Spatial Exploration Behavior in XR Learning: Toward Passive Assessment of Embodied Engagement

Yiyin Gu*
Purdue University

Yiqun Zhang†
Purdue University

Miguel A. Feijoo-Garcia‡
Purdue University

Alejandra J. Magana§
Purdue University

Bedrich Benes¶
Purdue University

Voicu Popescu‖
Purdue University

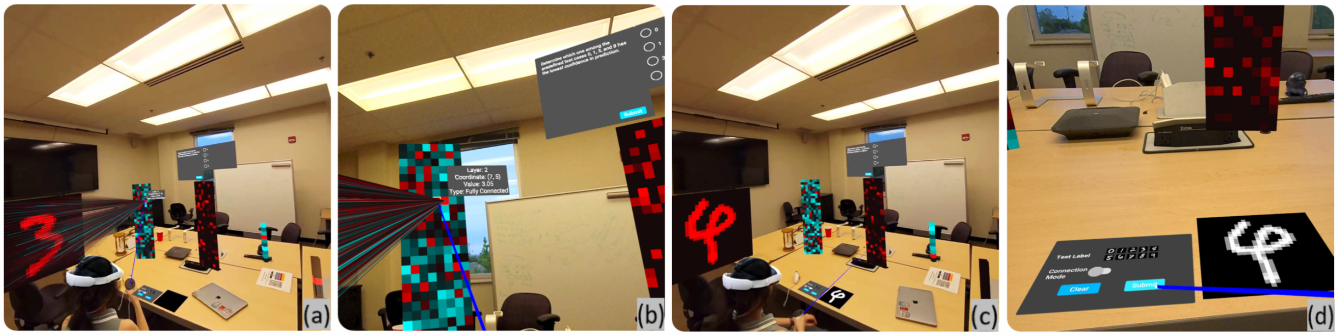

Figure 1: *ACHIEVE* XR lesson: learners explore a 3D neural network visualization through head movement and laser pointer interaction. The headset's video passthrough preserves visibility of the physical environment, allowing learners to see their surroundings, take notes, and interact with instructors.

Abstract

Extended Reality (XR) learning environments generate rich behavioral data through embodied interaction—head movements, gaze patterns, and spatial navigation—that could enable passive assessment without interrupting the learning experience. We investigate spatial exploration behavior in *ACHIEVE*, an XR environment for neural network visualization learning. In a between-subjects study ($N = 56$), XR participants exhibited significantly different spatial behavior than desktop users: over 14 times greater pointer movement ($M = 132.5$ m vs. $M = 9.5$ m), extensive head rotation ($M = 9{,}818°$), and $M = 22.5$ m of head translation during the learning session. We visualize individual exploration patterns through head and pointer trajectory traces, revealing substantial variation in how learners navigate the 3D content. These spatial metrics, automatically captured during learning, represent a promising avenue for passive assessment of embodied engagement—enabling educators to identify struggling learners, provide personalized feedback, and adapt content delivery without intrusive testing. Full learning outcomes and user experience metrics are reported in companion publications; here we focus on spatial behavior as a novel contribution toward spatialized learning analytics in education.

Index Terms: Extended Reality, Embodied Learning, Passive Assessment, Learning Analytics, Spatial Behavior.

*e-mail: gu251@purdue.edu
†e-mail: zhan2897@purdue.edu
‡e-mail: mfeijoog@purdue.edu
§e-mail: admagana@purdue.edu
¶e-mail: bbenes@purdue.edu
‖e-mail: popescu@purdue.edu

1 Introduction

Assessment in education traditionally interrupts the learning process—quizzes, tests, and surveys pause instruction to measure outcomes. This interruption is particularly problematic in immersive learning environments, where breaking presence can disrupt the flow state that makes embodied learning effective [4]. An alternative is passive assessment: evaluating learner engagement, attention, and potentially learning through behavioral data captured during instruction, without intrusive measurement instruments [6].

Extended Reality (XR) environments are uniquely positioned to enable passive assessment because they inherently generate rich behavioral data. XR headsets continuously track head position and orientation at high frequency; controllers or hand tracking capture pointer movements and gestures; and the spatial nature of the content means that navigation itself becomes a measurable learning behavior. This data stream is a natural byproduct of embodied learning—the physical movements that constitute the learning experience itself [14].

Unlike traditional metrics such as click logs or time-on-task, spatial exploration patterns capture how learners engage with 3D content: which areas they examine, how thoroughly they explore, their navigation strategies, and how their attention shifts over time. A learner who systematically examines each component of a 3D visualization demonstrates different engagement than one who fixates on a single region or moves erratically. These behavioral signatures could provide insight into learning processes without requiring learners to stop and answer questions.

