2026 ASEE ANNUAL CONFERENCE June 21–24, 2026 | Charlotte Convention Center, NC
Host Sponsor UNC Charlotte
Paper ID #50253

Work in Progress: From Textbooks to Interactive Visualization – 3D STEMulate: An AI-Powered Platform for STEM Education

Dr. Hamid S Timorabadi P.Eng., University of Toronto

Hamid Timorabadi received his B.Sc, M.A.Sc, and Ph.D. degrees in Electrical Engineering from the University of Toronto. He has worked as a project, design, and test engineer as well as a consultant to industry. His research interests include the applications in Engineering Education.

Jiayi Xu, University of Toronto
Kia Salehzadeh

©American Society for Engineering Education, 2026

Work in Progress: From Textbooks to Interactive Visualization – 3D STEMulate: An AI-Powered Platform for STEM Education

Abstract

Understanding spatially complex concepts remains a challenge in science, technology, engineering, and mathematics (STEM) education, particularly when inherently three-dimensional (3D) phenomena are presented using static two-dimensional (2D) representations. Students are often required to mentally reconstruct spatial relationships, which can increase cognitive load and make conceptual understanding more difficult, especially for abstract topics such as electromagnetism.

This paper presents 3D STEMulate, a work-in-progress, web-based platform that uses generative artificial intelligence (AI) to transform textbook diagrams and textual descriptions into interactive 3D visualizations. With minimal user input, the system generates manipulable 3D models that can be explored in real time through a standard web browser, allowing students to examine spatial relationships when needed. The platform also supports an optional low-cost, shared pseudo-3D holographic display for group viewing without requiring immersive head-mounted hardware.

The main contribution of this work is a practical workflow that lowers the barrier to generating and using interactive 3D visualizations within existing learning environments. Preliminary qualitative observations from an exploratory pilot case study in an undergraduate electromagnetism course suggest increased engagement with spatial aspects of the material. However, these findings are limited by a small sample size and are not intended to establish statistically significant learning gains. Ongoing work will focus on broader classroom deployment and more structured evaluation of AI-generated 3D visualization as a tool to support spatial reasoning across STEM disciplines.

1. Introduction

Students in science, technology, engineering, and mathematics (STEM) disciplines are routinely required to understand inherently three-dimensional (3D) phenomena—such as electromagnetic fields, molecular structures, mechanical assemblies, and biological systems—through static two-dimensional (2D) textbook representations. While 2D diagrams are efficient for publication and instruction, they place a significant burden on learners to mentally reconstruct depth, orientation, and spatial relationships that are not explicitly shown. In practice, this translation from 2D to 3D can become a barrier to conceptual understanding, particularly for students with underdeveloped spatial visualization skills.

Decades of research in cognitive science and engineering education have demonstrated that the ability to mentally manipulate and rotate objects in space is a foundational skill for learning in STEM fields [1]. In engineering education, spatial reasoning ability has been shown to correlate strongly with student performance and persistence, especially in courses that rely heavily on geometry, physical structures, and field-based representations [2]. When instructional materials demand high levels of mental rotation and spatial inference without sufficient external support, students may experience increased cognitive load, which can interfere with conceptual understanding and problem-solving [3], [4]. This challenge is particularly evident when students encounter new or abstract concepts for the first time, where both the underlying theory and its spatial representation must be constructed simultaneously.

A range of instructional approaches has been proposed to address these challenges, including physical manipulatives, computer-aided design (CAD) tools, animations, and immersive visualization technologies. Prior work has shown that visualization technologies can improve conceptual understanding by making spatial relationships explicit and interactive [5]. More recently, augmented reality (AR) and virtual reality (VR) systems have been explored as tools for enhancing spatial learning. However, despite their potential, these systems often introduce practical barriers—such as cost, required hardware, instructor training, and scalability—that limit their adoption in large undergraduate courses and resource-constrained environments [6]. As a result, many instructors continue to rely primarily on static textbook figures, even when interactive 3D representations may be beneficial.

