

Skill-Adaptive Ghost Instructors: Enhancing Retention and Reducing Over-Reliance in VR Piano Learning

Tzu-Hsin Hsieh
Computer Graphics and Visualization
Delft University of Technology
Delft, Netherlands
thhsieh@tudelft.nl

Cassandra Michelle
Stefanie Visser
Delft University of Technology
Delft, Netherlands
c.m.s.visser@outlook.com

Elmar Eisemann
Computer Graphics and Visualization
Delft University of Technology
Delft, Netherlands
e.eisemann@tudelft.nl

Ricardo Marroquim
Computer Graphics and Visualization
TU Delft
Delft, Netherlands
r.marroquim@tudelft.nl

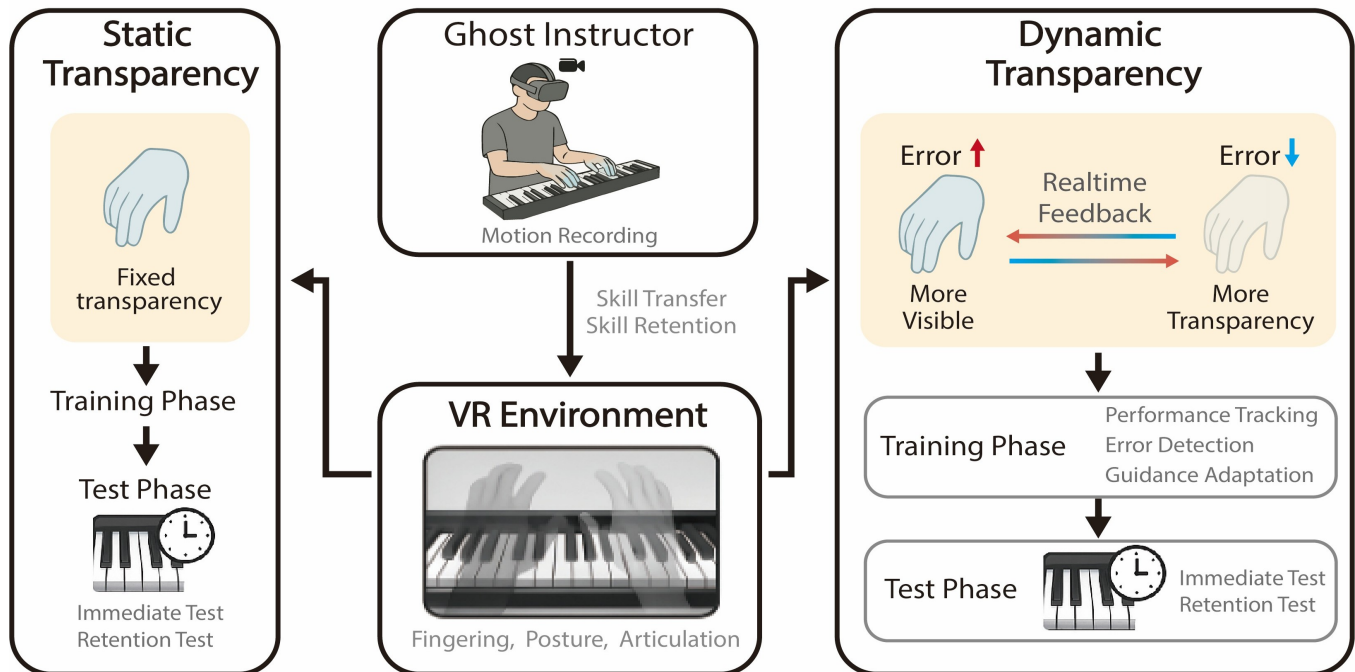


Figure 1: Skill-Adaptive Ghost Instructors: ghost-hand guidance with two modes: Static (fixed opacity) and Dynamic (opacity adapts to the learner’s performance).

Abstract

Motor-skill learning systems in XR rely on persistent cues. However, constant cueing can induce overreliance and erode memorization and skill transfer. We introduce a skill-adaptive, dynamically transparent ghost instructor whose opacity adapts in real time to learner performance. In a first-person perspective, users observe a ghost hand executing piano fingering with either a static or a performance-adaptive transparency in a VR piano training application. We conducted a within-subjects study (N=30), where learners

practiced with traditional *Static* (fixed-transparency) and our proposed *Dynamic* (performance-adaptive) modes and were tested without guidance immediately and after a 10-minute retention interval. Relative to *Static*, the *Dynamic* mode yielded higher pitch and fingering accuracy and limited error increases, with comparable timing. These findings suggest that adaptive transparency helps learners internalize fingerings more effectively, reducing dependency on external cues and improving short-term skill retention within immersive learning environments. We discuss design implications for motor-skill learning and outline directions for extending this approach to longer-term retention and more complex tasks.



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '26, Barcelona, Spain
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3791437>

CCS Concepts

• **Human-centered computing** → **Virtual reality**; *Empirical studies in HCI*; **Interaction devices**.

Keywords

Virtual Reality, Feedback-Aware, Skill Retention, Adaptive Learning, Passthrough, Hand Tracking

ACM Reference Format:

Tzu-Hsin Hsieh, Cassandra Michelle Stefanie Visser, Elmar Eisemann, and Ricardo Marroquim. 2026. Skill-Adaptive Ghost Instructors: Enhancing Retention and Reducing Over-Reliance in VR Piano Learning. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3772318.3791437>

1 Introduction

Immersive technologies such as Virtual Reality (VR) have demonstrated substantial potential for supporting *motor skill learning* across diverse domains, including sports [22, 58], medical training [75], and musical performance [44, 63]. These domains share common requirements: learners must coordinate spatial precision, temporal sequencing, and ergonomic control under time constraints. Piano playing is a representative case for VR-based music guidance systems, as it requires precise timing, pitch accuracy, and correct fingering. Prior research has demonstrated that VR training can support both synchronous remote instruction and in-situ assistance [25, 36], enabling learners to receive real-time guidance without the physical presence of an instructor. VR environments are particularly suited for such training because they can deliver embodied demonstrations and allow learners to observe and imitate expert actions from a first-person perspective [24, 43], thereby facilitating the acquisition and transfer of complex motor skills.

Despite these advantages, many existing VR-based music guidance systems continue to rely heavily on persistent visual cues [21]. This often results in learners focusing on reproducing notes accurately while neglecting the development of proper fingering strategies [68, 78]. Fingering accuracy is particularly important in piano performance, as it directly influences long-term skill transfer and technical fluency [16, 23, 26]. In addition, visual guidance in these systems is generally presented at a fixed visibility level, without adapting to the learner's performance [7, 8, 85]. This lack of adaptability can lead to over-reliance on guidance [78], reduced attention to self-correction [68], and ultimately weaker retention once the cues are removed, particularly for tasks that require both accuracy and correct technique [72]. This trade-off between *instructional support* and *learner autonomy* represents a broader HCI challenge in the design of adaptive feedback for immersive learning.

Taken together, the evidence motivates a shift from static to adaptive visual cueing. Recent research has shown that adaptive visibility can balance instructional support and learner autonomy [64, 78, 99]. While VR studies show that making avatar or object representations visible can improve motor performance and control [12, 76], evidence for such visibility manipulations in fine-motor, rhythm-dependent tasks remains limited [65, 85]. Piano training provides a representative case study for this problem: learners must coordinate pitch accuracy, timing precision, and ergonomic fingering, all while maintaining rhythmic flow. These characteristics make piano an ideal domain for investigating how adaptive visual guidance can both mitigate over-reliance and support retention.

To address the limitations of persistent visual cues, we present a **skill-adaptive ghost instructor** that dynamically adjusts its transparency according to learner performance. We examine these limitations in the context of VR piano learning. The system continuously estimates pitch accuracy, timing precision, and fingering correctness, mapping them into a performance signal that drives ghost visibility in real time, reducing guidance as proficiency increases. When performance is accurate, the ghost becomes faint, encouraging learner autonomy; when errors accumulate, the ghost reappears to provide additional support. We evaluate this approach in a within-subjects user study ($N = 30$) comparing two visibility modes (Static and Dynamic) across two melodies (A and B). Measured results include objective accuracy, subjective workload, and short-term retention.

Accordingly, we ask: *how can adaptive transparency balance instructional support and learner autonomy in VR-based motor-skill training, and what impact does it have on retention compared with static transparency?* In summary, this paper makes three contributions to immersive motor learning:

- (1) **Conceptual framing.** We identify over-reliance on persistent cues as a barrier to retention in VR-based motor skill training, highlighting the need for adaptive guidance that balances support with autonomy.
- (2) **Design intervention.** We introduce a skill-adaptive ghost instructor that modulates its visibility in real time based on learner performance. By framing adaptive transparency as a design principle, our approach addresses first-person occlusion and fingering correctness, advancing beyond static cues and conventional fading schemes.
- (3) **Empirical insights.** Through a controlled user study in VR piano training, we demonstrate that adaptive transparency improves fingering accuracy, reduces errors, lowers workload, and enhances short-term retention compared to static cues. We further distill design implications for extending adaptive visual guidance to other fine-motor training contexts such as sports practice and rehabilitation.

2 Related Work

2.1 XR-based Motor Learning Systems

2.1.1 General Motor Skill Training. Extended reality (XR) technologies are increasingly used to support motor-skill training that requires spatial precision, temporal sequencing, and ergonomic control. Critically, multiple domains report *transfer* to real-world performance or *error reduction*: in *sports*, training has shown on-field transfer in batting [33] and broader performance benefits across sport domains [59]; in *surgical training*, laparoscopy simulation improves operating-room performance [74]; in *clinical rehabilitation*, adding non-immersive VR to treadmill training reduced falls in high-risk older adults in a multicenter randomized controlled trial (RCT) [56]; and in *industrial assembly*, AR guidance shortens task time and lowers error rates compared with paper/video instructions [40, 42, 100]. Together, these results show that XR affords controlled, repeatable practice with task-aligned feedback and motivates skill-adaptive, fading assistance to mitigate cue dependence and promote retention across motor domains.

Observational learning is another robust pathway for acquiring motor skills: watching skilled or learning models can shape movement form, timing, and decision-making, particularly when demonstrations are structured and paired with concise cues [5, 11, 79]. From a design standpoint, the constraints-led perspective emphasizes aligning practice environments with performer, task, and environmental constraints to elicit functional coordination patterns [19]. Standard texts synthesize these principles into actionable guidance for feedback, practice variability, and progression. Foundational motor-learning work further motivates how feedback should be scheduled and framed: reduced or faded knowledge-of-results can yield better long-term retention than constant feedback [6, 86], contextual interference in practice schedules can benefit generalization [13], and an *external* focus of attention reliably improves performance and learning relative to an internal focus [92, 93]. Collectively, these findings highlight the need for XR systems that not only provide controlled, repeatable practice but also regulate feedback and adapt assistance to learner proficiency.

2.1.2 Piano and Music Training Systems. Within music education, piano is a representative case for XR-based guidance because it requires precise timing, pitch accuracy, and correct fingering. Prior work can be organized along two orthogonal axes: *pedagogical function* and *interface strategy*. On the pedagogical axis, four themes recur [21]: (i) synchronizing movement aids that help learners keep time (animated guides and mirrored or projected tutor hands [94–97]); (ii) sight-reading support that maps notation to the keyboard (colored lines, falling-note or piano-roll cues, and on-key highlights [8, 10, 15, 18, 37, 66, 67, 70, 98]); (iii) motivational designs that add curricula and gamified feedback [28, 57, 84]; and (iv) improvisation aids that suggest harmonies or recognize gestures (hand tracking or surface EMG) [20, 29, 38, 52]. On the interface axis, systems differ in how guidance is delivered: notation-assistance overlays project cues on or above the keys and can accelerate early learning [2, 61, 65, 84], HoloKeys superimposes target keys and supports demonstrations via AR with IoT key sensing [37], embodied demonstration/imitations foreground *how* to play [35, 48] with learner-friendly viewpoints in the MirrorFugue series [94–96], and remote MR/VR lesson rooms provide co-presence, hand tracking, and key-press visualization for bidirectional teaching [30, 71, 83]. Taken together, these archetypes cover *what* to play through cues, *how* to play through embodied examples, and *where* to learn through shared virtual studios.

Despite encouraging gains in note accuracy, timing, and engagement, gaps remain. Reviews of XR piano tutoring systems show that prior work largely employs *fixed-visibility guidance* and emphasizes *knowledge-of-results* (e.g., pitch correctness, rhythmic alignment), whereas *knowledge-of-performance* (e.g., fingering, posture, articulation) is comparatively under-supported [2, 21, 73, 85]. This imbalance can foster over-reliance on visual prompts during solo practice and weaker technique formation once cues are removed [54, 69], which in turn has been shown to impair skill retention in piano training [14]. Explicit occlusion handling is rarely addressed: most systems maintain visibility via fixed overlays rather than dynamically managing hand-ghost-keyboard occlusion [8, 85].

