Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Cerebellar interneurons execute a low-dimensional control policy for real-time kinematic calibration

Ming-Kai Pan

emorymkpan@ntu.edu.tw

National Taiwan University College of Medicine https://orcid.org/0000-0001-5535-2714

Jou-Yu Ho

National Taiwan University College of Medicine https://orcid.org/0009-0001-7624-1578

Po-Wei Chen

National Tsing Hua University

Shun-Chi Wu

National Tsing Hua University

Chih-An Lin

National Taiwan University College of Medicine

Jye-Chang Lee

National Taiwan University https://orcid.org/0000-0003-4374-482X

Article

Keywords:

Posted Date: November 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11446111/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Cerebellar interneurons execute a low-dimensional control policy for real-time kinematic calibration

Jou-Yu Ho[1,2,3]†, Po-Wei Chen[6,7]†, Chih-An Lin[3,4]†, Jye-Chang Lee[3,5], Shun-Chi Wu[6,7], Ming-Kai Pan[1,2,3,4,5,8,9]*

[1]Taiwan International Graduate Program in Interdisciplinary Neuroscience, National Taiwan University and Academia Sinica; Taipei, Taiwan

[2]Institute of Biomedical Sciences, Academia Sinica; Taipei, Taiwan

[3]Molecular Imaging Center, National Taiwan University; Taipei, Taiwan

[4]Department and Graduate Institute of Pharmacology, National Taiwan University College of Medicine; Taipei, Taiwan

[5]Department of Medical Research, National Taiwan University Hospital; Taipei, Taiwan

[6]Department of Engineering and System Science, National Tsing Hua University, Hsinchu, 30013 Taiwan

[7]Brain Research Center, National Tsing Hua University, Hsinchu, 30013 Taiwan

[8]Cerebellar Research Center, National Taiwan University Hospital, Yun-Lin Branch; Yun-Lin, Taiwan

[9]Research Center for Developmental Biology and Regenerative Medicine, National Taiwan University; Taipei, Taiwan

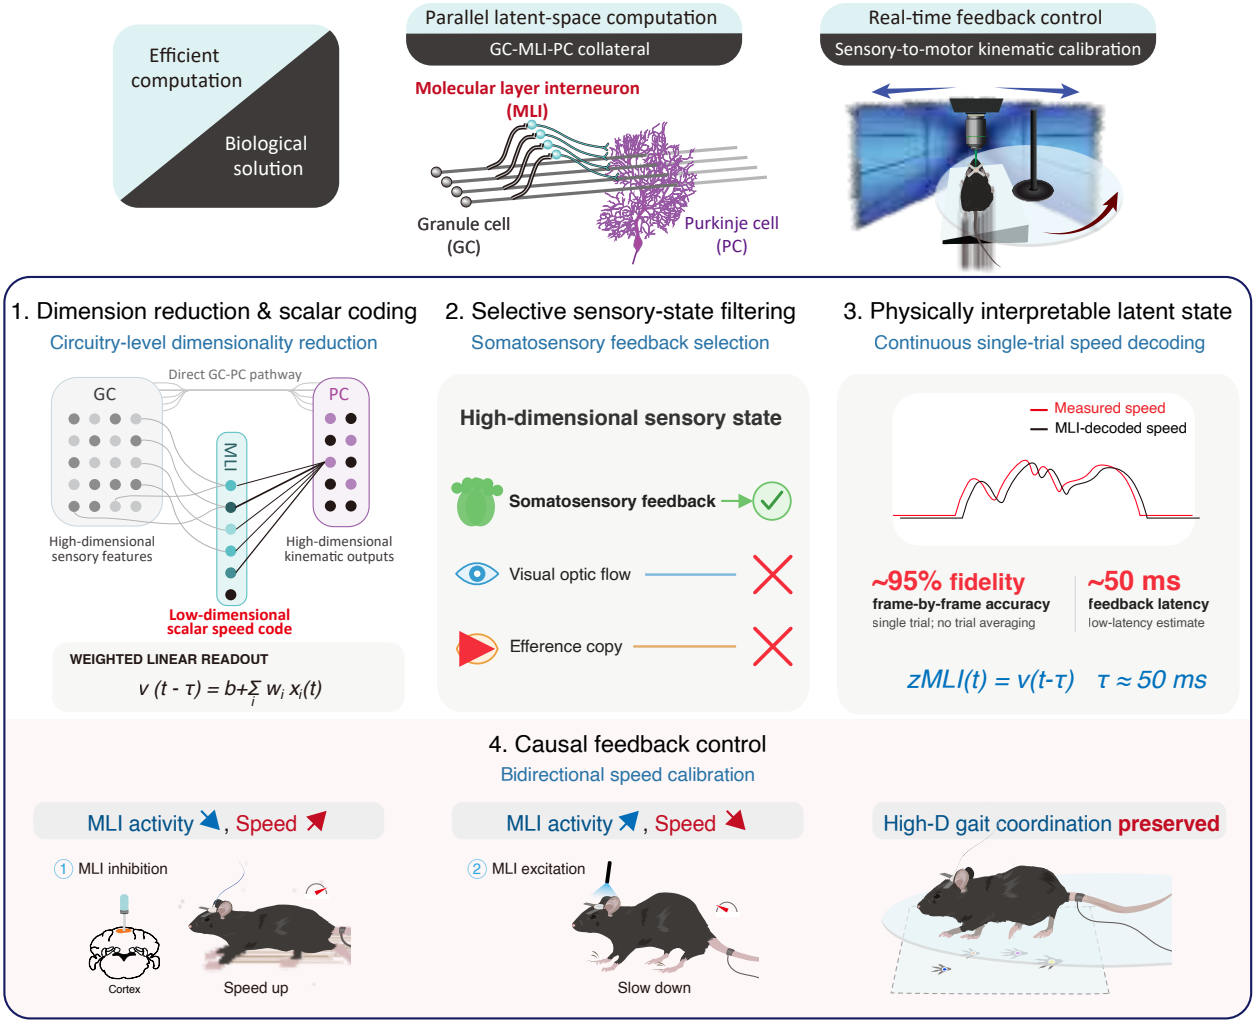
†These authors contributed equally to this work.

