

Fear and Anxiety Differentially Reweight Control Objectives: Evidence from Inverse Optimal Control

Tahmineh A. Koosha^{1,3}, Fabian Hahne^{2,3}, Alap Kshirsagar^{2,3}, Nick Augustat^{1,3}, Jan Peters^{1,3}, Dominik M. Endres^{1,3}

¹Theoretical Cognitive Science, Universität Marburg, Marburg, Germany

²Intelligent Autonomous Systems, Technische Universität Darmstadt, Darmstadt, Germany

³Center for Mind, Brain, and Behavior (CMBB), Philipps-Universität Marburg, Germany

Abstract

Exposure to height induces pronounced changes in human postural behavior, yet the computational mechanisms underlying these adaptations remain unclear. Previous studies report both reduced and increased postural sway under height exposure leading to seemingly contradictory interpretations. Here, we apply inverse optimal control (IOC) to infer how control objectives governing postural behavior are reweighted under threat. Participants performed quiet standing in a virtual reality environment under ground and height conditions while joint-level kinematics were recorded. Postural control was modeled as a multi-joint optimal control problem with fixed dynamics, and condition-specific cost weights on joint velocity and control effort were inferred. Height exposure was associated with increased penalization of joint angular velocity and altered effort weighting, consistent with more conservative control strategies. Substantial inter-individual variability was observed. Moreover, anxiety-related measures were linked to increased control variability, whereas fear of heights and heart-rate change were associated with motion-suppressive control. We thus provide computational evidence for a functional dissociation between fear of heights and anxiety in human balance control.

Keywords: inverse optimal control; postural control; fear of heights; anxiety; motor control; virtual reality

Introduction

Maintaining upright balance is a fundamental motor task that requires continuous integration of sensory information and adaptive control of the body's degrees of freedom. When balance is threatened, such as during exposure to heights, humans exhibit pronounced changes in posture and movement, including altered sway amplitude, increased muscle co-contraction, and heightened physiological arousal. These responses are commonly attributed to fear, which encompasses both perceptual distortions and autonomic reactions that dynamically influence motor behavior (Adkin et al., 2002; Wiederhold & Bouchard, 2014). Individuals with fear of heights often avoid edges, windows, or elevated environments even in situations with minimal physical risk, highlighting that perceived threat, rather than objective danger, plays a central role in shaping action.

Empirical studies investigating balance under height exposure have reported mixed findings. Some studies show reduced postural sway, consistent with a stiffening strategy aimed at minimizing movement (Cleworth et al., 2012; Spartakov et al., 2024), whereas others report increased sway at greater heights, possibly reflecting overcompensation or instability under extreme threat (Wuehr et al., 2019). Similar

variability has been observed across real-world settings, such as balconies or elevated platforms (Alpers & Adolph, 2008; Wuehr et al., 2014), as well as virtual reality (VR) environments that enable systematic manipulation of height and visual context (Bzdúšková et al., 2023; Krupić et al., 2021). Together, these findings challenge simple accounts in which fear uniformly destabilizes posture.

A key limitation of much of this literature is its reliance on descriptive sway metrics, such as center-of-pressure (CoP) amplitude or frequency. While informative, these measures do not reveal how fear alters the underlying control processes that generate behavior. From a cognitive and sensorimotor perspective, balance emerges from the integration of multisensory information to maintain an internal representation of body state (Horak, 2006; Mergner & Rosemeier, 1998). When this representation signals increased risk—such as an elevated likelihood of falling—control policies must be adjusted to restore safety, often at the expense of movement efficiency (Clark, 2013).

Optimal Feedback Control (OFC) theory provides a principled framework for formalizing these adjustments. OFC models the nervous system as selecting control policies that minimize a cost function trading off task-relevant state deviations against control effort (Todorov & Jordan, 2002). A central implication of this framework is that changes in behavior can arise from reweighting control objectives without altering the underlying dynamics. Applied to postural control, this opens up the possibility that threatening environments induce more conservative strategies, such as increased penalties on movement or velocity, resulting in stable but constrained posture.