To address this gap, we developed 3D STEMulate, a web-based platform that enables students and instructors to transform textbook diagrams and textual descriptions into interactive 3D visualizations using generative artificial intelligence (AI). By uploading an image or entering the name of a scientific concept, the system generates executable rendering code that produces a manipulable 3D model. Users can rotate, zoom, and explore the visualization in real time, allowing them to examine spatial relationships dynamically rather than infer them mentally. From an instructional standpoint, this approach aims to reduce the spatial reasoning demands placed on the learner, consistent with multimedia learning principles that emphasize the externalization of complex cognitive processes [3].

Beyond screen-based interaction, 3D STEMulate also supports projection onto a low-cost holographic display, providing a shared pseudo-3D viewing experience without the need for head-mounted displays or specialized eyewear. This multimodal design allows the platform to be used across individual study, small-group learning, and lecture-based instruction. Importantly, the primary educational contribution of this work is not the hardware itself, but the integration of AI-driven visualization into the learning process at the moment when students encounter conceptual difficulty. In this way, it may help reduce delays and minimize the need to search for external resources during study.

This paper presents 3D STEMulate as a work in progress and focuses on three contributions to engineering education:

1. Identifying limitations of static 2D representations for spatially complex STEM concepts [1], [2];
2. Introducing an AI-powered workflow that lowers the barrier to generating interactive 3D models from existing instructional materials [5]; and
3. Reporting preliminary observations from early classroom deployments that motivate a larger, formal evaluation study [6].

Rather than positioning the system as a replacement for traditional materials, this work explores how on-demand 3D visualization may complement existing instruction in spatially intensive STEM domains.

2. Background and Related Work

2.1 Spatial visualization and learning barriers in STEM

Many STEM concepts are inherently spatial, yet in practice instruction often relies on static two-dimensional (2D) representations that require learners to mentally infer depth, orientation, and structure. Foundational work in cognitive science established mental rotation as a core mechanism supporting spatial reasoning [1], and subsequent research in engineering education showed that spatial skills are predictive of student success and persistence in engineering programs [2]. Importantly, these studies also showed that spatial ability is not fixed and can be improved through targeted instructional interventions.

More recent research has reinforced and extended these findings. For example, a large-scale meta-analysis by Uttal et al. showed that spatial skills are malleable across age groups and instructional contexts, with training leading to measurable improvements that transfer to STEM learning tasks [7]. Similarly, studies examining learning with dynamic and interactive visualizations have shown that learners benefit when spatial relationships are made visible and can be manipulated, rather than inferred solely through mental rotation [8]. Overall, this body of work suggests that instructional tools that reduce the need for internal spatial reconstruction may be particularly helpful for novice learners and students encountering new or abstract concepts.

From a learning-sciences perspective, these challenges can be understood in terms of the limits of working memory. When learners must decode conceptual content while also constructing a three-dimensional (3D) representation from a 2D source, the combined demands can increase cognitive load and reduce learning efficiency, consistent with cognitive load theory [4] and multimedia learning principles [3]. Providing interactive visualizations that externalize spatial structure has therefore been proposed as one way to reduce unnecessary cognitive burden and support conceptual understanding.

2.2 Visualization technologies for STEM education

Visualization technologies have long been studied as tools for supporting conceptual understanding by allowing learners to inspect and manipulate representations rather than relying solely on static diagrams. In general, prior work suggests that interactive visual representations can enhance engagement and learning when they are well aligned with instructional goals and learner needs [5]. At the same time, this body of work highlights that effectiveness depends strongly on accessibility and how easily such tools can be integrated into existing instructional workflows.

In recent years, researchers have explored web-based and browser-delivered 3D visualization tools as more scalable alternatives to specialized software and laboratory-based platforms. These approaches aim to reduce technical barriers while still providing meaningful interactivity, enabling students to explore spatial structures directly within standard learning environments. Studies suggest that such tools can support conceptual exploration and reduce reliance on instructor-provided explanations, particularly when learners are able to manipulate models at their own pace [8]. Despite these advances, a persistent barrier remains the cost of authoring, as creating accurate, course-aligned 3D models typically requires time, expertise, or specialized tools, which limits widespread adoption.