*Corresponding author:
Ming-Kai Pan, M.D. Ph.D.
Email: emorymkpan@ntu.edu.tw
Department and Graduate Institute of Pharmacology
National Taiwan University College of Medicine
No. 1, Sec. 1, Ren-Ai Road, Taipei 100, Taiwan
Tel: (+886)-2-23123456 ext 88236

ABSTRACT

Artificial agents for continuous control currently struggle with the computational and latency overhead of processing high-dimensional sensory data. How biological neural networks perform similar feedback-control tasks with ultra-efficiency remains unclear. Here, we reveal a neurocomputational architecture in the mammalian cerebellum that resolves this bottleneck. Using in vivo two-photon calcium imaging in awake-behaving mice, we demonstrate that the cerebellar molecular layer interneuron (MLI) network acts as a native dimension-reduction layer. The MLI population distills high-dimensional somatosensory feedback into a continuous, linear scalar variable for locomotor speed. Operating without reliance on algorithmic trial-averaging, this network achieves 95% single-trial inference accuracy with 50-ms latency, yielding a strict mathematical control policy: $zMLI(t)=v(t-\tau)$. Bidirectional optogenetic modulations demonstrate that this *interpretable* biological latent space provides immediate, causal feedback calibration of motor output. These findings unearth a fundamental biological algorithm for real-time kinematic control, providing a deterministic blueprint for physics-informed neural networks and biomimetic robotic architectures.

GRAPHICAL ABSTRACT



The cerebellar MLI network implements a low-dimensional control policy for real-time kinematic calibration. Mammalian brains achieve accurate and reliable sensorimotor control under stringent computational and energetic constraints. Here, we identify a biological circuit solution that compresses high-dimensional sensory information into a low-dimensional state for rapid kinematic calibration. High-dimensional sensory features represented in the granule cell (GC) layer are transmitted directly to Purkinje cells (PCs) to support multidimensional kinematic outputs, whereas the parallel molecular layer interneuron (MLI) collateral performs circuitry-level dimensionality reduction and extracts a scalar representation of locomotor speed through weighted population decoding. The MLI network selectively processes somatosensory feedback while rejecting visual optic flow and an internally generated efference copy of motion. The resulting population code forms a physically interpretable latent state and achieves approximately 95% frame-by-frame decoding fidelity in individual self-paced trials without trial averaging. MLI population activity lags locomotor speed by approximately 50 ms, summarized as $zMLI(t) \approx v(t-\tau)$., where $\tau \approx 50$ ms. Pharmacological suppression of MLI-to-PC inhibitory transmission accelerates locomotion while preserving high-dimensional gait coordination, whereas optogenetic excitation of MLIs decelerates locomotion. Together, these findings identify the MLI latent space as a selective, low-latency feedback controller that bidirectionally calibrates the scalar speed policy without disrupting multidimensional motor coordination.