In this vein, recent work has proposed OFC as a unifying account of fear-related postural behavior, suggesting that apparently contradictory findings—such as sway reduction versus sway amplification—reflect distinct control regimes rather than inconsistent effects of fear (Spartakov et al., 2024). Complementary experimental results using indoor VR environments have shown that virtual height exposure can reduce sway amplitude while increasing sway frequency, consistent with a high-gain stabilization strategy (Koosha, Kshirsagar, et al., 2025). However, these studies remain largely descriptive and do not directly infer the control objectives responsible for observed behavior.

In the present study, we address this gap by applying in-

verse optimal control (IOC) to joint-level postural behavior. Rather than characterizing posture using aggregate sway metrics, we model standing balance as the outcome of a shared multi-joint control architecture acting on an inverted-pendulum-like body representation (Reimann & Schöner, 2017). Using joint kinematics recorded during quiet standing in VR under ground and height conditions, we infer condition-specific cost weights of a small number of control objectives while holding system dynamics and modeling assumptions fixed. Differences in inferred costs between conditions provide evidence that height exposure induces systematic changes in postural control strategy, linking fear- and anxiety-related behavioral adaptations to reweighting of these control objectives.

Materials and Methods

Experimental Design

This study builds on a previously reported virtual-reality balance experiment comparing postural control under ground and height conditions (Koosha, Hahne, et al., 2025; Koosha, Kshirsagar, et al., 2025). Experimental setup, preprocessing, and descriptive analyses followed prior work and are summarized briefly here. The present contribution introduces a novel inverse optimal control analysis of joint-level kinematics.

Experimental Task Participants performed quiet standing in a virtual environment while joint-level kinematics were recorded using a Microsoft Kinect v2 system (Microsoft, 2014). Each trial comprised a 20 s exploration phase, an anxiety rating (11-point Likert scale), and a 60 s fixation phase during upright standing.

Two conditions were tested: Ground (flat surface) and Height (virtual platform ~ 20 m above ground). The experiment consisted of seven trials, starting with a Ground trial, followed by two randomized Ground–Height blocks, and ending with a Height trial followed by a Ground trial.

Psychological and Physiological Measures

Heart Rate Change (HRC) Cardiac activity was recorded using a Polar H10 chest-strap sensor (Markspan, 2017). ECG signals were band-pass filtered, R-peaks were detected, and average heart rate segment’sted over the 60 s fixation phase. Heart rate change was defined as

$$\text{HRC} = \text{HR}_t - \text{HR}_{t-1},$$

and analyzed for consecutive Ground–Height trial pairs, with positive values indicating increased heart rate following height exposure.

Psychological Measures State anxiety was assessed using the cognitive and somatic subscales of the STICSA (Ree et al., 2008), alongside a single-item fear-of-heights rating. Additional questionnaire details have been reported previously (Koosha, Kshirsagar, et al., 2025).

Kinematic Representation Joint angles and center-of-mass (CoM) position were derived from Kinect joint positions using sagittal-plane geometry and a planar anthropometric model with standard segment proportions (Reimann & Schöner, 2017).

Biomechanical Model

Postural dynamics were modeled using a three-link sagittal-plane inverted pendulum representing the ankle, knee, and hip joints. The model was based on the multi-joint stance model of Reimann and Schöner (2017), adapted for use within an optimal control framework.

The state vector was defined as

$$x_t = [q_t \quad \dot{q}_t],$$

where $q_t \in \mathbb{R}^3$ denotes joint angles and $\dot{q}_t \in \mathbb{R}^3$ joint angular velocities. Control inputs were joint torques $u_t \in \mathbb{R}^3$.

System dynamics followed

$$M(q)\ddot{q} + C(q, \dot{q}) + G(q) + \tau_{\text{passive}}(q, \dot{q}) = u,$$

with M , C , and G denoting inertia, Coriolis/centrifugal, and gravitational terms, respectively. Passive joint torques were modeled as the sum of a constant offset and a linear spring–damper component defined relative to a reference posture. This reference posture was set to the subject-specific initial joint configuration at the beginning of each analyzed window, such that the linear spring–damper contribution was zero at this posture. All passive parameters and physical model parameters were held constant across experimental conditions.