We investigate spatial exploration behavior using *ACHIEVE* [13], an XR environment for neural network visualization learning (Fig. 1). In a between-subjects study (N = 56) comparing XR to desktop conditions, we recorded head tracking and pointer trajectory data throughout the learning session. Companion publications report the system design [13] and comprehensive learning outcomes with user experience metrics [9]. We report a quantitative analysis of spatial exploration behavior in XR versus desktop learning, revealing significant differences in

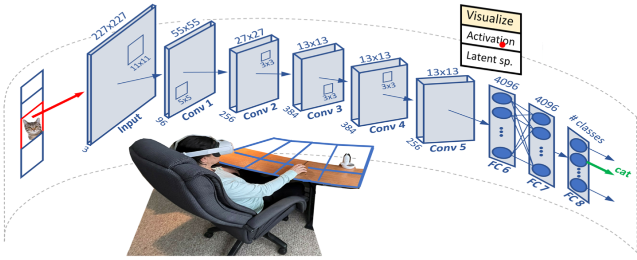

Figure 2: *ACHIEVE* interface: a convolutional neural network rendered in semi-cylindrical layout, spanning approximately 240° around the seated learner.

interaction patterns. Through trajectory visualizations, we show substantial individual variation in how learners navigate 3D content and discuss spatial behavior as a promising passive assessment modality for XR learning analytics.

2 Related Work

2.1 Embodied Learning in XR

Embodied cognition theory posits that learning is enhanced through physical interaction with the environment [26]. This perspective suggests that cognitive processes are grounded in bodily experience—we understand abstract concepts partly through physical metaphors and sensorimotor engagement [14]. XR environments leverage embodied cognition by enabling learners to navigate, manipulate, and explore 3D content through natural body movements rather than abstract mouse clicks [16].

Meta-analyses indicate VR's effectiveness for knowledge acquisition [21], with particular benefits for spatial understanding in STEM [24]. The sense of presence—feeling "there" in the virtual environment [4]—appears to enhance motivation and engagement. However, most studies measure learning outcomes through traditional assessments administered after the VR experience, missing the opportunity to leverage the behavioral data generated during the experience itself.

2.2 Learning Analytics and Passive Assessment

Learning analytics uses behavioral data to understand and optimize learning [25]. Traditional online learning generates click streams, time-on-page, and navigation patterns that can predict student performance and identify struggling learners [20]. In XR, richer multimodal data streams become available: gaze direction, head pose, hand movements, and locomotion patterns [2].

Eye tracking has been extensively used for attention analysis in VR, revealing where learners look and for how long [8]. Movement patterns have been used to infer cognitive load and emotional states [15]. However, most work has focused on individual modalities in isolation. The potential of spatial exploration patterns—the integrated record of how learners physically navigate and interact with 3D content—as a holistic indicator of engagement and learning strategy remains underexplored. Our work contributes to this gap by analyzing head and pointer trajectories as candidate metrics for passive learning assessment.

2.3 Neural Network Visualization for Education

Teaching neural networks is challenging due to their "black box" nature [12]. The high-dimensional, layered architectures with thousands of parameters resist intuitive understanding. 2D visualization tools like TensorFlow Playground [27] help illustrate network behavior through interactive diagrams, but they struggle to convey the true multidimensional structure and can impose high cognitive load