INTRODUCTION

Achieving stable, continuous motor control in dynamically changing environments is a central challenge in both systems neuroscience and artificial machine intelligence. In modern robotics and deep reinforcement learning, artificial agents rely on computationally heavy state-space estimation, where high-dimensional sensory streams must be continuously integrated to update a motor policy. This architectural demand introduces significant latency, model complexity, and energy costs. Conversely, biological neural networks (BNNs) effortlessly execute robust continuous behaviors, such as locomotion, with minimal-latency inference and extreme energy efficiency. A major frontier in neuromorphic engineering is identifying the specific hardware-level biological architectures that solve these bottlenecks, extracting their exact mathematical rules to inspire next-generation, physics-informed AI.

The cerebellum, also known as the "small brain", is a uniquely organized, parallel-processing architecture optimized to perform this massive sensory integration[1] (**Extended Data Fig. 1**). It receives dense, multidimensional streams of proprioceptive and exteroceptive information, establishing a highly expanded, high-dimensional state space in the granule cell (GC) layer [1-3]. Crucially, for such integrated sensory feedback to calibrate movement, it must converge onto neural circuits dedicated to precise kinematic control, a hallmark of cerebellar function [4-19]. In modern artificial neural networks, integrating a low-dimensional global state variable (like speed) into a high-dimensional feature space requires computationally expensive software algorithms, such as Squeeze-and-Excitation (SE) networks or Feature-wise Linear Modulation (FiLM). We hypothesized that the unique geometry of the cerebellar cortex provides a natively optimized, hardware-level solution to this exact problem—housing a dedicated neural population that bypasses the computational burden of deep integration to act solely as a continuous, low-dimensional "squeeze" readout of a kinematic state.

Within the cerebellar microcircuit, molecular layer interneurons (MLIs) provide dense, finely tuned inhibition [20,21] to Purkinje cells (PCs), the primary output layer of the cerebellar cortex for precise kinematic output [4-7,11,13]. Sensory-driven synchronization of MLIs occurs with millisecond precision [22]. This physiological motif establishes the necessary hardware for scalable, temporally precise modulation. We theorized that the MLI population functions as the biological equivalent of a latent decoding layer, physically reducing complex granular inputs into a continuous, scalar feedback signal for real-time kinematic calibration.

Here, we combined closed-loop virtual reality, in vivo two-photon calcium imaging, and bidirectional neuromodulation during self-paced mouse locomotion to investigate whether the MLI

network computes an explicit mathematical policy for feedback control. We demonstrate that population MLI activity continuously encodes an explicit, physically meaningful entity—locomotor speed—with frame-by-frame precision on a single-trial basis, bypassing visual integration to exclusively process motion-related somatosensory feedback. By natively solving the dimensionality-reduction problem at the hardware level, this circuit avoids the uninterpretable 'black box' of artificial latent spaces; instead, it yields a strict, linear control equation. We validate this biological policy through real-time optogenetic and pharmacological perturbations, demonstrating how the MLI latent space dynamically calibrates ongoing motor output without impairing local interlimb coordination. These findings decode a fundamental biological algorithm for real-time, low-latency continuous control, providing a deterministic blueprint for biomimetic robotics and neuromorphic architectures.

RESULTS

The MLI network functions as a continuous kinematic readout in closed-loop environments

To investigate whether the cerebellar cortex computes continuous mathematical representations of movement, we designed an experimental system to monitor real-time neural dynamics during active locomotion. Head-fixed mice were placed on a rotary disc [23,24] integrated with a closed-loop virtual reality (VR) environment, wherein visual optic flow was strictly coupled to ongoing running speed (**Fig. 1a**). Real-time kinematic states were recorded via a rotary encoder, while a cranial window over the motor cerebellum provided direct optical access to the underlying circuitry.

Using a dual-labeling genetic strategy (AAV-CaMKIIα-GCaMP6f in PCP2-tTA2-GCaMP6f mice), we performed simultaneous in vivo two-photon calcium imaging to compare the dynamics of the proposed latent layer of MLIs against the primary network output layer, the PC dendrites (**Fig. 1b-d and Supplementary Fig. 1**). We discovered that continuous MLI calcium dynamics robustly tracked the instantaneous kinematic state (locomotor speed) of the mice (**Fig. 1e; Supplementary Video 1**). In contrast, PC dendritic signals exhibited significantly weaker correlations with instantaneous speed, indicative of a more complex, high-dimensional output space (**Fig. 1f**). Pearson correlation analysis confirmed that MLIs encode the speed variable with significantly higher fidelity than PCs (**Fig. 1g**). Crucially, the predictive accuracy of this kinematic representation scaled dramatically when aggregating signals: population-averaged MLI activity consistently outperformed individual neuronal trajectories in predicting real-time velocity (**Fig. 1h**). This indicates that the explicit representation of locomotor speed is not a property of isolated nodes, but rather an emergent, continuous variable derived through population-level coding.