Optimal Control Formulation

Postural control was formulated as a finite-horizon deterministic optimal control problem, in which control signals minimize a cumulative cost subject to the nonlinear biomechanical dynamics. Rather than assuming trajectory tracking, this formulation assumes that observed behavior approximately optimizes a small number of task-relevant objectives.

The cost function was parameterized by weights $\phi = \{w_{\text{CoM}}, w_{\dot{q}}, w_u\}$:

$$J(\phi) = \sum_{t=0}^{T-1} \ell(x_t, u_t), \quad x_t = (q_t, \dot{q}_t).$$

The stage cost was defined as

$$\ell(x_t, u_t) = w_{\text{CoM}} \left(\frac{\text{CoM}(q_t) - \text{CoM}^*}{\sigma_{\text{CoM}}} \right)^2 + w_{\dot{q}} \left\| \frac{\dot{q}_t}{\sigma_{\dot{q}}} \right\|^2 + w_u \|u_t\|^2.$$

To ensure identifiability, the center-of-mass stabilization weight was fixed to $w_{\text{CoM}} = 1$ and served as a reference scale. Condition-dependent effects were quantified in the remaining control weights ($w_{\dot{q}}, w_u$).

The first term penalizes deviations of the center of mass from a reference position CoM^* , defined as the empirical

mean CoM position observed for each condition. This corresponds to a quasi-static equilibrium assumption, whereby balance is regulated around a condition-specific operating point inferred from behavior.

The joint-velocity term penalizes rapid movements, encouraging smooth, low-velocity motion, as commonly assumed in optimal control models of motor behavior. The control-effort term penalizes joint torques via an isotropic quadratic cost ($R = w_v I$), such that effort corresponds to the weighted sum of squared torques; only a single scalar effort weight was optimized.

State-related cost terms were normalized by fixed scaling constants ($\sigma_{\text{CoM}}, \sigma_{\dot{q}}$) to ensure dimensionless and comparable contributions. These constants were chosen based on empirical variability in pilot data and were held fixed across subjects and conditions ($\sigma_{\dot{q}} = 10^{-2}$ rad/s, $\sigma_{\text{CoM}} = 10^{-3}$ m).

No explicit joint-angle tracking term was included; joint configurations therefore emerged indirectly from CoM regulation, smoothness constraints, and the system dynamics.

Trajectory Optimization via iLQR

For fixed cost parameters, optimal trajectories were computed using the iterative Linear Quadratic Regulator (iLQR) algorithm (Li & Todorov, 2004). iLQR iteratively approximates the nonlinear dynamics with linear models and the cost with quadratic functions around a nominal trajectory, yielding a sequence of time-varying linear-quadratic subproblems.

Each iteration consisted of a backward pass to compute a locally optimal time-varying feedback policy, followed by a forward rollout with line search to update the nominal trajectory.

All simulations were performed over a 3 s horizon with a 3 ms integration timestep, providing sufficient temporal resolution to capture postural dynamics while remaining computationally tractable.

Inverse Optimal Control

Condition-specific cost parameters were inferred using a simulation-based inverse optimal control procedure. The goal was to identify the cost weights ϕ for which the resulting iLQR trajectories best matched the observed joint-angle behavior.

Discrepancy between simulated and empirical joint-angle trajectories was quantified as

$$\mathcal{L}(\phi) = \text{RMSE} \left(\frac{q_{\text{sim}}(\phi) - q_{\text{data}}}{\sigma_q} \right),$$

where $\sigma_q = 10^{-3}$ rad is the fixed normalization constant defined above, ensuring comparability across conditions and subjects. Only joint angles were used as fitting targets; center-of-mass position and joint velocities entered solely as objectives in the forward optimal control problem.

An outer-loop derivative-free optimizer (BOBYQA; Powell et al., 2009) was used to minimize $\mathcal{L}(\phi)$. Optimization was

performed independently for the Ground and Height conditions. Across conditions, system dynamics, simulation horizon, target definitions, and normalization constants were held fixed. Consequently, differences in inferred parameters reflect changes in the relative weighting of control objectives rather than changes in biomechanical structure or task dynamics.