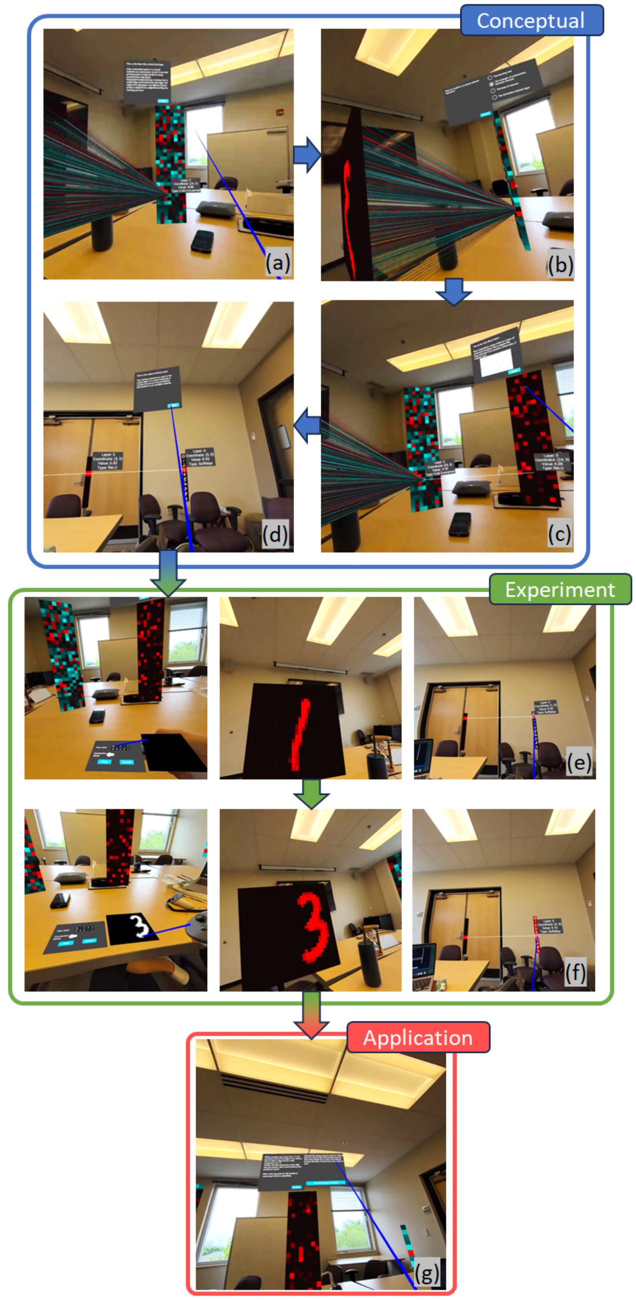

Figure 3: Lesson flow: (a–d) conceptual phase; (e–f) experiment phase with digit inputs; (g) application phase.

when learners must mentally reconstruct 3D relationships from 2D projections [5].

Prior work introduced *ACHIEVE*, an XR environment that renders neural networks as explorable 3D structures [13]. The system enables learners to physically navigate around and through network layers, observe activation patterns in real-time, and test the network with handwritten digit inputs. A controlled study comparing *ACHIEVE* to desktop visualization found that XR participants reported higher satisfaction and were more likely to recommend the experience, though learning outcomes were similar across conditions; full results are reported in [9]. The present paper focuses specifically on the spatial behavior data collected during that study, examining what head tracking and pointer trajectories reveal about how learners engage with 3D educational content.

3 System and Study Overview

ACHIEVE [13] is an XR environment where learners explore a 3D neural network visualization rendered in a semi-cylindrical layout around them (Fig. 2). Running on Meta Quest 3 [22] with video passthrough, learners navigate by turning their head and select network elements with a laser pointer. The lesson covers neural network concepts using MNIST digit recognition, with learners able to draw digits and observe real-time inference.

Throughout each session, we logged head position and rotation (6DOF) and laser pointer position at 30Hz. For desktop participants, we logged mouse cursor position. This continuous recording enables comparison of spatial exploration behavior between modalities without interrupting the learning experience.

4 Study Design

We conducted a between-subjects study (N = 56) comparing spatial exploration in XR vs. desktop. Full methodology, learning outcomes, and user experience metrics are reported in [9]; here we summarize methods relevant to spatial behavior analysis.

Participants: 56 participants (29 XR, 27 desktop), mean age 20.4 years. Most had limited XR experience (51.7% never used XR).

Conditions: XR participants used *ACHIEVE* on Meta Quest 3; desktop participants used a 2D version with identical content. The lesson followed three phases (Fig. 3): conceptual exploration, experimentation with digit inputs, and reflective application.

Spatial Metrics: For XR, we computed: (1) total laser pointer movement distance, (2) head translation distance, and (3) cumulative head rotation. For desktop, we computed mouse cursor movement distance. All participants also completed the SSQ [17] (XR only) and the session was timed.

5 Spatial Exploration Results

5.1 Quantitative Comparison

Table 1 summarizes spatial exploration metrics. There was a significant difference in pointer/cursor movement between conditions. EC participants moved their laser pointer an average of $M = 132.5$ m during the session, compared to $M = 9.5$ m of mouse movement for CC participants ($p < 0.001$, Cliff's $\delta = 0.93$, power $= 0.91$). This 14-fold difference reflects the fundamentally different interaction modalities and indicates that EC participants engaged with a much larger area of the interface. The high statistical power (0.91) provides confidence that this effect would reliably replicate.

EC participants also exhibited substantial head movement: mean translation of $M = 22.5$ m and cumulative rotation of $M = 9{,}818^\circ$ (equivalent to approximately 27 full rotations) during the approximately 12-minute session. This indicates active physical engagement with the 240° visualization space, suggesting that learners leveraged the expanded spatial affordances of the XR environment.

