learning experience (P34).

5.4.4 Knowledge tracing provided valuable practice opportunities to learners. Participants found the knowledge tracing functionalities of Rubikon to be valuable. Several participants pointed out that the knowledge tracing algorithm was able to track the knowledge components they did not master well. *"I felt like I did struggle with those low scores skills. (P4)"* Participants (P13, P34) also found the skillometer functionality to be useful. P34 mentioned, *"I think it was useful to know which skills I still need practice with so I could pay more attention during the tutorial."*

6 Discussion and Future Work

6.1 Generalizability of Rubikon

The study revealed that utilizing the proposed AR setup to generate new Rubik's Cube configurations led to more comprehensive coverage of knowledge components during practice, resulting in enhanced learning. For example, the system generates configurations of the Rubik's Cube where a target square is placed at varying positions so that learners could practice mapping the square back in more diverse situations than what they would have seen from a video tutorial. Learners experienced a reduced "preparation cost" for engaging in meaningful practice, meaning that they did not need to spend time aligning their physical Rubik's Cube to a virtual cube in the video tutorial. We want to emphasize that Rubikon is a Rubik's Cube-specific solution, and the AR setup and the behavior tracking techniques may not generalize beyond the task of learning to solve a Rubik's Cube. On the other hand, this work investigates the feasibility and pedagogical value of automatically generating targeted and personalized practice opportunities for physical task learning while retaining the haptic experience. We consider the concept of physical task reconfiguration applicable in other mixed reality-based physical task tutoring and guidance systems. We will discuss future work directions next.

6.2 Beneficial Situations to Repeatedly Generate Configurations of Physical Objects

Based on the case of solving a Rubik's Cube, we lay out some properties of a 3D physical task where the generation of new configurations of physical objects could be helpful.

- (1) **The spatial layout is important for learning.** For tasks where the spatial layout is important, e.g., solving a Rubik's

Cube, playing chess, etc., generating new configurations that represent different spatial layouts could help learning.

- (2) **Learners want to jump to intermediate states in a procedure to practice.** When it is critical for learners to be able to jump to any intermediate state of a procedure to practice, the capability to reset the physical task to any state would be beneficial. For example, in the case of learning drawing, learners can work on half-completed drawings (rendered digitally or by AR) to deliberately practice a skill instead of having to start from scratch every time.
- (3) **One solution path does not necessarily contain all the knowledge components encompassed in the task.** When the task is complex and has multiple solution paths, learners may not easily get exposed to all the required knowledge components if they just practice on their own. For example, when someone is learning to use a physical interface that has diverse paths, using AR to project new configurations on a simulated physical interface could help people more easily exhaust the possible scenarios.
- (4) **The task requires complex or expensive configurations of physical objects.** When the task requires complex or expensive configurations of physical objects, AR-enabled reconfiguration could save both time and resources.

6.3 Implications for Physical Task Tutoring

In this work, we present evidence on the benefit of automatically configuring a Rubik's Cube in order for learners to more efficiently practice and learn to solve it. In particular, a main reason Rubikon is more helpful than traditional tutorials is that learners can easily align the state of the physical object they are manipulating with the one in the tutorial. We observed struggles and frustrations from the users when they could not match the Rubik's Cube in their hands with the one in the video tutorial. When using AR to present different configurations of the physical objects required in the physical tasks, it also offers an opportunity to provide extra information for the users which is inaccessible without AR. In Rubikon, we showed users the hidden views of a Rubik's Cube and used simple visualizations to make them more comprehensible. We believe this notion can be extended to other domains, even though they may require a different AR setup from Rubikon (e.g., ARUco markers) for tracking and display that fits the nature of the task.

Although there has been a plethora of work in both the field of intelligent tutoring systems [2, 4, 67], and AR-based tutorial systems for physical tasks [6, 18, 22, 41, 63], there has been limited work integrating both, and limited empirical studies examining the learning benefit of such tutors. One example is Origami Sensei [12], which uses computer vision and real-time projection to guide learners through structured, planar origami tasks. Although it showcases the potential of AI-augmented MR systems, its scope is restricted to relatively linear tasks with fixed action sequences. We carefully designed Rubikon following the design principles of cognitive tutors, which has witnessed huge success in disciplines such as math [32, 34, 43, 53] and programming [50, 55] over the past decade. In our study, the knowledge and model tracing algorithms were applauded by the participants for offering valuable feedback and directing them to useful practice opportunities. We encourage

more work to deeply integrate MR technologies with theory-driven instructional design to produce effective tutoring systems for 3D physical task learning.