MLI population dynamics encode delayed sensory feedback rather than feedforward predictions

Biological motor coordination requires the continuous integration of efferent control signals with incoming sensory state updates. To determine whether the MLI network functions as a predictive (feedforward) controller or a state-updating (feedback) estimator, we analyzed the temporal phase relationship between the MLI network dynamics and the real-time kinematic output (**Fig. 2a,b**). Theoretically, if the MLI network generates feedforward motor commands, its latent dynamics should precede physical gait execution; conversely, if it processes sensory state updates, the network activity must strictly lag the physical movement with a measurable processing delay (**Fig. 2c**).

We extracted the temporal trajectories from individual MLI nodes (**Fig. 2d**) and their aggregated population mean (**Fig. 2e**), computing cross-correlations against the real-time velocity stream (**Fig. 2f,g**). Individual MLI nodes predominantly exhibited delayed peak correlations relative to the physical kinematic state (**Fig. 2h**). Critically, at the population level, the unified MLI trajectory consistently lagged behind running speed by 150–500 ms (**Fig. 2i and Supplementary Fig. 2**). This temporal lag confirms that the network operates on delayed sensory state-space inputs rather than generating predictive efferent commands. Following temporal alignment (latency correction), individual MLIs showed improved correlations with the kinematic state (mean $r = 0.51 \pm 0.17$, $n =$ 9,521 MLI traces across 107 trials). However, the aggregated population dynamic consistently yielded vastly superior decoding accuracy ($r = 0.76 \pm 0.094$, $n = 107$ trials; **Fig. 2j**). These results define the MLI network as a delayed, feedback-driven kinematic estimator, where population aggregation suppresses noise to enhance coding precision. Crucially, this robust estimation was achieved on single-trial, self-paced trajectories without the need for trial-averaging, highlighting the network's capacity for real-time, online state computation during continuous behavior.

The MLI latent space selectively filters visual state-spaces to exclusively process somatosensory feedback

To rigorously test the robustness of this biological state estimator, we investigated whether the MLI network relies on multidimensional visual state-spaces or selectively filters them to isolate physical movement data. In the VR environment, we decoupled visual optic flow from physical locomotion by presenting preprogrammed visual motion while immobilizing the rotary disc to minimize voluntary movement (**Extended Data Fig. 2a**). Under these purely visual conditions, the MLI population fails to correlate with optic flow dynamics, and response correlations remained at baseline levels across paradigms (**Extended Data Fig. 2b-h**). Instead, residual MLI activity reliably tracked tail kinematics—a proxy for movement-related somatosensory input (**Supplementary Fig. 3;**

Supplementary Fig. 4a-d). Cross-correlation analysis confirmed that MLI activity lagged tail kinematics, and correlations between MLI activity and tail motion exceeded those between MLI activity and optic flow (**Supplementary Fig. 4e-g**).

To establish whether the network operates exclusively on somatosensory channels, we removed all external visual input, requiring head-fixed mice to navigate in complete darkness relying strictly on proprioceptive and tactile feedback (**Extended Data Fig. 3a**). The MLI latent space continued to robustly encode locomotor speed throughout the episodes (**Extended Data Fig. 3b-d**). The aggregated population dynamics uniformly lagged the physical kinematic state, confirming a feedback-driven computation (**Extended Data Fig. 3e-g**). Delay-corrected correlations confirmed that this coding was more accurate at the population level than at the level of individual neurons (**Extended Data Fig. 3h,i**). Furthermore, MLI activity exhibited weak correlations with individual gait phases (**Supplementary Fig. 5**), indicating that the network does not simply relay primary tactile sensations, but actively distills them into a transformed, continuous representation of locomotor speed.

To definitively uncouple these variables and exclude the possibility of artificial dissociation, we introduced a dynamic visual-motor gain mismatch, progressively altering the speed of optic flow relative to the physical locomotor speed reported by the animal's somatosensation. In the modified condition, visual gain gradually increased over 90 seconds (**Extended Data Fig. 4a,b**). Direct comparisons revealed that both individual and population MLI activity tracked locomotor speed more closely than optic flow dynamics (**Extended Data Fig. 4c-f**). A complementary gain-down manipulation produced similar results (**Extended Data Fig. 5**).

A hallmark of robust artificial control systems is the ability to selectively filter multidimensional sensory data to isolate a target state variable. Together, these results demonstrate that the MLI population performs selective state-space filtering, explicitly rejecting deceptive visual variables to compute an uncorrupted, real-time estimate of true kinematic momentum.

