Human Learning Dynamics During Teleoperation of Surgical Robots

Yinqi Huang (corresponding author)[1], Yilin Cai[2], Mengyao Li[1], Yue Chen[3], Robert C. Wilson[1]

[1]School of Psychology, Georgia Institute of Technology, Atlanta, Georgia, USA.
[2]School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, Georgia, USA.
[3]Department of Biomedical Engineering, Georgia Institute of Technology/Emory University, Atlanta, Georgia, USA.

Contributing authors: yinqi.huang@gatech.edu; yilincai@gatech.edu; mengyao.li@gatech.edu; yue.chen@bme.gatech.edu; rwilson337@gatech.edu;

Abstract

Robot-assisted minimally invasive surgery requires operators to learn precise motor control. In this paper, we examined how task complexity and the human–robot motion scaling ratio affect teleoperation learning in a ring-through-rail task implemented on da Vinci Research Kit. Human participants completed training, retention, and transfer phases across different rail shapes and motion scaling ratios. In parallel, a linear–quadratic–Gaussian (LQG) controller performed the same task by combining state estimation with optimal feedback control. Humans and the model showed qualitatively similar patterns of improvement with practice and generalization across complexity. These results suggest that a simple control theory model can capture features of human teleoperation learning, providing a mechanistic model to predict how motor skill transfers and how robotic interface settings influence motor learning.

Keywords: robot-assisted minimally invasive surgery (RAMIS), teleoperation, motor learning, embodied cognition

1 Introduction

Robot-assisted minimally invasive surgery (RAMIS) has become an important component of modern clinical practice, extending the capabilities of conventional minimally invasive procedures. Systems such as the da Vinci Surgical System (Intuitive Surgical, Sunnyvale, CA) exemplify this class of technology and have demonstrated how teleoperated robotic platforms can enhance geometric accuracy, reduce operator fatigue, and support smaller incisions that can improve patient recovery and outcomes [1, 2]. RAMIS also advances conventional minimally invasive surgery by offering capabilities such as motion scaling and tremor reduction, which enable precise manipulation of delicate surgical instruments [3].

Despite these benefits, mastering RAMIS presents unique challenges for surgeons. During teleoperation, surgeons must perform complex eye-hand-foot coordination while working at a remote console, relying on video-mediated 3D endoscopic imagery instead of direct line-of-sight. The absence of tactile feedback, the altered visual-motor mappings, and the need to integrate multiple streams of sensory information all contribute to a steep teleoperation learning curve [4–6]. These features also make robotic teleoperation a complex embodied learning task, in which operators must coordinate physical action, visual feedback, spatial reasoning, and limited cognitive resources. Recent work integrating embodied cognition and cognitive load theory suggests that physical action can support learning, but the learning outcomes depend on whether task demands are structured in a way that keeps cognitive load manageable [7, 8]. As RAMIS becomes widely used, there is an increasing need for structured, evidence-based RAMIS training protocols that can accelerate skill acquisition while maintaining patient safety. To develop such protocols, it is essential to understand the specific factors that affect motor learning during teleoperation.

This study aims to understand how humans learn to teleoperate surgical robots by disentangling two types of factors that affect the learning curve: (i) task complexity, and (ii) robot-related sensorimotor adaptation.

First, the complexity of the training task directly influences learning. Tasks that require more precise motor control, involve complex bimanual coordination, or impose higher cognitive workload are more difficult and are expected to produce slower learning, whereas simpler tasks may allow for rapid improvement. Beyond immediate performance gains, task complexity is also widely believed to affect how well acquired skills are retained over time and transferred to novel tasks. Although simulation-based and proficiency-based robotic surgery curricula are increasingly used, how task complexity influences skill retention and transfer in the context of RAMIS training have not been systematically analyzed [9, 10].

Second, skill acquisition depends on the surgeon's ability to adapt to the robotic interface. Successful teleoperation requires surgeons to construct an internal model of the teleoperated system, adapt to altered visual-motor mappings and compensate for the absence of tactile cues. Surgeons must understand how the robot responds to their control inputs and how changes in the motion scaling ratio alter the sensitivity and stability of movements [11]. Early in training, unfamiliarity with the robot's interface, such as the gripper inertia and responsiveness, can lead to jerky, imprecise movements and frequent errors. With practice, surgeons update their internal model, allowing them to predict the outcome of their motor commands more accurately and to stabilize the end-effector along the desired trajectory. This process of sensorimotor adaptation is independent from task complexity per se, yet the two together affect overall teleoperation performance.

The goal of our study is to disentangle the contributions of task complexity and motion scaling to human learning during surgical teleoperation. Drawing on cognitive science and control theory, we model the human teleoperator as a linear–quadratic–Gaussian (LQG) controller with an internal estimate of the motion scaling ratio that is updated over training. This modeling choice is motivated by Bayesian and optimal-control accounts of sensorimotor behavior, which emphasize that effective action requires representing uncertainty arising from sensory and motor noise, as well as ambiguity about the environment. In these frameworks, the nervous system constructs internal models of sensorimotor transformations and uses prediction errors, weighted by their reliability, to guide inference and learning [12–14]. Consistent with this view, research in sensorimotor control has shown that many features of human movement, including goal-directed reaching and flexible responses to perturbations can be explained by optimal feedback control, in which participants tried to balance minimizing task error and control effort under uncertainty [13, 15–18]. Related work has also proposed computational accounts of human learning using Bayesian models [19] and Markov decision processes [20]. Building on these complementary perspectives, we treat surgical teleoperation learning as adaptive control under uncertainty: operators must infer the human–robot motion scaling while generating feedback commands that trade off accuracy and effort. This motivates our use of an LQG controller with an internal mapping estimate that is updated during training. Compared with simpler controllers such as PD or first-order models, LQG explicitly separates state estimation from feedback control. This distinction is important for the proposed task because participants performed continuous teleoperation under noisy sensory observations and altered visuomotor mappings. Although the controlled plant can be represented using second-order position-velocity dynamics, the LQG controller goes beyond a simple second-order model by specifying how noisy observations are combined with internal predictions to estimate the current state and generate feedback control. In this model, the Kalman gain provides an interpretable parameter governing how internal state estimates are updated across training, retention, and transfer. Simpler controllers can reproduce position tracking when tuned appropriately, but they do not directly model sensory uncertainty or belief updating. Reinforcement learning approaches are powerful for reward-driven sequential decision-making when system dynamics or cost functions are unknown [21]; however, our focus is on continuous closed-loop sensorimotor control under sensory noise, where the system dynamics are defined. LQG therefore provides a mathematically tractable and interpretable framework for modeling adaptive teleoperation behavior in this study.

In our study, both humans and the LQG model perform the same set of ring-through-rail tasks, under the same task complexity and mapping-ratio conditions. Their behaviors were then compared qualitatively. By having humans and the LQG controller perform the same task, we aim to (i) characterize how task complexity and motion scaling ratio affects learning, retention, and transfer, and (ii) assess to what extent a simple LQG model can reproduce the patterns observed in human teleoperation. Ultimately, this work is intended to inform the design of surgical training protocols and robotic interfaces that better support efficient skill acquisition in RAMIS. More broadly, the study motivates viewing teleoperation as a human-machine interaction problem, in which operator skill and robotic interface design jointly affect how limited perceptual, motor, and cognitive resources are allocated during learning [22].

Task complexity during training would non-linearly influence motor learning, which we assess using retention and transfer tests. Retention test evaluates the persistence of performance after a delay when participants are re-tested under the trained condition. Transfer test evaluates the ability to generalize learned skills to a novel motion scaling ratio different from that used during training [23]. Therefore, we hypothesized an inverted U-shaped relationship between task complexity and motor learning outcomes, consistent with the Challenge Point Hypothesis [24–26]. This prediction also aligns with evidence that learning is optimized at intermediate levels of challenge, when task demands produce an intermediate error rate rather than being either too easy or too demanding [27]. Accordingly, we predicted that intermediate levels of task complexity would yield optimal performance in both retention and transfer tasks, whereas very low or very high complexity would be less effective. We designed three rails that differed in geometric complexity, thereby increasing task complexity along multiple dimensions simultaneously: (1) a circular rail constrained to a single level (mean curvature $\bar{\kappa} = 0.031$, $\kappa_{\max} = 0.061$), (2) a rail with two primary turns and two elevation levels ($\bar{\kappa} = 0.037$, $\kappa_{\max} = 0.211$), and (3) a rail with three turns and three elevation levels ($\bar{\kappa} = 0.060$, $\kappa_{\max} = 0.409$). Here $\bar{\kappa}$ is the mean curvature of the rail and $\kappa_{\max}$ is the maximum curvature. Higher curvature requires more precise and continuous end-effector corrections, elevation changes introduce an additional spatial degree of freedom that must be actively monitored, and additional turns increase the number of rotational transitions that must be anticipated and executed.

Beyond task performance, we also hypothesized that subjective cognitive workload would decrease with learning. From a cognitive-load perspective, more complex rails increase task difficulty because participants must simultaneously coordinate trajectory tracking, orientation control, collision avoidance, depth perception, and adaptation to the teleoperation mapping, and these components cannot be managed independently [28]. As geometric complexity increases, the cognitive load imposed by the task places greater demands on working memory and attentional resources. For novice teleoperators, these component processes have not yet been automatized and therefore compete for limited cognitive resources during performance [8]. As skill develops and these component processes become more automatic, the same objective task complexity may impose lower subjective difficulty, thereby freeing cognitive resources for higher-order control refinement. We therefore also predicted that subjective cognitive workload would decrease with learning as participants progressively developed more efficient sensorimotor control strategies and internal representations of the teleoperation system. Perceived workload was assessed using NASA-TLX [29].

We further examined whether the motion scaling ratio would produce asymmetric transfer of learning. The motion scaling ratio is an adjustable parameter in the robotic interface, defined as the ratio of operator hand motion to robot end-effector motion. For example, a ratio of 0.2 maps a 1-unit human input to a 0.2-unit robotic movement, which improves precision and stability due to downscaled operator motion. Conversely, larger ratios (e.g., 0.6) amplify operator motion and required finer human control to achieve the same level of robotic precision. Thus, scaling ratio is a potential factor that may influence how learning generalizes across mappings. Specifically, we hypothesized that training under a higher ratio (0.6) would transfer more effectively to a moderate (0.4) than training under a lower ratio (0.2) would transfer to 0.4. This asymmetric transfer was hypothesized because higher motion scaling ratios might impose greater control, and therefore encourage a more robust internal model that transfers to less demanding conditions. This hypothesis was consistent with findings that asymmetric transfer occurred when training under highly complex conditions exposed individuals to the full range of movement required for harder tasks [30]. However, it is also plausible that participants rapidly adapt for scaling differences through feedback, such that scaling primarily affects movement efficiency (e.g., path length and duration) rather than aiming-related outcomes (e.g., drops and collisions). Thus, the present study evaluates whether scaling ratio meaningfully modulates transfer performance, and if so, which behavioral metrics are most sensitive to this effect.

2 Results

2.1 Human studies of performance improves over training.

Across training, robust improvements were observed in translational path length, trial duration, number of drops, and the number of handovers, whereas rotational path length remained largely unchanged. Below, we report paired t-test between Block 1 and the Retention block, followed by linear mixed-effects (LME) models assessing block-wise learning effects from Block 1 to Retention.

We observed a reliable reduction in gripper's translational path length from Block 1 to Retention. The paired t-test comparison between Block 1 and Retention reveals a statistically significant decrease in path length, paired t-test: $t(23) = 2.840$, $p = 0.009$; Cohen's $d = 0.580$; mean percentage change (Retention vs. Block 1) $= -18.8\%$. The corresponding LME with Block as a fixed effect and random intercepts for participants revealed a significant main effect of Block on path length, $F(4, 115) = 5.162$, $p = 0.001$, indicating a reduction in translational distance across training (Table 1; Fig. 2).