6.4 Limitations

Our work has several limitations. Firstly, despite the comparable NASA TLX score, Rubikon's on-screen display for rendering AR objects may still increase cognitive load, as some users reported. We plan to explore head-mounted displays in the future, expecting improved coordination between learners' views and hand operations, potentially reducing cognitive load further. Secondly, our evaluation study compares Rubikon with conventional approaches to physical task learning. However, for the Rubik's Cube, there are web-based platforms allowing practice through GUI-based interactions, as well as smart cubes that track the cube's state (as mentioned in Section 2.1). We did not include a baseline where learners use such platforms, though the benefits of Rubikon's AR-enabled task reconfiguration remain applicable. Thirdly, due to time constraints in the user study, we evaluated Rubikon's performance only on solving the first layer of the Rubik's Cube. The current system only specifies the task model for the first layer. However, we believe our findings are generalizable to learning other parts of the Rubik's Cube. Future studies could adopt longer instruction and assessment time spans to examine the effects of intelligent tutoring on long-term skill retention.

7 Conclusion

We presented Rubikon, an intelligent tutoring system for learning to solve the Rubik's Cube. The foundational design of Rubikon is an AR setup, where learners manipulate a physical cube while seeing an AR-rendered cube. This allows Rubikon to automatically generate configurations of the Rubik's Cube to target learners' weaknesses and help them exercise diverse knowledge components. In the evaluation study, we found that the users learned more in Rubikon in comparison to a business-as-usual baseline, where they watched video tutorials while manipulating a physical cube. We found that learners in Rubikon exercised more knowledge components and had lower "preparation cost" to engage in meaningful practice. Beyond solving a Rubik's Cube, we identified situations when repeatedly generating configurations of physical objects could help with the learning of physical tasks. This includes tasks that have an emphasis on spatial layout, require learners to easily jump to intermediate states, have multiple diverse solution paths, and require complex or expensive configuration of physical objects.