2.3 AR/VR and immersive learning

Immersive technologies such as augmented reality (AR) and virtual reality (VR) have received significant attention as potential solutions for spatial learning challenges in STEM education. Multiple systematic reviews and meta-analyses suggest learning-related benefits [6], while also highlighting both opportunities and practical constraints [10]. However, these studies also report substantial variability in learning outcomes, as well as challenges related to cost, classroom logistics, and instructor training.

More recent experimental studies grounded in multimedia learning theory have shown that immersive environments do not automatically lead to improved learning and may, in some cases, introduce additional cognitive demands if instructional design is not carefully considered [11]. These findings suggest that immersive technologies should be evaluated based on how they support learning processes, rather than on novelty alone. As a result, there is growing interest in approaches that provide some benefits of spatial immersion while avoiding the scalability and hardware constraints associated with head-mounted displays, particularly in large-class or shared-viewing contexts.

2.4 Generative AI for visualization and 3D content creation

Recent advances in generative artificial intelligence (AI), particularly large language models (LLMs), have expanded the possibilities for automating the creation of visual and interactive representations. In the visualization research community, recent surveys highlight growing interest in using generative AI to support visualization workflows, pointing to opportunities to reduce authoring effort while also noting challenges related to controllability, reliability, and evaluation [12], [13]. Much of this work builds on earlier research in natural language interfaces for visualization, where rule-based or intent-driven systems translated user descriptions into visual outputs. More recent LLM-based approaches extend these ideas toward more complete, end-to-end visualization pipelines.

In parallel, rapid progress in text-to-3D and image-to-3D generation reflects a broader trend toward lowering the technical barrier to producing 3D assets [14]. These systems aim to automate aspects of 3D model creation that previously required specialized expertise, often focusing on visual fidelity and scalability for graphics, design, or entertainment applications.

However, most generative AI systems for visualization focus on data analytics, dashboard generation, or high-fidelity content creation, rather than on reconstructing conceptual scientific representations for instructional use. Relatively little work has examined how generative AI can be integrated directly into classroom learning workflows, particularly in ways that support rapid, lightweight model generation aligned with textbook material. In educational contexts, practical considerations such as speed, interpretability, and ease of use may be more important than photorealism.

From an engineering education perspective, there remains a need for systems that leverage existing instructional materials—such as textbook diagrams—and generate manipulable 3D representations quickly enough to support learning during study or instruction, without requiring specialized hardware or instructor-authored models. This work explores this design space by positioning generative AI as a tool for supporting spatial reasoning within everyday learning environments, rather than as a standalone content-creation technology.

3. Technical Implementation

The 3D STEMulate platform is designed as an end-to-end pipeline that enables the transformation of static instructional content into interactive three-dimensional (3D) visualizations. The system architecture consists of four main components: (i) a user-facing web interface, (ii) a middleware processor for handling image and text input, (iii) a backend processing pipeline responsible for interpretation and model generation, and (iv) an optional physical visualization output for shared viewing.

Figure 1 shows the overall system workflow. The process begins at the frontend, where users provide input through the web application. This input is processed in the middleware and then passed to the backend, where concept interpretation and rendering code generation take place. The resulting 3D model is displayed on an interactive web-based canvas and, when desired, projected through an acrylic pyramid display to support shared pseudo-3D visualization.

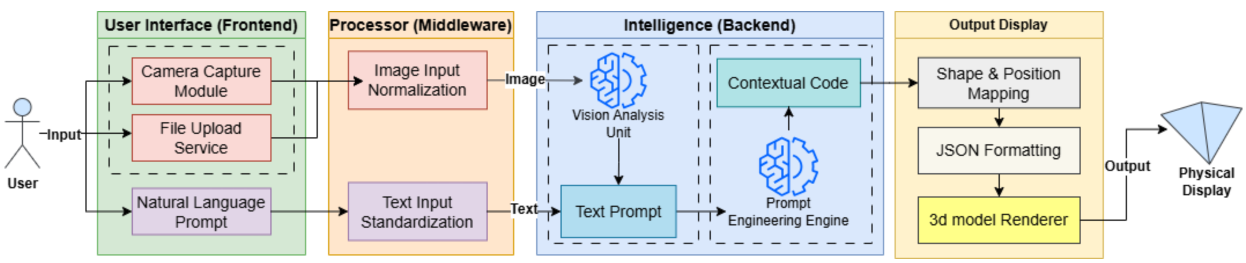

Figure 1. Block diagram of the system architecture and data flow

3.1 Web Application Interface

The web application is designed to be simple and easy to use, allowing students to focus on the visualization itself. Users can interact with the system by capturing a live photo of a textbook diagram, uploading an image, or entering the name of a scientific concept. Once a request is submitted, the generated 3D model is rendered on a full-screen canvas. A caching mechanism is used to reduce response time for commonly requested concepts, helping improve performance during repeated use.