In contrast, rotational path length did not show a Block 1-to-Retention reduction at the group level (Fig. 2b). The paired t-test comparison between Block 1 and Retention is non-significant, paired t-test: $t(23) = 1.020$, $p = 0.320$; Cohen's $d = 0.210$; mean percentage change (Retention vs. Block 1) $= 10.2\%$. Consistently, the LME with Block (1–5) as a fixed effect did not show a reliable Block effect, $F(4, 115) = 0.706$, $p = 0.590$.

Here we observe a dissociation between translational and rotational control (Fig. 3). Translational path length decreased over training, reflecting more direct and efficient end-effector movement, whereas rotational path length did not show a consistent reduction. This pattern is consistent with the hypothesis that participants continued to purposefully explore the arm's joint redundancy during learning [16, 31, 32]. One possible interpretation is that participants reduced inefficient translational corrections to lower overall task demands, thereby freeing attentional and working-memory resources for exploration of joint-level coordination strategies. Recent work from cognitive load theory suggest that motor control and cognitive processing rely on integrated and limited attentional resources rather than fully independent systems [8]. Under this view, improving translational efficiency may allow learners to allocate more cognitive and motor resources toward exploring alternative rotational and joint configurations that support motor skill acquisition. Together, these results suggest that participants learned to stabilize end-effector translational movement while continuing to explore alternative strategies to adapt to task complexity.

For trial duration, participants became significantly faster across training. The paired t-test between Block 1 and Retention showed a robust reduction in trial duration, paired t-test: $t(23) = 3.770$, $p = 0.001$; Cohen's $d = 0.770$; mean percentage change (Retention vs. Block 1) $= -34.7\%$. The LME analysis confirmed a significant main effect of Block on trial duration across Blocks 1–5, $F(4, 115) = 10.381$, $p < 0.001$, with progressively shorter trial times from Block 1 through Retention (Fig. 2).

The frequency of participants dropping the ring also decreased significantly across training. The paired comparison between Block 1 and Retention showed a robust reduction in the number of drops, paired t-test: $t(23) = 4.260$, $p < 0.001$; Cohen's $d = 0.870$; mean percentage change (Retention vs. Block 1) $= -50.4\%$. The LME analysis also revealed a significant Block effect, $F(4, 115) = 8.564$, $p < 0.001$, indicating significant reductions in drop frequency across training and into retention (Fig. 2).

Similarly, the frequency of collisions between the ring and the rail decreased across training. The paired comparison between Block 1 and Retention showed a reduction in the number of drops, paired t-test: $t(23) = 2.740$, $p = 0.012$; Cohen's $d = 0.560$; mean percentage change (Retention vs. Block 1) $= -14.8\%$. The LME analysis also revealed a significant Block effect, $F(4, 115) = 6.036$, $p < 0.001$, indicating significant reductions in collision frequency across training and into retention (Fig. 2).

Lastly, the collision duration also decreased significantly across training. The paired comparison between Block 1 and Retention showed a robust reduction in collisions duration, paired t-test: $t(23) = 3.530$, $p = 0.002$; Cohen's $d = 0.720$; mean percentage change (Retention vs. Block 1) $= -47.6\%$. The LME analysis revealed a significant Block effect, $F(4, 115) = 10.271$, $p < 0.001$, indicating significant reductions in collisions duration across training and into retention (Fig. 2).

Taken together, these analyses show that participants learned to perform the task more efficiently and with fewer errors over training, and that these performance were retained at the delayed retention test. Translational path length, trial duration, collisions duration, number of collisions, and number of drops all showed significant improvements from Block 1 to Retention, with medium-to-large effect sizes and substantial percentage reductions. In contrast, rotational path length did not show a reliable decrease, suggesting that participants stabilized end-effector transport while continuing to explores rotational strategies.

To obtain a measurement of the overall task performance, we applied principal component analysis (PCA) to the metrics defined above: translational path length (x_1), rotational path length (x_2), trial duration (x_3), collisions duration (x_4), number of collisions (x_5), and number of drops (x_6).

The PCA-derived performance score increased across the four training blocks and retention block (Fig. 4). A linear mixed-effects model (LME) with random intercepts for subjects revealed a significant main effect of Block,

$$F(4, 115) = 16.53, \quad p < 0.001,$$

indicating reliable improvement across the one-hour training session.

Subjective workload, assessed by NASA-TLX, decreased consistently across blocks (Fig. 4c). An LME with random subject intercepts confirmed a significant main effect of Block,

$$F(3, 92) = 6.20, \quad p < 0.001,$$

indicating reduced perceived workload across training.

To assess the relationship between participants between subjective workload and performance, we calculated Pearson correlations between NASA-TLX scores and performance scores in training blocks for each subject. One participant was excluded due to no variance in NASA-TLX responses. Across the remaining participants ($N = 23$), the mean within-subject correlation was

$$\bar{r} = -0.504, \quad \mathrm{SD} = 0.447.$$

A one-sample t-test against zero showed that the negative correlation was significant,

$$t(22) = -5.405, \quad p < 0.001,$$

indicating that lower subjective workload was reliably associated with higher performance score. Therefore, the human results support Hypothesis 2.

This negative relationship between subjective workload and performance is consistent with cognitive-load accounts of complex skill learning, in which practice reduces task-irrelevant processing demands and allows greater cognitive resources to be allocated toward task-relevant control processes [33]. In the context of teleoperation learning, these results suggest that participants progressively developed more stable internal sensorimotor representations of the robotic interface, thereby reducing cognitive effort while improving movement efficiency and overall task performance.

2.2 LQG studies of performance improves over training.

To compare participants' performance with the LQG controller, we provided the controller with the centerline of each rail and let it generate simulated gripper trajectories. Note that this differs from the visual information available to human participants, who relied on stereo endoscopic feedback rather than direct access to the rail centerline. Thus, the LQG controller should be interpreted as a baseline model for rail-tracking behavior, rather than as a full perceptual model of teleoperation. As shown in Fig. 1, the controller qualitatively reproduced several trajectory features observed in human performance.

For each condition, the LQG controller was initialized with the same incorrect estimate of the motion scaling ratio gain g. It then completed 12 training trials. Every three trials were averaged into one block, resulting into four blocks in total. We compared translational path length, rotational path length, number of drops, and trial duration based on the simulated trajectory. After the controller had adapted to either $g = 0.6$ or $g = 0.2$, it was tested on a transfer trial with $g = 0.4$ to mirror the human experiment. As shown in Fig. 5, the controller was able to improve over training similar to humans.

2.3 Human studies of retention and transfer across complexity levels and ratios

To examine how task complexity and motion scaling ratio influenced the generalization of learning, we modeled retention and transfer PCA composite scores using linear mixed-effects (LME) models. These analyses evaluated whether post-training performance improvements depended more strongly on complexity imposed by rail geometry or on the trained motion scaling ratio. Because the number of participants per condition was modest ($n = 4$), we also reported effect sizes and confidence intervals in addition to statistical significance.

As shown in Fig. 6, performance normalized by pre-training performance differed across complexity levels in both retention and transfer, suggesting that rail geometry imposes persistent constraints on performance. Similar complexity-dependent performance patterns were also observed in absolute performance metrics without normalization (Fig. 7).

For absolute transfer performance, all human participants were tested at the intermediate motion scaling ratio of 1:0.4 (Fig. 7). The LME revealed a significant main effect of Complexity Level,

$$F(2, 18) = 3.858, \quad p = 0.040, \quad \eta^2_{\text{partial}} = 0.300, \quad f = 0.655,$$

indicating a large effect size despite the modest sample size.

Mean transfer improvements averaged across Scaling Ratio are summarized in Table 2. Transfer was strongest at the lowest complexity level, with Level 1 showing the largest improvement ($M = 9.24$), followed by Level 2 ($M = 7.01$) and Level 3 ($M = 5.75$). The wider confidence intervals for Levels 1 and 3 suggest greater participant-wise variance.

Similarly for retention, we also observed a significant main effect of Complexity Level,

$$F(2, 18) = 5.331, \quad p = 0.015, \quad \eta^2_{\text{partial}} = 0.372, \quad f = 0.770.$$

which indicates a large effect size.

Retention improvements averaged across Scaling Ratio are summarized in Table 3. Retention improvement increased with task complexity, with the largest improvement observed at Level 3 ($M = 11.74$), followed by Level 2 ($M = 5.74$) and Level 1 ($M = 3.32$). However, the wider confidence interval at the highest difficulty level suggests higher participant-wise variance.

In contrast to transfer, the largest retention improvement was observed at the highest complexity level. This indicates that more complex training conditions produced stronger within-condition learning gains, likely because the increased geometric demands provided greater opportunity for improvement across practice. However, these gains did not translate into the strongest transfer performance, suggesting a dissociation between retaining performance improvements on the trained condition and generalizing those improvements to a new task condition.

Together, these results do not support Hypothesis 1 as originally stated. Specifically, we did not observe an inverted-U relationship between task complexity and post-training performance. Instead, transfer performance showed a generally decreasing pattern with increasing complexity, whereas retention improvement showed the opposite pattern in human participants. This dissociation suggests that retention and transfer were influenced differently by geometric task complexity. Higher-complexity rails may have provided greater opportunity for within-condition improvement, leading to stronger retention, but the same complexity may have reduced generalization to the transfer condition. However, given the limited sample size, this result should be interpreted cautiously.

Task complexity explained a considerable proportion of the observed variance in transfer and retention, with partial effect sizes of $\eta^2_{\text{partial}} = 0.372$ and 0.300, respectively. In contrast, the observed effect of motion scaling ratio was small, and we did not observe a clear asymmetric transfer pattern in human. Thus, the human results do not support Hypothesis 3 as originally stated. Specifically, transfer performance did not show a strong directional dependence on the trained motion scaling ratio. This suggests participants were able to partially compensate for moderate scaling changes through visual feedback and sensorimotor adaptation.

Combined with the LQG simulations shown in Fig. 8, these findings suggest that motion scaling ratio primarily affected movement efficiency, such as path length and trial duration, rather than task-success measures such as drops and collisions. In contrast, geometric task complexity imposed by rail geometry was the dominant factor influencing how learning outcome was retained and transferred across conditions. However, given the limited sample size, these findings should be interpreted cautiously.

2.4 LQG studies of retention and transfer across complexity levels and ratios

To investigate how task complexity and motion scaling affect teleoperation, we simulated transfer performance using an LQG controller with an internal mapping estimate g_{est} that differed from the true mapping g_{true}. We considered two mismatch scenarios, underestimation ($g_{\text{est}} < g_{\text{true}}$) and overestimation ($g_{\text{est}} > g_{\text{true}}$), and asked whether these conditions produce reliable differences in transfer when the number of participants was matched in human and simulation. We also examined the evolution of the LQG feedback gain K, the control command u, and the estimated mapping gain g_{est} over time (see Supplementary 1.2).

As shown in Fig. 8, the controller exhibited performance similar to that observed in humans (Fig. 7) with mismatch gain. We analyzed simulated transfer performance using LME (Table 4). In the LQG simulations, transfer revealed significant differences in both task complexity and motion scaling ratio. Specifically, translational and rotational path lengths showed significant main effects of Complexity and Ratio, and rotational path length additionally showed a significant Complexity × Ratio interaction, indicating that the effect of motion scaling depended on rail geometry. In contrast, contact-related metrics showed a different pattern: collisions and collision duration both showed significant main effects of Complexity, whereas their Ratio effects and interactions were not significant, suggesting that contact dynamics during transfer were largely insensitive to the mapping ratio in this LQG setting. The number of drops that were not analyzed transferred trials contained many zeros (i.e. the LQG controller rarely dropped after training) and the duration of the trials were not analyzed because it is effectively invariant by the design of LQG.