These observations motivate systems that adapt both the *visibility* and the *content* of guidance in real time to balance support

with autonomy. While AR/MR is often cited as ideal for in-situ deployment, prior research highlights VR as the preferred modality for motor-skill training that requires tight experimental control, precise feedback timing, and rigorous logging [41, 82]. We leverage VR as a controllable testbed to isolate the effects of our adaptive transparency mechanism, avoiding the confounds of variable lighting or background clutter. This choice is supported by Zahabi et al. [102], who find that adaptive training systems rely on fully virtual environments to standardize stimuli and isolate adaptive control strategies, and by Lui et al. [53], who show that immersive VR is uniquely capable of supporting precise, performance-contingent adjustments during training. Consequently, our system is VR-based: we use headset passthrough for a brief calibration and alignment step, then train the target piano skills in a fully virtual scene. This setup employs embodied demonstrations with *adaptive instructor visibility* to reduce cue dependence, improve transfer to the physical keyboard, and enhance retention once guidance is removed.

2.2 Ghost Instructors and Visual Perspectives

Perspective shapes how learners interpret demonstration visuals. A first-person perspective tends to increase embodiment and imitation for fine motor actions, while a third-person perspective reduces occlusion and supports global inspection [32, 80]. Comparative studies report complementary benefits for the two views, which motivate the use of perspective blending or shared viewpoints between the instructor and learner. A blended-view approach [45] merges real-world camera feeds of both users into a single composite for head-mounted display viewing. While this reduced positional error compared to most view configurations, showing only the partner’s view produced better velocity control, suggesting that no single perspective is optimal for all performance metrics. Augmented reality telepresence has also been explored for remote guidance. The Vishnu system [50] lets a distant expert control two virtual arms attached to the learner’s shoulders, manipulating virtual proxies aligned to real objects. This method significantly outperformed a monitor-based baseline for complex tasks, highlighting the value of spatially registered, embodied instruction without co-location.

Within these perspectives, ghost-style instruction systems in VR overlay a semi-transparent, co-located tutor avatar directly onto the learner’s body to guide motion [17]. In Just Follow Me [99], the tutor’s motion was visualized as a first-person ghost aligned to the learner’s body, with swept-volume trails and acceleration vectors to convey path and dynamics. In controlled motion learning tasks, first-person, co-located guidance delivered training and transfer effects comparable to or better than indirect video imitation. Other approaches explore virtual co-embodiment [24], where teacher and student jointly control a single avatar via a weighted average of their movements. Comparative studies [47] show that co-embodiment and first-person ghost metaphors (1PPGM) accelerate learning more than no guidance. However, both approaches suffer from performance drops when guidance is removed, with the largest drop observed in co-embodiment, a smaller drop in 1PPGM, and the smallest drop with no guidance. Beyond perspective, tuning visibility itself affects performance: in navigational obstacle-avoidance tasks, transparent or invisible hands with opaque objects yielded faster completion and fewer collisions, and are preferred by users

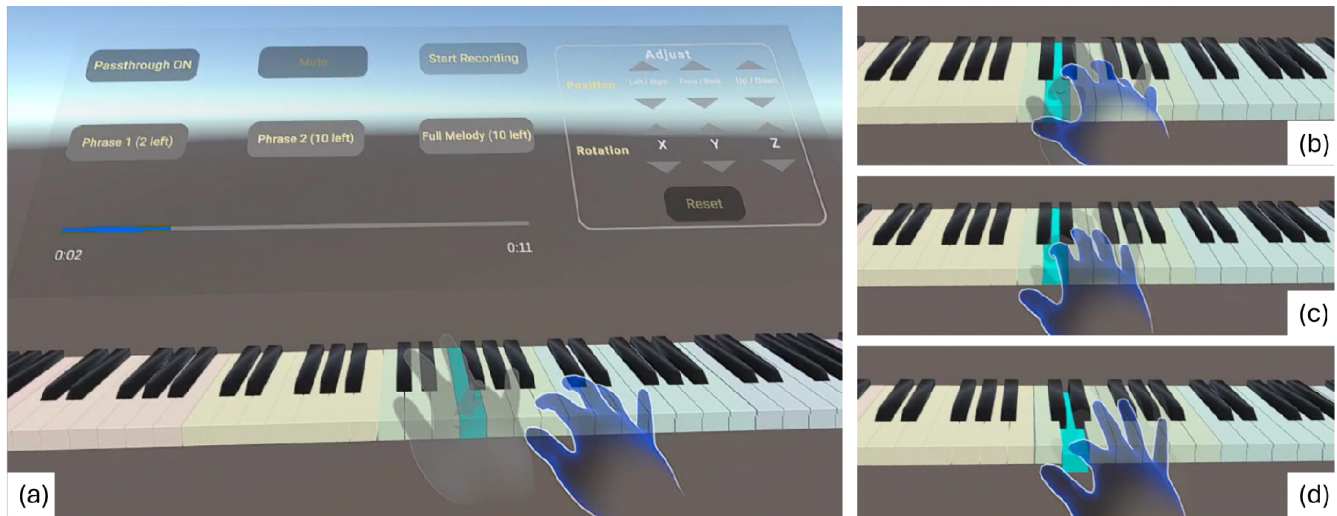


Figure 2: System interface and dynamic-transparency feedback. The blue hand denotes the *participant*, and the gray hand denotes the *ghost instructor*. (a) Controls. Passthrough desk alignment; ghost audio mute/unmute; record button. The second row selects *Phrase 1*, *Phrase 2*, or *Full Melody* and shows the remaining-trial count (default 10). The *Adjust* panel fine-tunes the virtual piano rig to the physical desk (translate/rotate for alignment). (b–d) Adaptive feedback. Ghost-hand opacity is driven by performance in real time: higher visibility under poorer performance, and fainter visibility as pitch, timing, and fingering stabilize.

[39]. Notably, prior ghost-based systems fixed the tutor’s visibility, leaving the potential of adaptive transparency unexplored.

2.3 Adaptive Support Strategies for Motor-Skill Learning

Adaptive training refers to closed-loop personalization, continuously adjusting task challenge or scheduling feedback based on a learner’s ongoing performance [34, 46]. For a system to be adaptive, three core components are required: (1) performance measurement, (2) an adaptive variable, and (3) adaptive logic [102]. Prior work on motor-skill tutoring typically adapts either *how* assistance is delivered or *who* provides guidance, which can be categorized into two broad streams: *solo practice* and *collaboration-guided* support.

In solo practice, most systems adapt either *task difficulty* or the *feedback modality*. *BACH* [101] increases musical complexity based on cognitive workload to maintain an optimal challenge zone, whereas *Adapt2Learn* [81] maps performance to a running score and bi-directionally raises or lowers difficulty when progress plateaus. Other work, such as *Hapticus* [51], explores *modality selection and ordering* (EMS, linkage exoskeleton, vibrotactile gloves) through pre-specified haptic sequences that improve in-the-moment performance but do not provide closed-loop, within-practice adaptation. However, beyond hardware complexity, haptic guidance (HG) faces significant efficacy challenges. Specifically, HG can hinder learning if the guidance reduces the learner’s active involvement in error correction. Marchal-Crespo et al. [55] demonstrated that while HG is beneficial for gross motor tasks such as tennis or rowing, visual feedback can be more effective for internalizing complex spatial topologies, as haptic assistance may induce a “passive learning state” where users fail to actively correct their own errors. Given

that piano playing requires extreme fine-motor precision, complex haptic devices can often constrain natural movement and interfere with the ergonomic flow of performance. Consequently, our approach leverages dynamic transparency as a lightweight visual scaffold that guides spatial positioning without interfering with physical touch or the piano’s natural tactile feedback.

Collaboration-guided approaches emphasize making demonstrations and instructions accessible and clear. *Digitusync* [60] mechanically couples a teacher’s and learner’s hands through a passive exoskeleton to improve imitation fidelity during co-located instruction. *ShowMe* and *ARTEMIS* [3, 27] support immersive tele-gestures and spatial annotations to strengthen shared attention in remote collaboration. These methods primarily optimize *who provides guidance*, rather than algorithmically adapting support based on performance. While effective for following expert demonstrations, they can inadvertently foster over-reliance, making it unclear whether learners can perform autonomously once such support is withdrawn.

Evidence further shows that the timing and direction of adaptation critically shape learning effectiveness and efficiency [34], while guidance-hypothesis and faded-feedback literature warns that excessive or continuous guidance can induce dependency [4, 68, 88]. Addressing this dependency requires regulating the salience of guidance over time. Sigrist et al. [77] employed error-based transparency in rowing training. While effective for gross-motor trajectory learning, their visibility modulation is primarily driven by a single scalar spatial error signal, without integrating multi-dimensional performance cues. In contrast, our Ghost Instructor targets the high-dimensional requirements of piano performance,

and its transparency controller integrates pitch correctness, rhythmic timing, and fingering accuracy into a unified performance signal. Moreover, we introduce an asymmetric adaptation mechanism that fades guidance out quickly as performance stabilizes while reintroducing guidance more cautiously when errors increase, aiming to encourage retrieval practice and mitigate over-reliance on continuous visual cues.

Building on this work, our study extends adaptive training theory in three ways. First, we introduce *adaptive transparency* as a simple, purely visual adaptive variable that can be implemented on commodity XR devices, complementing hardware-intensive systems that require worn devices to adapt task difficulty or haptic actuation and may constrain users' natural movement. Second, we instantiate a continuously corrected control strategy that maps *multi-dimensional and fine-grained* piano performance signals, including not only pitch accuracy but also timing precision and finger-specific correctness, to the moment-to-moment visibility of a virtual tutor. This allows assistance to become more salient when needed and to gradually fade when the learner is performing well, while performance is continuously monitored in real time. Third, we explicitly operationalize and evaluate *dependency* by comparing immediate and retention performance once guidance is removed, addressing an outcome that is rarely measured in prior motor-skill tutoring work.

2.4 Summary

Prior work in motor learning highlights three key requirements for XR-based training: (i) integrate observation and task-representative practice to reduce representational gaps; (ii) regulate the *frequency, timing, and framing* of feedback to preserve error-driven learning while avoiding dependence; and (iii) adapt assistance dynamically to the learner's moment-to-moment proficiency, balancing instructional support with autonomy [89, 92].

Existing XR-based piano learning systems integrate notation overlays, embodied demonstrations, and shared lesson spaces. However, they often emphasize knowledge of results more strongly than knowledge of performance, which can lead to visual over-reliance and limit skill retention. Research on instructional perspectives shows that first-person ghost overlays and co-embodiment can accelerate motor learning. However, both approaches experience performance drops when guidance is removed and have so far only employed fixed-visibility avatars.

In response to these limitations, we propose *skill-adaptive ghost instructors* that modulate virtual-ghost transparency in real time based on the learner's performance in note accuracy, timing precision, and fingering correctness. Within adaptive training theory, our design instantiates a different adaptive variable: rather than changing task difficulty or switching feedback modalities, we regulate the *visual availability* of a tutor. Conceptually, we treat ghost visibility as a limited guidance resource whose salience must be matched to the learner's current proficiency to avoid reliance on assistance. By tailoring visibility to moment-to-moment performance while holding the task constant, our approach aims to reduce assistance dependence, maintain engagement, and improve retention within a VR piano setting.

3 Research Question and Hypotheses

RQ1. What effects does performance-adaptive transparency, compared with static transparency, have on objective learning outcomes (transfer and short-term retention) in VR-based motor-skill training?

RQ2. How does performance-adaptive transparency, compared with static transparency, shape the balance between instructional support and learner autonomy?

3.1 Hypotheses

Guided by motor-learning theory on guidance dependence and feedback scheduling, we hypothesized the following effects of *Dynamic* (performance-adaptive transparency) versus *Static* (fixed transparency):

H1 (Skill transfer). Dynamic transparency will promote better transfer when switching to a new melody and condition. Participants who trained under Dynamic first will start a new melody and condition at a higher initial performance than those who trained under Static first.

H2 (Retention). Dynamic transparency will improve short-term retention relative to static transparency. The decrease in the composite performance score from the Immediate Test to the Retention Test will be smaller for Dynamic than for Static.

H3 (Training convergence). Across training, both conditions will show positive practice effects, but Dynamic transparency will lead to faster convergence than Static. Participants in the Dynamic condition will show larger gains in early loops and smaller marginal improvements in later loops, or will reach a predefined performance criterion in fewer loops, compared with Static.

4 Methodology

4.1 System Overview

Our VR piano training system (Figure 1) is designed to improve fingering, posture, and articulation, with the goal of enhancing both skill transfer and short-term skill retention. The setup uses a ghost instructor that appears as an overlaid virtual hand performing the target fingering. Two guidance modes are supported: a static transparency mode with fixed visibility, and a performance-adaptive mode where the ghost's transparency changes in real time according to the learner's fingering accuracy and timing performance.