Direct activation of ascending hardware validates the somatosensory origin of the speed computation

In artificial control systems and biological motor control, kinematic representations can arise either from real-time sensory state updates or from predictive internal models (such as an efference copy of a motor command) [25-30]. To distinguish sensory-driven signals from an internal efference copy, we isolated the ascending sensory hardware. We transfected the dorsal column nuclei with AAV-ChrimsonR, enabling targeted optogenetic excitation of ascending somatosensory afferents in the lower brainstem (**Fig. 3a and Supplementary Fig. 6**).

Validating the activation of this sensory pathway, we recorded reliable evoked potentials in the primary somatosensory cortex following photostimulation (**Fig. 3a-e**). Crucially, optogenetic activation of these ascending afferents elicited immediate, robust responses within the MLI latent space, occurring entirely in the absence of voluntary movement and preceding any overt motor behavior (**Fig. 3f-j**). Consequently, direct perturbation of the biological sensor hardware is sufficient to drive the network. Combined with the visual-filtering data, this confirms that the MLI computation is not an internally generated predictive model, but a dedicated estimator driven explicitly by ascending somatosensory state-space inputs.

Weighted population coding generates a highly precise, deterministic kinematic policy

In artificial neural networks, latent dimension reduction relies on weighted integration rather than uniform pooling. Although simple population averaging improved kinematic state decoding (**Fig. 2 and Extended Data Fig. 3**), it assumes an equal contribution from each node and does not reflect the physiological reality of weighted synaptic integration. To test whether the MLI network natively executes a weighted dimensionality reduction, we applied mathematical linear regression to the population activity. Weighted regression substantially enhanced the decoding of locomotor speed relative to simple population means (**Fig. 4a**), indicating that the kinematic state is extracted via a distributed, non-uniform architecture across the MLI population.

To evaluate the scalability of this biological algorithm, we randomly sampled subsets of the network and computed their decoding performance (**Fig. 4b**). Repeated sampling (100 iterations per subset size) stabilized performance estimates and minimized selection bias (**Fig. 4c**). Expanding the population size progressively improved decoding across simple averaging, ranked linear regression (recruiting neurons with weaker single-cell correlations first), and randomly recruited linear regression (**Fig. 4d**). Crucially, weighted population coding consistently outperformed simple averaging, achieving ~95% frame-by-frame formulaic precision without relying on computational trial averaging. The similarity between random and ranked recruitment strategies indicates that the kinematic parameter is broadly distributed across the latent space rather than confined to a privileged subset of nodes. Across continuous self-paced trajectories, simple population means failed to achieve 90% decoding accuracy in 89.3% of trials, whereas the weighted population architectures failed in only 2 of 122 trials (1.64%) (**Fig. 4e**). Further analysis revealed that engaging just 20–30% of the weighted MLI network was sufficient to achieve robust decoding accuracy (**Fig. 4f**).

Together, these data demonstrate that the MLI population natively utilizes weighted coding to extract locomotor speed with high formulaic precision on a single-trial basis.

Weighted decoding unmasks ultra-low latency real-time feedback processing

In real-time control architectures, accurate latency estimation depends critically on the signal-to-noise ratio of the latent readout. We therefore re-evaluated the processing delay between the MLI network state and the physical locomotor speed using cross-correlation analysis. Initial estimates derived from simple uniform pooling suggested substantial processing lags of 300–500 ms (**Figs. 2i, 4g, and Extended Data Fig. 3g**). However, because simple averaging assigns equal weight across the network, it is disproportionately biased by nodes with lower decoding fidelity and unphysiological delays (**Fig. 4g**).

By contrast, implementing weighted decoding substantially enhanced signal fidelity and minimized this bias, unmasking an ultra-low temporal lag of approximately 50 ms (**Fig. 4g**). This enhanced decoding precision allowed for a more accurate derivation of the network's processing latency, revealing that the MLI state estimator tracks physical kinematics within tens of milliseconds. While the scale of the dataset optimized temporal estimation, further resolution is currently bounded by the intrinsic hardware sampling limit (30.3 frames per second, fps) of the two-photon imaging configuration, constraining the true biological feedback delay to between 33 and 66 ms.

Together, these data indicate that the cerebellar MLI network computes locomotor speed through natively weighted population dynamics, functioning as an ultra-low latency state estimator optimized for real-time, continuous kinematic control. This network distills multidimensional somatosensory inputs into a strict, deterministic kinematic policy, describable by the explicit neural formula:

$$zMLI(t) = v(t-\tau)$$

where zMLI indicates weighted MLI population activity, with a $\tau \approx 50$ ms.

The MLI network natively executes massive dimensionality reduction to isolate a scalar speed variable.