EC participants also spent significantly more time on task ($M = 723.7$ s vs. $M = 578.5$ s, $p = 0.01$, Cohen's $d = 0.72$, power $= 0.76$). This medium effect size suggests meaningful differences in engagement duration, though increased time could reflect either deeper engagement with the material or additional overhead of XR interaction. Click counts did not differ significantly between conditions ($M = 117.7$ vs. $M = 91.9$, $p = 0.09$, Cliff's $\delta = 0.26$). This lack of difference, despite the large differences in spatial movement, suggests that XR promotes a different, more continuous style of interaction rather than simply more frequent discrete actions—learners in EC explored more space through smooth physical movements rather than additional point-and-click operations. Cybersickness was minimal (SSQ $M = 19.4$), well within acceptable ranges for educational use, and notably all 29 EC participants completed the session [1].

5.2 Individual Exploration Patterns

Beyond aggregate statistics, the trajectory data reveals substantial individual variation in exploration strategies. The standard deviations in Table 1 are notable: pointer distance $SD = 60.6$ m for EC (46% of the mean) and head rotation $SD = 3{,}897^\circ$ (40% of the mean), indicating that learners adopted quite different approaches to navigating the same content. We visualize these differences using space-time cubes [18], where horizontal axes encode spatial position on the cylindrical display (angle and height) and the vertical axis encodes time. Horizontal planes partition the visualization into the three lesson phases. Fig. 4 shows head rotation traces for two EC participants with contrasting exploration patterns.

The high-coverage participant (Fig. 4a) systematically explored the entire 240° visualization, with traces spreading broadly across each phase plane and evidence of revisitation as attention returned to earlier content regions. The low-coverage participant (Fig. 4b) progressed more linearly through the lesson, with tighter trace clusters that rise vertically through the cube. Similar variation appears in laser pointer traces (Fig. 5): one participant (a) shows a clustered pattern with focused interaction within each phase, while another (b) exhibits a distributed pattern with broader coverage across the display throughout all phases.

This variation suggests that spatial exploration patterns could differentiate learning behaviors and engagement strategies. The range of pointer movement across EC participants (from approximately 70 m to over 250 m) represents meaningful behavioral differences that would be invisible in traditional assessments. Whether broader exploration correlates with better learning outcomes, or whether different strategies suit different learners, remains an open question for future investigation.

6 Discussion: Toward Passive Assessment

6.1 Spatial Behavior as Learning Analytics

The spatial exploration data we collected—head tracking, pointer trajectories, and temporal patterns—represents a form of passive assessment that emerges naturally from embodied XR interaction. Unlike surveys or quizzes that interrupt learning and break presence, this behavioral data is captured continuously without learner awareness or effort. The learner simply engages with the content; the assessment happens invisibly in the background.

Several characteristics make spatial behavior particularly promising for learning analytics:

Richness: Spatial trajectories encode not just where learners looked but how they explored—systematically or randomly, thoroughly or superficially, with focused attention or broad scanning. The 14-fold difference in pointer movement between XR and desktop conditions suggests that XR elicits fundamentally different—and potentially more informative—exploration behavior.

Temporal dynamics: The color-coded traces (Figs. 4–5) reveal exploration patterns over time. A learner who revisits earlier layers after seeing later ones may be building connections between concepts; one who never returns may be passively progressing without integration. These temporal signatures could distinguish active learning from passive exposure.

Individual differences: The substantial variation between high and low-coverage participants demonstrates that spatial metrics capture individual learning styles or engagement levels. Some learners systematically survey the entire visualization and revisit earlier content; others progress linearly through the lesson without looking back. Whether these strategies correlate with different learning outcomes is a question for future investigation.

Non-intrusiveness: Because spatial data is a byproduct of normal XR interaction, it requires no additional instrumentation or

Metric	Condition	Mean	Median	SD	Test	p	Effect Size	Power
Pointer/Cursor Distance (m)	EC (XR)	132.5	120.7	60.6	Mann-Whitney U	$< .001$	Cliff's $\delta = 0.93$	0.91
	CC (Desktop)	9.5	7.6	5.2				
Head Translation (m)	EC (XR)	22.5	19.9	9.2	–	–	–	–
Head Rotation (°)	EC (XR)	9818.7	9594.0	3897.2	–	–	–	–
Session Time (s)	EC (XR)	723.7	731.0	179.8	Independent t-test	.01	Cohen's $d = 0.72$	0.76
	CC (Desktop)	578.5	584.0	212.2				
Click Count	EC (XR)	117.7	99.0	59.5	Mann-Whitney U	.09	Cliff's $\delta = 0.26$	0.15
	CC (Desktop)	91.9	83.0	38.4				

Table 1: Spatial exploration and interaction metrics comparing EC (XR, $n = 29$) and CC (Desktop, $n = 27$).