Parameter Inference and Data Analysis

Analyses were performed on data from nine participants for whom complete motion-capture recordings were available for the analyzed time window in both conditions. Data from additional participants were excluded due to missing or incomplete kinematic recordings within this window.

Condition-wise Inference and Performance Metrics Inverse optimal control analyses focused on the first Height trial and the immediately preceding Ground trial for each participant. This pairing isolates acute effects of height exposure while minimizing habituation and order effects. All group-level analyses used this trial pairing.

To facilitate interpretation and cross-subject comparison, condition effects were quantified as within-subject log-differences between inferred control weights,

$$\Delta \log w = \log w_{\text{Height}} - \log w_{\text{Ground}},$$

which correspond to log-ratios of the optimized parameters. This representation emphasizes relative, condition-dependent shifts in control strategy while avoiding reliance on absolute parameter magnitudes.

Model performance was evaluated using root-mean-square error (RMSE) between simulated and observed joint trajectories. RMSE was computed for joint angles and joint angular velocities after aligning simulated trajectories to the data sampling rate, and summarized across subjects to characterize reconstruction accuracy and inter-individual variability.

To illustrate the modeling approach, we first report results for a single representative participant using a 3 s analysis window (101 samples at 33 Hz). This example serves as a qualitative demonstration of the inferred controller, whereas all subsequent analyses quantify group-level effects and inter-individual variability.

Correlation Analysis To examine whether condition-dependent changes in inferred control parameters relate to psychological and physiological measures, we computed within-subject differences for each model-derived metric m as

$$\Delta m = m_{\text{Height}} - m_{\text{Ground}}.$$

Pearson correlation coefficients were then computed between each Δm —including ΔRMSE measures, $\Delta \log w_{\dot{q}}$, $\Delta \log w_u$, and differences in control signal statistics—and subject-level measures of fear of height, cognitive and somatic anxiety, and heart-rate change (HRC).