The ability to compute locomotor speed through weighted decoding implies that this kinematic variable is encoded within a highly compressed, low-dimensional subspace. To formalize the latent architecture of this network, we assessed the dimensionality of the calcium dynamics using principal component analysis (PCA). Visualization of the first three principal components revealed that the variance within the MLI population was predominantly constrained to the first one or two dimensions (**Fig. 4h**). In contrast, the variance of populational PCs, the primary output layer, remained highly

distributed across a broader, multidimensional manifold. Unbiased dimensionality quantification confirmed that the MLI latent space occupies a significantly lower-dimensional regime compared to the PC layer (**Fig. 4i-j**). Indeed, the first principal component alone accounted for 40–60% of the total activity variance in the MLI population (**Fig. 4k**).

Together, these data reveal a fundamentally optimized neurocomputational architecture: the low-dimensional MLI network operates as a critical biological information bottleneck, situated directly between the massively expanded, high-dimensional feature space of the granule cell layer and the complex, multidimensional output of the PCs (**Fig. 4l**). Thus, the MLI population acts as a continuous, low-dimensional collateral branch, natively executing extreme dimensionality reduction to distill massive granular inputs into a unified, explicit kinematic variable for the continuous calibration of PC output.

Bidirectional actuation of the MLI latent space exclusively calibrates the scalar speed policy without disrupting high-dimensional coordination

Having mathematically defined the MLI network as a low-dimensional state estimator, we tested its causal role as a feedback controller. In artificial closed-loop systems, perturbing the latent estimate of a kinematic variable dynamically forces the primary network to recalibrate its output to compensate for the perceived error. We hypothesized that if the MLI network continuously reports a scalar speed state, artificially suppressing this signal should induce a state-estimation error (an underestimation of true locomotor speed), thereby triggering compensatory acceleration.

To test this prediction, we infused bicuculline, a γ-aminobutyric acid type A ($GABA_A$) receptor antagonist, into the cerebellar cortex to the parallel MLI-to-PC inhibitory collateral (**Fig. 5a and Supplementary Fig. 7a,b**). During self-paced locomotion (**Fig. 5b**), this targeted suppression produced a marked, continuous acceleration in locomotor speed (**Fig. 5c,d; Supplementary Video 2**). Notably, the probability of voluntary stopping epochs remained unchanged (10–20%; **Fig. 5e, red arrow**), confirming that state transitions remained self-initiated rather than forcibly driven by the manipulation.

We next examined whether updating this scalar policy corrupted the broader multidimensional motor output. High-speed videography (120 fps) enabled rigorous quantification of high-dimensional limb kinematics, including stride length, stride duration, step length, and step phase (**Fig. 5f-l and Supplementary Fig. 7c-j**). Despite the profound acceleration in locomotor speed, all analyzed limb kinematic parameters were preserved (**Supplementary Video 2**). These findings indicate that the

MLI latent space exclusively updates the scalar speed variable without corrupting the high-dimensional coordination of the primary network.

To contrast this highly selective dimensional modulation with a global network disruption, we infused bicuculline bilaterally into the deep cerebellar nuclei (DCN) to suppress primary PC-to-DCN transmission, the kinematic output. This manipulation produced pronounced ataxia and a catastrophic failure of high-dimensional gait coordination (**Supplementary Fig. 8**). Rotarod and open-field assays further confirmed widespread locomotor deficits (**Supplementary Fig. 9**; **Supplementary Video 3**). Thus, while disrupting the primary multidimensional output node compromises global coordination, selective manipulation of the MLI latent space cleanly isolates and recalibrates the speed setpoint while preserving gait structure.

To definitively establish bidirectional causal control over this deterministic policy, we optogenetically actuated the MLI network in PV-Ai203 mice (**Fig. 6a**), which express the blue-light-sensitive cation channel ChroME in MLIs [31]. Acute optogenetic excitation of the MLIs produced a rapid, reversible deceleration of locomotion (**Fig. 6b-h; Supplementary Video 4**). Together, pharmacological suppression and optogenetic actuation demonstrate robust bidirectional control: attenuating the latent speed estimate drives acceleration (**Fig. 5**), whereas amplifying it induces compensatory braking (**Fig. 6**). Collectively, these results prove that the MLI network functions as an interpretable biological feedback controller, continuously applying a strict, low-dimensional kinematic calibration to high-dimensional motor outputs in real time.