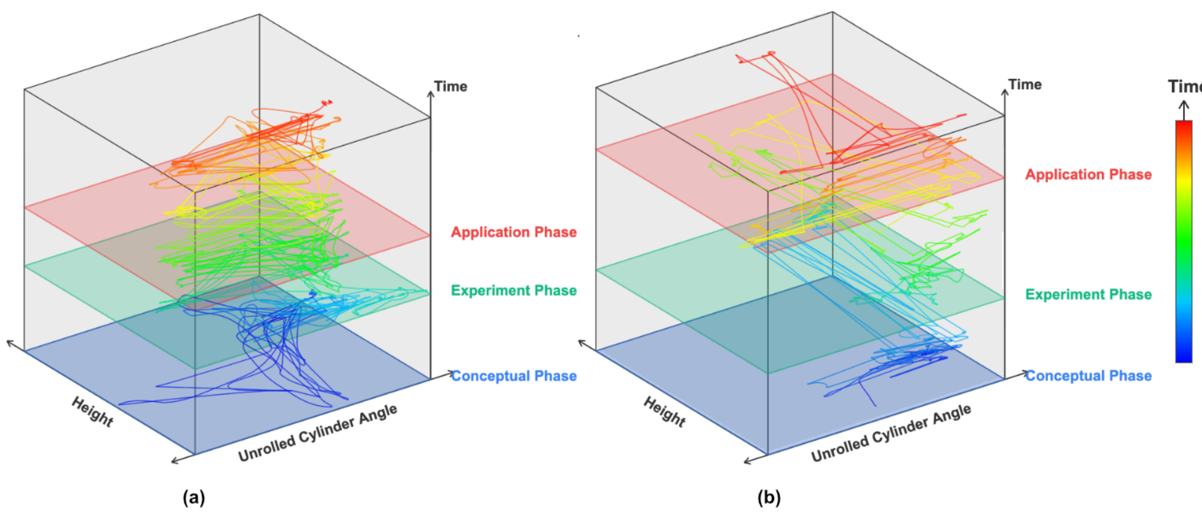

Figure 4: Space-time cube visualization of head rotation traces. The horizontal axes represent spatial position (unrolled cylinder angle and height); the vertical axis represents time. Horizontal planes delineate lesson phases: Conceptual (blue), Experiment (green), and Application (red). (a) High-coverage participant with broad exploration across the display surface. (b) Low-coverage participant with more focused, linear progression. Color gradient indicates time (blue→red).

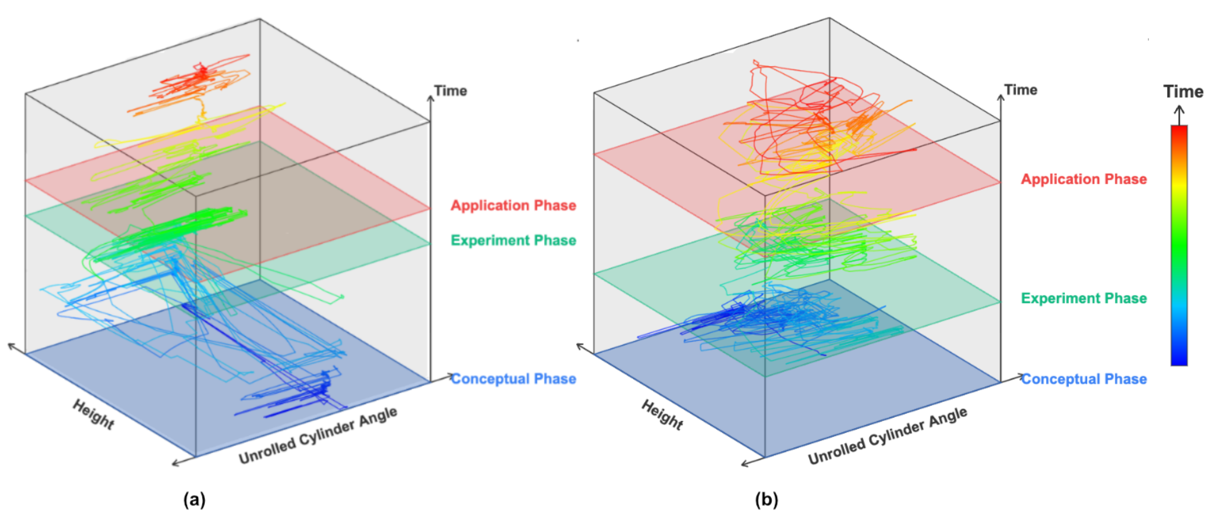

Figure 5: Space-time cube visualization of laser pointer traces. Same axes and phase planes as Fig. 4. (a) Clustered exploration pattern with focused, localized interaction within each phase. (b) Distributed exploration pattern with broader coverage spanning the display. The contrasting trace densities reveal different interaction strategies.

learner cooperation. This contrasts with physiological sensors (which require wearing additional hardware) or self-reports (which require interruption). Importantly, spatial behavior data can be collected remotely without supervision, enabling scalable assessment of learners using XR content independently—an advantage for distributed or self-paced learning scenarios.