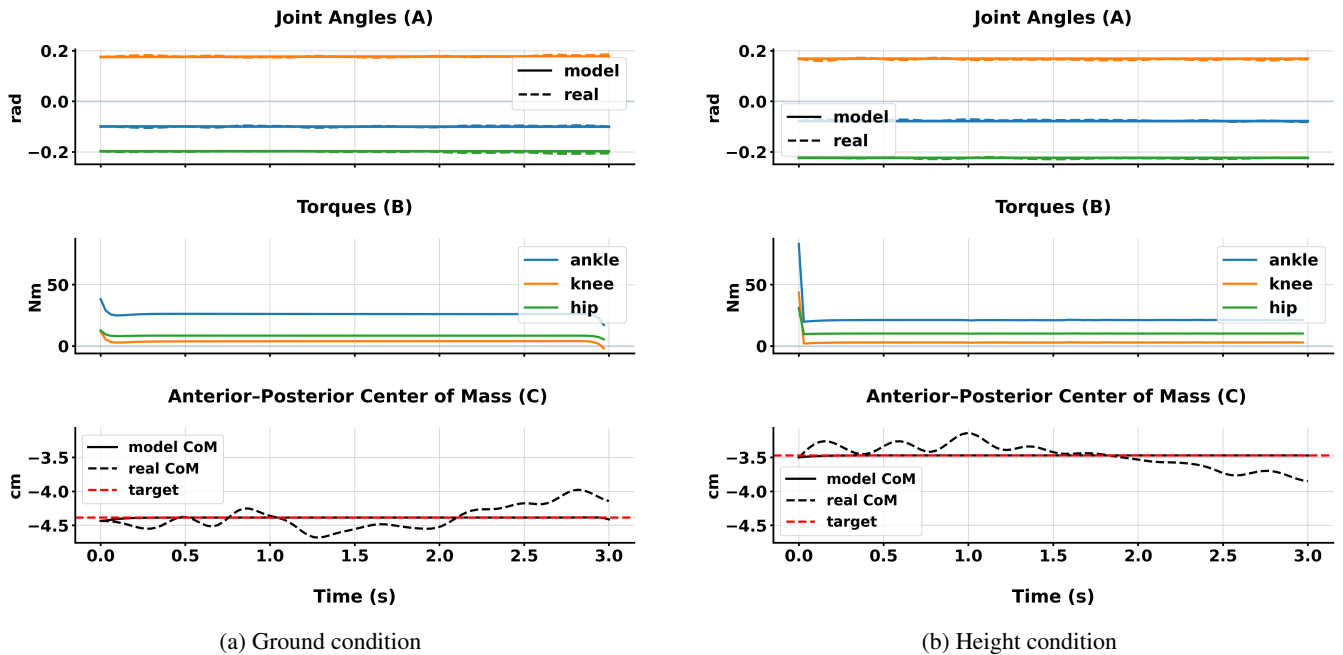


Figure 1: iLQR model fits to experimental postural sway data. Solid lines indicate model predictions and dashed lines indicate experimental data. Across both conditions, the iLQR controller accurately reproduces joint angle trajectories and average center-of-mass behavior.

Results

Inverse Optimal Control Reveals Condition-Dependent Control Strategies

Inverse optimal control was applied to joint-level postural data to infer control cost parameters separately for the ground and height conditions, while holding system dynamics, state representation, and optimization procedures fixed. Optimization converged reliably for all analyzed participants.

A Representative Participant Figure 1 illustrates model fits for a representative participant over the 3 s analysis window. In both Ground and Height conditions, the iLQR controller accurately reproduced joint-angle trajectories and the average center-of-mass (CoM) trajectory, despite the absence of explicit joint-angle tracking. This indicates that key features of postural control emerged from CoM regulation, smoothness constraints, and the system dynamics.

For this participant, the inferred control weights differed markedly between conditions. Relative to the Ground condition ($w_{\dot{q}} = 0.8$, $w_u = 10^{-3}$), the Height condition was characterized by a higher penalty on joint angular velocity ($w_{\dot{q}} = 1.1$) and a substantially reduced penalty on control effort ($w_u = 10^{-6}$). This pattern reflects tighter suppression of joint motion under height exposure, combined with less penalized corrective control actions.

Inter-individual Variability in Control Reweighting

To characterize group-level effects, we analyzed within-subject log-differences in inferred control weights between

the Height and Ground conditions ($\Delta \log w = \log w_{\text{Height}} - \log w_{\text{Ground}}$).

The distributions reveal pronounced inter-individual variability in how control objectives are reweighted under height exposure, see Figure 2. Because system dynamics and optimization settings were identical across conditions, this variability reflects differences in inferred control strategies rather than estimation noise. For example, participants in the lower-right quadrant (Q4) increased their velocity penalty while reducing the control effort penalty. In our sample, this pattern was observed in the majority of participants. In contrast, participants in Q2 reduced their velocity penalty while increasing effort penalization, a pattern consistent with more passive sway. Participants in Q3 decreased both penalties, suggesting a more active control strategy. The absence of participants in Q1 may reflect the infeasibility of simultaneously increasing both movement suppression and effort penalization within this task. These results indicate that height exposure does not induce a uniform reweighting of control objectives across individuals, but instead amplifies subject-specific control strategies.

Trajectory Reconstruction Accuracy

Model performance was quantified using root-mean-square error (RMSE) between simulated and empirical trajectories for joint angles (q), joint angular velocities (\dot{q}), and center-of-mass (CoM) position (Table 1). RMSE values are reported as mean \pm standard deviation across subjects for the Ground and Height conditions.

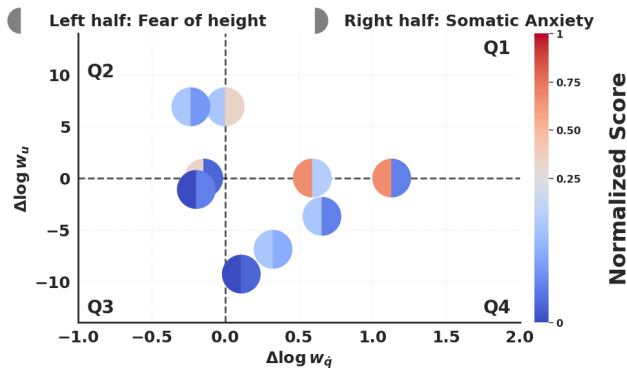


Figure 2: Subject-wise changes in inferred control weights between the height and ground conditions. Each point represents one subject, positioned by the difference in log-weight for joint-velocity cost ($\Delta \log w_{\dot{q}}$) and control-effort cost ($\Delta \log w_u$). Dashed lines indicate no change between conditions. The four quadrants correspond to distinct patterns of control reweighting, including concurrent increases or decreases in both costs, as well as trade-offs between joint-velocity regulation and control effort.