DISCUSSION

Our findings identify MLIs as a specialized cerebellar feedback circuit that continuously reports the outcome of massive sensory dimensionality reduction of locomotor speed to PCs, in parallel with the high-dimensional kinematic control in the GC-PC axis. While classical motor control models often emphasize complex, high-dimensional population dynamics for motor execution and learning, we demonstrate that the nervous system also employs explicit, low-dimensional mathematical equations, akin to engineered control systems, for continuous feedback calibration. MLIs natively express movement-related somatosensory inputs as a continuous, linear sensory reference for locomotor speed, independent of visual optic flow. Through weighted population coding, the MLI network achieves frame-by-frame precision on a single-trial basis without relying on post-hoc algorithmic extraction, yielding a strict systems-level formulation: $zMLI(t) = v(t-\tau)$. This quantitative relationship provides a deterministic framework for predicting the behavioral consequences of

perturbing the feedback signal. Increasing MLI activity above the level corresponding to actual speed induces immediate compensatory slowing, whereas reducing MLI signaling produces acceleration, confirming the dynamic, formula-driven adjustment of a locomotor speed setpoint. The findings of this study are summarized in **Extended Data Fig. 6**.

The bidirectional speed control also highlights the separation between global kinematic calibration and detailed motor coordination, in which a compact state variable can influence an ongoing multidimensional motor pattern without specifying its full kinematic structure. This separation addresses a fundamental challenge in continuous control: extracting compact, task-relevant state variables while preserving the high-dimensional representations required for flexible behaviour. Artificial locomotor systems can often achieve robust control from a restricted subset of informative feedback variables rather than from the full sensory space [32], and complex observations can often be described by substantially lower-dimensional intrinsic dynamics [33,34]. The cerebellar circuit studied here suggests a biological implementation of a related principle. Classical theories proposed that mossy-fiber inputs are expanded through the granule-cell layer to form a rich basis for downstream computation [35-37], and population recordings support high-dimensional coding in granule-cell axons [1]. Our results raise the possibility that this expanded representation is distributed into pathways with different representational geometries: a comparatively rich PC representation coexists with a lower-dimensional MLI representation of locomotor speed. Expansion and compression may therefore operate in parallel within the same microcircuit, rather than forming successive stages of a single processing stream. A related balance between representational expansion and compression also support the observations in artificial recurrent networks optimized for task performance [38].

The sensory and temporal properties of the MLI signal help define the role of this compact representation. MLI activity followed rather than preceded locomotion, remained robust in darkness, preferentially tracked physical movement when visual and locomotor signals were placed in conflict, and was driven by ascending somatosensory input. These observations are more consistent with an online representation of the achieved physical state than with a purely prospective motor command. Such a signal could complement predictive cerebellar computations: forward models have been proposed to estimate the consequences of motor commands, whereas closed-loop control also requires information about the state that was actually reached [39,40]. The MLI population may therefore provide a rapidly updated sensory reference that links ongoing cerebellar computation to the physical consequences of movement.

Its position within the cerebellar circuit provides a possible mechanism for using this state information. MLIs receive parallel-fiber excitation and directly inhibit PCs, placing the compact speed representation in contact with the principal output layer of the cerebellar cortex. Recent work shows that MLI populations possess structured inhibitory and disinhibitory connectivity and can be rapidly synchronized by sensory input [20,22,41]. One possibility is that this organization supports state-dependent modulation of PC processing: MLIs may not encode the detailed motor command itself, but instead alter how the PC population responds to its inputs according to the current locomotor state. Conceptually, this resembles conditional modulation in artificial networks, where compact contextual signals regulate processing within larger feature representations, as in squeeze-and-excitation and feature-wise modulation [42,43]. Here, however, the compact state signal and the richer representation are carried by distinct but directly connected neural populations. Direct measurements of PC input–output transformations during selective MLI manipulation will be required to determine whether this interaction implements gain modulation or another form of state-dependent control.

The causal dissociation further constrains this interpretation. MLI perturbation changed locomotor speed bidirectionally while leaving much of the spatial and temporal organization of stepping intact, whereas disruption of cerebellar nuclear processing produced broader deficits in motor coordination. This suggests that the MLI pathway does not simply scale cerebellar output globally, but preferentially influences a kinematic dimension of an already coordinated movement. Previous work similarly indicates that cerebellar circuits can regulate specific movement parameters and organize multiple kinematic variables within structured population activity [44,45]. More generally, behaviorally relevant variables can occupy restricted dimensions within larger neural population spaces [46,47]. In the present case, such dimensional specialization is associated with an anatomically defined inhibitory population and remains accessible through a relatively simple weighted population readout.