References

- [1] Ghodai Abdelrahman, Qing Wang, and Bernardo Nunes. 2023. Knowledge tracing: A survey. *Comput. Surveys* 55, 11 (2023), 1–37.
- [2] Nour N. AbuEloun and Samy S. Abu Naser. 2017. Mathematics Intelligent Tutoring System. *International Journal of Advanced Scientific Research* 2, 1 (2017), 11–16.
- [3] Forest Agostinelli, Mihir Mavalankar, Vedant Khandelwal, Hengtao Tang, Dezhi Wu, Barnett Berry, Biplav Srivastava, Amit Sheth, and Matthew Irvin. 2021. Designing Children's New Learning Partner: Collaborative Artificial Intelligence for Learning to Solve the Rubik's Cube. In *Proceedings of the 20th Annual ACM Interaction Design and Children Conference* (Athens, Greece) (IDC '21). Association for Computing Machinery, New York, NY, USA, 610–614. <https://doi.org/10.1145/3459990.3465175>
- [4] Bashar G. Al-Bastami and Samy S. Abu Naser. 2017. Design and Development of an Intelligent Tutoring System for C# Language. *European Academic Research* 4, 10 (2017).
- [5] Vincent Alevan. 2010. Rule-Based Cognitive Modeling for Intelligent Tutoring Systems. In *Advances in Intelligent Tutoring Systems*, Janusz Kacprzyk, Roger Nkambou, Jacqueline Bourdeau, and Riichiro Mizoguchi (Eds.). Vol. 308. Springer Berlin Heidelberg, Berlin, Heidelberg, 33–62. https://doi.org/10.1007/978-3-642-14363-2_3 Series Title: Studies in Computational Intelligence.
- [6] Fraser Anderson, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2013. YouMove: Enhancing Movement Training with an Augmented Reality Mirror. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) (UIST '13). Association for Computing Machinery, New York, NY, USA, 311–320. <https://doi.org/10.1145/2501988.2502045>
- [7] 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.
- [8] Christoph Bandelow. 2012. *Inside Rubik's Cube and Beyond*. Springer Science & Business Media, New York, USA.
- [9] Build-its. 2018. The Rubik's Contraption. <https://build-its-inprogress.blogspot.com/2018/03/the-rubiks-contraption.html> Build-its in Progress Blog.
- [10] Yuanzhi Cao, Anna Fuste, and Valentin Heun. 2022. MobileTutAR: a Lightweight Augmented Reality Tutorial System using Spatially Situated Human Segmentation Videos. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. ACM, New Orleans LA USA, 1–8. <https://doi.org/10.1145/3491101.3519639>
- [11] Yuanzhi Cao, Xun Qian, Tianyi Wang, Rachel Lee, Ke Huo, and Karthik Ramani. 2020. An Exploratory Study of Augmented Reality Presence for Tutoring Machine Tasks. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376688>
- [12] Qiyu Chen, Richa Mishra, Dina El-Zanfaly, and Kris Kitani. 2023. Origami Sensei: mixed reality AI-assistant for creative tasks using hands. In *Companion Publication of the 2023 ACM Designing Interactive Systems Conference*. 147–151.
- [13] Ziwei Chen and Lin Liu. 2019. Analysis the Influencing Factors of Intelligent Cube Interaction Design. In *Proceedings of the 3rd International Conference on Computer Science and Application Engineering* (Sanya, China) (CSAE '19). Association for Computing Machinery, New York, NY, USA, Article 72, 5 pages. <https://doi.org/10.1145/3331453.3360955>
- [14] Pei-Yu Chi, Sally Ahn, Amanda Ren, Mira Dontcheva, Wilmo Li, and Björn Hartmann. 2012. MixT: Automatic Generation of Step-by-Step Mixed Media Tutorials. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (Cambridge, Massachusetts, USA) (UIST '12). Association for Computing Machinery, New York, NY, USA, 93–102. <https://doi.org/10.1145/2380116.2380130>
- [15] Subramanian Chidambaram, Hank Huang, Fengming He, Xun Qian, Ana M Villanueva, Thomas S Redick, Wolfgang Stuerzlinger, and Karthik Ramani. 2021. ProcessAR: An Augmented Reality-Based Tool to Create in-Situ Procedural 2D/3D AR Instructions. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference* (Virtual Event, USA) (DIS '21). Association for Computing Machinery, New York, NY, USA, 234–249. <https://doi.org/10.1145/3461778.3462126>
- [16] 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.
- [17] Daniel Eckhoff, Christian Sandor, Christian Lins, Ulrich Eck, Denis Kalkofen, and Andreas Hein. 2018. TutAR: Augmented Reality Tutorials for Hands-Only Procedures. In *Proceedings of the 16th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry* (Tokyo, Japan) (VRCAI '18). Association for Computing Machinery, New York, NY, USA, Article 8, 3 pages. <https://doi.org/10.1145/3284398.3284399>
- [18] Mehrad Faridan, Bheesha Kumari, and Ryo Suzuki. 2023. ChameleonControl: Teleoperating Real Human Surrogates through Mixed Reality Gestural Guidance for Remote Hands-on Classrooms. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 203, 13 pages. <https://doi.org/10.1145/3544548.3581381>
- [19] Markus Funk, Thomas Kosch, and Albrecht Schmidt. 2016. Interactive Worker Assistance: Comparing the Effects of in-Situ Projection, Head-Mounted Displays, Tablet, and Paper Instructions. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Heidelberg, Germany) (UbiComp '16). Association for Computing Machinery, New York, NY, USA, 934–939. <https://doi.org/10.1145/2971648.2971706>
- [20] GiiKER. 2024. SUPERCUBE® i3S. <https://giiker.com/products/supercube-electronic-bluetooth-speed-cube-3x3> Accessed: 2024-11-13.
- [21] Arthur C. Graesser, Patrick Chipman, Brian C. Haynes, and Andrew Olney. 2005. AutoTutor: An intelligent tutoring system with mixed-initiative dialogue. *IEEE Transactions on Education* 48, 4 (2005), 612–618.
- [22] Ankit Gupta, Dieter Fox, Brian Curless, and Michael Cohen. 2012. DuploTrack: a real-time system for authoring and guiding duplo block assembly. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*.