Together, these results suggest that rail geometry (Complexity Level) affects transfer performance by increasing the frequency and magnitude of corrective actions required to remain aligned with the rail. While the controller can adapt its internal mapping to compensate for motion scaling mismatch, ratio effects remain detectable in efficiency, particularly in rotational path length, where the significant interaction suggests geometry-dependent amplification of reorientation demands. In contrast, contact-related outcomes appear to be driven primarily by task complexity, consistent with the interpretation that tighter curvature and elevation increase control effort and amplify the consequences of estimation noise and execution variability. Overall, the transfer simulations support the hypothesis that task complexity affect performance through geometric constraints, whereas motion scaling ratio effects are observed most clearly in efficiency.

The asymmetric convergence mechanism derived in Supplementary 1.2 provides a possible explanation for why motion scaling ratio effects were more detectable in the LQG simulations than in the human behavioral results. In the model, underestimation and overestimation of the motion scaling ratio produce different adaptation rates because the gain adaptation update scales with the squared control input u_k^2. Underestimation generates larger control commands and therefore faster adaptation, whereas overestimation produces smaller commands and slower convergence. This asymmetry produced measurable differences in trajectory efficiency metrics in the simulations, particularly rotational path length. In contrast, the human participants did not exhibit a strong asymmetric transfer pattern, suggesting that humans may compensate for scaling mismatches using additional mechanisms not represented in the simplified LQG controller, such as rapid sensorimotor adaptation. Thus, the simulations suggest that asymmetric convergence may exist as a latent control-theoretic mechanism, but its behavioral expression in humans may be partially masked by higher-level compensatory processes.

3 Discussion

By having humans and an LQG controller perform the same ring-through-rail task under matched experimental conditions, this study examined how task complexity and motion scaling ratio influence learning, retention, and transfer during surgical robot teleoperation. Although the controller is a simplified, idealized learner, humans and the model converged on several similar qualitative patterns, suggesting that important components of teleoperation learning can be captured by a control-theoretic model in which operators update an internal estimate of the human-robot mapping while generating feedback commands that trade off accuracy and effort.

First, both humans and the LQG controller showed learning across training, expressed as more efficient movement execution. In humans, performance improved with training, reflected by shorter translational path lengths, reduced trial durations, and fewer drops. The model showed a parallel pattern: as its internal mapping estimate approached the true motion scaling ratio, trajectories became more direct and required fewer corrective actions and fewer drops. Together, these results support the interpretation that teleoperation skill acquisition can be captured by prediction error-driven updates of control policies.

Second, both humans and the model were sensitive to task complexity. In both retention and transfer assessments, the effects related to complexity were more robust than the effects of the scaling-ratio, which also suggests that the task complexity was driven primarily by the constraints imposed by the geometry of the rails. This pattern is consistent with the idea that learning is optimized when task complexity generate informative errors without overwhelming the learner. From a control perspective, higher complexity increases the frequency and magnitude of corrective actions required to maintain contact-free tracking, and therefore amplifies the behavioral consequences of sensorimotor noise and imperfect internal models. These results highlight that training outcomes depend not only on robotic interface parameters (e.g., motion scaling), but also on how task structure affects the control problem faced by the operator.

Third, comparing humans with the LQG controller clarifies what a simple cognitive-level control model can explain and where it remains incomplete. The LQG controller provides an interpretable mechanistic hypothesis for how learning and generalization depend on both task constraints and motion scaling. At the same time, the model does not incorporate several cognitive factors known to affect real-world skill acquisition, including strategic exploration, fatigue, attention [34]. The model should therefore be viewed as a baseline model of optimal feedback control rather than a complete description of human learning in surgical settings. Humans and the LQG controller also differed in the information available for control. Human participants relied on stereo endoscopic visual feedback, whereas the LQG controller was given the rail centerline as an explicit reference trajectory. Therefore, the model should be interpreted as an idealized control-theoretic baseline, not a full perceptual model of teleoperation. Future work could incorporate noisy visual rail estimates, depth uncertainty, occlusion, and feedback delay.

Several limitations of the present work motivate future directions. First, our participant population consisted of general university students rather than surgical trainees or surgeons. This choice is appropriate for probing general motor learning processes, such as error correction, adaptation to altered motion scaling, and the emergence of stable control policies. However, surgical training also involves domain-specific expertise, including procedural knowledge, anticipatory control, and bimanual strategies, which may alter learning dynamics and skill generalization. Prior work has shown substantial differences between novice and expert performance in robotic surgery tasks, including movement efficiency, coordination strategies, and workload management [35].

In addition, although we used NASA-TLX to assess subjective workload, future studies involving surgical trainees or surgeons may benefit from domain-specific workload measures such as the Surgery Task Load Index (SURG-TLX), which was developed specifically for surgical performance assessment [36]. Future work should therefore test whether the task complexity-dependent learning patterns and mapping-related transfer effects generalize to clinically relevant populations and whether they predict training outcomes relevant to surgical competency.

Second, the task itself focused on unimanual control. Although implemented on a da Vinci system, participants primarily used a single gripper to guide the ring along a rail. In contrast, real robotic surgery requires complex bimanual coordination, including role differentiation between hands (e.g., traction with the non-dominant hand while manipulating or suturing with the dominant hand), continuous reallocation of effort, and interaction with deformable tissue [37]. A growing body of work shows that expert surgeons differ from novices not only in single-hand smoothness, but also in bimanual efficiency and synergy, and that meaningful skill assessment often requires metrics that capture inter-hand coordination rather than treating the two hands independently [38, 39]. Consistent with this view, recent work on bimanual motor learning suggests that, compared with unimanual task, naturalistic bimanual tasks may rely more strongly on online feedback control and inter-hand coordination [40]. Thus, future studies should therefore extend this paradigm to bimanual tasks that more directly approximate surgical workflows (e.g., suturing, knot tying, traction-and-dissection sequences) and quantify bimanual learning using measures that capture inter-hand coordination and closed-loop control.

Third, our LQG model is a cognitive model that intentionally simplifies human kinematics. While this simplification is useful for isolating how internal mapping estimates and optimal feedback control can generate qualitative learning patterns, it does not represent biomechanical constraints or capture aspects of joint-level coordination, such as joint redundancy resolution and fine-grained hand and arm dynamics that underlie real surgical manipulation. A future direction is to integrate control-theoretic learning with richer kinematic models of the human upper limb and hand (such as those proposed in [41, 42]), including both high-dimensional representations (e.g., 27-DOF hand models that capture finger articulations) and simplified but tractable models (e.g., 7-DOF upper-extremity models that capture joint dynamics). Additionally, the current model does not incorporate cognitive-load mechanisms, even though teleoperation performance likely depends not only on optimal feedback control but also on working-memory demands, attentional allocation, and movement-related cognitive processing [8, 33]. Incorporating both kinematic and cognitive-load components could support more realistic predictions of movement efficiency, coordination strategies, bimanual role allocation, and workload-dependent adaptation, and would naturally support comparisons to emerging approaches in human-robot collaborative bimanual manipulation [43] and autonomous robotic bimanual manipulation [44]. Importantly, adopting such models would also require richer data collection: beyond endpoint trajectories and task-level outcomes (drops and collisions), future experiments would need richer kinematic measurements (joint angles, joint velocities, end-effector orientation dynamics, and interlimb coupling metrics) to constrain model parameters and evaluate coordination-level predictions.

Despite these limitations, this study provides evidence that a simple LQG model can serve as a useful baseline for human teleoperation learning. By linking behavioral improvements to an interpretable internal-model framework, this approach provides descriptive performance metrics and motivates hypotheses about how task complexity and robotic interface parameters affect skill acquisition in RAMIS.

4 Methods

4.1 Teleoperation Task

The teleoperation task was performed using a da Vinci Research Kit system (dVRK; Intuitive Surgical, Inc., Sunnyvale, CA) [45]. Participants interacted with the system at the surgeon console using joystick handles to control the movements of patient-side robotic arms in real time. Visual feedback was provided via a stereo endoscopic camera, delivering 3D depth perception and visual guidance similar to that of clinical robotic surgery (Fig. 9).

The behavioral task was a ring-through-rail teleoperation task. The ring-through-rail task is a standard, validated exercise in robot-assisted surgery included in the Fundamentals of Robotic Surgery (FRS) curriculum and has been widely used in surgical skills assessment [46]. During the task, participants controlled the robotic grippers to grasp a ring and guide it along a fixed 3D rail while minimizing contact with the rail. This task serves as an abstraction of surgical scenarios, rather than a realistic simulation of a specific surgical procedure. This abstraction allowed us to systematically manipulate geometric features, including curvature, elevation change, and spatial precision, while maintaining experimental control and reproducibility [47]. In this study, geometric manipulations were designed to increase task complexity along multiple dimensions simultaneously. Higher curvature requires more precise and continuous end-effector corrections, elevation changes introduce an additional spatial degree of freedom that must be actively controlled, and additional turns increase the number of rotational transitions that must be anticipated and executed.

To manipulate task complexity, we designed three rails that varied in geometric complexity:

- **Level 1 (circular rail):** A circular rail constrained to a single elevation, with mean curvature $\bar{\kappa} = 0.031$ and maximum curvature $\kappa_{\max} = 0.061$.
- **Level 2 (two-turn, two-level rail):** A rail with two primary turns and two elevation levels ($\bar{\kappa} = 0.037$, $\kappa_{\max} = 0.211$).
- **Level 3 (three-turn, three-level rail):** A rail with three turns and three elevation levels, which has the highest curvature and complexity ($\bar{\kappa} = 0.060$, $\kappa_{\max} = 0.409$).

Here, κ denotes local geometric curvature estimated from discrete 3D spline samples along the rail.

The dVRK is designed to scale down human movement to achieve precise and stable control of the patient-side arms. As such, 1:1 motion scaling (equal magnitudes between hand and robot motion) is not typically used in clinical practice, because it does not leverage the benefits of robotic mediation. Therefore, three motion scaling ratios of 1:0.2, 1:0.4, and 1:0.6 (human movement : robot movement) were selected:

- **Low ratio (1:0.2):** 1 unit of human input produces a 0.2-unit robotic movement, high precision and stability. Used in training.
- **Intermediate ratio (1:0.4):** 1 unit of human input produces a 0.4-unit robotic movement. Used in pre-training and post-training tests.
- **High ratio (1:0.6):** 1 unit of human input produces a 0.6-unit robotic movement, requires finer motor control to achieve the same precision. Used in training.

These ratios allowed us to systematically test the asymmetric transfer hypothesis, with training at 1:0.2 or 1:0.6 and transfer tested at 1:0.4.

To test the hypotheses proposed above, we designed a 2×3 between-subjects factorial design with the following factors:

- **Motion Scaling Ratio (2 levels):** 1:0.2 and 1:0.6.
- **Task Complexity (3 levels):** circular rail (Level 1), rail with two turns and two levels (Level 2), and rail with three turns and three levels (Level 3).

Each participant was randomly assigned to one of the six Ratio × Complexity combinations and trained only on that combination.

4.2 Experimental Procedure

The experiment consists of three main phases: pre-training test, training, and post-training test (which including skill transfer and retention).

Upon arrival, all participants began with a **pre-training test** consisting of three trials. During this baseline assessment, participants controlled the robotic arm using a fixed motion scaling ratio of $1 : 0.4$ across all three task complexity levels: the circular rail (Level 1), the two-turn rail (Level 2), and the three-turn rail (Level 3).

Participants were then randomly assigned to one of the six **training** conditions, each defined by a fixed complexity level and a fixed motion scaling ratio (e.g., Level 1: circular rail at a ratio of $1 : 0.2$). The assigned complexity level and motion scaling ratio remained the same throughout this phase. Each participant completed 12 training trials, where each three trials forms one block. At the end of each training block, participants completed the NASA Task Load Index (NASA-TLX) survey to measure and track changes in perceived task load over the course of the training.