The interface is intentionally minimal, providing only essential controls to avoid distracting from the visualization (Figure 2a). Interaction follows standard conventions for 3D environments and supports both mouse- and touch-based input. Users can rotate, zoom, and translate the model using familiar gestures, allowing them to explore spatial relationships at their own pace (Figure 2b). This interaction design supports exploratory learning by enabling direct manipulation of spatial structures rather than passive observation.

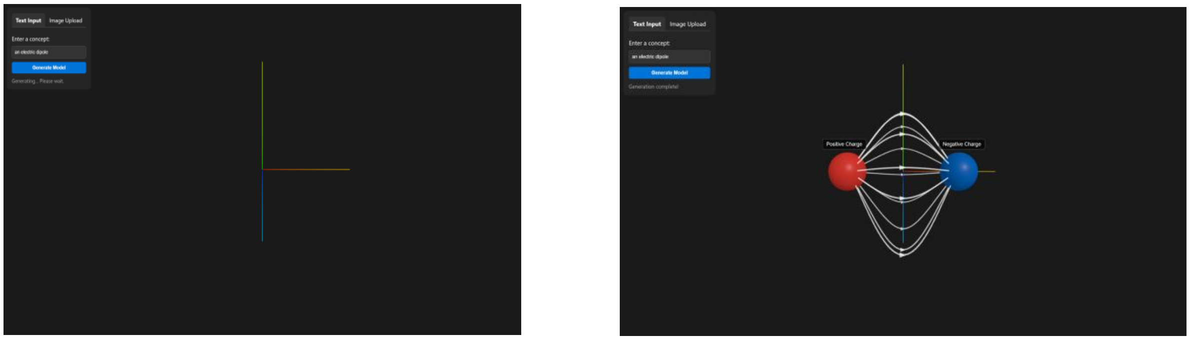

(a) **(b)**

Figure 2. Example of the web application interface during model interaction where *(a)* shows the landing page of the web application and *(b)* shows the interactive output environment where the AI-generated model is rendered in full-screen.

3.2 Backend Processing Pipeline

The backend processing logic implements a two-stage AI-driven pipeline that converts visual or textual input into executable 3D rendering code. In the first stage, the system performs a concept interpretation step, where an AI model analyzes the user input to identify the underlying scientific concept and its key geometric characteristics. This stage focuses on determining structure rather than generating a high-fidelity mesh.

In the second stage, the interpreted concept description is passed to a code-generation model that has been pre-prompted with a library of basic geometric primitives and rendering templates. Using building blocks such as spheres, cylinders, and rectangular prisms, the system generates modular code that assembles a corresponding 3D representation. This approach prioritizes interpretability and real-time performance over photorealism, allowing the system to generate manipulable models quickly enough for use during study or instruction.

The system emphasizes conceptual correctness by constructing models from well-defined geometric primitives, rather than aiming for high visual fidelity. As a result, it is intended to support conceptual exploration rather than replace formal analytical or simulation-based tools.

By relying on compositional primitives instead of static model retrieval, the backend supports a wide range of concepts and user-specified variations, such as highlighting specific components or adjusting visual attributes. This design aligns with the educational goal of enabling flexible exploration rather than producing visually perfect replicas.

3.3 Holographic Display for Shared Visualization

To provide an alternative to standard flat displays, the system optionally supports a pseudo-3D holographic viewing mode using an acrylic pyramid placed above a conventional LCD screen, as shown in Figure 3 and in line with prior pseudo-holographic visualization approaches [15]. This setup presents the same 3D model from four synchronized viewpoints, creating the impression of a volumetric object visible from multiple sides.