Expert demonstrations are recorded using hand and finger tracking to capture joint positions and movements in 3D space, along with key press events and audio. These recordings are then replayed in a fully virtual piano scene aligned to the learner's real keyboard. During practice, the ghost instructor provides real-time, moment-by-moment guidance. The adaptive transparency further encourages learners to focus on self-correction when performance is accurate. After training, users perform an Immediate Test and a Retention Test without any ghost guidance to evaluate skill transfer and retention.

4.2 Motion and Input Capture

4.2.1 Hand and Finger Tracking. We employ the Meta XR Hands API for 26 joint skeletal tracking of the learner's right hand in real

time. Joint positions and rotations are streamed at the device’s native update rate in the headset’s local coordinate system, and then transformed into the calibrated piano coordinate frame. This alignment allows us to directly compare the learner’s finger positions with the target ghost-hand motion.

The tracking data serves two main purposes: (1) detecting fingering correctness in real time, and (2) providing the basis for performance-aware transparency adaptation. Fingering correctness is evaluated by mapping tracked fingertip joints to pressed keys and comparing them with the intended fingering from the ghost demonstration. In addition, per-frame hand-posture data (joint positions and rotations) are recorded for subsequent analysis in Section 6.3.1. Timing precision and note accuracy are computed from synchronized key event logs, enabling integrated assessment of pitch, rhythm, and fingering to drive adaptive feedback.

4.2.2 Motion Recording. Building on the tracking setup, we log the 26-joint skeleton at a fixed 30fps with per-frame timestamps, storing joint rotations and positions in the calibrated piano frame. For animation export, positions are recorded internally. However, emitting position curves is optional, allowing a trade-off between playback fidelity and file size.

For each included joint, we store a timestamped local pose (relative to the hand root) and later serialize the sequence into a Unity AnimationClip. Keyframe curves are compressed using rotation and position slope thresholds to reduce file size without losing perceptible motion fidelity.

The motion recording module serves two purposes in our system:

- (1) **Ghost Instructor Demonstrations** – We first record our own expert performances, producing smooth, reusable ghost-hand animations that users follow during the training phase.
- (2) **User Performance Capture** – During the test phase, users’ hand motions are recorded in the same format. These recordings are not used in real-time feedback but are stored for *post-hoc* analysis of timing, fingering correctness, and articulation patterns.

This dual-use design standardizes instructor and learner motion data, enabling accurate playback, direct comparison, and kinematic analysis. In this proof-of-concept, instructor motions were author-recorded and kept identical across all participants and conditions. Therefore, any stylistic bias applies equally across participants and conditions, and cannot be used to explain condition effects. Future work will include multiple expert performers.

4.2.3 Audio and Key Event Logging. We developed a logger to capture every press–release cycle on the virtual keyboard. For each key event, the system records the key identifier (pitch), onset time, press duration, and finger index. A dictionary of active keys tracks onset times until release, when the duration is computed and the event is stored. Recording can be toggled on/off.

Two modes match the study workflow: demonstration recordings capture the author’s ghost–hand performances, while user recordings capture participants’ play during the test phase. All sessions are saved as JSON files with versioned names, time–aligned with motion recordings to support post-hoc analysis of note accuracy, timing precision, sustain characteristics, and fingering correctness.

4.3 Playback and Visualization

The playback system replays pre-recorded ghost-hand motions in the VR environment to guide learners during practice. Learners view the scene from a first-person perspective to promote embodiment and imitation. The ghost hand is displayed as a semi-transparent grey model, overlaid on the virtual piano and spatially aligned so that each animated press matches the learner’s corresponding virtual key. To clearly differentiate between the instructor and the learner, the participant’s own tracked hands are rendered in semi-transparent blue. Additionally, a key highlight system provides note-specific visual cues by changing the color of pressed keys to light blue; thus indicating both the precise timing and location of each note.

To structure practice, the interface provides three playback modes: *Phrase 1*, *Phrase 2*, and *Full Melody*, where *Phrase 1* and *Phrase 2* together constitute the *Full Melody*. Each mode replays its corresponding ghost-hand performance ten times consecutively, with a two-second countdown shown between repetitions to allow participants to prepare. This segmentation encourages focused practice on smaller sections before attempting the full piece, while ensuring consistent repetition and exposure to each segment.

4.4 Adaptive Transparency Control

To mitigate over-reliance on visual guidance while still supporting learners who struggle, we implemented an adaptive transparency controller that dynamically adjusts the ghost hand’s visibility based on real-time performance metrics. For each frame during playback, the system evaluates the learner’s input against the reference melody and ghost avatar motion, producing normalized sub-scores for pitch correctness, timing precision, and fingering correctness, denoted s_p , s_t , and s_f , respectively, with $s \in [0, 1]$.

The overall performance score is a weighted sum

$$S = w_p s_p + w_t s_t + w_f s_f, \quad (1)$$

Calculating Sub-scores. These sub-scores are calculated in real time based on the user’s input. **Pitch correctness** (s_p) assigns full credit when the pressed key’s pitch matches the target and the onset falls within a temporal window $\pm\tau$ of the reference time; otherwise $s_p = 0$. In other words, a note is treated as correct only if both pitch and timing are within this acceptance window. **Timing precision** (s_t) is computed as $1 - \min(\frac{|\Delta t|}{\tau}, 1)$ where Δt is the onset deviation (ms) relative to the reference time and $\tau = 120$ ms (approximately 0.20–0.25 beats at 100–120 BPM), a musically meaningful window that also accommodates VR latency. **Fingering correctness** (s_f) uses semantic finger labels (Thumb–Index–Middle–Ring–Pinky) from the hand-tracking API. We apply a two-level partial-credit rule: $s_f = 1$ for the intended finger, $s_f = 0.5$ if the used finger is adjacent to the intended one on the same hand, and $s_f = 0$ otherwise. To robustly handle individual differences in hand size, comparisons rely on these semantic labels rather than absolute 3D joint coordinates.

The corresponding error that drives transparency is

$$E = 1 - S, \quad (2)$$

with $S, E \in [0, 1]$. Here, E denotes the instantaneous error signal (higher means worse performance and thus higher target opacity).

Our adaptive transparency score is a weighted sum of pitch, timing, and fingering correctness. Prior work indicates that pitch, timing, and fingering are co-learned and mutually reinforcing in beginner training. Within this interplay, curricula and assessments often place early emphasis on pitch accuracy, while allowing some timing flexibility and introducing fingering strategies progressively [1, 9, 31, 62]. Based on this rationale, we adopted the configuration $w_p = 0.7$, $w_t = 0.2$, $w_f = 0.1$, which prioritizes note accuracy while still incorporating rhythm stability and ergonomic fingering. While this weighting is grounded in prior pedagogical research, the optimal balance between pitch, timing, and fingering may vary across learners and tasks.

From error to opacity. We first stabilize frame-to-frame fluctuations with an asymmetric exponential moving average (EMA) on the error:

$$\hat{E}_t = \lambda_t E_t + (1 - \lambda_t) \hat{E}_{t-1}, \quad \lambda_t = \begin{cases} \lambda_{\uparrow}, & \text{if } E_t > \hat{E}_{t-1} \\ \lambda_{\downarrow}, & \text{otherwise} \end{cases} \quad (3)$$

E_t is the per-frame value of E , and \hat{E}_t is its temporally smoothed version. Here $\lambda_t \in (0, 1)$ is the smoothing factor: larger values make \hat{E}_t follow E_t more quickly, while smaller values yield heavier smoothing. We choose a smaller rise rate λ_{\uparrow} and a larger decay rate λ_{\downarrow} so that the ghost fades quickly when the learner improves and reappears more cautiously after occasional slips. This reduces cue dependence and promotes attention to pitch accuracy, rhythmic alignment, and proper fingering. Importantly, this EMA introduces only brief within-stream smoothing (≈ 0.5 – 1 s) to stabilize visuals. This short latency operates within a single performance stream and is therefore distinct from the multi-trial feedback delays (terminal feedback) typically associated with enhanced retention in motor-learning studies.

The filtered error is then mapped linearly to opacity and clamped to the display range:

$$\tilde{\alpha}_t = \alpha_{\min} + (\alpha_{\max} - \alpha_{\min}) \hat{E}_t, \quad (4)$$

$$\alpha_t = \min\{\max(\tilde{\alpha}_t, \alpha_{\min}), \alpha_{\max}\}. \quad (5)$$

To ensure legible yet unobtrusive guidance, we clamp opacity to $\alpha_{\min} = 0.08$ and $\alpha_{\max} = 0.8$. Thus, a low filtered error \hat{E}_t yields a faint ghost that does not occlude the keyboard, whereas a high error makes the ghost clearly visible without becoming fully opaque. We also apply a small per-frame visual easing when updating the material to avoid shader jitter; this does not alter the linear mapping in Eqs. 4–5. In our study, we compare a *Static* condition (constant $\alpha = 0.5$) with a *Dynamic* condition (adaptive opacity as described above).

4.5 Implementation Details

We developed the system in Unity 2022.3.44f1 (LTS) targeting Meta Quest 3. For rapid iteration, the headset was tethered to the development PC via Oculus Link over USB-C, enabling in-headset Play-in-Editor without producing Android builds during prototyping. The project uses the Meta XR All-in-One SDK (v69.0.1, October 2024). From this SDK, we employed the camera-rig prefab for HMD tracking, the interaction building blocks (e.g., buttons and UI interactors) for interface elements, and the hand-tracking modules to capture and drive hand-based interactions. This configuration

provided a stable, up-to-date integration path for Quest features while keeping iteration times low.

5 User Studies

5.1 Experiment Design

We employed a within-subject design with two experimental conditions:

- **Static Ghost Hand:** a semi-transparent ghost hand with fixed visibility.
- **Dynamic Ghost Hand:** a ghost hand whose transparency is adapted in real time according to the learner’s performance (pitch accuracy, timing precision, fingering correctness).

Each participant practiced and performed two short piano melodies under both conditions. The order of conditions (Static first vs. Dynamic first) and melody assignment (Melody A vs. Melody B) were counterbalanced across participants using a Latin square design, resulting in four groups (SA/DB, SB/DA, DA/SB, DB/SA) (see Figure 3). Each session began with a short familiarization phase, followed by practice and performance trials.

After completing both conditions, participants filled out the NASA-TLX workload questionnaire and additional Likert-scale questions. Objective performance metrics (note accuracy, timing deviations, and fingering correctness) were logged automatically during the trials.

Design Rationale. We acknowledge that the Dynamic condition represents a composite intervention consisting of performance-adaptive transparency and a sub-second stabilization latency (≈ 0.5 – 1 s) introduced by the asymmetric EMA. This delay was intentionally incorporated to prevent visual jitter and to implement a "cautious reappearance" strategy: when errors occur, the instructor ghost returns gradually rather than abruptly, encouraging the learner to attempt self-correction before relying on the visual prompt. Consequently, the observed learning benefits in the Dynamic condition likely stem from the synergistic effect of both transparency adaptation and this brief feedback delay.

5.2 Participants

A total of **30 participants (18 male, 12 female; age 22–39 years, $M = 27.6$, $SD = 4.33$)** took part in the study. A power analysis ($\alpha = .05$, two-sided) indicated that with $N = 30$ our within-subjects design has approximately 80% power to detect medium-to-large effects ($d_z \approx 0.55$; $f \approx 0.50$, $\eta_p^2 \approx .20$). This sample size provides adequate power for medium-to-large effects; detecting smaller effects would require a larger sample. Half of the participants ($n = 15$) reported having prior piano experience (self-rated score ≥ 3 on a 5-point scale), while the other half ($n = 15$) had little or no experience (< 3). All participants had prior experience with Virtual Reality (VR) and were comfortable using VR headsets (self-rated ≥ 3). All had normal or corrected-to-normal vision. The study was approved by an institutional review board committee.

Participants were assigned to one of four counterbalanced groups following a Latin square design (Figure 3). Each group included both experienced and inexperienced pianists, though the distribution was not perfectly even (e.g., SA/DB: 2 inexperienced, 6 experienced; SB/DA: 5 inexperienced, 3 experienced). A chi-square test confirmed

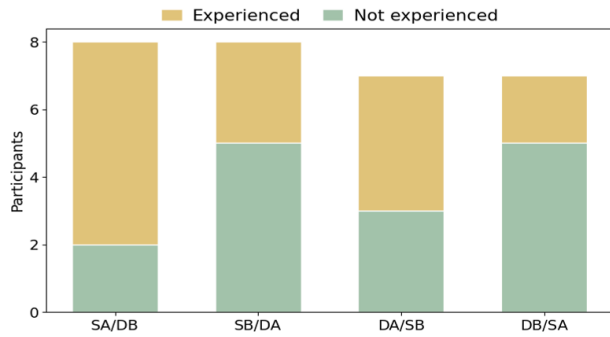


Figure 3: Participant distribution across Latin-square conditions and piano experience. S = *Static*, D = *Dynamic*; each condition pairs a visibility mode with a melody A/B.

that the proportion of experienced vs. inexperienced participants did not differ significantly across groups ($\chi^2(3, N = 30) = 3.93, p = .27$).