Table 1: RMSE (mean \pm SD) for joint angles (q), joint velocities (\dot{q}), and center of mass (CoM) under Ground and Height conditions.

Metric	Ground	Height
RMSE $_q$ (rad)	0.00281 \pm 0.00099	0.00404 \pm 0.00197
RMSE $_{\dot{q}}$ (rad/s)	0.01500 \pm 0.00809	0.03301 \pm 0.05456
RMSE $_{CoM}$ (m)	0.00175 \pm 0.00101	0.00136 \pm 0.00065

Across both conditions, the inverse optimal control model achieved low reconstruction error at the joint-angle level, with mean RMSE $_q$ less than 10^{-2} radians. Reconstruction error was slightly higher in the Height condition compared to Ground, consistent with increased postural variability under height exposure. A similar pattern was observed for joint angular velocities.

In contrast, CoM reconstruction error remained comparable across conditions, indicating that the model consistently captured whole-body balance dynamics despite differences in inferred control strategies. Notably, increased q reconstruction error under Height did not coincide with degraded CoM accuracy.

Associations with Psychological and Physiological Measures

Figure 3 suggests a dissociation between anxiety-related and fear-related affective measures in how postural control strategies are reweighted under height exposure. Cognitive and somatic anxiety scores (STICSA) were associated with increases in control effort weighting, control signal variability, and joint-velocity reconstruction error, consistent with a more reactive and variable control strategy. In contrast, fear of

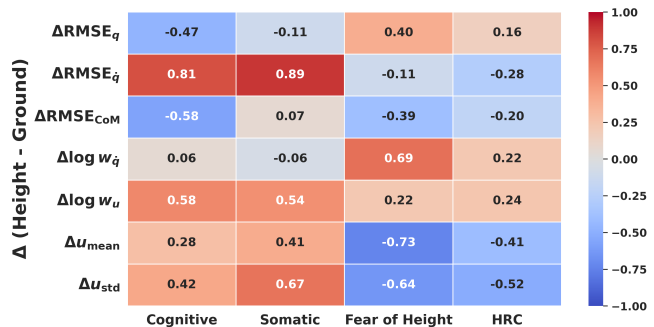


Figure 3: Pearson correlations between psychological/physiological measures and within-subject condition differences in model-derived metrics. For each subject and metric m , differences were computed as $\Delta m = m_{\text{Height}} - m_{\text{Ground}}$. Rows show $\Delta RMSE$ of joint angles, joint velocities, and center of mass, $\Delta \log w_{\dot{q}}$, $\Delta \log w_u$, and differences in control signal statistics (Δu_{mean} , Δu_{std}). Columns show fear, cognitive and somatic anxiety scores, and heart-rate reactivity (HRC). Correlations are reported descriptively; no correction for multiple comparisons was applied.

height was primarily associated with increased penalization of joint velocity and reduced mean and variability of control output, consistent with a conservative, motion-suppressive strategy. Heart-rate change showed a partially overlapping pattern with fear of height, relating primarily to reductions in control output. Together, these patterns suggest that anxiety and fear of heights, although related, map onto distinct computational adaptations of postural control.

Discussion

The present study used inverse optimal control to investigate how exposure to height alters the computational objectives underlying human postural control. Rather than relying on descriptive sway measures, we inferred condition-specific reweighting of control costs while holding biomechanical dynamics and task structure fixed. This approach allowed us to interpret fear-related postural adaptations as changes in control strategy rather than changes in system stability or mechanics.

For a representative participant, height exposure was associated with an increased penalty on joint angular velocity and a strongly reduced penalty on control effort. Within an optimal feedback control framework, this combination implies tighter regulation of joint motion together with permissive use of corrective torques. Such a policy naturally leads to suppressed movement variability while retaining the ability to respond rapidly to perturbations. This pattern is consistent with behavioral reports of reduced postural sway under height exposure, often interpreted as a stiffening strategy (Cleworth et al., 2012).