The holographic display is particularly suited to lecture halls and group learning environments, where multiple students can observe and discuss the same visualization at the same time. Unlike head-mounted AR or VR systems, this approach does not require individual hardware, calibration, or user isolation, making it more practical for shared instructional settings. To support different use cases, both large-format displays for classroom deployment and smaller, portable units for individual study have been developed.

Importantly, the holographic display is an optional output modality. The core educational functionality of the system remains accessible through standard web browsers.

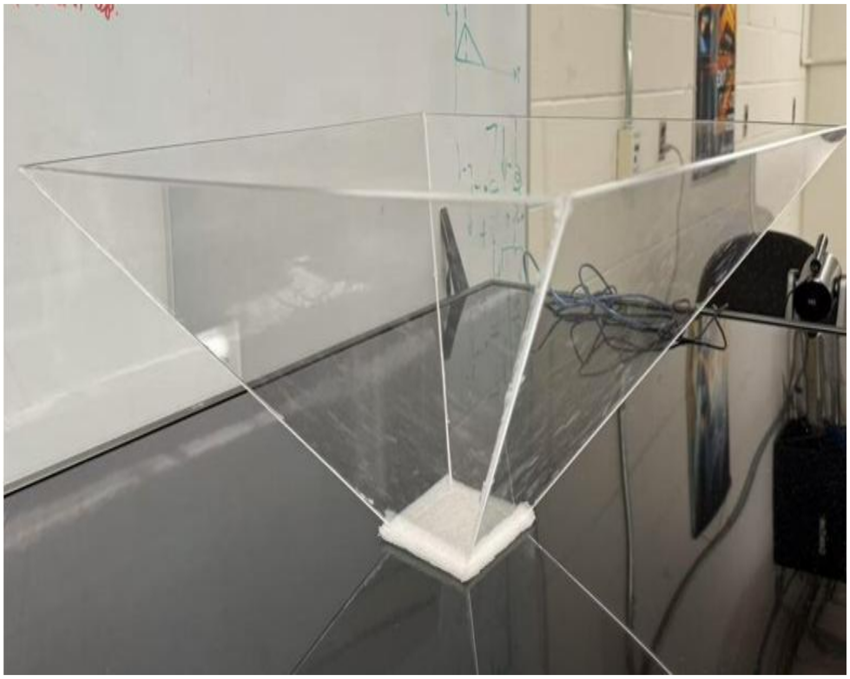

Figure 3. The Acrylic Pyramid Display Setup

4. Educational Impact

Traditional STEM learning often involves a gap between encountering a spatially complex concept and accessing an appropriate three-dimensional (3D) representation to support understanding. When students struggle to interpret a two-dimensional (2D) diagram, they may need to wait for a scheduled laboratory session to interact with a physical model or spend additional time searching for external resources, such as online videos or simulations. These delays can interrupt study and shift attention away from conceptual reasoning toward locating supplementary materials.

The 3D STEMulate platform is intended to reduce this gap by enabling students to generate interactive 3D visualizations at the point of conceptual difficulty. By allowing learners to externalize spatial structure through direct manipulation of a model, the system supports exploratory interaction without requiring prior preparation, specialized software, or instructor intervention. This capability may be particularly useful for electromagnetism topics that are difficult to represent in two dimensions, such as visualizing electric field lines around multiple charges or understanding magnetic field orientation around current-carrying conductors.

Providing immediate access to a manipulable representation may help students focus on spatial relationships rather than searching for additional resources. Unlike pre-authored simulations or visualization repositories, 3D STEMulate enables on-demand generation of interactive representations directly from existing instructional materials. This workflow can be applied across multiple STEM domains.

From an instructional perspective, this capability may complement existing teaching practices by providing on-demand visualization during self-study, group discussion, or live instruction. Rather than replacing traditional representations, the platform augments them by offering a way to transition from static diagrams to interactive spatial exploration. As a result, students may engage with complex structures at their own pace, which could support conceptual understanding, particularly when encountering new or abstract material.

5. Preliminary Results and Work In Progress

The potential educational value of 3D STEMulate lies in its ability to provide a scalable and accessible way to introduce interactive three-dimensional visualization into existing classroom workflows. To explore how the platform might be used in an instructional setting, we conducted an initial pilot case study within an undergraduate electromagnetism course.