5.3 Task Representation

Participants were asked to learn and perform two short piano melodies in VR. Each melody was approximately 12 seconds in length, contained *single notes only*, and was performed with the *right hand only*. During practice, the system rendered a ghost hand co-located with the participant’s virtual hand to illustrate intended finger placement and motion.

Each session began with a brief familiarization, followed by two phases: (1) a *training* phase, during which participants followed the ghost hand to rehearse the melody, and (2) a *test* phase, in which they were asked to play as accurately and steadily as possible without the ghost hand. The system automatically logged key events, per-finger joint trajectories, loop indices, and timing information. During training, the system computed the per-loop, per-note sub-scores (Eq. 1) and the aggregated error (Eq. 2) in real time to drive ghost transparency, and logged these values for offline analysis.

5.4 Procedure

Each session lasted approximately 50–60 minutes per participant (pre-study 5 min; two blocks × [training 10 min + immediate test (2 trials) + 10-min retention interval + retention test (2 trials) + questionnaires 5 min]).

Pre-study. After providing informed consent, participants completed a brief pre-study questionnaire assessing piano experience and VR/AR experience and comfort, followed by headset fitting and a short calibration/familiarization to align the virtual keyboard slightly above the physical desk height, ensuring keys could be played on the tabletop rather than in midair.

Experimental blocks. Each participant then completed two counterbalanced blocks, one per condition (Static and Dynamic). In each block, the training phase comprised three progressive parts with the condition’s ghost hand visualization active: Phrase 1 (first subphrase; 10 loops), Phrase 2 (second subphrase; 10 loops), and Full Melody (Phrase 1 + Phrase 2; 10 loops). The training lasted about

10 minutes. Immediately after training, an Immediate Test was administered consisting of two trials; the better trial was retained for analysis. About 10 minutes later, participants completed a Retention Test, again consisting of two trials, where the better trial was retained. During the interval, participants solved an expert-level Sudoku as a distractor to discourage rehearsal or phone use. Following the retention test of each block, participants completed the NASA TLX assessment as well as the post-condition Likert items. A short break was offered before switching to the other block.

Post-study. After the completion of both blocks, participants filled out a post-study survey capturing overall preference (Static vs. Dynamic), comparative ratings, and open-ended feedback.

5.5 Measures

We collected both objective performance measures and subjective responses.

Performance Measures. We use the same pitch, timing, and fingering definitions as in Section 4.4. Briefly, pitch accuracy counts a note as correct only if the pressed key pitch matches the target and the onset falls within the temporal window $\pm\tau$; timing precision is computed as $1 - \min(|\Delta t|/\tau, 1)$; and fingering correctness follows the partial-credit rule over semantic finger labels (intended vs. adjacent vs. other fingers). During the Training Phase, we compute these sub-scores per loop to analyze learning curves; during the Test Phase, we aggregate them per trial for inferential statistics. The error rate E is the signal that drives adaptive transparency and is defined as $E = 1 - S$, where $S = w_p s_p + w_t s_t + w_f s_f$ (weights as in Section 4.4). In the *Training Phase*, we logged pitch accuracy, finger accuracy, timing accuracy, and error rate for each practice loop to analyze the learning curve. In the *Test Phase*, we recorded (a) pitch, timing, and fingering via the audio and key event logging system (Section 4.2.3) to compute accuracy and error measures, and (b) 3D hand motion trajectories via motion capture (Section 4.2.2) for exploratory motion similarity analysis. Unless otherwise specified, all error bars in the figures indicate standard error of the mean (SEM).

Subjective Measures. Participants rated statements on 5-point Likert scales (1 = strongly disagree, 5 = strongly agree). During the pre-study, we collected demographics and musical background. Post-study measures included: *Learning Experience* (7 items), *Visual Design & Transparency* (4), and *Retention* (7). We also included a *Comparative evaluation* with two Likert items (“easy to learn from” for *Static* and *Dynamic*) and three forced-choice questions (Static/Dynamic/Not sure) on overall preference, memorization support, and naturalness. Perceived workload was assessed with the *NASA-TLX*.

6 Results

6.1 Performance Results

This section presents objective performance results. Analyses of Static vs. Dynamic, effects of prior experience, and Immediate vs. Retention use the participants’ *performance test data* (Immediate/Retention), while learning-curve progression is based on *training phase* practice logs.

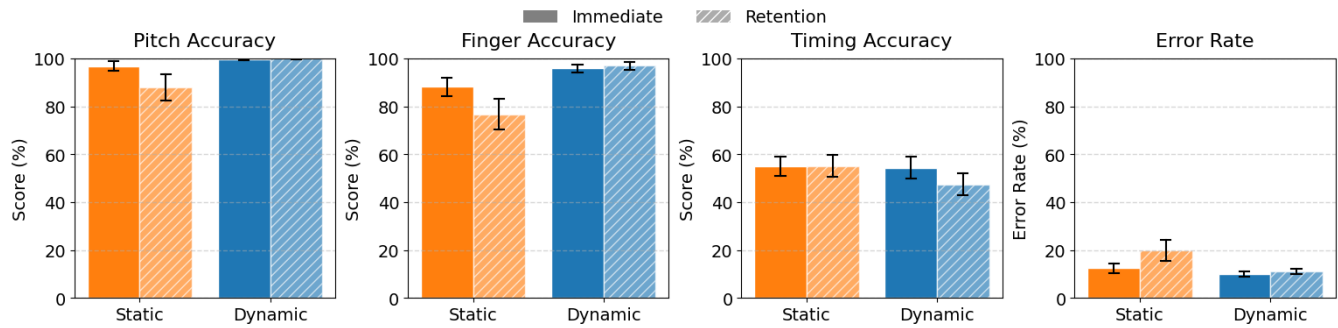


Figure 4: Performance comparison between Static and Dynamic ghost hand conditions across pitch accuracy, finger accuracy, timing accuracy, and error rate. Error bars represent the standard error of the mean (SEM).

6.1.1 Static vs. Dynamic. As shown in Figure 4, we compared participants’ performance between the *Static ghost hand* and *Dynamic ghost hand* conditions across four objective measures: pitch accuracy, finger accuracy, timing accuracy, and error rate. A two-way repeated-measures ANOVA with factors *Condition* (Static vs. Dynamic) and *Test Type* (Immediate vs. Retention) was conducted for each measure.

Pitch Accuracy. There was a significant main effect of Condition ($F(1,29) = 5.04, p = 0.03$), with higher pitch accuracy in the Dynamic condition compared to the Static. As shown in Table 1, participants achieved nearly ceiling-level accuracy in the Dynamic condition (99.1% in the Immediate test, 99.6% in the Retention test), whereas accuracy dropped under the Static condition, particularly in the Retention test (96.8% Immediate vs. 87.8% Retention). This indicates that the advantage of Dynamic guidance was consistent across both Immediate and Retention tests.

Finger Accuracy. A strong main effect of Condition was found ($F(1,29) = 9.68, p < 0.01$), showing that Dynamic guidance significantly outperformed Static. Numerically, finger accuracy was stable and high in the Dynamic condition (95.8–96.5%), but decreased in the Static condition from 88.2% in the Immediate test to 76.7% in the Retention test (Table 1).

Timing Accuracy. No significant effects were observed for Condition ($F(1,29) = 1.42, p = 0.24$), Test Type ($F(1,29) = 0.88, p = 0.36$), or their interaction ($F(1,29) = 1.04, p = 0.32$). As shown in Table 1, timing performance remained comparable across conditions and test phases, ranging between 46.8% and 55.0%.

Error Rate. A significant main effect of Condition was found ($F(1,29) = 4.29, p = .047$), with Dynamic yielding fewer errors. Specifically, error rates in the Dynamic condition remained low (9.9% Immediate, 11.6% Retention), while errors increased more strongly in the Static condition, particularly in the Retention test (12.4% Immediate vs. 19.8% Retention; Table 1).

Table 1: Mean performance scores (in %) for Static vs. Dynamic ghost hand conditions in Immediate and Retention tests.

Condition	Test Type	PitchAcc	FingerAcc	TimingAcc	ErrorRate
Static	Immediate	96.79	88.16	55.00	12.43
	Retention	87.84	76.71	55.04	19.83
Dynamic	Immediate	99.11	95.84	54.36	9.85
	Retention	99.56	96.52	46.81	11.61

Summary. Across performance measures, the Dynamic ghost hand consistently outperformed the Static condition. Specifically, Dynamic guidance significantly improved pitch accuracy and finger accuracy, while also reducing overall error rates. Timing accuracy did not statistically differ between conditions. As reported above for timing, the Condition \times Test Type interaction was not significant, indicating that the Dynamic–Static difference did not vary between Immediate and Retention.

6.1.2 Immediate Test vs. Retention Test. To assess memory for the learned material, we computed a *retention score* for each metric as the difference between the Retention and Immediate tests (Retention–Immediate). Scores near zero indicate stable performance, negative scores indicate decline, and positive scores indicate improvement. We compared retention scores between the Static and Dynamic conditions using paired *t*-tests for each metric. As shown in Figure 5, retention patterns differed by condition.

For *pitch accuracy*, Static showed a marked decline ($\Delta = -9.0 \pm 4.8\%$), whereas Dynamic was essentially preserved ($\Delta = +0.2 \pm 0.5\%$). *Finger accuracy* followed the same pattern: Static dropped substantially ($\Delta = -11.4 \pm 5.7\%$), while Dynamic remained close to baseline ($\Delta = +1.1 \pm 2.1\%$). *Timing accuracy* revealed the opposite tendency, with Static showing no reliable change ($\Delta = 0.0 \pm 5.6\%$) and Dynamic decreasing ($\Delta = -7.0 \pm 4.5\%$). Finally, *error rate* increased in both conditions, but the rise was larger for Static ($\Delta = +7.4 \pm 4.0$ percentage points) than for Dynamic ($\Delta = +1.1 \pm 1.1$ points).

Overall, Dynamic ghost hand guidance not only yielded higher immediate performance but also supported better retention for pitch and fingering accuracy, and limited the growth of errors over time, whereas Static guidance was associated with larger declines and a greater increase in errors.

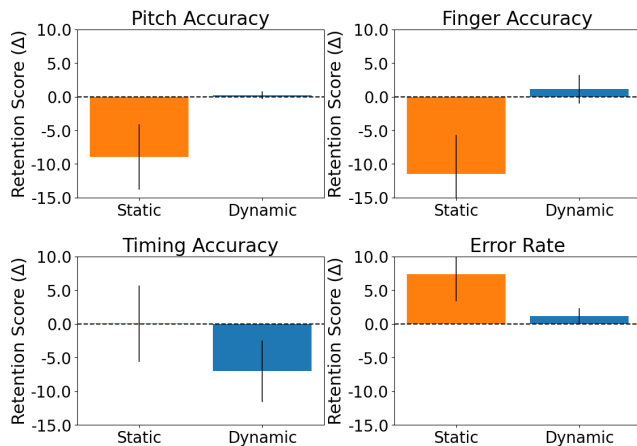


Figure 5: Retention scores (Retention – Immediate) for pitch, finger, and timing accuracy as well as error rate, comparing Dynamic and Static ghost hand conditions. Values near zero indicate stable performance, with negative values showing decline and positive values showing improvement.




Figure 6: Interaction plots showing performance across Static and Dynamic conditions for experienced and non-experienced participants.

6.1.3 Effect of Experience. We tested whether prior piano experience moderated performance using a two-way design with factors *Condition* (Static vs. Dynamic) and *Experience* (experienced vs. non-experienced). As shown in Figure 6, both groups follow the same pattern across metrics: Dynamic yields higher pitch and finger accuracy and lower error rates than Static, whereas timing differences are small for experienced learners but more variable for inexperienced learners. Experienced participants are slightly better overall, and the lines are roughly parallel across metrics except for timing accuracy. This may reflect that experienced musicians often have a stronger sense of rhythm.

Across these measures, ANOVAs confirmed a main effect of *Condition* (Dynamic > Static; e.g., pitch $F(1, 116) = 4.74, p = .032$) and no reliable main effect of *Experience* nor *Condition*×*Experience* interaction (all $p \geq .12$). Thus, the benefit of Dynamic guidance is robust regardless of prior musical background.

6.1.4 Learning curve. We modeled performance across ten practice loops in the training phase with a linear mixed-effects model (LMM) including fixed effects of *Condition* (Static vs. Dynamic), *Loop* (mean-centered), and their interaction, while controlling for *Melody* (A/B) and *Phrase* (Phrase1, Phrase2, Full Melody). Participant-specific random intercepts and random slopes for *Loop* were included.