- ACM, Cambridge Massachusetts USA, 389–402. <https://doi.org/10.1145/2380116.2380167>
- [23] Joseph Hale. 2022. DIY Smartcube: Tracking the Face Turns of a Rubik's Cube Using Embedded Speakers. (2022).
- [24] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, Amsterdam, 139–183.
- [25] How to solve a Rubik's Cube. 2015. How to solve a Rubik's Cube | The ultimate beginner's guide. <https://www.learnhowtosolvearubikscube.com/how-to-solve-a-rubiks-cube-solution-overview>.
- [26] Gaoping Huang, Xun Qian, Tianyi Wang, Fagun Patel, Maitreya Sreeram, Yuanzhi Cao, Karthik Ramani, and Alexander J. Quinn. 2021. AdapTutAR: An Adaptive Tutoring System for Machine Tasks in Augmented Reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 417, 15 pages. <https://doi.org/10.1145/3411764.3445283>
- [27] Lauri Jokinen. 2024. SOLVING A RUBIK'S CUBE WITH COMPUTER VISION. (2024).
- [28] David Joyner. 2008. *Adventures in Group Theory: Rubik's Cube, Merlin's Machine, and Other Mathematical Toys*. JHU Press, Baltimore, MD, USA.
- [29] WD Joyner. 1996. Mathematics of the Rubik's cube. *Spring semester 7* (1996), 1–275.
- [30] Awais Hameed Khan, Stephen Snow, and Ben Matthews. 2023. Participatory Design Tools: Leveraging Materiality and Familiarity to Adapt Unconventional Materials into Design Tools. In *Proceedings of the 15th Conference on Creativity and Cognition* (Virtual Event, USA) (C&C '23). Association for Computing Machinery, New York, NY, USA, 399–412. <https://doi.org/10.1145/3591196.3593339>
- [31] Emily C King, Max Benson, Sandra Raysor, Thomas A Holme, Jonathan Sewall, Kenneth R Koedinger, Vincent Aleven, and David J Yaron. 2022. The Open-Response Chemistry Cognitive Assistance Tutor System: Development and Implementation.
- [32] Kenneth R Koedinger and Vincent Aleven. 2007. Exploring the assistance dilemma in experiments with cognitive tutors. *Educational Psychology Review* 19 (2007), 239–264.
- [33] Kenneth R Koedinger and John R Anderson. 1998. Illustrating principled design: The early evolution of a cognitive tutor for algebra symbolization. *Interactive Learning Environments* 5, 1 (1998), 161–179.
- [34] Kenneth R Koedinger, John R Anderson, William H Hadley, Mary A Mark, et al. 1997. Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education* 8, 1 (1997), 30–43.
- [35] Kenneth R Koedinger, Albert Corbett, and others. 2006. *Cognitive tutors: Technology bringing learning sciences to the classroom*.
- [36] Kenneth R Koedinger, Albert T Corbett, and Charles Perfetti. 2012. The Knowledge-Learning-Instruction framework: Bridging the science-practice chasm to enhance robust student learning. *Cognitive science* 36, 5 (2012), 757–798.
- [37] Junhan Kong, Dena Sabha, Jeffrey P Bigham, Amy Pavel, and Anhong Guo. 2021. TutorialLens: Authoring Interactive Augmented Reality Tutorials Through Narration and Demonstration. In *Proceedings of the 2021 ACM Symposium on Spatial User Interaction* (Virtual Event, USA) (SUI '21). Association for Computing Machinery, New York, NY, USA, Article 16, 11 pages. <https://doi.org/10.1145/3485279.3485289>
- [38] Ziyi Liu, Zhengzhe Zhu, Enze Jiang, Feichi Huang, Ana M Villanueva, Xun Qian, Tianyi Wang, and Karthik Ramani. 2023. InstruMentAR: Auto-Generation of Augmented Reality Tutorials for Operating Digital Instruments Through Recording Embodied Demonstration. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 32, 17 pages. <https://doi.org/10.1145/3544548.3581442>