This study was exploratory in nature and limited by a small sample size (four students per group). The pilot involved three small groups of students. Two groups consisted of students with no prior exposure to the target concept. One of these groups used traditional instructional resources, while the other used a blended approach that incorporated the 3D STEMulate platform alongside standard materials. A third group consisted of former students who had previously completed the course and revisited the same concept using 3D STEMulate. This setup was intended to explore how interactive 3D visualization might support both first-time exposure and revisiting of concepts.

During the pilot, observations focused on how students interacted with the visualizations, how they used the models in their reasoning and discussions, and how easily they were able to interpret spatial relationships. In groups that used 3D STEMulate, students appeared to engage more with the spatial aspects of the concepts and spent more time exploring the representations.

These observations are qualitative and limited by the small sample size, and should not be interpreted as evidence of statistically significant learning gains. The trends are reflected in the preliminary student perception survey results shown in Figure 4.

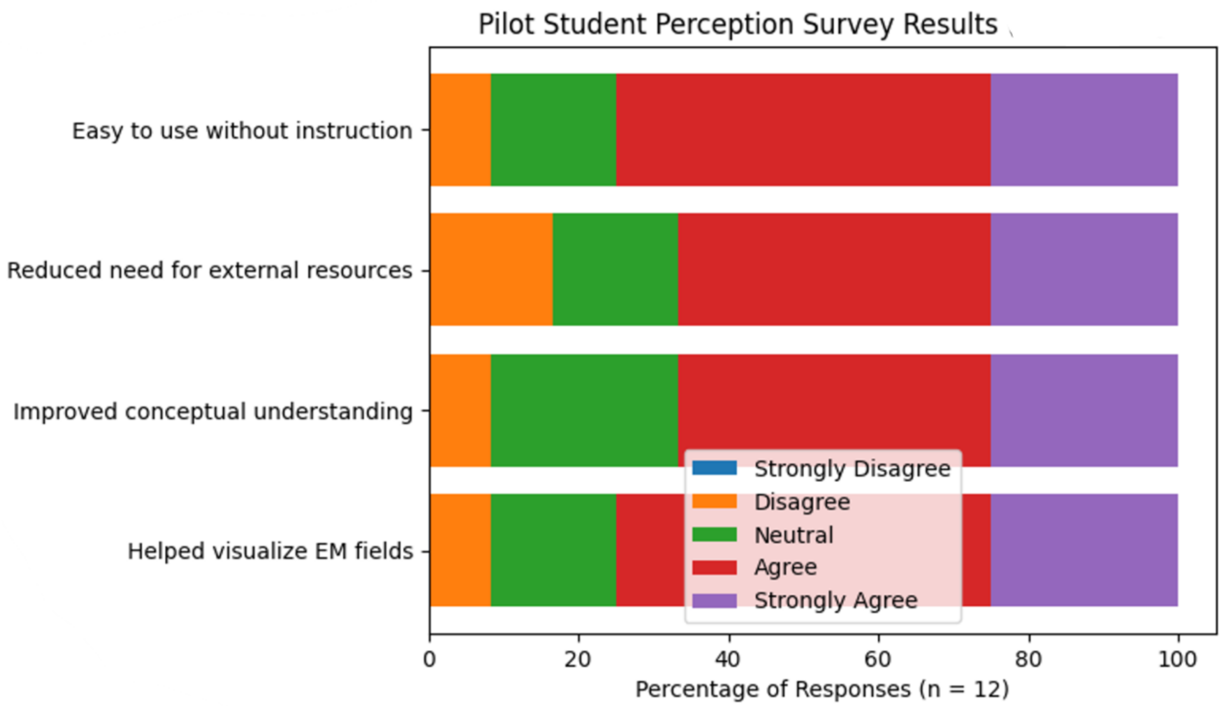

Figure 4. Stacked bar chart of responses to a five-point Likert-scale survey following pilot use of 3D STEMulate in an undergraduate electromagnetism course (n = 12). Results are illustrative and intended to inform future study design, not establish statistically significant outcomes.