6.2 Potential Applications

Similar applications have been explored in 2D learning environments, but often face limitations: click-stream data provides limited insight into attention and engagement, while eye tracking—though more informative—raises privacy concerns as learners may be reluctant to have their faces recorded [19], and requires specialized hardware. XR spatial behavior offers a middle ground: richer behavioral signals than mouse tracking, without the privacy concerns of facial video. If spatial behavior proves predictive of learning outcomes, several applications become possible:

Real-time intervention: Following the model of intelligent tutoring systems that monitor student behavior and provide adaptive feedback [29], an XR learning system could detect when a learner is not exploring thoroughly and prompt them to examine neglected areas, or offer additional scaffolding when exploration patterns suggest confusion.

Adaptive content: Similar to how 2D learning environments adapt based on click patterns and time-on-task [3], the system could adjust content difficulty or pacing based on exploration behavior—slowing down for learners who need more time to explore, or advancing for those who have thoroughly examined the current material.

Instructor dashboards: As demonstrated in learning analytics dashboards for online courses [30], aggregated spatial behavior could help instructors identify which content areas receive insufficient attention across a class, informing lesson redesign.

Learning research: Spatial behavior data provides a window into cognitive processes during embodied learning that traditional assessments cannot capture, extending the use of multimodal log data for understanding learning processes [7] to the embodied XR domain.

6.3 On the Pedagogical Value of Neural Network Visualization

A fundamental question underlying this work is whether spatial visualization of neural networks actually helps students understand them, independent of the delivery medium. Neural networks are mathematical abstractions—compositions of linear transformations and nonlinear activations operating in high-dimensional spaces. Rendering them as 3D node-and-edge structures provides an intuitive visual metaphor, but this metaphor may also introduce misconceptions or oversimplifications [12].

Prior work suggests that interactive visualizations can help learners develop intuitions about neural network behavior. TensorFlow Playground [27], for instance, has been widely adopted in education precisely because it makes abstract concepts (decision boundaries, hidden representations, overfitting) perceptible. However, the relationship between visualization and understanding is not straightforward: learners may develop "illusions of understanding" from engaging visualizations without acquiring transferable knowledge [23]. The fact that XR and desktop participants in our study achieved similar learning outcomes [9] despite vastly different exploration behavior raises the question of whether the spatial affordances of XR meaningfully enhance conceptual understanding or primarily enhance engagement and motivation.

We do not resolve this question here—it requires dedicated comparative studies that isolate visualization approach from delivery medium. However, we note that the value of spatial exploration data for learning analytics does not depend on resolving it. Even if 3D visualization provides limited pedagogical advantage over 2D alternatives, the behavioral data it generates may still enable valuable insights into learner engagement and attention that support adaptive instruction.

6.4 Open Questions and Future Work

This work raises questions we cannot yet answer but that define a research agenda:

Correlation with learning: Does broader exploration lead to better understanding? A companion study [9] reports learning outcomes, but correlating specific spatial patterns (coverage, revisitation, temporal dynamics) with assessment performance requires further analysis. We plan to investigate whether trajectory features predict quiz scores or performance on novel problems that require applying learned concepts to new situations.

Disambiguating exploration types: High movement could reflect engaged, purposeful exploration or confused, aimless searching. Distinguishing these requires combining spatial data with other modalities (e.g., task performance, verbal responses, physiological signals). Future work should develop classification methods that can identify productive vs. unproductive exploration.

Generalizability: These patterns emerged in neural network visualization; whether similar dynamics apply to other XR learning contexts (e.g., anatomy, chemistry, historical sites) remains unknown. The specific meaning of spatial patterns may depend on content structure and learning objectives.

Individual differences: Do exploration patterns reflect stable individual traits or situational factors (content difficulty, fatigue)? In particular, might spatial exploration patterns correlate with preferred learning styles [10]? Visual learners might exhibit broader scanning patterns to take in the overall structure, while kinesthetic or read/write learners might focus more narrowly on interactive elements or text annotations. If spatial behavior maps onto learning style preferences, it could enable automatic adaptation of content presentation. Longitudinal studies could reveal whether learners exhibit consistent spatial signatures across sessions.

6.5 The Ambiguity of Embodied Engagement Metrics

A critical challenge for interpreting spatial exploration data is that similar movement patterns can arise from fundamentally different cognitive states. High physical activity in an XR learning environment does not necessarily equate to high learning or deep engagement—it may equally signal confusion, frustration, or a "search for clarity" as the learner struggles to locate relevant information or make sense of the content [28].

Consider two learners who both exhibit extensive head rotation and pointer movement: one may be systematically surveying the 3D neural network visualization, building a coherent mental model by examining each layer in relation to others; another may be moving erratically, unable to identify which elements are important or how they connect. The raw spatial metrics—total distance traveled, coverage area, movement velocity—could be nearly identical despite representing opposite ends of productive engagement.