The **post-training test** was conducted immediately following the training phase and consisted of two parts: a transfer task and a retention task. The transfer task was designed to assess the generalization of skills to untrained conditions. This test used a motion scaling ratio of $1 : 0.4$ (identical to the pre-training test) and included all three complexity levels. Performance in this phase was compared to the pre-training test results to quantify the overall improvement attributable to training. To evaluate skill retention, participants completed three additional trials using the specific complexity level and motion scaling ratio they were trained on. Performance in these retention trials was compared to their performance in the training phase to determine how well the acquired motor skills were retained over time.

4.3 Measurements

All user data were recorded in real-time and synchronized using Robot Operating System (ROS), and collected as multi-dimensional time-series data for each trial. The kinematic data include motion information from both the participants and the robotic arm's tip, which includes 3D position (x, y, z), rotation (θ, ψ, ϕ), translational velocity, and angular velocity.

Perceived workload was assessed using the NASA Task Load Index (NASA-TLX), a multidimensional rating scale designed to measure workload across six dimensions: mental demand, physical demand, temporal demand, performance, effort, and frustration [29].

4.4 Data Analysis

All analyses and simulations were performed in MATLAB R2024b.

To obtain a data-driven measure of task performance, we applied principal component analysis (PCA) to the six behavioral metrics: translational path length (x_1), rotational path length (x_2), trial duration (x_3), collisions duration (x_4), number of collisions (x_5), and number of drops (x_6). Here, collisions refer specifically to physical contact events between the ring and the rail during task execution, not robot self-collision parameters.

To control for condition-level variability (e.g., longer rails naturally inducing longer durations), each metric was first z-scored within its respective complexity level and mapping-ratio group. Prior to PCA, metrics for which lower values indicate better performance were sign-aligned (i.e., multiplied by -1) so that larger values consistently corresponded to better performance. PCA was then performed on the standardized and sign-aligned data. In our dataset, the first principal component (PC1) accounted for 55.3% of the total variance (Fig. 10b) and was used to compute the composite performance score.

PC1 contains the following signed loadings: translational path length ($r = 0.483$), rotational path length ($r = -0.439$), trial duration ($r = 0.488$), number of drops ($r = 0.263$), number of collisions ($r = 0.242$), and collision duration ($r = 0.457$).

To construct a composite performance score, we normalized the loadings to obtain a non-negative contribution of each metric:

$$w_j = \frac{|v_{1j}|}{\sum_{k=1}^{p} |v_{1k}|},$$

where v_{1j} is the loading of metric j on PC1 and p is the number of selected metrics. As shown in Fig. 10c, the resulting weights were: translational path length ($w = 0.204$), rotational path length ($w = 0.185$), trial duration ($w = 0.206$), drops ($w = 0.111$), collisions ($w = 0.102$), and collision duration ($w = 0.193$).

Each metric was then rescaled into the $[0, 1]$ range, with the direction of improvement taken into account. For metrics where lower is better (translational path length, trial duration, drops, and collision-related metrics), we defined

$$d_{ij} = 1 - \frac{x_{ij}}{x_j^{\max}},$$

and for metrics where higher is better (rotational path length), we define

$$d_{ij} = \frac{x_{ij}}{x_j^{\max}},$$

where $x_j^{\max}$ is the maximum observed value of metric j.

Finally, the composite performance score for each trial was computed as:

$$\text{Score}_i = 100 \sum_{j=1}^{p} w_j \, d_{ij}, \qquad 0 \le \text{Score}_i \le 100.$$

This yields a normalized, PCA-derived measure in which higher values reflect better overall task performance, and the weights w_j correspond to the multivariate structure captured by PC1. For example, a composite score of 73 indicates that, averaged across all weighted metrics, the participant performed at the 73rd percentile of the observed performance range, with 0 representing the worst and 100 representing the best performance observed in this dataset.

To evaluate learning, for each of the behavioral metric and the performance score, we compared Block 1 and the Retention block using paired t-tests and computed Cohen's d and percentage change. To assess block-wise improvement over the entire training session (four training blocks and retention), we fit linear mixed-effects (LME) models with Block as a fixed effect and random intercepts for participants:

$$\text{Metric} \sim \text{Block} + (1|\text{Subject}).$$

To evaluate retention, we defined it as the relative change between performance in the Retention block and Block 1:

$$\text{Rate of Improvement} = \frac{\text{Retention} - \text{Training Block 1}}{\text{Training Block 1}}$$

To evaluate transfer, we defined it as the relative change between performance in the Transfer block and the pre-training baseline:

$$\text{Rate of Transfer} = \frac{\text{Transfer} - \text{Pre-training}}{\text{Pre-training}}$$

We then modeled these rates using LME models with Complexity and Ratio as fixed effects and random intercepts for participants:

$$\text{Rate} \sim \text{Ratio} \times \text{Complexity} + (1|\text{Subject}).$$

Similarly, for the LQG-simulated transfer metric, we fit the same LME models:

$$\text{Metric} \sim \text{Ratio} \times \text{Complexity} + (1|\text{Subject}).$$

Additionally, we conducted a Pearson's correlation to examine the relationship between the NASA-TLX and the training performance score.

4.5 Optimal Control Model of Human Teleoperation

Beyond the composite performance score as an outcome measure, we also want to understand human control by analyzing the gripper's movement throughout each trial. We begin by mathematically defining the human's internal model of human-robot interaction, followed by the dynamics of learning. Our goal is to describe how performance evolves throughout the one-hour training.

As shown in Fig. 11, we model the teleoperation behavior of the participant as an optimal linear–quadratic–Gaussian (LQG) controller acting on a rail-centric discrete-time plant. Each trial simulates the interaction between a *true* physical plant, parameterized by the actual teleoperation motion scaling gain g_{true}, and the human's *internal* model, parameterized by an internal estimate g_{est}, which may differ from g_{true}.

The system state is defined as $X = [x, y, z, \theta, \psi, \phi, \dot{x}, \dot{y}, \dot{z}, \dot{\theta}, \dot{\psi}, \dot{\phi}] \in \mathbb{R}^{12}$, where x, y, z denote the operator's hand positions, and θ, ψ, ϕ denote the rotation, followed by their corresponding first order derivatives (velocities). The control input $u_k = [a_{t,k}^\top, a_{r,k}^\top]^\top \in \mathbb{R}^6$ contains translational and rotational accelerations on velocities commanded by the operator, and y_k is the measurement.

The true system dynamics are approximated by a discrete-time linear time-invariant (LTI) model

$$X_{k+1} = A\,X_k + B\,(g_{\text{true}} u_k) + w_k, \tag{1}$$

$$y_k = C\,X_k + v_k, \tag{2}$$

where g_{true} is the true teleoperation motion scaling ratio implemented by the robot. w_k and v_k represent process disturbance and measurement noise, respectively. Both noise terms are modeled as independent, identically distributed zero-mean Gaussian random variables with covariance matrices W and V:

$$w_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, W), \qquad v_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, V) \tag{3}$$

By defining zero-order hold on control inputs and a full state measurement, the state space realization is given by

$$A = \begin{bmatrix} I_{6\times6} & dt\, I_{6\times6} \\ \mathbf{0}_{6\times6} & I_{6\times6} \end{bmatrix} \in \mathbb{R}^{12\times12},\ B = \begin{bmatrix} \mathbf{0}_{6\times6} \\ dt\, I_{6\times6} \end{bmatrix} \in \mathbb{R}^{12\times6},\ C = I_{12\times12} \tag{4}$$

The matrix A describes discrete-time integration of translational and rotational velocities into positions with time step dt. The matrix B indicates that the control inputs act as accelerations on the corresponding velocity states, scaled by the true teleoperation motion scaling ratio g_{true} over each time step. Since our constructed rail-frame state x_t is directly observable, we take $C = I$.

The human teleoperator is modeled as an infinite-horizon H_2-optimal controller [48] with an internal estimate g_{est} of the motion scaling ratio. The internal design model mirrors the plant structure, but replaces g_{true} with the internal estimate g_{est}.

Because the true system state is observed through noisy measurements, the human's internal state estimate $\hat{X}_k$ is updated via a Kalman filter. The estimator dynamics are given by

$$\hat{X}_{k+1} = A\hat{X}_k + B(g_{\text{est}}u_k) + L(y_{k+1} - C(A\hat{X}_k + B(g_{\text{est}}u_k))). \tag{5}$$

where L denotes the steady-state Kalman gain and is given by:

$$L = PC^\top(CPC^\top + V)^{-1}, \tag{6}$$

where P is the solution to the discrete-time algebraic Riccati equation (DARE):

$$P = APA^\top - APC^\top(CPC^\top + V)^{-1}CPA^\top + W. \tag{7}$$

The optimal control law is:

$$u_k = -K\big(\hat{X}_k - r_k\big), \tag{8}$$

where r_k denotes the reference state at time step k, and K is the optimal state-feedback gain given by

$$K = (B^\top SB + R)^{-1}B^\top SA, \tag{9}$$

where S is the solution to the discrete-time algebraic Riccati equation (DARE)

$$S = A^\top SA - A^\top SB(B^\top SB + R)^{-1}B^\top SA + Q. \tag{10}$$

The control objective is to track a rail-aligned reference trajectory r_k at each time step k and minimize the deviation from reference $e_k \triangleq \hat{X}_k - r_k$. In the current implementation, this reference trajectory was generated from the known rail centerline, which provides an idealized task reference for the controller rather than a camera-based perceptual estimate available to human participants. Here $r_k = [p_k^{\text{ref}\top}, v_k^{\text{ref}\top}, \psi_k^{\text{ref}\top}, \omega_k^{\text{ref}\top}]^\top \in \mathbb{R}^{12}$, where p_k^{ref} and v_k^{ref} denote the translational reference position and velocity along the rail, and ψ_k^{ref} and ω_k^{ref} denote the corresponding rotational reference and angular velocity.

The human operator is modeled as an H_2-optimal controller that minimizes an infinite-horizon quadratic cost on state tracking error and control effort.

$$J = \mathbb{E}\sum_{k=0}^{\infty}\left[\, e_k^\top Qe_k + u_k^\top Ru_k \,\right], \tag{11}$$

where $Q = Q^\top \geq 0$ and $R = R^\top > 0$ are the weighting matrices of state tracking errors and control effort, respectively.

To compare the model with human teleoperators, we simulated the LQG controller acting on the same rails with adaptable g_{est}, under the same motion scaling ratios, and with the same experimental procedure as in the human experiment. For each Complexity × Ratio condition, the controller was initialized with an incorrect motion scaling ratio estimate g_{est} and executed 12 training trials, grouped into four blocks, followed by a transfer trial with a different motion scaling ratio (e.g., 1:0.4). For each simulated trial, we extracted the same trajectory metrics as in the human data (translational path length, rotational path length, trial duration, drops, collisions, collisions duration). These simulations were compared qualitatively with human data, to see whether a simple control-theoretic model reproduces patterns of human teleoperation learning, including complexity-dependent performance costs and asymmetric transfer across mapping ratios.

To further verify that the LQG model can be adjusted through interpretable parameters rather than arbitrary trajectory fitting, we performed a sensitivity analysis of non-gain LQG parameters, including the tracking-effort ratio, observation and process uncertainty, command limits, and feedforward contribution (Supplementary 1.3). The analysis showed that the tracking-effort ratio Q/R produced the largest trajectory changes, while the command limit primarily modulated control effort, supporting the interpretation that different behavioral regimes can be generated by structured changes in accuracy, effort, uncertainty, and movement-conservatism parameters.

4.6 Model comparison

To evaluate whether the LQG controller provides advantages over simpler controller, we compared the LQG against two classical controllers: a PD controller and a first-order controller. These additional models were implemented using the same rail geometry. However, unlike the LQG controller, the simpler controllers did not include latent state estimation, internal uncertainty representations, or optimal feedback control under noisy observations.