6. Conclusion and Future Work

This paper presented 3D STEMulate, a work-in-progress platform that transforms static textbook diagrams and textual descriptions into interactive three-dimensional (3D) visualizations using generative artificial intelligence. By combining AI-driven model generation with web-based interaction and an optional shared pseudo-3D display, the system demonstrates that on-demand spatial visualization can be provided without specialized software or immersive head-mounted hardware. A key contribution is an end-to-end workflow that lowers the barrier for students and instructors to access interactive 3D representations at the point of conceptual difficulty.

Rather than replacing traditional instructional materials, 3D STEMulate is intended to complement existing 2D representations by enabling exploratory inspection of spatial structure. It supports individual study, group discussion, and lecture-based instruction while remaining compatible with typical classroom constraints.

Future work will focus on both technical refinement and educational evaluation. Ongoing development includes expanding supported scientific representations and improving the robustness and controllability of AI-generated models, as well as exploring more compact display configurations. Planned classroom studies will examine how students interact with these visualizations and how they are incorporated into learning workflows. An important direction is understanding how students naturally integrate such tools into their study practices. These efforts aim to clarify when and how AI-assisted visualization can effectively support conceptual learning in STEM education.

References

[1] R. N. Shepard and J. Metzler, “Mental rotation of three-dimensional objects,” *Science*, vol. 171, no. 3972, pp. 701–703, 1971.

[2] S. A. Sorby, “Educational research in developing 3-D spatial skills for engineering students,” *International Journal of Science Education*, vol. 31, no. 3, pp. 459–480, 2009.

[3] R. E. Mayer, *Multimedia Learning*, 2nd ed. Cambridge, U.K.: Cambridge University Press, 2009.

[4] J. Sweller, “Cognitive load during problem solving: Effects on learning,” *Cognitive Science*, vol. 12, no. 2, pp. 257–285, 1988.

[5] H.-K. Wu, S. W.-Y. Lee, H.-Y. Chang, and J.-C. Liang, “Visualization technologies in STEM education,” *Educational Psychology Review*, vol. 25, no. 1, pp. 1–35, 2013.

[6] G. Makransky, G. B. Petersen, and the Immersive Learning Group, “Immersive virtual reality and learning: A meta-analysis,” *Educational Psychology Review*, vol. 31, no. 3, pp. 693–730, 2019.

[7] D. H. Uttal, N. G. Meadow, E. Tipton, L. L. Hand, A. R. Alden, C. Warren, and N. S. Newcombe, “The malleability of spatial skills: A meta-analysis of training studies,” *Psychological Bulletin*, vol. 145, no. 1, pp. 1–26, 2019.

[8] M. Stieff, J. R. Bateman, and D. Uttal, “Learning with dynamic visualizations: A review of instructional features,” *Educational Psychology Review*, vol. 30, no. 3, pp. 1–24, 2018.

[9] J. Buckley, N. Seery, and D. Canty, “A heuristic framework of spatial ability: A review and synthesis,” *International Journal of Technology and Design Education*, vol. 28, no. 3, pp. 1–23, 2018.

[10] J. Radianti, T. A. Majchrzak, J. Fromm, and I. Wohlgenannt, “A systematic review of immersive virtual reality applications for education,” *Computers & Education*, vol. 147, p. 103778, 2020.

[11] J. Parong and R. E. Mayer, “Learning science in immersive virtual reality,” *Journal of Educational Psychology*, vol. 113, no. 4, pp. 749–765, 2021.

[12] S. Liu, B. Kim, J. Hullman, and Z. Wang, “Generative AI for visualization: Opportunities and challenges,” *IEEE Computer Graphics and Applications*, vol. 43, no. 3, pp. 8–18, 2023.

[13] Y. Wu, X. Ma, J. Li, and W. Chen, “AI-assisted visualization authoring: A survey,” *IEEE Transactions on Visualization and Computer Graphics*, early access, 2024.

[14] J. Chen, Y. Zhang, R. Zhang, and H. Li, “Text-to-3D and image-to-3D generation: A survey,” *ACM Computing Surveys*, early access, 2024.

[15] Q. Wang, L. Zhang, and Y. Chen, “Pseudo-holographic visualization systems for education and training,” *IEEE Access*, vol. 8, pp. 180000–180015, 2020.