- [39] Marsha C. Lovett. 1998. Cognitive Task Analysis in Service of Intelligent Tutoring System Design: A Case Study in Statistics. In *Intelligent Tutoring Systems*, Barry P. Goettl, Henry M. Halff, Carol L. Redfield, and Valerie J. Shute (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 234–243.
- [40] Bill Moggridge and Bill Atkinson. 2007. *Designing interactions*. Vol. 17. MIT press Cambridge, Cambridge, MA, USA.
- [41] Kyzyl Monteiro, Ritik Vatsal, Neil Chulpongsoatorn, Aman Parnami, and Ryo Suzuki. 2023. Teachable Reality: Prototyping Tangible Augmented Reality with Everyday Objects by Leveraging Interactive Machine Teaching. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 459, 15 pages. <https://doi.org/10.1145/3544548.3581449>
- [42] Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. 2024. Using an LLM to Help With Code Understanding. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering* (Lisbon, Portugal) (ICSE '24). Association for Computing Machinery, New York, NY, USA, Article 97, 13 pages. <https://doi.org/10.1145/3597503.3639187>
- [43] Jennifer K. Olsen, Daniel M. Belenky, Vincent Aleven, and Nikol Rummel. 2014. Using an Intelligent Tutoring System to Support Collaborative as well as Individual Learning. In *Intelligent Tutoring Systems*, Stefan Trausan-Matu, Kristy Elizabeth Boyer, Martha Crosby, and Kitty Panourgia (Eds.). Springer International Publishing, Cham, 134–143.
- [44] 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.
- [45] Jaebum Park and Changhoon Park. 2016. Augmented reality based guidance for solving rubik's cube using HMD. In *HCI International 2016—Posters' Extended Abstracts: 18th International Conference, HCI International 2016 Toronto, Canada, July 17–22, 2016 Proceedings, Part II* 18. Springer, 524–529.
- [46] Jocelyn Parong and Richard E Mayer. 2018. Learning science in immersive virtual reality. *Journal of Educational Psychology* 110, 6 (2018), 785.
- [47] Particula. 2024. GoCube. <https://particula-tech.com/products/gocube> Accessed: 2024-11-13.
- [48] Kayur Patel, Jeremy N Bailenson, Sang Hack-Jung, Rosen Diankov, and Ruzena Bajcsy. 2006. The Effects of Fully Immersive Virtual Reality on the Learning of Physical Tasks. In *Proceedings of the 9th Annual International Workshop on Presence, Ohio, USA*. 87–94.
- [49] Philip I Pavlik, Hao Cen, and Kenneth R Koedinger. 2009. Performance Factors Analysis—A New Alternative to Knowledge Tracing. In *Artificial Intelligence in Education*. IOS Press, Amsterdam, Netherlands, 531–538.
- [50] Thomas W. Price, Yihuan Dong, and Dragan Lipovac. 2017. ISnap: Towards Intelligent Tutoring in Novice Programming Environments. In *Proceedings of the 2017 ACM SIGCSE Technical Symposium on Computer Science Education* (Seattle, Washington, USA) (SIGCSE '17). Association for Computing Machinery, New York, NY, USA, 483–488. <https://doi.org/10.1145/3017680.3017762>
- [51] Iulian Radu. 2014. Augmented reality in education: a meta-review and cross-metanalysis. *Personal and ubiquitous computing* 18 (2014), 1533–1543.
- [52] Shwetha Rajaram and Michael Nebeling. 2022. Paper trail: An immersive authoring system for augmented reality instructional experiences. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–16.
- [53] Martina A Rau, Vincent Aleven, and Nikol Rummel. 2009. Intelligent Tutoring Systems with Multiple Representations and Self-Explanation Prompts Support Learning of Fractions. In *Artificial Intelligence in Education*. IOS Press, Amsterdam, Netherlands, 441–448.