For comparison, human trajectories were transformed into a rail-centered coordinate frame before computing deviation metrics. The left and right gripper positions measured Cartesian position signals were read, and the active hand at each time point was identified using the gripper signals. When one hand was active, the corresponding gripper position was used as the ring-center trajectory. The resulting hand trajectory was then aligned to the rail centerline using an iterative closest-point procedure. At each iteration, trajectory samples were matched to their nearest points on the spline-fitted rail centerline, and a rigid Kabsch transform was estimated and applied. This procedure used rotation and translation only, without scaling, so that subsequent distance-based metrics remained in physical units. Translational deviation was computed as the Euclidean distance between the aligned ring-center trajectory and the nearest point on the rail centerline at each normalized time point. Rotational deviation was computed using tangent-based trajectory orientation relative to the local rail tangent.

Figure 12 shows the resulting trajectory-level comparisons across the human data, the LQG controller, the PD controller, and the first-order controller. The first row shows the aligned 3D trajectories relative to the rail centerline. The second row shows translational deviation from the rail centerline across normalized movement time. The third row shows tangent-based rotatinal deviation relative to the local rail direction. The simulation parameters are summarized in Table 6.

To provide a quantitative comparison, we computed lagged cross-correlations and cumulative mismatch values for translational and rotational deviations relative to the human trajectories. Cumulative mismatch was computed as the integrated absolute difference between the model and human deviation profiles across normalized time.

To account for potential timing differences in corrective movements, we computed the maximum lagged cross-correlation between the human and model deviation profiles within a bounded window of ± 0.50, corresponding to at most half of the normalized trial duration (Table 5). Thus, the analysis allowed moderate temporal shifts in corrective behavior. We additionally computed cumulative mismatch as the integrated absolute deviation between the model and human trajectories across normalized movement time. The reported correlations therefore correspond to the maximum lagged correlation values within the bounded window, whereas the mismatch values reflect cumulative trajectory differences over the full movement trajectory.

The LQG controller showed the strongest similarity to human translational correction dynamics, with a high maximum lagged translational correlation ($r = 0.8958$), exceeding both the PD controller ($r = 0.4171$) and the first-order controller ($r = 0.6770$). This suggests that, when allowing for timing differences in corrective movements, the LQG controller captured a temporal structure of translational deviation that was more similar to human teleoperation behavior than the simpler baseline controllers. Although the PD and first-order controllers achieved lower cumulative translational mismatch, these simpler models primarily behaved as smooth geometric rail followers. In contrast, the LQG controller produced multi-phase corrective behavior, curvature-dependent deviations, transient overshoot, and changes in rotational correction near high-curvature segments of the rail.

Thus, the primary advantage of the LQG controller is not that it minimizes cumulative trajectory error under every metric, but that it provides a mechanistically richer account of teleoperation as feedback control under uncertainty. Specifically, the LQG framework incorporates state estimation, sensory noise, and adaptive correction dynamics, whereas the simpler rail-tracking baselines lack these mechanisms and therefore fail to capture several qualitative characteristics of human movement observed in the task.

4.7 Participants

This study was approved by the Institutional Review Board at the Georgia Institute of Technology (IRB number: IRB2025-410). All participants provided informed consent prior to participation. Based on a power analysis for an expected medium-to-large effect size [26], we recruited 24 participants (10 female) from the Georgia Institute of Technology student population through the SONA recruitment system. The participants included 11 undergraduate and 13 postgraduate students (Master's and PhD). Their age ranged from 18 to 30 years (mean age = 22.2 years, SD = 4.7 years) and reported normal or corrected-to-normal vision before participation.

5 Data availability

All data used for analyses are publicly available on the Open Science Framework (https://osf.io/acuyr/).

6 Code availability

Custom MATLAB code used for analyses are publicly available on Open Science Framework (https://osf.io/acuyr/).

7 Acknowledgements

The authors thank Tianyang Hu for helpful discussions.

8 Author Contributions

Yinqi Huang: Conceptualization; Methodology; Software; Formal analysis; Investigation; Data curation; Project administration; Visualization; Writing - Original Draft; Writing - Review & Editing. Yilin Cai: Conceptualization; Methodology; Software; Validation; Investigation; Resources; Data curation; Writing - Original Draft; Writing - Review & Editing. Mengyao Li: Supervision; Writing - Original Draft. Yue Chen: Resources; Writing - Review & Editing. Robert C. Wilson: Supervision; Writing - Original Draft; Writing - Review & Editing.

9 Competing Interests

The authors declare no competing interests.

10 Funding

This research received no external funding.

References

[1] DiMaio, S., Hanuschik, M., Kreaden, U.: The da vinci surgical system. Surgical robotics: systems applications and visions, 199–217 (2011)

[2] Satava, R.M.: Future directions in robotic surgery. In: Rosen, J., Hannaford, B., Satava, R.M. (eds.) Surgical Robotics: Systems Applications and Visions, pp. 3–11. Springer, Boston, MA (2011). https://doi.org/10.1007/978-1-4419-1126-1_1

[3] Lanfranco, A.R., Castellanos, A.E., Desai, J.P., Meyers, W.C.: Robotic surgery: a current perspective. Annals of surgery **239**(1), 14–21 (2004)

[4] Klein, M.I., Lio, C.H., Grant, R., Carswell, C.M., Strup, S.: A mental workload study on the 2d and 3d viewing conditions of the da vinci surgical robot. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting, vol. 53, pp. 1186–1190 (2009). SAGE Publications Sage CA: Los Angeles, CA

[5] Culjat, M.O., Bisley, J.W., King, C.-H., Wottawa, C., Fan, R.E., Dutson, E.P., Grundfest, W.S.: Tactile feedback in surgical robotics. Surgical Robotics: Systems Applications and Visions, 449–468 (2011)

[6] Gao, Y., Wang, S., Li, J., Li, A., Liu, H., Xing, Y.: Modeling and evaluation of hand–eye coordination of surgical robotic system on task performance. The International Journal of Medical Robotics and Computer Assisted Surgery **13**(4), 1829 (2017)

[7] Zou, L., Zhang, Z., Mavilidi, M., Chen, Y., Herold, F., Ouwehand, K., Paas, F.: The synergy of embodied cognition and cognitive load theory for optimized learning. Nature Human Behaviour **9**(5), 877–885 (2025) https://doi.org/10.1038/s41562-025-02152-2 . Accessed 2026-05-09

[8] Sepp, S., Howard, S.J., Tindall-Ford, S., Agostinho, S., Paas, F.: Cognitive load theory and human movement: towards an integrated model of working memory. Educational Psychology Review **31**(2), 293–317 (2019) https://doi.org/10.1007/s10648-019-09461-9 . Accessed 2026-05-09

[9] Basile, G., Gallioli, A., Diana, P., Gallagher, A., Larcher, A., Graefen, M., Harke, N., Traxer, O., Tilki, D., Van Der Poel, H., Emiliani, E., Angerri, O., Wagner, C., Montorsi, F., Wiklund, P., Somani, B., Buffi, N., Mottrie, A., Liatsikos, E., Breda, A.: Current standards for training in robot-assisted surgery and endourology: a systematic review. European Urology **86**(2), 130–145 (2024) https://doi.org/10.1016/j.eururo.2024.04.008 . Accessed 2026-05-09

[10] Puliatti, S., Mazzone, E., Dell'Oglio, P.: Training in robot-assisted surgery. Current Opinion in Urology **30**(1), 65–72 (2020) https://doi.org/10.1097/MOU.0000000000000687 . Accessed 2026-05-09

[11] Prasad, S.M., Prasad, S.M., Maniar, H.S., Chu, C., Schuessler, R.B., Damiano, R.J.: Surgical robotics: impact of motion scaling on task performance. Journal of the American College of Surgeons **199**(6), 863–868 (2004) https://doi.org/10.1016/j.jamcollsurg.2004.08.027

[12] Orbán, G., Wolpert, D.M.: Representations of uncertainty in sensorimotor control. Current Opinion in Neurobiology **21**(4), 629–635 (2011) https://doi.org/10.1016/j.conb.2011.05.026

[13] Todorov, E.: Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. Neural Computation **17**(5), 1084–1108 (2005) https://doi.org/10.1162/0899766053491887

[14] Säfström, D.: The multisensory control of sequential actions. Experimental Brain Research **243**(1), 13 (2024) https://doi.org/10.1007/s00221-024-06962-0

[15] Flash, T., Hogan, N.: The coordination of arm movements: an experimentally confirmed mathematical model. The Journal of Neuroscience **5**(7), 1688–1703 (1985) https://doi.org/10.1523/JNEUROSCI.05-07-01688.1985 . Accessed 2026-01-23

[16] Jagacinski, R.J., Flach, J.M.: Control Theory for Humans: Quantitative Approaches to Modeling Performance, transferred to digital print edn. Erlbaum, Mahwah, NJ (2009)

[17] Todorov, E., Jordan, M.I.: Optimal feedback control as a theory of motor coordination. Nature Neuroscience **5**(11), 1226–1235 (2002) https://doi.org/10.1038/nn963 . Accessed 2026-01-21

[18] Franklin, D., Wolpert, D.: Computational mechanisms of sensorimotor control. Neuron **72**(3), 425–442 (2011) https://doi.org/10.1016/j.neuron.2011.10.006 . Accessed 2025-02-08

[19] Ullman, T.D., Tenenbaum, J.B.: Bayesian models of conceptual development: learning as building models of the world. Annual Review of Developmental Psychology **2**(1), 533–558 (2020) https://doi.org/10.1146/annurev-devpsych-121318-084833 . Accessed 2026-01-21

[20] Tian, R., Tomizuka, M., Dragan, A., Bajcsy, A.: Towards modeling and influencing the dynamics of human learning. arXiv. arXiv:2301.00901 (2023). https://doi.org/10.48550/arXiv.2301.00901 . http://arxiv.org/abs/2301.00901 Accessed 2025-12-03

[21] Bertsekas, D.P., Bertsekas, D.P.: Approximate Dynamic Programming, Fourth edition, updated printing edn. Dynamic programming and optimal control / Dimitri P. Bertsekas, Massachusetts Institute of Technology, vol. volume 2. Athena Scientific, Belmont, Massachusetts (2018)

[22] Wang, P., Liu, H., Zou, L., Paas, F.: Overloaded minds and machines: a cognitive load framework for human-AI symbiosis. Artificial Intelligence Review **59**(3), 109 (2026) https://doi.org/10.1007/s10462-026-11510-z . Accessed 2026-05-09

[23] Adams, J.A.: Historical review and appraisal of research on the learning, retention, and transfer of human motor skills. Psychological Bulletin **101**(1), 41–74 (1987) https://doi.org/10.1037/0033-2909.101.1.41

[24] Guadagnoli, M.A., Lee, T.D.: Challenge point: a framework for conceptualizing the effects of various practice conditions in motor learning. Journal of Motor Behavior **36**(2), 212–224 (2004) https://doi.org/10.3200/JMBR.36.2.212-224

[25] Guadagnoli, M., Morin, M.-P., Dubrowski, A.: The application of the challenge point framework in medical education. Medical Education **46**(5), 447–453 (2012) https://doi.org/10.1111/j.1365-2923.2011.04210.x

[26] Akizuki, K., Ohashi, Y.: Measurement of functional task difficulty during motor learning: What level of difficulty corresponds to the optimal challenge point? Human Movement Science **43**, 107–117 (2015) https://doi.org/10.1016/j.humov.2015.07.007 . Accessed 2025-03-04

[27] Wilson, R.C., Shenhav, A., Straccia, M., Cohen, J.D.: The eighty five percent rule for optimal learning. Nature communications **10**(1), 4646 (2019)

[28] Chen, O., Paas, F., Sweller, J.: A cognitive load theory approach to defining and measuring task complexity through element interactivity. Educational Psychology Review **35**(2), 63 (2023) https://doi.org/10.1007/s10648-023-09782-w . Accessed 2026-05-09