Inter-Individual Variability Reveals Multiple Control Regimes

Despite consistent group-level trends, we observed pronounced inter-individual variability in how control weights were reweighted under height exposure. Some participants exhibited concurrent increases in velocity penalization and reductions in effort cost, a pattern consistent with stiffening strategies. Others showed reduced velocity penalties or increased effort penalization, suggesting alternative control regimes that may permit greater passive sway or rely more heavily on biomechanical stability.

This variability offers a principled explanation for the mixed findings reported in the height-exposure literature, where both sway reduction and sway amplification have been observed across studies and individuals (Alpers & Adolph, 2008; Wuehr et al., 2019). From an optimal control perspective, these patterns need not be contradictory: they can arise from different solutions to the same control problem under uncertainty and perceived threat.

By explicitly modeling control objectives, our results support recent proposals that fear-related postural behavior reflects transitions between distinct control regimes rather than a single monotonic effect of fear (Spartakov et al., 2024). In this sense, inverse optimal control provides a unifying computational framework for understanding heterogeneity in fear-induced motor behavior.

Dissociable Contributions of Anxiety, Fear of Heights, and Physiological Arousal

A key contribution of the present study is the dissociation between anxiety-related and fear-related affective measures in their association with control reweighting. Cognitive and somatic anxiety scores were primarily associated with increased control effort weighting, increased variability of control output, and increased joint-velocity reconstruction error. This pattern is consistent with a more reactive and variable control strategy, potentially reflecting heightened monitoring, reduced confidence in state estimation or increased motor noise.

In contrast, fear of heights showed a markedly different signature. Higher fear of heights was associated with increased penalization of joint velocity and reduced mean and variability of control output, consistent with a conservative, motion-suppressive strategy. These findings suggest that anxiety and fear of heights, although related at a phenomenological level, map onto distinct computational adaptations of postural control.

Heart-rate change following height exposure showed a partially overlapping pattern with fear of heights rather than anxiety. Specifically, increased heart-rate change was associated with reductions in control output magnitude and variability. This link suggests that autonomic arousal during height exposure is more closely tied to conservative stabilization strategies than to reactive or variable control. Rather than simply reflecting nonspecific stress, heart-rate reactivity may index a shift toward safety-oriented motor policies under perceived

threat.

Implications, Limitations, and Future Work

These results support an optimal control account of fear-related postural behavior in which height exposure induces adaptive reweighting of control objectives rather than uniform changes in stability or motor performance. Within this framework, both reduced and increased sway can arise from different cost structures, providing a unified explanation for previously mixed findings in the literature.

Several limitations constrain the present study. First, inverse optimal control was applied to a short 3 s analysis window, enabling a proof-of-concept demonstration. Extending the framework to longer windows and time-varying inference will be important for assessing the temporal stability of control objectives. Second, the biomechanical model was limited to a three-link sagittal-plane representation and did not include arm or trunk dynamics, which may contribute to balance under threat (cf. Koosha, Kshirsagar, et al. (2025)). Third, the model assumed deterministic dynamics and did not explicitly account for sensory or motor noise, despite evidence that fear and anxiety modulate uncertainty and feedback gain. Finally, the focus on sagittal-plane dynamics excludes mediolateral control, and the modest sample size limits statistical inference.

Future work should extend the model to higher-dimensional, stochastic control formulations, incorporate additional degrees of freedom, and test the robustness of inferred control strategies in larger samples.

Acknowledgments

This work was supported by the Research Cluster “The Adaptive Mind”, funded by the Excellence Program of the HMWK, the Deutsche Forschungsgemeinschaft (German Research Foundation, DFG) under Germany’s Excellence Strategy (EXC 3066/1 “The Adaptive Mind”, Project No. 533717223), the DFG GRK-RTG 2271 ‘Breaking Expectations’ project number 290878970 and the DFG CRC/TRR 135, project number 222641018, subproject C6.