- [54] Liam Rigby, Burkhard C. Wünsche, and Alex Shaw. 2021. PiARno - An Augmented Reality Piano Tutor. In *Proceedings of the 32nd Australian Conference on Human-Computer Interaction* (Sydney, NSW, Australia) (OzCHI '20). Association for Computing Machinery, New York, NY, USA, 481–491. <https://doi.org/10.1145/3441000.3441039>
- [55] Kelly Rivers and Kenneth R Koedinger. 2017. Data-driven hint generation in vast solution spaces: a self-improving python programming tutor. *International Journal of Artificial Intelligence in Education* 27 (2017), 37–64.
- [56] Rubik's. 2020. How To Solve A Rubik's Cube | Full Tutorial. https://www.youtube.com/watch?v=qrSOC_o-Fds.
- [57] Rubik's. 2020. YOU CAN DO THE Rubik's Cube Solution Guide. <https://rubiks.com/en-US/solve-it>.
- [58] RubiksCu. 2023. Online Rubik's Cube - Simulator, Solver, Tutorial, Timer. <https://rubikscu.be/>.
- [59] Ceylan Sen, Zeynep Sonay Ay, and Seyit Ahmet Kiray. 2021. Computational thinking skills of gifted and talented students in integrated STEM activities based on the engineering design process: The case of robotics and 3D robot modeling. *Thinking Skills and Creativity* 42 (2021), 100931.
- [60] David Singmaster. 1981. Notes on Rubik's magic cube. (*No Title*) (1981).
- [61] Arthur Tang, Charles Owen, Frank Biocca, and Weimin Mou. 2003. Comparative Effectiveness of Augmented Reality in Object Assembly. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA) (CHI '03). Association for Computing Machinery, New York, NY, USA, 73–80. <https://doi.org/10.1145/642611.642626>
- [62] Michael Thees, Sebastian Kapp, Martin P Strzys, Fabian Beil, Paul Lukowicz, and Jochen Kuhn. 2020. Effects of augmented reality on learning and cognitive load in university physics laboratory courses. *Computers in Human Behavior* 108 (2020), 106316.
- [63] Balasaravanan Thoravi Kumaravel, Fraser Anderson, George Fitzmaurice, Björn Hartmann, and Tovi Grossman. 2019. Loki: Facilitating Remote Instruction of Physical Tasks Using Bi-Directional Mixed-Reality Telepresence. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 161–174. <https://doi.org/10.1145/3332165.3347872>
- [64] Jesslyn Valerie. 2020. *Supporting middle school students' spatial skills through Rubik's Cube play*. Ph.D. Dissertation. University of Minnesota.
- [65] Kurt VanLehn. 2006. The behavior of tutoring systems. *International journal of artificial intelligence in education* 16, 3 (2006), 227–265.
- [66] Matt Whitlock, George Fitzmaurice, Tovi Grossman, and Justin Matejka. 2020. AuthAR: concurrent authoring of tutorials for AR assembly guidance. In *Graphics Interface* 2020.
- [67] Masahiro Yamaguchi, Shohei Mori, Peter Mohr, Markus Tatzgern, Ana Stanescu, Hideo Saito, and Denis Kalkofen. 2020. Video-Annotated Augmented Reality Assembly Tutorials. In *Proceedings of the 33rd Annual ACM Symposium on User*

- Interface Software and Technology*. ACM, Virtual Event USA, 1010–1022. <https://doi.org/10.1145/3379337.3415819>
- [68] Saelyne Yang, Anh Truong, Juho Kim, and Dingzeyu Li. 2025. VideoMix: Aggregating How-To Videos for Task-Oriented Learning. In *Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI '25)*. Association for Computing Machinery, New York, NY, USA, 1564–1580. <https://doi.org/10.1145/3708359.3712144>
- [69] Da-Xing Zeng, Ming Li, Juan-Juan Wang, Yu-Lei Hou, Wen-Juan Lu, and Zhen Huang. 2018. Overview of Rubik's cube and reflections on its application in mechanism. *Chinese Journal of Mechanical Engineering* 31 (2018), 1–12.