This ambiguity has several implications for the development of XR learning analytics:

Movement alone is insufficient. Spatial exploration metrics must be combined with other indicators to support valid inferences about learning [31]. Task performance data (e.g., accuracy on in-lesson interactions), temporal patterns (e.g., systematic versus erratic trajectories), and physiological signals could help disambiguate purposeful exploration from confused searching.

Context shapes interpretation. The same movement pattern may signal engagement in one lesson phase and confusion in another. During initial orientation, broad scanning is expected and likely productive; during focused problem-solving, it may indicate difficulty. Analytics systems must account for the instructional context when interpreting spatial behavior.

Caution in automated assessment. Before spatial metrics can be used for consequential decisions (e.g., identifying struggling learners, adapting content difficulty), validation studies must establish when high movement predicts positive versus negative outcomes [11]. Premature deployment of movement-based assessment could misidentify engaged learners as struggling or vice versa.

We acknowledge this fundamental ambiguity as a limitation of embodied engagement metrics and emphasize that the spatial behavior data we present should be understood as a *candidate* signal requiring validation, not a confirmed indicator of learning or engagement.

6.6 Limitations

Our sample ($N = 56$) limits pattern generalization and statistical power for subgroup analyses. The comparison between XR pointer movement and desktop mouse movement conflates modality differences with engagement differences—it is possible that XR simply requires more movement to accomplish the same cognitive work. We did not manipulate content or task difficulty to probe how spatial behavior changes under different conditions. Finally, this analysis is exploratory; confirming the predictive value of spatial metrics requires replication with held-out data.

7 Conclusion

We presented an analysis of spatial exploration behavior in *ACHIEVE*, an XR environment for neural network visualization learning [13]. XR participants exhibited dramatically greater spatial exploration than desktop users, with substantial individual variation in exploration patterns. These findings, combined with full learning outcomes reported in [9], suggest that spatial behavior data could serve as a passive assessment modality for XR learning environments. Future work will investigate correlations between exploration patterns and learning outcomes, toward automated detection of engagement and comprehension in embodied learning.

Acknowledgments

This work was supported by the National Science Foundation under Grant Numbers 2212200, 2219842, 2309564, 2318657 and 2417510. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

[1] P. Bimberg, T. Weissker, and A. Kulik. On the usage of the simulator sickness questionnaire for virtual reality research. In *2020 IEEE conference on virtual reality and 3D user interfaces abstracts and workshops (VRW)*, pp. 464–467. IEEE, 2020. 3

[2] P. Blikstein and M. Worsley. Multimodal learning analytics and education data mining: using computational technologies to measure complex learning tasks. *Journal of Learning Analytics*, 3(2):220–238, Sept. 2016. doi: 10.18608/jla.2016.32.11 2

[3] P. Brusilovsky. Methods and techniques of adaptive hypermedia. *User Modeling and User-Adapted Interaction*, 6(2–3):87–129, July 1996. doi: 10.1007/bf00143964 5

[4] O. I. Caldas, M. Mauledoux, O. F. Aviles, and C. Rodriguez-Guerrero. Breaking presence in immersive virtual reality toward behavioral and emotional engagement. *Computer Methods and Programs in Biomedicine*, 248:108124, May 2024. doi: 10.1016/j.cmpb.2024.108124 1, 2

[5] P. Chandler and J. Sweller. Cognitive load theory and the format of instruction. *Cognition and instruction*, 8(4):293–332, 1991. 2

[6] M. A. A. Dewan, M. Murshed, and F. Lin. Engagement detection in online learning: a review. *Smart Learning Environments*, 6(1), Jan. 2019. doi: 10.1186/s40561-018-0080-z 1

[7] D. Di Mitri, J. Schneider, M. Specht, and H. Drachsler. From signals to knowledge: A conceptual model for multimodal learning analytics. *Journal of Computer Assisted Learning*, 34(4):338–349, July 2018. doi: 10.1111/jcal.12288 5

[8] A. Duchowski. *Eye Tracking Methodology: Theory and Practice*. Methods in molecular biology. Springer London, 2007. 2

[9] M. Feijoo-Garcia, Y. Zhang, Y. Gu, A. Magana, B. Benes, and V. Popescu. Exploring extended reality (xr) in teaching ai: A comparative study of xr and desktop environments. In *Proceedings of the 20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, p. 472–482. SCITEPRESS - Science and Technology Publications, 2025. doi: 10.5220/0013141000003912 1, 2, 3, 5, 6

[10] N. Fleming and C. Mills. Not another inventory, rather a catalyst for reflection. *To improve the academy*, 11, 06 1992. doi: 10.1002/j.2334-4822.1992.tb00213.x 5