Group learning curves. Figure 7 shows group-averaged learning curves for the *Full Melody* across *Pitch*, *Finger*, *Timing*, and *Error* under *Static* and *Dynamic* guidance (lines = participant means; shaded bands = 95% CIs). During practice, the two modalities follow essentially the same trajectory: *Pitch* and *Finger* increase, *Error* decreases, and *Timing* remains approximately stable. A mixed-effects model ($score \sim condition * loop_c + melody + phase + (1 + loop_c | user)$, REML) confirms improvement across loops ($\beta = 0.014, SE = 0.003, z = 4.64, p < .001$), comparable learning rates (*Condition*×*Loop*: $\beta = -0.003, SE = 0.004, z = -0.90, p = .368$), and only a small average *Dynamic* offset ($\beta = 0.028, SE = 0.011, z = 2.57, p = .010$). Melody B exceeds A, and later phases are associated with slightly lower learning scores Phrase 1 (Phrase 2: $\beta = -0.044, SE = 0.013, p < .001$; Full Melody: $\beta = -0.053, SE = 0.013, p < .001$).

Individual differences. Figure 8(a) shows per-participant *error-rate* learning slopes (**lower is better**), color-coded by training order (blue = *Dynamic-first*, orange = *Static-first*; lines connect within subject; white squares = mean ±95% CI). More negative values indicate faster error reduction. Slopes were negative in both conditions (Static: $-0.009, SE = 0.002, 95\% CI [-0.0138, -0.0049]$; Dynamic: $-0.017, SE = 0.003, 95\% CI [-0.024, -0.010]$), indicating decreasing error over practice, with a steeper (more negative) decline under *Dynamic*. Figure 8(b) shows within-subject differences $\Delta = \text{Dynamic} - \text{Static}$. The mean Δ was negative and its 95% CI excluded zero (mean = $-0.0077, SE = 0.0036, 95\% CI [-0.015, -0.0003], n = 30$); Cohen’s $d_z \approx 0.40, p = .038$). Critically, the effect depended on order: *Dynamic-first* ($n = 14$) showed a clear negative shift (mean $\Delta = -0.019, 95\% CI [-0.031, -0.006], p = .006$), whereas *Static-first* ($n = 16$) centered near zero (mean $\Delta = +0.002, 95\% CI [-0.005, 0.008], p = .642$). The order groups differed (Welch $t = 3.08, p = .006$), indicating that adaptive transparency accelerates error reduction primarily when experienced first.

Figure 9 summarizes the block switch from the first to the second block. Figure 9(a) indicates that, at the switch to the second block (new melody and condition), participants who trained with *Dynamic* first began the new melody and condition with higher initial performance, operationalized as the mean of the first two loops, than those who trained with *Static* first. Using a linear mixed-effects model that includes melody (A/B) as a covariate, and a participant random intercept ($InitialScore \sim ConditionFirst + Melody + (1 | Participant)$), we estimate an adjusted *Dynamic-first* advantage of $\beta = +0.073 (p = .050)$. Melodies differed in difficulty, with higher




Figure 7: Group learning curves for Full Melody: Pitch, Finger, Time, and Error.

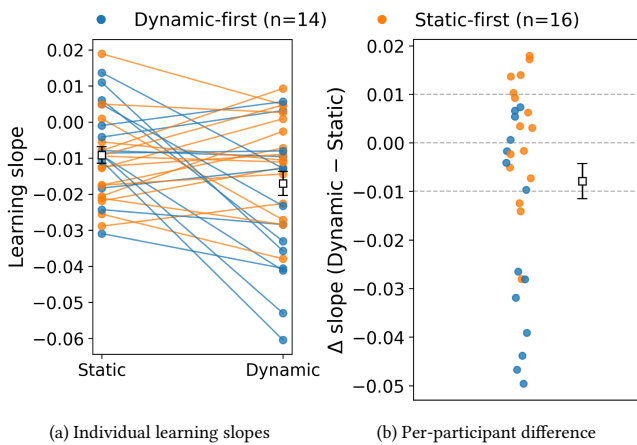


Figure 8: Per-participant *error-rate* learning slopes by training order. white squares = mean \pm 95% CI. Lower (more negative) is better. (a) Individual learning slopes for Static vs. Dynamic; lines link the same participant; (b) Within-subject differences Δ = Dynamic - Static grouped by order.

scores on melody B ($\beta = +0.123$, $p = .006$). Critically, the Dynamic-first advantage holds after controlling for melody and is also visible within each melody (melody A: 0.336 vs. 0.266; melody B: 0.461 vs. 0.385 for Dynamic-first vs. Static-first). By the end of the second block, computed as the mean of the last two loops, the *Dynamic*-first group still performs better, with more individuals exceeding the 0.60 criterion, whereas *Static*-first rarely does (Figure 9(b)). Taken together, these distributions indicate a transfer advantage at the switch that persists over subsequent practice, while also revealing substantial between-participant variability in how quickly learners improve within the second block.

6.2 Subjective Results

We measured participants' subjective workload with NASA-TLX after each condition. Post-study, participants provided learning-experience ratings for *Static/Dynamic*, completed a Visual Design & Transparency survey for the *Dynamic* ghost, and a comparative evaluation. All items used a 5-point Likert scale.

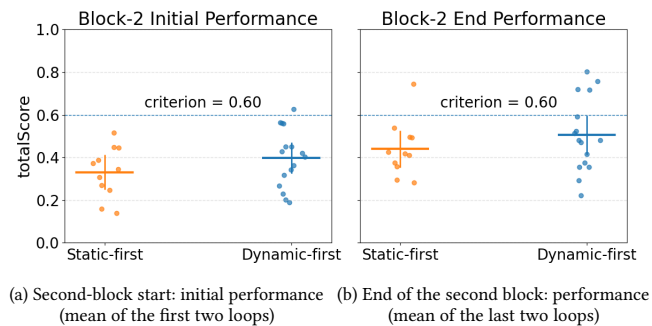


Figure 9: Summarizes the block switch from the first to the second block. Dots represent participants; crosshairs indicate group means; and the dashed line marks the 0.60 criterion. Orange = *Static*-first, blue = *Dynamic*-first.




Figure 10: The NASA-TLX results for the *Static* and *Dynamic* conditions (lower indicates lower workload).

6.2.1 NASA-TLX. We compared perceived workload between *Static* and *Dynamic* conditions with the NASA Task Load Index (NASA-TLX), which comprises six dimensions: *Mental demand*, *Physical demand*, *Temporal demand*, *Performance*, *Effort*, and *Frustration* (see Figure 10). Lower scores indicate less workload. Overall workload was lower with *Dynamic* than with *Static* ($\Delta_{\text{DYN-STA}} = -4.5$ on the

Table 2: Hand similarity per user. Static/Dynamic \times Immediate/Retention.

User Index	Static (Imm.) μ	Dynamic (Imm.) μ	Static (Ret.) μ	Dynamic (Ret.) μ
1	0.641	0.775	0.18	0.523
2	0.349	0.652	0.315	0.584
3	0.689	0.741	0.569	0.934
⋮	⋮	⋮	⋮	⋮
29	0.558	0.657	0.486	0.703
30	0.590	0.800	0.549	0.862
Overall t-Test	0.465	0.610	0.414	0.611
Cohen's d_z		$p = 0.000743$		$p = 0.000021$
		0.688		0.925

0–100 scale, 95% CI [−5.9, −3.2]; Holm-adjusted $p < .001$; large within-subject effect $d_z = 1.30$).

Across the individual NASA-TLX dimensions, *Dynamic* reduced perceived workload relative to *Static*. The largest reductions were in *Performance* ($\Delta_{\text{Dyn-Sta}} = -8.5$) and *Effort* ($\Delta_{\text{Dyn-Sta}} = -8.2$), followed by *Frustration* ($\Delta_{\text{Dyn-Sta}} = -5.8$) and *Mental demand* ($\Delta_{\text{Dyn-Sta}} = -4.2$). These effects remained significant after Holm adjustment ($p \leq .013$) and were medium to large in magnitude ($d_z = 0.57$ – 1.16). In contrast, differences in *Temporal demand* ($\Delta_{\text{Dyn-Sta}} = -0.7$) and *Physical demand* ($\Delta_{\text{Dyn-Sta}} = +0.1$) were small and not reliable ($p \geq .14$).

Together, adaptive transparency primarily reduced perceived performance difficulty/success cost, effort, and frustration, with smaller effects on mental demand and no change in physical demand.

6.2.2 Post-study Questionnaires. Reliability. All questionnaires used 5-point Likert scales. Internal consistency was acceptable for Learning Experience in *Static* ($\alpha = .646$) and good in *Dynamic* ($\alpha = .885$); the Visual Design & Transparency scale (Dynamic-only) was also good ($\alpha = .805$). Retention reliabilities were lower (*Static* $\alpha = .627$, *Dynamic* $\alpha = .587$), so retention phase estimates should be interpreted with caution.

Learning experience and Retention. Participants rated *Dynamic* higher than *Static* during practice ($t(29) = 3.834$, $p = .0006$, $d_z = 0.700$; $\Delta = 0.529$, 95% CI [0.258, 0.799]) and during the retention phase ($t(29) = 4.317$, $p = .0002$, $d_z = 0.788$; $\Delta = 0.452$, 95% CI [0.247, 0.658]).

Comparative evaluation and Preference. The comparative Likert items were numerically higher for *Dynamic* (3.867 vs. 3.533; $t(29) = 1.980$, $p = .057$, $d_z = 0.361$; $\Delta = 0.334$). Preference responses strongly favored *Dynamic*: across three items ($30 \times 3 = 90$ responses), 37 were marked “Not sure” and excluded, leaving $N = 53$ decisive responses; of these, 47 (88.7%) preferred *Dynamic* (Wilson 95% CI [0.774, 0.947]; binomial $z = 5.632$, $p < .001$). Facet-level results were consistent with the overall finding: *Dynamic* exceeded *Static* on Overall, Memorization, and Naturalness.

6.3 Exploratory Analysis

6.3.1 Hand-Ghost Motion Similarity. We aligned users’ right-hand joint trajectories to the ghost reference via key-event-based linear time fitting and resampled them onto the reference timeline. Comparisons used the user-ghost joint overlap in wrist-anchored coordinates: (i) position cosine; (ii) velocity cosine on derivatives;

and (iii) wrist→joint direction (frame-wise cosine; median angular error). As an alignment-robust check, DTW mean step cost is computed on position, direction, or velocity signals. Condition effects (Dynamic vs. Static) were tested within subjects (Immediate/Retention) using paired t -tests with Cohen’s d_z ; per-user means (μ) are summarized in Table 2.

In a within-subjects comparison ($N=30$), the *Dynamic* transparency mode outperformed *Static* both immediately and at retention. Immediate test: $M_{\text{DYN}} = 0.610$ vs. $M_{\text{STAT}} = 0.465$, $\Delta = 0.144$, 95% CI = [0.069, 0.216], paired t -test $p = 0.000743$, Cohen’s $d_z = 0.688$. Retention test: $M_{\text{DYN}} = 0.611$ vs. $M_{\text{STAT}} = 0.414$, $\Delta = 0.197$, 95% CI = [0.124, 0.272], $p = 0.000021$, $d_z = 0.925$. Effects are medium to large for the retention test, indicating better short-term learning retention under *Dynamic*.

7 Discussion

7.1 Quantitative Findings

Our central finding is that *Dynamic* guidance yields significantly better retention than *Static*, with higher pitch accuracy, higher fingering accuracy, and lower error rates—directly addressing over-reliance on persistent cues in VR motor-skill training [78]. These outcomes align with **H2 (Retention)** and provide clear empirical support for the *guidance hypothesis*, which holds that excessive or continuous feedback, characteristic of our *Static* condition, suppresses error processing during practice and degrades later performance [68, 72]. Prior ghost/co-embodiment studies have similarly noted that fixed-visibility guidance can accelerate early learning yet incurs a performance drop once guidance is removed [47, 99]. By strategically fading visibility during proficient passages and restoring it after slips, our *Dynamic* controller mitigates this drop, yielding stable retention across pitch, fingering, and overall errors. Furthermore, surveys of XR piano systems [21, 85] highlight that prior work predominantly uses *fixed-visibility* cues (e.g., [2, 65]). We directly address this gap by adapting *knowledge-of-performance* cues—specifically, fingering visibility—and demonstrate significant gains over *Static* on retention-oriented outcomes.

Unlike prior adaptive-support approaches that primarily vary *how* it is delivered (e.g., haptic modality [51] or task-difficulty actuation [81, 101]), or *who* provides guidance (e.g., remote or dyadic co-embodiment [60]), our method varies *how much* guidance is provided at any moment. Treating tutor visibility as a continuously allocatable resource allows us to scale support in real time, preserving learner autonomy during practice while still enabling expert demonstration when needed. In doing so, we directly mitigate the over-reliance that co-embodiment systems frequently encounter.