[29] Hart, S.G., Staveland, L.E.: Development of NASA-TLX (Task load index): Results of empirical and theoretical research. In: Human Mental Workload. Advances in psychology, 52, pp. 139–183. North-Holland, Oxford, England (1988). https://doi.org/10.1016/S0166-4115(08)62386-9

[30] Kwon, O.-S., Zelaznik, H.N., Chiu, G., Pizlo, Z.: Human motor transfer is determined by the scaling of size and accuracy of movement. Journal of Motor Behavior **43**(1), 15–26 (2010) https://doi.org/10.1080/00222895.2010.520359 . Accessed 2025-03-04

[31] Singh, P., Jana, S., Ghosal, A., Murthy, A.: Exploration of joint redundancy but not task space variability facilitates supervised motor learning. Proceedings of the National Academy of Sciences **113**(50), 14414–14419 (2016) https://doi.org/10.1073/pnas.1613383113 . Accessed 2026-01-22

[32] Buzzi, J., De Momi, E., Nisky, I.: An uncontrolled manifold analysis of arm joint variability in virtual planar position and orientation telemanipulation. IEEE Transactions on Biomedical Engineering **66**(2), 391–402 (2019) https://doi.org/10.1109/TBME.2018.2842458 . Accessed 2026-01-23

[33] Paas, F., Van Merriënboer, J.J.G.: Cognitive-load theory: methods to manage working memory load in the learning of complex tasks. Current Directions in Psychological Science **29**(4), 394–398 (2020) https://doi.org/10.1177/0963721420922183 . Accessed 2026-05-09

[34] VanLehn, K.: Cognitive skill acquisition. Annual Review of Psychology **47**(1), 513–539 (1996) https://doi.org/10.1146/annurev.psych.47.1.513 . Accessed 2026-01-26

[35] Siyar, S., Azarnoush, H., Rashidi, S., Winkler-Schwartz, A., Bissonnette, V., Ponnudurai, N., Del Maestro, R.F.: Machine learning distinguishes neurosurgical skill levels in a virtual reality tumor resection task. Medical & Biological Engineering & Computing **58**(6), 1357–1367 (2020) https://doi.org/10.1007/s11517-020-02155-3 . Accessed 2026-05-09

[36] Wilson, M.R., Poolton, J.M., Malhotra, N., Ngo, K., Bright, E., Masters, R.S.W.: Development and validation of a surgical workload measure: the surgery task load index(Surg-tlx). World Journal of Surgery **35**(9), 1961–1969 (2011) https://doi.org/10.1007/s00268-011-1141-4 . Accessed 2026-05-09

[37] Shafiei, S.B., Shadpour, S., Mohler, J.L., Attwood, K., Liu, Q., Gutierrez, C., Toussi, M.S.: Developing surgical skill level classification model using visual metrics and a gradient boosting algorithm. Annals of Surgery Open **4**(2), 292 (2023) https://doi.org/10.1097/AS9.0000000000000292 . Accessed 2026-01-23

[38] Alzahrani, T., Haddad, R., Alkhayal, A., Delisle, J., Drudi, L., Gotlieb, W., Fraser, S., Bergman, S., Bladou, F., Andonian, S., Anidjar, M.: Validation of the da Vinci Surgical Skill Simulator across three surgical disciplines: A pilot study. Canadian Urological Association Journal = Journal De l'Association Des Urologues Du Canada **7**(7-8), 520–529 (2013) https://doi.org/10.5489/cuaj.419

[39] Soleymani, A., Tavakoli, M., Aghazadeh, F., Ou, Y., Rouhani, H., Zheng, B., Li, X.: Hands collaboration evaluation for surgical skills assessment: an information theoretical approach. IEEE Transactions on Medical Robotics and Bionics **6**(4), 1490–1501 (2024) https://doi.org/10.1109/TMRB.2024.3464110 . Accessed 2026-01-23

[40] Varghese, R., Rossi, C., Malone, L.A., Bastian, A.J.: Task demands shift motor learning from adaptation to feedback control in a naturalistic bimanual task. bioRxiv, 2025–0501651653 (2026) https://doi.org/10.1101/2025.05.01.651653 . Accessed 2026-05-22

[41] Saul, K.R., Hu, X., Goehler, C.M., Vidt, M.E., Daly, M., Velisar, A., Murray, W.M.: Benchmarking of dynamic simulation predictions in two software platforms using an upper limb musculoskeletal model. Computer Methods in Biomechanics and Biomedical Engineering **18**(13), 1445–1458 (2015) https://doi.org/10.1080/10255842.2014.916698 . Accessed 2026-01-23

[42] Fischer, F., Bachinski, M., Klar, M., Fleig, A., Müller, J.: Reinforcement learning control of a biomechanical model of the upper extremity. Scientific Reports **11**(1), 14445 (2021) https://doi.org/10.1038/s41598-021-93760-1 . Accessed 2026-01-23

[43] Wen, R., Rouxel, Q., Mistry, M., Li, Z., Tiseo, C.: Collaborative bimanual manipulation using optimal motion adaptation and interaction control. arXiv. arXiv:2206.00528 (2023). https://doi.org/10.48550/arXiv.2206.00528 . http://arxiv.org/abs/2206.00528 Accessed 2026-01-23

[44] Grotz, M., Shridhar, M., Asfour, T., Fox, D.: Peract2: benchmarking and learning for robotic bimanual manipulation tasks. arXiv. arXiv:2407.00278 (2024). https://doi.org/10.48550/arXiv.2407.00278 . http://arxiv.org/abs/2407.00278 Accessed 2026-01-23

[45] Kazanzides, P., Chen, Z., Deguet, A., Fischer, G.S., Taylor, R.H., DiMaio, S.P.: An open-source research kit for the da vinci® surgical system. In: 2014 IEEE International Conference on Robotics and Automation (ICRA), pp. 6434–6439 (2014). IEEE

[46] Satava, R.M., Stefanidis, D., Levy, J.S., Smith, R., Martin, J.R., Monfared, S., Timsina, L.R., Darzi, A.W., Moglia, A., Brand, T.C., Dorin, R.P., Dumon, K.R., Francone, T.D., Georgiou, E., Goh, A.C., Marcet, J.E., Martino, M.A., Sudan, R., Vale, J., Gallagher, A.G.: Proving the effectiveness of the fundamentals of robotic surgery (Frs) skills curriculum: a single-blinded, multispecialty, multi-institutional randomized control trial. Annals of Surgery **272**(2), 384–392 (2020) https://doi.org/10.1097/SLA.0000000000003220 . Accessed 2026-05-09

[47] Huang, Y.: Understanding as designed projection. Computational Brain & Behavior (2026) https://doi.org/10.1007/s42113-026-00299-3 . Accessed 2026-05-09

[48] Hespanha, J.P.: Linear Systems Theory, Second edition edn. Princeton University Press, Princeton (2018)

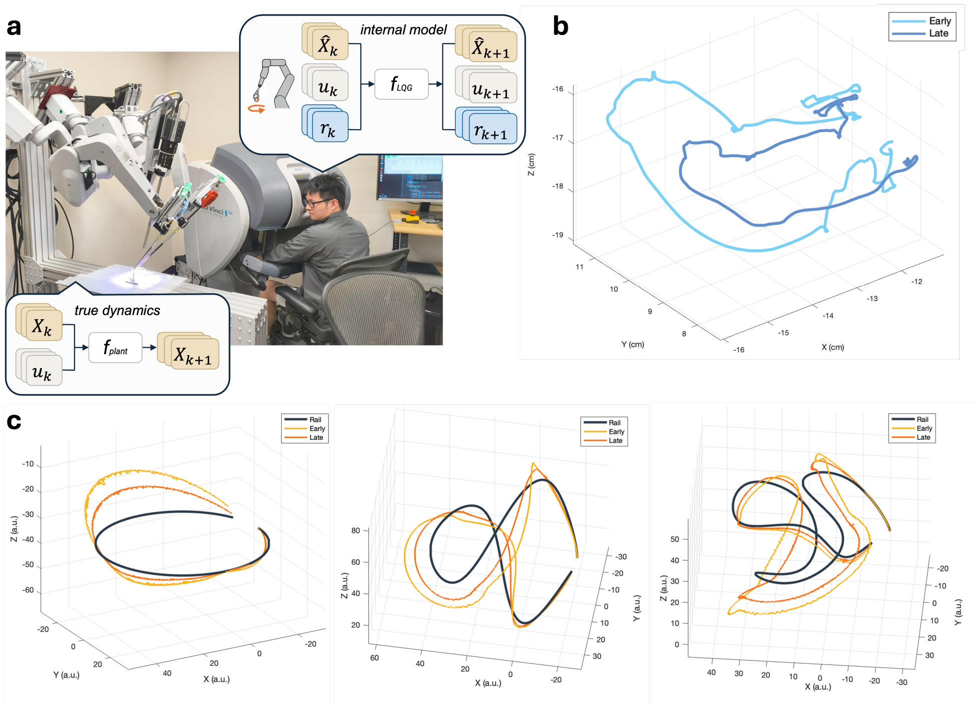

Fig. 1 Human teleoperation model and simulation. (a) A participant teleoperates a da Vinci Research Kit (dVRK) robotic system to move a ring along a 3D rail. Through repeated trials, they update their internal model by acting and observing the outcomes of their own motor commands, progressively becoming a more accurate teleoperator. The true plant dynamics f_{plant} describe how the physical system evolves from state X_k to X_{k+1} under input u_k, while the internal model f_{LQG} predicts $\hat{X}_{k+1}$ from $\hat{X}_k$, r_k, and u_k. (b) Gripper trajectories from early (light blue) and late (dark blue) trials on Level 1 show that with practice, participants' movements become more direct and efficient. The smoother paths with reduced overall length reflects improved sensorimotor calibration and more accurate internal prediction of system dynamics. (c) LQG simulated gripper translational trajectories relative to the rail centerline across the three rail complexity levels. The model captures characteristic deviations around curved segments: near high curvature, the simulated controller overshoots due to delayed correction, and in segments with rapid elevation changes, the trajectory lags behind the reference.

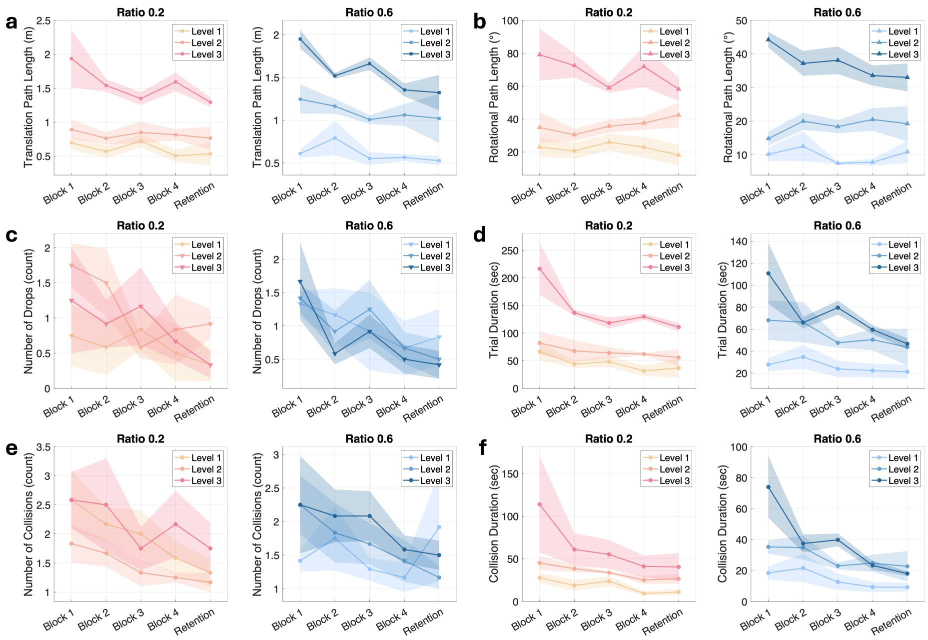

Fig. 2 **Human performance improves across training.** (a-f) Across Blocks 1-4, participants showed improved performance, reflected by reduced translational path length, number of drops, trial duration, number of collisions, and collision duration. Rotational path length did not show a consistent reduction. Shaded regions represent $\pm$SEM.