[11] D. Gašević, S. Dawson, and G. Siemens. Let's not forget: Learning analytics are about learning. *TechTrends*, 59(1):64–71, Dec. 2014. doi: 10.1007/s11528-014-0822-x 6

[12] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. http://www.deeplearningbook.org. 2, 5

[13] Y. Gu, M. A. Feijoo-Garcia, Y. Zhang, A. J. Magana, B. Benes, and V. Popescu. An xr environment for ai education: Design and first implementation. In *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, p. 157–162. IEEE, Mar. 2024. doi: 10.1109/vrw62533.2024.00032 1, 2, 3, 6

[14] S. Hrach. *Minding Bodies: How Physical Space, Sensation, and Movement Affect Learning*. Teaching and learning in higher education. West Virginia University Press, 2021. 1, 2

[15] C. Ilyas, R. Nunes, K. Nasrollahi, M. Rehm, and T. Moeslund. Deep emotion recognition through upper body movements and facial expression. In *Proceedings of the 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*. SCITEPRESS - Science and Technology Publications, 2021. doi: 10.5220/0010359506690679 2

[16] S. Jang, J. M. Vitale, R. W. Jyung, and J. B. Black. Direct manipulation is better than passive viewing for learning anatomy in a three-dimensional virtual reality environment. *Computers & Education*, 106:150–165, Mar. 2017. doi: 10.1016/j.compedu.2016.12.009 2

[17] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The international journal of aviation psychology*, 3(3):203–220, 1993. 3

[18] M.-J. Kraak. The space-time cube revisited from a geovisualization perspective. In *21st international Cartographic Conference, ICC 2003: Cartographic renaissance*, pp. 1988–1996. International Cartographic Association, 2003. 3

[19] J. L. Kröger, O. H.-M. Lutz, and F. Müller. *What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking*, p. 226–241. Springer International Publishing, 2020. doi: 10.1007/978-3-030-42504-3_15 5

[20] L. P. Macfadyen and S. Dawson. Mining lms data to develop an "early warning system" for educators: A proof of concept. *Computers & Education*, 54(2):588–599, Feb. 2010. doi: 10.1016/j.compedu.2009.09.008 2

[21] Z. Merchant, E. T. Goetz, L. Cifuentes, W. Keeney-Kennicutt, and T. J. Davis. Effectiveness of virtual reality-based instruction on students' learning outcomes in k-12 and higher education: A meta-analysis. *Computers & education*, 70:29–40, 2014. 2

[22] Meta. Quest 3: New Mixed Reality Headset. https://www.meta.com/quest/quest-3/. Accessed: 2024-9-6. 3

[23] F. Paas, J. E. Tuovinen, H. Tabbers, and P. W. M. Van Gerven. Cognitive load measurement as a means to advance cognitive load theory. *Educational Psychologist*, 38(1):63–71, Jan. 2003. doi: 10.1207/s15326985ep3801_8 5

[24] J. Radianti, T. A. Majchrzak, J. Fromm, and I. Wohlgenannt. A systematic review of immersive virtual reality applications for higher education: Design elements, lessons learned, and research agenda. *Computers & Education*, 147:103778, 2020. doi: 10.1016/j.compedu.2019.103778 2

[25] G. Siemens. Learning analytics: The emergence of a discipline. *American Behavioral Scientist*, 57(10):1380–1400, Aug. 2013. doi: 10.1177/0002764213498851 2

[26] A. Skulmowski and G. D. Rey. Embodied learning: introducing a taxonomy based on bodily engagement and task integration. *Cognitive research: principles and implications*, 3:1–10, 2018. 2

[27] D. Smilkov, S. Carter, D. Sculley, F. B. Viégas, and M. Wattenberg. Direct-manipulation visualization of deep networks. *arXiv preprint arXiv:1708.03788*, 2017. 2, 5

[28] J. Sweller, P. Ayres, and S. Kalyuga. *Cognitive Load Theory*. Springer New York, 2011. doi: 10.1007/978-1-4419-8126-4 5

[29] K. VanLEHN. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4):197–221, Oct. 2011. doi: 10.1080/00461520.2011.611369 5

[30] K. Verbert, S. Govaerts, E. Duval, J. L. Santos, F. Van Assche, G. Parra, and J. Klerkx. Learning dashboards: an overview and future research opportunities. *Personal and Ubiquitous Computing*, 18(6):1499–1514, 2014. doi: 10.1007/s00779-013-0751-2 5

[31] M. Worsley and P. Blikstein. A multimodal analysis of making. *International Journal of Artificial Intelligence in Education*, 28(3):385–419, Nov. 2017. doi: 10.1007/s40593-017-0160-1 5