Our system contributes a simpler, and purely visual method that achieves the same conceptual goal: balancing instructional support and learner autonomy, as conceptualized by Guadagnoli et al. in the Challenge Point framework [34]. The success of our approach is further reinforced by the subjective perception results. Participants perceived the *Dynamic* condition as having significantly lower workload, particularly in dimensions of perceived performance, effort, and frustration. This convergence of improved subjective experience and objective retention (H2) indicates that adapting visibility is a highly effective method for regulating instructional support. During the training phase, while group-average

learning rates were similar, the per-participant slope analysis revealed that *Dynamic* guidance produced significantly steeper error-reduction slopes. This supports **H3 (Training convergence)**. Moreover, learners who started with *Dynamic* began the second block (new melody) at a higher initial performance level, supporting **H1 (Skill transfer)**. This suggests that the benefits of faded feedback are not limited to retention, but also promote the generalization of self-correction skills, a known failing of continuous guidance which can degrade transfer [72].

By contrast, we found no reliable timing advantage. This boundary condition suggests that visibility adaptation alone does not provide sufficient temporal scaffolding. Two plausible interpretations are that the controller emphasized correctness cues more than temporal structuring, and that the interface lacked explicit temporal cues (e.g., a beat indicator or metronome). It is also possible that learners relied on their own internal timing once note targets became familiar, reducing the marginal benefit of visibility adaptation for rhythm. Consistent with the specificity-of-practice principle [54], our controller prioritized pitch ($w_p=0.7$) over timing ($w_t=0.2$), and qualitative feedback frequently requested explicit rhythmic aids. To address timing guidance, we outline targeted visual and rhythmic additions in Limitations and Future Work (Section 7.5).

A notable diagnostic pattern in the *Static* condition helps explain part of the pitch gap: several participants (e.g., **P1**, **P5**, **P15**) produced consistent *octave-shift* errors (playing one octave too high) in retention tests. As shown in Table 2, **P1**'s similarity score dropped sharply from the *Static Immediate* to the *Static Retention* phase, reflecting the impact of the octave transposition on positional alignment. This kind of systematic transposition inflates pitch error even when relative finger sequences are otherwise preserved. By contrast, *Dynamic* guidance appears to stabilize spatial anchoring on the keyboard and to reinforce the intended fingering trajectory; this is consistent with its higher finger accuracy and lower error rates. The adaptively transparent exemplar reduces external guidance during stable passages and restores it after slips, providing just-in-time spatial landmarks (where to land, which finger to use) and temporal hints (when to move). Together, these help prevent drift into a transposed register and facilitate reinstatement of the learned pattern, thereby increasing retention.

7.2 Hypotheses Revisited

Based on the results, we return to our hypotheses.

H1 (Skill transfer). Evidence supports H1. As shown in the individual differences in Section 6.1.4, at the switch to the second block (new melody and condition), participants who trained under *Dynamic* first began at a higher initial performance (mean of the first two loops). This advantage held after controlling for melody and was visible within each melody; it also persisted to the end of the block with a larger share of participants exceeding the 0.60 criterion. These patterns indicate a transfer advantage that carries over to the new task setting.

H2 (Retention). Largely supported. *Dynamic* showed smaller *Immediate* → *Retention* declines, especially for pitch and fingering, while timing remained comparable. This implies that correctness benefits stem from visibility adaptation, whereas rhythm likely requires explicit temporal cues.

H3 (Training convergence). Partially supported. Group learning curves (Sec. 6.1.4) showed overall improvement under both conditions and a small *Dynamic* advantage, but no reliable *Condition*×*Loop* interaction, implying similar average learning rates. Crucially, at the individual level, the per-participant slope analysis (Figure 8(b)) reveals a statistically reliable shift toward steeper error-rate reductions under *Dynamic*: the mean Δ slope (*Dynamic*–*Static*) < 0 with a 95% CI entirely below 0. Together with the larger share of *Dynamic*-first learners exceeding the 0.60 criterion by the end of the second block (Figure 9(b)), these results indicate a modest but consistent convergence advantage for *Dynamic*. The absence of a group-level interaction likely reflects learner heterogeneity and non-linear/ceiling effects that mask a subtle individual-level difference in the aggregate.

7.3 Qualitative Feedback

We analyzed open-ended responses from the post-study questionnaire (26/30 participants provided at least one comment). Three themes emerged.

What participants liked. The most frequent positives were *finger positioning clarity* and *followability*. Participants highlighted that the ghost hand made finger placement easy to imitate and track (11 mentions; e.g., “*The position of the fingers is easy to follow.*”—**P10**; “*... useful for practicing both timing and finger positions.*”—**P8**). Eight comments praised the overall clarity/guidance (“*clear and easy to follow*”—**P2**). A few explicitly mentioned help with memory/location (2) and rhythm/timing (2), and some noted that the dynamic transparency felt natural (3; e.g., “*looking back I actually like the concept and it felt natural*”—**P7**) and fun/motivating (2).

What was confusing or difficult. **General issues**. The most common issue was *hand-tracking on ring/pinky fingers* (6; e.g., “*The ring finger didn't function too well.*”—**P2**; “*... didn't register properly when played with ring- or pinky finger.*”—**P8**). Some reported *accidental keypresses* / touching neighboring keys (3; **P1**, **P7**, **P16**), *audio/rhythm confusion* when both sources were audible (1; **P9**), and *physical fatigue* (1; **P1**).

Condition-specific. In the *Static* condition, participants mentioned *occlusion/visibility* of ghost fingers by their own hands (1; **P4**) and *afterimages* (1; **P1**). By contrast, several comments favored *Dynamic* for being more intuitive and easier to follow, with fewer visibility/occlusion issues and less conflicting cues (e.g., **P21**, **P11**, **P19**, **P24**); some also noted that visibility recovered when dynamic fading became too light (e.g., **P28**).

Improvements suggested. Participants asked for (i) *better tracking for ring/pinky* (3; **P8**, **P10**), (ii) *alternative views / camera options* to improve *viewpoint/line-of-sight* and judging contact (4; **P4**, **P9**), (iii) *haptic/physical keyboard support* (3; **P10**, **P16**), (iv) *beat/metronome indicators* (1; **P7**), and (v) *notation/labels on keys* (1; **P1**). Several of these directly target the issues above—tracking, viewpoint/line-of-sight for contact judgment (with occlusion largely mitigated by *Dynamic*), rhythm guidance, and input robustness.

7.4 Implications and Generalizability

Our findings suggest that dynamically modulated transparency reduces occlusion and focuses attention on *where* to land and *when*

to move in visuomotor tasks that share three properties: (i) clear spatial targets, (ii) temporal sequencing, and (iii) intermittent self-occlusion. This is consistent with motor-learning evidence that adaptive or faded augmented feedback mitigates guidance dependence and supports retention [4, 87]. Potential application areas include *sports posture and rhythm drills* with AR/VR coaching [90], *handwriting and calligraphy* with guided stroke trajectories [91], and *upper-limb rehabilitation* where VR/AR adjuncts support functional gains [49]. However, *domain-specific validation is required* before drawing conclusions.

We propose the following design guidelines: (1) Fade with proficiency, recover after slips. Treat guidance as a scaffold, not a crutch: use an opacity controller that decreases guidance visibility as the learner's performance stabilizes and increases visibility after detectable errors; (2) Keep guidance visible. Composite the instructor above the learner's representation with adaptive transparency, rendered without depth buffering, thereby keeping guidance visible while preventing occlusion by the environment; (3) Structure practice and probe retention. Support controlled, repeatable practice with task-aligned feedback: provide phrase-level chunking and looped replays, and follow with no-guidance probes to assess retention. These guidelines operationalize adaptive transparency as a strategy that reduces visual occlusion and cue dependence while enabling repeatable, task-aligned practice. These guidelines cover tasks that make use of clear spatial targets and temporal sequencing; assessing their validity in other domains remains future work.

7.5 Limitations and Future Work

Algorithm Personalization. Our study employed a fixed weighting scheme (e.g., pitch 70%, timing 20%, fingering 10%) to calculate the performance score that drives transparency. While this weighting is grounded in pedagogical principles prioritizing note accuracy for beginners, a "one-size-fits-all" approach may not be optimal for every learner. For instance, a user who masters pitch quickly but struggles with rhythm might benefit from a higher weight on timing accuracy. Future work should investigate dynamic and personalized weighting schemes. This could involve allowing the system to automatically adjust weights based on a learner's recurring error patterns, or even providing users direct control to customize the feedback's focus according to their personal learning goals.

Tracking and input robustness. Participants noted ring and pinky finger tracking errors and occasional stray key presses. In our current prototype, key presses are triggered solely by virtual contact between the tracked hand and the virtual keys; there is no physical keyboard or force-sensitive dynamics. As a result, learners cannot rely on tactile cues to judge when a key has been fully depressed, and small tracking jitters can lead to missed or spurious key events. In addition, both the learner's tracked hand and the tutor's ghost hand use the platform's generic hand model with fixed proportions that do not exactly match each user's hand size and span, which reduces ergonomic realism. Promising directions for future work include improving distal-joint tracking (e.g., with an auxiliary optical hand-tracking camera), adding light *debounce/false-touch* filtering, and personalizing hand scaling via a brief calibration step. Another avenue is to deploy the same adaptive transparency controller on top of a physical keyboard and real hands in AR/MR,

where tactile feedback and visual appearance more closely match real-world piano study.

View control and rhythm support. The current system offers a single first-person view and minimal tempo cues. Because timing differences between conditions were not reliable and participants suggested adding view options and rhythm aids, provide flexible view presets to keep the active region and the instructor visible; useful presets include snap to keyboard center, follow the instructor, and fit the active phrase region. Add rhythm cues such as an auditory metronome with a subtle visual pulse or a tempo strip, a one bar count-in, gradual tempo ramps, and optional beat subdivisions to support timing and improve visibility as well as temporal guidance.

Instructor motion provenance. In this study, the instructor trajectories were recorded by the authors for proof of concept. Such recordings may encode idiosyncratic fingering and micro-timing that differ from expert pedagogy and across styles. Future work will capture demonstrations from professional piano instructors, compare expert- vs. self-recorded guidance, and explore aggregating or synthesizing instructor motions to provide style-appropriate cues.

Generalizability and longevity. Our lab study focused on right-hand, monophonic, short phrases (≈ 12 s) in a controlled setting with immediate and 10-minute retention tests. This simplification was intentional to isolate the effect of adaptive visibility and ensure reliable hand-pose and key-press sensing in a within-subjects design, but it limits ecological validity relative to real piano learning (bimanual coordination, polyphony, expressive timing, longer phrases). It could be extended to bimanual/polyphonic repertoire, longer excerpts, and variable tempi/dynamics in authentic practice settings. Because the controller is performance-contingent (fade when stable, restore after slips) rather than score-specific, it should transfer naturally to these scenarios. In such settings, issues arising from self-occlusion and visual clutter are likely to be more pronounced, and we expect that our dynamic method would mitigate these effects more effectively than a static condition.

8 Conclusion

We introduced a skill-adaptive, dynamically transparent ghost instructor for VR piano training that scales its visibility to the learner's moment-to-moment performance. The controller aggregates pitch, timing, and fingering accuracy into a weighted performance signal. This is mapped to instructor opacity with an asymmetric exponential moving average combined with a clamped visibility range, aiming to reduce cue dependence, and improve learning and short-term retention.

Across a within-subjects study, the *Dynamic* mode consistently outperformed a *Static* baseline. Participants achieved higher pitch and fingering accuracy and lower error rates under guidance, and these advantages persisted when cues were removed. *Dynamic* also limited performance decline from immediate to retention, while timing showed no reliable condition effect. Learning-curve analyses indicated steady improvement for both modes with a small but reliable advantage for *Dynamic*. Subjectively, workload was lower and preferences favored *Dynamic*.

Overall, *Skill-Adaptive Ghost Instructors* helps learners internalize finger placement and motion patterns without fostering over-reliance on on-screen cues. Practically, we recommend fading assistance after correct behavior and reintroducing it after slips, capping visibility to preserve scene readability, and weighting feedback toward task-critical dimensions while monitoring technique. Beyond piano, dynamically transparent guidance is promising for skills combining spatial accuracy, rhythmic control, and ergonomics, including other instruments, upper-limb rehabilitation, and sports technique.

9 Use of Large Language Models

We used large language models solely for English language editing (grammar and wording). All study designs, analyses, and claims are the authors' own, and no participant data were uploaded to any LLM services.