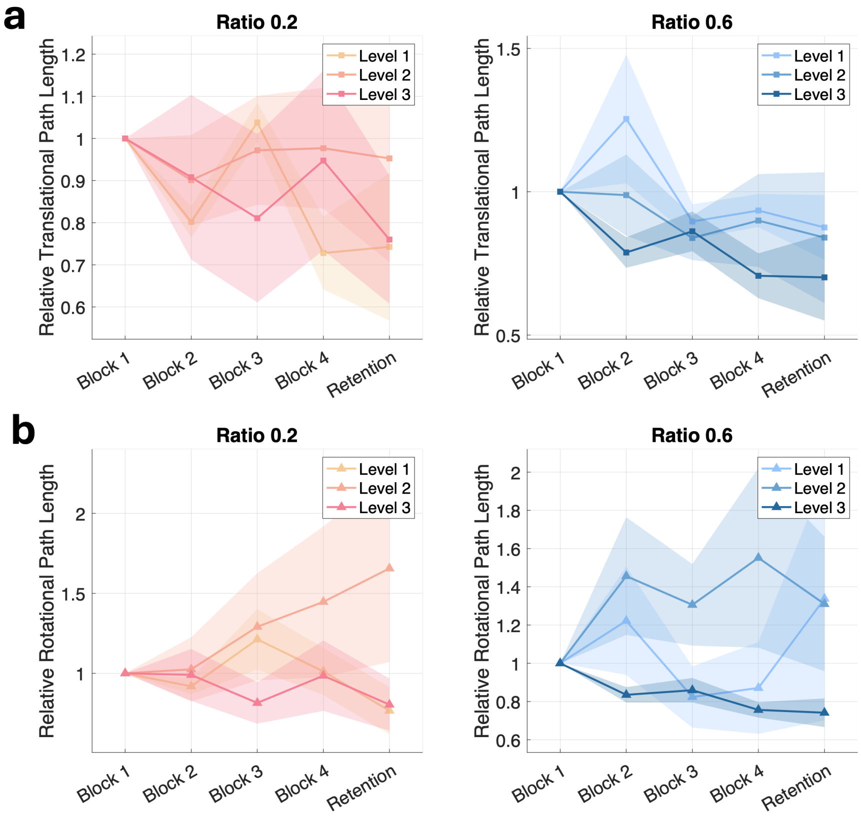

Fig. 3 Relative changes in translational and rotational path length. Translational path length, normalized to Block 1, decreased across training blocks and continued to decrease during retention, indicating more direct and efficient trajectories. In contrast, rotational path length, also normalized to Block 1, did not show a consistent reduction, indicating exploration of wrist angles during training. Error bars represent $\pm$SEM.

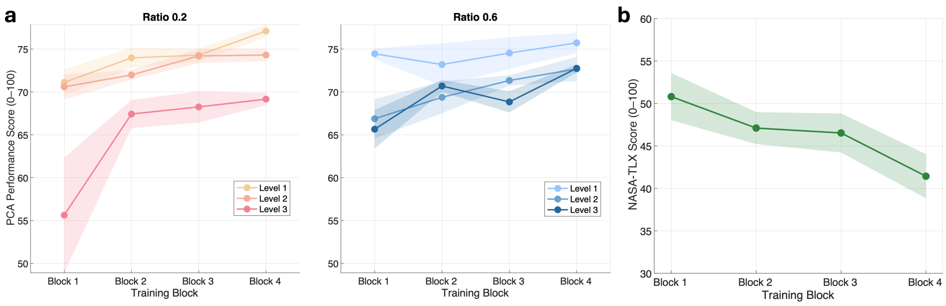

Fig. 4 (a) Performance improvement across training. The x-axis denotes training blocks in chronological order. The y-axis shows mean performance scores (0–100), with shaded regions representing the standard error of the mean (SEM). (b) NASA-TLX workload reduced across training. Shaded regions represent SEM.

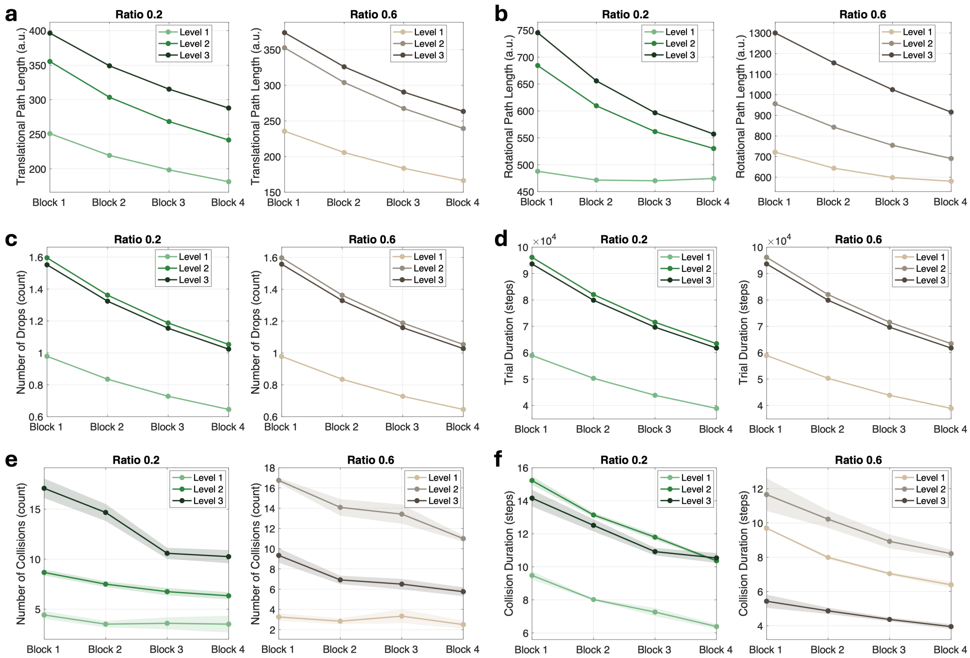

Fig. 5 LQG performance improves across training. (a-f) Across Blocks 1-4, LQG showed reduced translational path length, number of drops, trial duration, number of collisions, and collision duration. Rotational path length did not show a consistent reduction. Shaded regions indicate $\pm$ SEM.

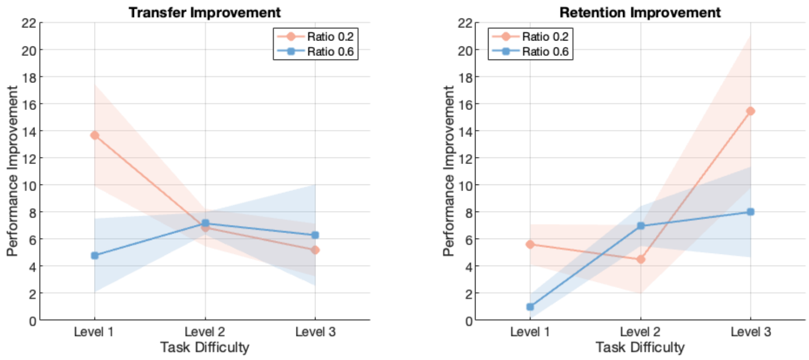

Fig. 6 Relative retention and transfer improvement in human. Composite performance was computed by applying PCA to six behavioral metrics (translational path length, rotational path length, drops, trial duration, collisions, and collision duration). Relative improvements for the individual metrics are shown in Fig. S1. The resulting PCA score was used as a single measure of task performance. **(a)** Retention, defined as the change in PCA composite performance score from the first training block to the retention block. Positive values indicate improved performance relative to initial training (higher values correspond to greater retention). **(b)** Transfer, defined as the change in PCA composite performance score from the pre-training test to the transfer test. Positive values indicate improved performance relative to pre-training (higher values correspond to greater transfer). Shaded regions indicate $\pm$standard deviation (SD) across participants.

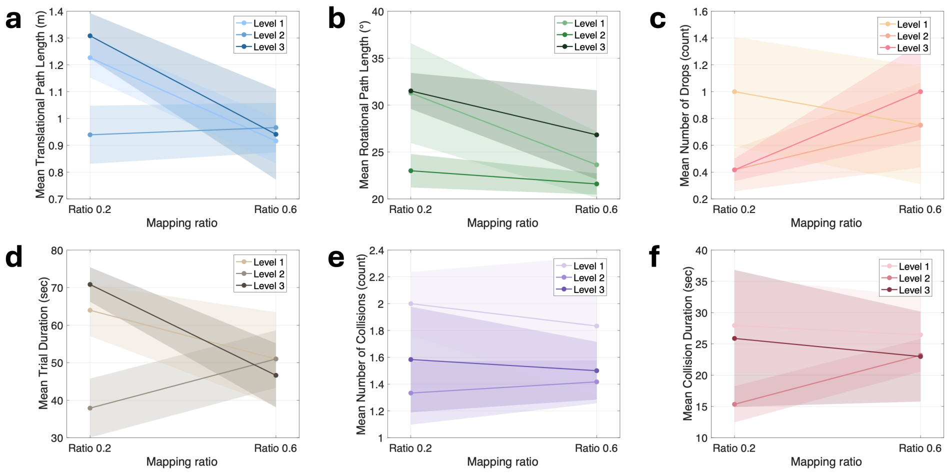

Fig. 7 Absolute transfer performance in human. Behavioral metrics measured at the transfer block (without normalized by pre-training performance). This figure shows whether absolute transfer performance differs across conditions (different from the previous transfer score, which is normalized by pre-training performance). Only translational path length showed a significant difference between scaling ratios, whereas rotational path length, drops, trial duration, collisions, and collision duration did not show reliable scaling-ratio effects. Shaded regions indicate $\pm$SD across participants.

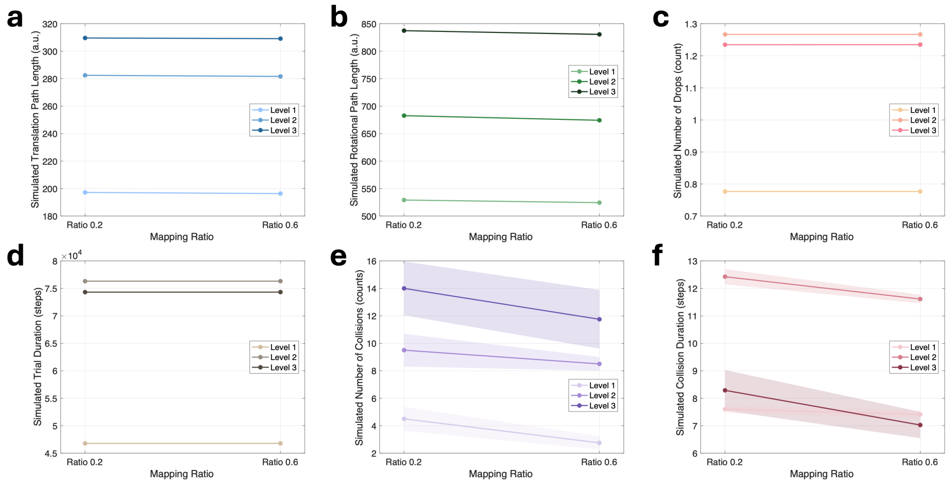

Fig. 8 Absolute transfer performance simulated by LQG. Simulated metrics at the transfer block for two training ratios ($g = 0.2$ vs 0.6) across three rail complexity levels. Shaded regions indicate $\pm$SD across four simulated subjects per condition (matching the human sample size). Consistent with the human results, rail geometry induces changes in trajectory-based outcomes, whereas the scaling-ratio effect is weak.