References

- Adkin, A. L., Frank, J. S., Carpenter, M. G., & Peysar, G. W. (2002). Fear of falling modifies anticipatory postural control. *Experimental Brain Research*, *143*(2), 160–170.
- Alpers, G. W., & Adolph, D. (2008). Height intolerance and visual exploration of heights in virtual environments. *Journal of Anxiety Disorders*, *22*(7), 1207–1213.
- Bzdúšková, D., Marko, M., Hirjaková, Z., Riečanský, I., & Kimijanová, J. (2023). Fear of heights shapes postural responses to vibration-induced balance perturbation at virtual height. *Frontiers in Human Neuroscience*, *17*, 1229484.
- Clark, J. E. (2013). From the ground up: Rethinking human postural development. *Motor Control*, *17*(1), 1–20.
- Cleworth, T. W., Horslen, B. C., & Carpenter, M. G. (2012). Influence of real and virtual heights on standing balance. *Gait & Posture*, *36*(2), 172–176.

- Horak, F. B. (2006). Postural orientation and equilibrium. *Age and Ageing*, 35(suppl.2), ii7–ii11.
- Koosha, T. A., Hahne, F., Kshirsagar, A., Augustat, N., Melzig, C. A., Bremmer, F., Peters, J., & Endres, D. M. (2025). Inferring height-induced changes in postural control via inverse optimal control. *Proceedings of the Conference on Cognitive Computational Neuroscience (CCN)*.
- Koosha, T. A., Kshirsagar, A., Augustat, N., Hahne, F., Mühl, D., Melzig, C. A., Bremmer, F., Peters, J., & Endres, D. M. (2025). Staring down the elevator shaft: Postural responses to virtual heights in an indoor environment. *Proceedings of the Annual Meeting of the Cognitive Science Society*.
- Krupić, D., Grewenig, S., Volbracht, A., & Bühlhoff, H. H. (2021). Walking on a plank at high virtual heights: Effects on gait and balance. *Scientific Reports*, 11, 17304.
- Li, W., & Todorov, E. (2004). Iterative linear quadratic regulator design for nonlinear biological movement systems. *First International Conference on Informatics in Control, Automation and Robotics*, 2, 222–229.
- Markspan. (2017). *PolarBand2Isl* (version 1.0.0) [source code].
- Mergner, T., & Rosemeier, T. (1998). Interaction of vestibular, somatosensory and visual signals for postural control. *Brain Research Reviews*, 28(1–2), 118–135.
- Microsoft. (2014). Microsoft Kinect v2.
- Powell, M. J., et al. (2009). The bobyqa algorithm for bound constrained optimization without derivatives. *Cambridge NA Report NA2009/06*, University of Cambridge, Cambridge, 26, 26–46.
- Ree, M. J., French, D., MacLeod, C., & Locke, V. (2008). Distinguishing cognitive and somatic dimensions of state and trait anxiety: Development and validation of the state-trait inventory for cognitive and somatic anxiety (sticsa). *Behavioural and Cognitive Psychotherapy*, 36(3), 313–332.
- Reimann, H., & Schöner, G. (2017). A multi-joint model of quiet, upright stance accounts for the “uncontrolled manifold” structure of joint variance. *Biological Cybernetics*, 111, 389–403.
- Spartakov, R., Kshirsagar, A., Mühl, D., Schween, R., Endres, D. M., Bremmer, F., Melzig, C. A., & Peters, J. (2024). Balancing on the edge: Review and computational framework on the dynamics of fear of falling and fear of heights in postural control. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 46.
- Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11), 1226–1235.
- Wiederhold, B. K., & Bouchard, S. (2014). *Advances in virtual reality and anxiety disorders*. Springer.
- Wuehr, M., Breitkopf, K., Decker, J., Ibarra, G., Huppert, D., & Brandt, T. (2019). Fear of heights in virtual reality saturates 20 to 40 m above ground. *Journal of Neurology*, 266(Suppl 1), 80–87.
- Wuehr, M., Kugler, G., Schniepp, R., Eckl, M., Müller, F., Brandt, T., & Jahn, K. (2014). Balance control and anti-gravity muscle activity during the experience of fear at heights. *Physiology & Behavior*, 138, 277–286.