Acknowledgments

Research was supported by the e-DIPLOMA project (101061424)¹, funded by the European Union. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

References

- [1] Emma Allingham, Renee Timmers, and George Waddell. 2022. Slow practice and tempo-management strategies in instrumental learning. *Psychology of Music* 50, 7 (2022), 2694–2712. doi:10.1177/03057356211073481
- [2] V. Amm et al. 2024. Mixed reality strategies for piano education. *Frontiers in Virtual Reality* (2024). <https://www.frontiersin.org/articles/10.3389/frvir.2024.1397154/full>
- [3] Judith Amores, Xavier Benavides, and Pattie Maes. 2015. Showme: A remote collaboration system that supports immersive gestural communication. In *Proceedings of the 33rd annual ACM conference extended abstracts on human factors in computing systems*. 1343–1348.
- [4] David I Anderson, Richard A Magill, Hiroshi Sekiya, and Greg Ryan. 2005. Support for an explanation of the guidance effect in motor skill learning. *Journal of motor behavior* 37, 3 (2005), 231–238.
- [5] Mathieu Andrieux and Luc Proteau. 2016. Observational Learning: Tell Beginners What They Are about to Watch and They Will Learn Better. *Frontiers in Psychology* 7, 51 (Jan 2016). doi:10.3389/fpsyg.2016.00051
- [6] Yoichiro Aoyagi, Eri Ohnishi, Yoshinori Yamamoto, Naoki Kado, Toshiaki Suzuki, Hitoshi Ohnishi, Nozomi Hokimoto, and Naomi Fukaya. 2019. Feedback protocol of 'fading knowledge of results' is effective for prolonging motor learning retention. *Journal of Physical Therapy Science* 31, 8 (2019), 687–691.
- [7] Mariano Banquero, Gracia Valdeolivas, David Ramón, and M-Carmen Juan. 2024. A color Passthrough mixed reality application for learning piano. *Virtual Reality* 28, 2 (2024), 67.
- [8] Mariano Banquero, Gracia Valdeolivas, Sergio Trincado, Natasha Garcia, and M-Carmen Juan. 2023. Passthrough mixed reality with oculus quest 2: A case study on learning piano. *IEEE MultiMedia* 30, 2 (2023), 60–69.
- [9] Simone Dalla Bella and Caroline Palmer. 2011. Rate Effects on Timing, Key Velocity, and Finger Kinematics in Piano Performance. *PLoS ONE* 6, 6 (2011), e20518. doi:10.1371/journal.pone.0020518
- [10] Amare Birhanu and Stefan Rank. 2017. KeynVision: exploring piano pedagogy in mixed reality. In *Extended abstracts publication of the annual symposium on computer-human interaction in play*. 299–304.
- [11] Yannick Blandin, Léna Lhuisset, and Luc Proteau. 1999. Cognitive Processes Underlying Observational Learning of Motor Skills. *The Quarterly Journal of Experimental Psychology Section A* 52, 4 (Nov 1999), 957–979. doi:10.1080/713755856
- [12] Malte Borgwardt, Jonas Boueke, María Fernanda Sanabria, Michael Bonfert, and Robert Porzel. 2023. Vrisbee: How hand visibility impacts throwing accuracy and experience in virtual reality. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–7.
- [13] Arnaud Boutin and Yannick Blandin. 2010. Cognitive underpinnings of contextual interference during motor learning. *Acta psychologica* 135, 2 (2010), 233–239.
- [14] John J Buchanan and Chaoyi Wang. 2012. Overcoming the guidance effect in motor skill learning: feedback all the time can be beneficial. *Experimental Brain Research* 219, 2 (2012), 305–320.
- [15] Minya Cai, Muhammad Alfian Amrizal, Toru Abe, and Takuo Suganuma. 2019. Design and implementation of ar-supported system for piano learning. In *2019 IEEE 8th Global Conference on Consumer Electronics (GCCE)*. IEEE, 49–50.
- [16] Matteo Candidi, Lucia Maria Sacheli, Ilaria Mega, and Salvatore Maria Aglioti. 2014. Somatotopic mapping of piano fingering errors in sensorimotor experts: TMS studies in pianists and visually trained musically naives. *Cerebral cortex* 24, 2 (2014), 435–443.
- [17] Winyu Chinthammit, Troy Merritt, Scott Pedersen, Andrew Williams, Denis Visentin, Robert Rowe, and Thomas Furness. 2014. Ghostman: augmented reality application for telerehabilitation and remote instruction of a novel motor skill. *BioMed research international* 2014, 1 (2014), 646347.
- [18] Jonathan Chow, Haoyang Feng, Robert Amor, and Burkhard C Wünsche. 2013. Music education using augmented reality with a head mounted display. In *Proceedings of the Fourteenth Australasian User Interface Conference-Volume 139*. 73–79.
- [19] Matthew E Clark, Kayla McEwan, and Candice J Christie. 2019. The effectiveness of constraints-led training on skill development in interceptive sports: A systematic review. *International Journal of Sports Science & Coaching* 14, 2 (2019), 229–240.
- [20] Palle Dahlstedt. 2015. Mapping strategies and sound engine design for an augmented hybrid piano. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 271–276.
- [21] Jordan Aiko Deja, Sven Mayer, Klen Čopič Pucihar, and Matjaž Kljun. 2022. A survey of augmented piano prototypes: Has augmentation improved learning experiences? *Proceedings of the ACM on Human-Computer Interaction* 6, ISS (2022), 226–253.
- [22] Peter Dürking, Hans-Christer Holmberg, and Billy Sperlich. 2018. The Potential Usefulness of Virtual Reality Systems for Athletes: A Short SWOT Analysis. *Frontiers in Physiology* 9 (2018), 128. <https://www.frontiersin.org/articles/10.3389/fphys.2018.00128/full>
- [23] Annerose Engel, Marc Bangert, David Horbank, Brenda S Hijmans, Katharina Wilkens, Peter E Keller, and Christian Keysers. 2012. Learning piano melodies in visuo-motor or audio-motor training conditions and the neural correlates of their cross-modal transfer. *NeuroImage* 63, 2 (2012), 966–978.
- [24] Rebecca Fribourg, Nami Ogawa, Ludovic Hoyet, Ferran Argelaguet, Takuji Narumi, Michitaka Hirose, and Anatole Lécuyer. 2020. Virtual co-embodiment: evaluation of the sense of agency while sharing the control of a virtual body among two individuals. *IEEE Transactions on Visualization and Computer Graphics* 27, 10 (2020), 4023–4038.
- [25] Kodai Fuchino, Mohammed Al-Sada, and Tatsuo Nakajima. 2023. T2Remoter: A remote table tennis coaching system combining VR and robotics. In *2023 IEEE 29th International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*. IEEE, 275–276.
- [26] Shinichi Furuya, Ayumi Nakamura, and Noriko Nagata. 2013. Transfer of piano practice in fast performance of skilled finger movements. *BMC neuroscience* 14, 1 (2013), 133.
- [27] Danilo Gasques, Janet G Johnson, Tommy Sharkey, Yuanyuan Feng, Ru Wang, Zhuoqun Robin Xu, Enrique Zavala, Yifei Zhang, Wanze Xie, Xinning Zhang, et al. 2021. Artemis: A collaborative mixed-reality system for immersive surgical telementoring. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–14.
- [28] Lynda Gerry, Sofia Dahl, and Stefania Serafin. 2019. ADEPT: exploring the design, pedagogy, and analysis of a mixed reality application for piano training. In *16th sound and music computing conference*. Sound and Music Computing Network, 241–249.
- [29] Seth Glickman, Byunghwan Lee, Fu Yen Hsiao, and Shantanu Das. 2017. Music Everywhere—Augmented Reality Piano Improvisation Learning System. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 511–512.
- [30] Naska Goagoses, Heike Winschiers-Theophilus, Selma Auala, Nicolas Pope, Erkki Rötönen, Helvi Itenge, Calkin Suero Montero, Tomi Suovuo, and Erkki Sutinen. 2024. Teachers and students envisioning mixed reality remote learning: A qualitative exploration on fostering academic engagement. *Technology, Knowledge and Learning* (2024), 1–25.
- [31] Werner Goebel and Caroline Palmer. 2013. Temporal Control and Hand Movement Efficiency in Skilled Music Performance. *Journal of Experimental Psychology: Human Perception and Performance* (2013).
- [32] Geoffrey Gorisse, Olivier Christmann, Etienne Armand Amato, and Simon Richir. 2017. First- and third-person perspectives in immersive virtual environments:

¹<https://e-diplomaproject.eu/>

- presence and performance analysis of embodied users. *Frontiers in Robotics and AI* 4 (2017), 33.
- [33] Rob Gray. 2017. Transfer of training from virtual to real baseball batting. *Frontiers in psychology* 8 (2017), 317328.
- [34] Mark A Guadagnoli and Timothy D Lee. 2004. Challenge point: a framework for conceptualizing the effects of various practice conditions in motor learning. *Journal of motor behavior* 36, 2 (2004), 212–224.
- [35] Ruoxi Guo, Jiahao Cui, Wanru Zhao, Shuai Li, and Aimin Hao. 2021. Hand-by-hand mentor: An AR based training system for piano performance. In *2021 IEEE conference on virtual reality and 3D user interfaces abstracts and workshops (VRW)*. IEEE, 436–437.
- [36] Shlomi Haar, Guhan Sundar, and A Aldo Faisal. 2021. Embodied virtual reality for the study of real-world motor learning. *Plos one* 16, 1 (2021), e0245717.
- [37] Dominik Hackl and Christoph Anthes. 2017. HoloKeys—an augmented reality application for learning the piano.. In *Forum media technology*. 140–144.
- [38] Lamtharn Hantrakul and Konrad Kaczmarek. 2014. Implementations of the Leap Motion in sound synthesis, effects modulation and assistive performance tools.. In *ICMC*.
- [39] Amal Hatira, Zeynep Ecem Gelmez, Anil Ufuk Batmaz, and Mine Sarac. 2024. Effect of Hand and Object Visibility in Navigational Tasks Based on Rotational and Translational Movements in Virtual Reality. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, 115–125.
- [40] Steven Henderson and Steven Feiner. 2011. Exploring the Benefits of Augmented Reality Documentation for Maintenance and Repair. *IEEE Transactions on Visualization and Computer Graphics* 17, 10 (2011), 1355–1368. doi:10.1109/TVCG.2010.245
- [41] Melynda Hoover and Eliot Winer. 2021. Designing adaptive extended reality training systems based on expert instructor behaviors. *IEEE Access* 9 (2021), 138160–138173.
- [42] Lei Hou, Xiangyu Wang, and Martijn Truijens. 2015. Using Augmented Reality to Facilitate Piping Assembly: An Experiment-Based Evaluation. *Journal of Computing in Civil Engineering* 29, 1 (2015), 04014007. doi:10.1061/(ASCE)CP.1943-5487.0000344
- [43] Felix Hülsmann, Cornelia Frank, Irene Senna, Marc O Ernst, Thomas Schack, and Mario Botsch. 2019. Superimposed skilled performance in a virtual mirror improves motor performance and cognitive representation of a full body motor action. *Frontiers in Robotics and AI* 6 (2019), 43.
- [44] David Johnson, Daniela Damian, and George Tzanetakis. 2020. Evaluating the effectiveness of mixed reality music instrument learning with the theremin. *Virtual Reality* 24, 2 (2020), 303–317.
- [45] Hiroki Kawasaki, Hiroyuki Iizuka, Shin Okamoto, Hideyuki Ando, and Taro Maeda. 2010. Collaboration and skill transmission by first-person perspective view sharing system. In *19th international symposium in robot and human interactive communication*. IEEE, 125–131.
- [46] Charles R Kelley. 1969. What is adaptive training? *Human Factors* 11, 6 (1969), 547–556.
- [47] Daiki Kodama, Takato Mizuho, Yuji Hatada, Takuji Narumi, and Michitaka Hirose. 2023. Effects of collaborative training using virtual co-embodiment on motor skill learning. *IEEE Transactions on Visualization and Computer Graphics* 29, 5 (2023), 2304–2314.
- [48] Katerina Labrou, Cagri Hakan Zaman, Arda Turkyasar, and Randall Davis. 2023. Following the master’s hands: Capturing piano performances for mixed reality piano learning applications. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [49] Kate E Laver, Stacey George, Susie Thomas, Judith E Deutsch, and Maria Crotty. 2011. Virtual reality for stroke rehabilitation. *Cochrane database of systematic reviews* 9 (2011).
- [50] Morgan Le Chénéchal, Thierry Duval, Valérie Gouranton, Jérôme Royan, and Bruno Arnaldi. 2016. Vishnu: virtual immersive support for helping users an interaction paradigm for collaborative remote guiding in mixed reality. In *2016 IEEE third VR international workshop on collaborative virtual environments (3DCVE)*. IEEE, 9–12.
- [51] Kyungyeon Lee, Daniel S Yang, Kriti Singh, and Jun Nishida. 2025. Hapticus: Exploring the Effects of Haptic Feedback and its Customization on Motor Skill Learning: Tactile, Haptic, and Somatosensory Approaches. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [52] Hui Liang, Jin Wang, Qian Sun, Yong-Jin Liu, Junsong Yuan, Jun Luo, and Ying He. 2016. Barehanded music: real-time hand interaction for virtual piano. In *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. 87–94.
- [53] Lasse F Lui, Unnikrishnan Radhakrishnan, Francesco Chinello, and Konstantinos Koumaditis. 2025. The efficacy of adaptive training in immersive virtual reality for a fine motor skill task. *Virtual Reality* 29, 1 (2025), 20.
- [54] Isabelle Mackrous and Luc Proteau. 2007. Specificity of practice results from differences in movement planning strategies. *Experimental brain research* 183, 2 (2007), 181–193.
- [55] Laura Marchal-Crespo, Mark van Raai, Georg Rauter, Peter Wolf, and Robert Riener. 2013. The effect of haptic guidance and visual feedback on learning a complex tennis task. *Experimental brain research* 231, 3 (2013), 277–291.
- [56] Anat Mirelman, Lynn Rochester, Inbal Maidan, Silvia Del Din, Lisa Alcock, Freek Nieuwhof, Marcel Olde Rikkert, Bastiaan R Bloem, Elisa Pelosin, Laura Avanzino, et al. 2016. Addition of a non-immersive virtual reality component to treadmill training to reduce fall risk in older adults (V-TIME): a randomised controlled trial. *The Lancet* 388, 10050 (2016), 1170–1182.
- [57] Will Molloy, Edward Huang, and Burkhard C Wünsche. 2019. Mixed reality piano tutor: A gamified piano practice environment. In *2019 International conference on electronics, information, and communication (ICEIC)*. IEEE, 1–7.
- [58] David L. Neumann, Robyn L. Moffitt, Patrick R. Thomas, Katrina Loveday, et al. 2018. A Systematic Review of the Application of Interactive Virtual Reality to Sport. *Frontiers in Psychology* 9 (2018), 2159. <https://research-repository.griffith.edu.au/items/25a2bc11-d7a6-590f-bab3-ab8622679ec8>
- [59] David L Neumann, Robyn L Moffitt, Patrick R Thomas, Kylie Loveday, David P Watling, Chantal L Lombard, Simona Antonova, and Michael A Tremer. 2018. A systematic review of the application of interactive virtual reality to sport. *Virtual Reality* 22, 3 (2018), 183–198.
- [60] Jun Nishida, Yudai Tanaka, Romain Nith, and Pedro Lopes. 2022. DigiTusync: A dual-user passive exoskeleton glove that adaptively shares hand gestures. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–12.
- [61] Honghu Pan, Xingxi He, Hong Zeng, Jia Zhou, and Sai Tang. 2018. Pilot study of piano learning with AR smart glasses considering both single and paired play. In *International Conference on Human Aspects of IT for the Aged Population*. Springer, 561–570.
- [62] Richard Parmcutt, John A. Sloboda, Eric F. Clarke, Mikko Raekallio, and Peter Desain. 1997. An Ergonomic Model of Keyboard Fingering for Melodic Fragments. In *Proceedings of the International Society for Music Information Retrieval (or related venue)*. Foundational model linking fingering to ergonomic efficiency.
- [63] Matevž Pesek, Nejc Hirci, Klara Žnidarsič, and Matija Marolt. 2024. Enhancing music rhythmic perception and performance with a VR game. *Virtual Reality* 28, 2 (2024), 118.
- [64] Sathiya Kumar Renganayagalu, Steven C Mallam, and Salman Nazir. 2021. Effectiveness of VR head mounted displays in professional training: A systematic review. *Technology, Knowledge and Learning* 26, 4 (2021), 999–1041.
- [65] Liam Rigby, Burkhard C Wünsche, and Alex Shaw. 2020. piARno—an augmented reality piano tutor. In *Proceedings of the 32nd Australian conference on human-computer interaction*. 481–491.
- [66] Luiz Rodrigues, Paula T Palomino, Armando M Toda, Ana CT Klock, Wilk Oliveira, Anderson P Avila-Santos, Isabela Gasparini, and Seiji Isotani. 2021. Personalization improves gamification: Evidence from a mixed-methods study. *Proceedings of the ACM on Human-Computer Interaction* 5, CHI PLAY (2021), 1–25.
- [67] Katja Rogers, Amrei Röhlig, Matthias Weing, Jan Gugenheimer, Bastian Könings, Melina Klepsch, Florian Schaub, Enrico Rukzio, Tina Seufert, and Michael Weber. 2014. Piano: Faster piano learning with interactive projection. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces*. 149–158.
- [68] Alan W Salmoni, Richard A Schmidt, and Charles B Walter. 1984. Knowledge of results and motor learning: a review and critical reappraisal. *Psychological bulletin* 95, 3 (1984), 355.
- [69] Sean Sanford, Brian Collins, Mingxiao Liu, Sophie Dewil, and Raviraj Nataraj. 2022. Investigating features in augmented visual feedback for virtual reality rehabilitation of upper-extremity function through isometric muscle control. *Frontiers in virtual reality* 3 (2022), 943693.
- [70] Marc Ericson C Santos, Angie Chen, Takafumi Taketomi, Goshiro Yamamoto, Jun Miyazaki, and Hirokazu Kato. 2013. Augmented reality learning experiences: Survey of prototype design and evaluation. *IEEE Transactions on learning technologies* 7, 1 (2013), 38–56.
- [71] Ruben Schlagowski, Dariia Nazarenko, Yekta Can, Kunal Gupta, Silvan Mertes, Mark Billinghurst, and Elisabeth André. 2023. Wish you were here: Mental and physiological effects of remote music collaboration in mixed reality. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–16.
- [72] Richard A Schmidt and Gabriele Wulf. 1997. Continuous concurrent feedback degrades skill learning: Implications for training and simulation. *Human factors* 39, 4 (1997), 509–525.
- [73] Ben Sellers and G Waddell. 2024. Augmented reality and in-person piano tuition: project report. (2024).
- [74] Neal E. Seymour, Anthony G. Gallagher, Sanziana A. Roman, Michael K. O’Brien, Vipin K. Bansal, Dana K. Andersen, and Richard M. Satava. 2002. Virtual reality training improves operating room performance: results of a randomized, double-blinded study. *Annals of Surgery* 236, 4 (2002), 458–463. doi:10.1097/0000658-200210000-00008
- [75] Neal E Seymour, Anthony G Gallagher, Sanziana A Roman, Michael K O’Brien, Vipin K Bansal, Dana K Andersen, and Richard M Satava. 2002. Virtual reality

- training improves operating room performance: results of a randomized, double-blinded study. *Annals of surgery* 236, 4 (2002), 458–464.
- [76] Yongmin Shin, Jaeseo Lim, Yonggwan Kim, Deog-Gyu Seo, and Jungjoon Ihm. 2022. Effects of virtual body-representation on motor skill learning. *Scientific Reports* 12, 1 (2022), 15283.
- [77] Roland Sigrist, Georg Rauter, Laura Marchal-Crespo, Robert Riener, and Peter Wolf. 2015. Sonification and haptic feedback in addition to visual feedback enhances complex motor task learning. *Experimental brain research* 233, 3 (2015), 909–925.
- [78] Roland Sigrist, Georg Rauter, Robert Riener, and Peter Wolf. 2013. Augmented visual, auditory, haptic, and multimodal feedback in motor learning: a review. *Psychonomic bulletin & review* 20, 1 (2013), 21–53.
- [79] Diane M Ste-Marie, Barbi Law, Amanda M Rymal, O Jenny, Craig Hall, and Penny McCullagh. 2012. Observation interventions for motor skill learning and performance: an applied model for the use of observation. *International Review of Sport and Exercise Psychology* 5, 2 (2012), 145–176.
- [80] Juan Trelles Trabucco, Andrea Rottigni, Marco Cavallo, Daniel Bailey, James Patton, and G Elisabeta Marai. 2019. User perspective and higher cognitive task-loads influence movement and performance in immersive training environments. *BMC Biomedical Engineering* 1, 1 (2019), 21.
- [81] Dishita G Turakhia, Andrew Wong, Yini Qi, Lotta-Gili Blumberg, Yoonji Kim, and Stefanie Mueller. 2021. Adapt2learn: A toolkit for configuring the learning algorithm for adaptive physical tools for motor-skill learning. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*. 1301–1312.
- [82] Neil Vaughan, Bodgan Gabrys, and Venketesh N Dubey. 2016. An overview of self-adaptive technologies within virtual reality training. *Computer Science Review* 22 (2016), 65–87.
- [83] Xizi Lucy Wang. 2020. *Ppvr: Teaching and learning to play the piano in virtual reality*. Ph. D. Dissertation.
- [84] Matthias Weing, Amrei Röhligh, Katja Rogers, Jan Gugenheimer, Florian Schaub, Bastian Könings, Enrico Rukzio, and Michael Weber. 2013. P.I.A.N.O.: enhancing instrument learning via interactive projected augmentation. In *Proceedings of the 2013 ACM conference on Pervasive and ubiquitous computing adjunct publication*. 75–78.
- [85] Katie Wilson and Phillip E Pfeiffer. 2023. Feedback in augmented and virtual reality piano tutoring systems: a mini review. *Frontiers in Virtual Reality* 4 (2023), 1207397.
- [86] Carolee Joyce Winstein. 1988. *Relative frequency of information feedback in motor performance and learning*. University of California, Los Angeles.
- [87] Carolee J Winstein, Patricia S Pohl, and Rebecca Lewthwaite. 1994. Effects of physical guidance and knowledge of results on motor learning: support for the guidance hypothesis. *Research quarterly for exercise and sport* 65, 4 (1994), 316–323.
- [88] Carolee J Winstein and Richard A Schmidt. 1990. Reduced frequency of knowledge of results enhances motor skill learning. *Journal of Experimental Psychology: Learning, memory, and cognition* 16, 4 (1990), 677.
- [89] C. J. Winstein and R. A. Schmidt. 1990. Reduced Frequency of Knowledge of Results Enhances Motor Skill Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16, 4 (1990), 677–691. doi:10.1037/0278-7393.16.4.677
- [90] Kerstin Witte, Dan Bürger, and Stefan Pastel. 2025. Sports training in virtual reality with a focus on visual perception: a systematic review. *Frontiers in Sports and Active Living* 7 (2025), 1530948.
- [91] Xiaodi Wu, Qiang Zhang, Jun Qiao, Nan Chen, and Xie Wu. 2022. Calligraphy-based rehabilitation exercise for improving the upper limb function of stroke patients: protocol for an evaluator-blinded randomised controlled trial. *BMJ open* 12, 5 (2022), e052046.
- [92] Gabriele Wulf. 2013. Attentional focus and motor learning: a review of 15 years. *International Review of sport and Exercise psychology* 6, 1 (2013), 77–104.
- [93] Gabriele Wulf and Rebecca Lewthwaite. 2016. Optimizing performance through intrinsic motivation and attention for learning: The OPTIMAL theory of motor learning. *Psychonomic bulletin & review* 23, 5 (2016), 1382–1414.
- [94] Xiao Xiao, Paula Aguilera, Jonathan Williams, and Hiroshi Ishii. 2013. MirrorFugue iii: conjuring the recorded pianist. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. 2891–2892.
- [95] Xiao Xiao and Hiroshi Ishii. 2010. MirrorFugue: communicating hand gesture in remote piano collaboration. In *Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction*. 13–20.
- [96] Xiao Xiao and Hiroshi Ishii. 2011. Duet for solo piano: MirrorFugue for single user playing with recorded performances. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*. 1285–1290.
- [97] Xiao Xiao, Basheer Tome, and Hiroshi Ishii. 2014. Andante: Walking figures on the piano keyboard to visualize musical motion. In *Proceedings of the international conference on New interfaces for musical expression*. 629–632.
- [98] Qi Yang and Georg Essl. 2013. Visual associations in augmented keyboard performance. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 252–255.
- [99] Ungyeon Yang and Gerard Jounghyun Kim. 2002. Implementation and evaluation of “just follow me”: An immersive, VR-based, motion-training system. *Presence* 11, 3 (2002), 304–323.
- [100] Zhen Yang, Jinlei Shi, Wenjun Jiang, Yuexin Sui, Yimin Wu, Shu Ma, Chunyan Kang, and Hongting Li. 2019. Influences of Augmented Reality Assistance on Performance and Cognitive Loads in Different Stages of Assembly Task. *Frontiers in Psychology* 10 (2019), 1703. doi:10.3389/fpsyg.2019.01703
- [101] Beste F Yuksel, Kurt B Oleson, Lane Harrison, Evan M Peck, Daniel Afergan, Remco Chang, and Robert JK Jacob. 2016. Learn piano with BACH: An adaptive learning interface that adjusts task difficulty based on brain state. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 5372–5384.
- [102] Maryam Zahabi and Ashiq Mohammed Abdul Razak. 2020. Adaptive virtual reality-based training: a systematic literature review and framework. *Virtual Reality* 24, 4 (2020), 725–752.