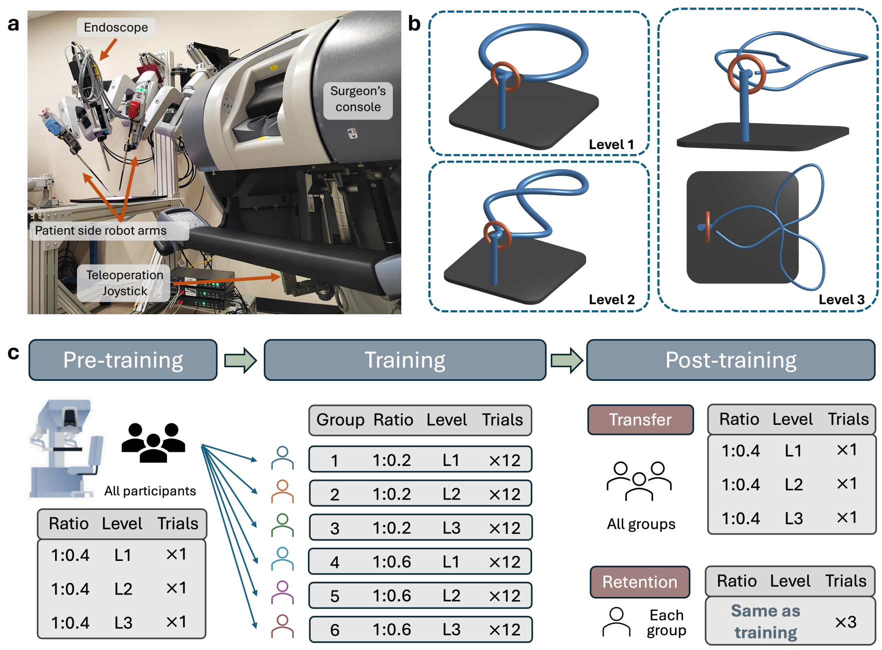

Fig. 9 Experimental setup and procedure for surgical robot teleoperation. (a) dVRK robotic system used for surgical teleoperation. (b) Ring-through-rail apparatus with three different complexity levels. (c) Experimental procedure.

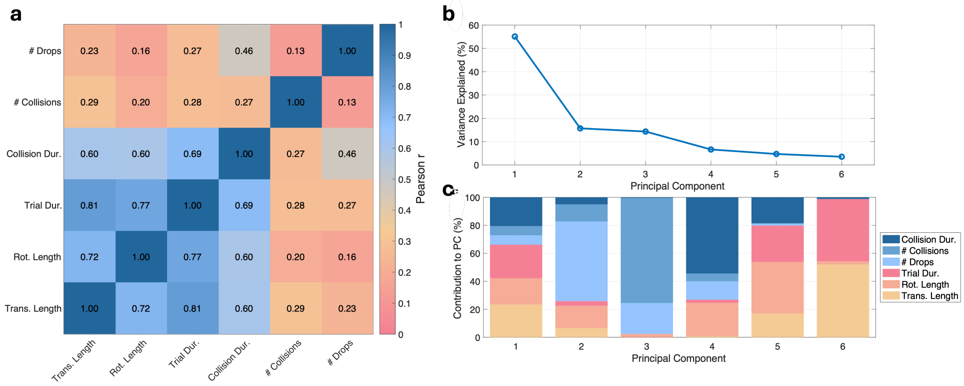

Fig. 10 Principal component analysis (PCA) of trial-wise behavioral metrics. (a) Pairwise correlation matrix (Pearson r) across all six behavioral metrics, computed over all trials. (b) Scree plot showing the percentage of variance explained by the first six principal components. (c) Bar plots showing the relative contribution of each behavioral metric to each principal component.

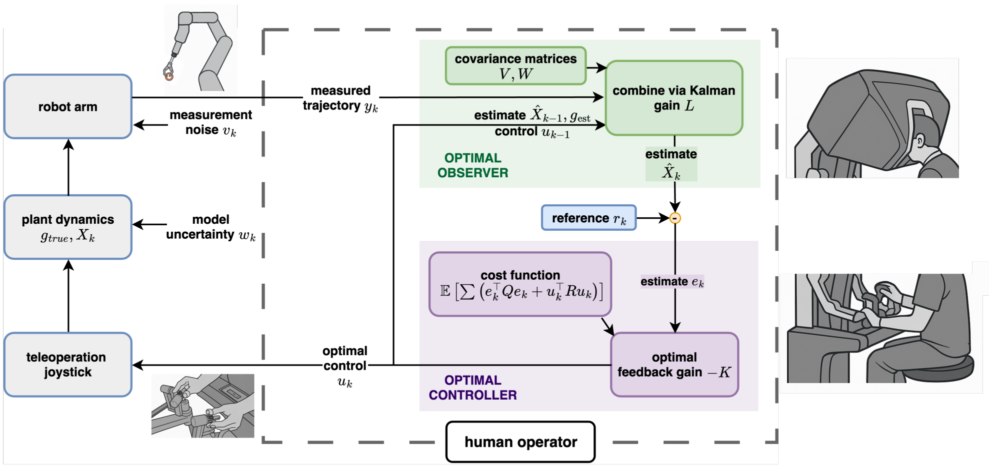

Fig. 11 **Optimal control model of human teleoperation.** The human operator is modeled as an LQG controller with state estimation via a Kalman filter and optimal feedback control for reference tracking under measurement noise and model uncertainty.

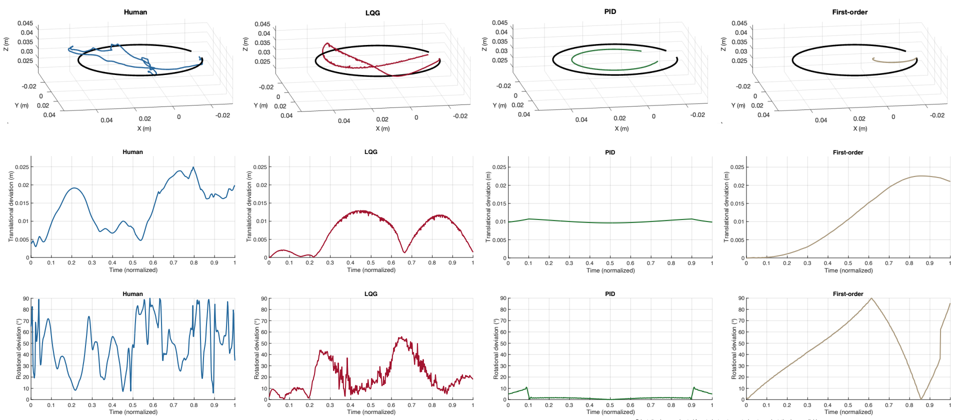

Fig. 12 Comparison between human teleoperation trajectory and controller model trajectories in the level 1 circular rail tracking task. Columns correspond to the human trajectory, the LQG controller, the PD controller, and the first-order controller. Top row: 3D translational trajectories relative to the rail centerline. Middle row: translational deviation from the rail across normalized movement time. Bottom row: tangent-based rotational deviation relative to the local rail direction. Although the simpler PD and first-order controllers produced smoother geometric rail-following behavior, the LQG controller qualitatively reproduced several characteristic features observed in the human trajectories, including curvature-dependent deviations, overshoot, and changing rotational corrections near high-curvature regions.

Table 1 Summary of statistical tests for training-related changes across metrics.

Metric	Paired t	Cohen's d	% Change
Translational length	$t(23) = 2.840$, $p = 0.009$	0.580	-18.8%
Rotational length	$t(23) = 1.020$, $p = 0.320$	0.210	$+10.2\%$
Trial duration	$t(23) = 3.770$, $p = 0.001$	0.770	-34.7%
Collision duration	$t(23) = 3.530$, $p = 0.002$	0.720	-47.6%
Collisions	$t(23) = 2.740$, $p = 0.012$	0.560	-14.8%
Drops	$t(23) = 4.260$, $p < 0.001$	0.870	-50.4%
LME Block Effect (Blocks 1–Retention)			
Translational length	$F(4, 115) = 5.162$, $p = 0.001$		
Rotational length	$F(4, 115) = 0.706$, $p = 0.590$		
Trial duration	$F(4, 115) = 10.381$, $p < 0.001$		
Collision duration	$F(4, 115) = 10.271$, $p < 0.001$		
Collisions	$F(4, 115) = 6.036$, $p < 0.001$		
Drops	$F(4, 115) = 8.564$, $p < 0.001$		

Table 2 Transfer improvement across complexity levels.

Complexity Level	**Mean (M)**	**95% CI**	n
Level 1	9.24	$[2.79, 15.70]$	8
Level 2	7.01	$[5.25, 8.77]$	8
Level 3	5.75	$[1.11, 10.38]$	8

Table 3 Retention improvement across difficulty levels.

Complexity Level	**Mean (M)**	**95% CI**	n
Level 1	3.32	$[0.52, 6.12]$	8
Level 2	5.74	$[2.34, 9.13]$	8
Level 3	11.74	$[3.82, 19.66]$	8

Table 4 Linear mixed-effects model (LME) tests on LQG transfer performance.

Metric	**Effect**	**Test statistic**
Translational path length	Level	$F(2, 18) = 132880.000,\ p < 0.001$
	Ratio	$F(1, 18) = 14.424,\ p = 0.001$
	Level × Ratio	$F(2, 18) = 0.879,\ p = 0.432$
Rotational path length	Level	$F(2, 18) = 79552.000,\ p < 0.001$
	Ratio	$F(1, 18) = 35.519,\ p < 0.001$
	Level × Ratio	$F(2, 18) = 9.429,\ p = 0.002$
Collision duration	Level	$F(2, 18) = 43.220,\ p < 0.001$
	Ratio	$F(1, 18) = 0.117,\ p = 0.736$
	Level × Ratio	$F(2, 18) = 0.352,\ p = 0.708$
Collisions	Level	$F(2, 18) = 23.190,\ p < 0.001$
	Ratio	$F(1, 18) = 1.167,\ p = 0.294$
	Level × Ratio	$F(2, 18) = 2.389,\ p = 0.120$

Table 5 Cross-correlation and cumulative mismatch between model and human deviations.

Model	**Metric**	**Corr. / Lag**	**Mismatch**
LQG	Translation	$r = 0.8958$, lag $= -0.3186$	0.0091
	Rotation	$r = 0.4807$, lag $= -0.2926$	30.4442
PD	Translation	$r = 0.4171$, lag $= -0.0200$	0.0059
	Rotation	$r = 0.5800$, lag $= 0.3888$	45.7314
First-order	Translation	$r = 0.6770$, lag $= -0.2725$	0.0056
	Rotation	$r = 0.3733$, lag $= -0.1804$	24.3457

Table 6 Parameters used in model comparison.

Model	Parameter	Value
LQG	Translational process noise W_t	1×10^{-5}
LQG	Rotational process noise W_r	4×10^{-6}
LQG	Translational measurement noise V_t	2×10^{-5}
LQG	Rotational measurement noise V_r	1×10^{-5}
LQG	Position cost q_P	2.0×10^{5}
LQG	Velocity cost q_V	2.5×10^{2}
LQG	Rotation cost q_R	5.0×10^{2}
LQG	Angular velocity cost q_W	1.0×10^{1}
LQG	Translational control cost r_T	1.5×10^{-2}
LQG	Rotational control cost r_R	4.0×10^{-3}
PD	K_p	0.80
PD	K_i	0
PD	K_d	0.02
PD	Damping	0.94
PD	Command limit	0.004
First-order	α	0.002
First-order	Maximum step	0.0005