These features suggest a circuit motif of potential relevance to machine intelligence. Contemporary continuous-control architectures typically transform high-dimensional observations before extracting compact state or contextual variables needed for action. The cerebellar organization examined here suggests a complementary strategy in which richer and compressed representations are maintained in parallel, with the compact feedback pathway positioned to directly modulate the richer output pathway. We describe this as factorized processing across circuit elements. Related work in machine learning and neuromorphic engineering has emphasized selective feedback-state representations, task-dependent changes in representational dimensionality, and the benefits of assigning different computational functions to interacting pathways [32,48,49]. Whether the cerebellar organization

identified here provides analogous computational or energetic advantages remains to be tested. Several biological questions also remain unresolved, including the true temporal precision of the MLI signal, where along the ascending pathway locomotor speed is first constructed, and how distinct MLI subtypes contribute to the population representation. Nevertheless, our results suggest that efficient continuous control may depend not simply on compressing information, but on maintaining rich and compact representations in parallel and allowing a task-relevant feedback state to calibrate the richer control pathway.

MATERIALS and METHODS

Animals

All experimental procedures were approved by the Institutional Animal Care and Use Committee (IACUC) of National Taiwan University (Protocol No. NTU-112-EL-00068) and conducted in accordance with institutional and national guidelines. Mice were housed in the central and satellite animal facilities of National Taiwan University under a reversed 12-hour light/dark cycle, with food and water provided freely. Both male and female mice aged 2 to 6 months were used. *Pcp2-Cre* mice (JAX #004146) crossed with Ai148 (TIT2L-GC6f-ICL-tTA2; JAX #030328) were used for two-photon calcium imaging and optogenetic manipulations involving the dorsal column nuclei. *PV*-Ai203 (*PV-Cre* × Ai203; JAX #008069, #037939) mice were used for MLI-targeted optogenetic excitation via a cranial window. Wild-type C57BL/6J mice (JAX #000664) were used for pharmacological manipulations. Additional surgical procedures were performed as required for each experimental paradigm.

Viral vectors

All adeno-associated viral (AAV) vectors were obtained from Addgene (https://www.addgene.org/) and stored at -80 °C until use. For simultaneous calcium imaging of Purkinje cells (PCs) and molecular layer interneurons (MLIs), *AAV9-CaMKIIα-GCaMP6f* (Addgene #100834-AAV9; $\geq 1\times10^{13}$ vg/mL) was injected into the cerebellar cortex of *Pcp2-Cre* × Ai148 mice. In these animals, Cre-dependent GCaMP6f expression selectively labeled PCs, whereas *CaMKIIα* promoter–driven viral expression labeled MLIs. This strategy enabled concurrent visualization of PC dendrites and MLIs within the same two-photon imaging field. For sensory optogenetic stimulation, *pAAV-Syn-ChrimsonR-tdTomato* (Addgene #59171-AAV9; $\geq 1\times10^{13}$ vg/mL) was bilaterally injected into the dorsal column nuclei to express ChrimsonR in ascending somatosensory afferents. Viral aliquots were thawed once and injected at calibrated volumes using defined stereotaxic coordinates.

Stereotaxic surgery for viral injection

Mice were anesthetized with 1-2% isoflurane at a flow rate of 1 L/min and secured in a stereotaxic frame (Model 930, Kopf Instruments, Los Angeles, CA, USA) for all surgical procedures. All viral injections were performed using a syringe pump system (Hamilton Company, Reno, NV, USA) with either a stainless-steel injection needle or a pulled-glass pipette, depending on the target site. For two-photon calcium imaging, *AAV9-CaMKII-GCaMP6f-WPRE-SV40* was injected into lobule V of the cerebellar vermis (2.5 μL per site; AP: -6.2 mm, ML: 0 mm, DV: -0.3 and –0.15 mm from dura) using a 6 mm stainless-steel injection needle (33 gauge, Small Hub Removable Needle (RN), PN: 7803-05, Hamilton) at 0.5 μL/min. For optogenetic activation of ascending somatosensory afferents, *AAV9-Syn-ChrimsonR-tdTomato* was bilaterally injected into the dorsal column nuclei (0.25 μL per site; AP:

obex level, ML: ±0.25 mm, DV: -0.15 mm from surface) using a pulled-glass pipette at 0.075 µL/min. Injections were performed sequentially on each side to avoid major blood vessels. After injection, pipettes or needles were left in place for 10 min to allow viral diffusion before withdrawal. To confirm functional activation of ascending sensory pathways, stainless-steel miniature screws (Antrin Miniature Specialties; 0.089″ diameter, 0.0625″ length) were implanted epidurally over bilateral primary somatosensory cortex (S1; AP: -0.02 mm, ML: ±1.75 mm) and secured with dental cement (Super Bond, Sun Medical Co., LTD). An optical fiber (200 µm core, FT-200-UMT, Thorlabs) was inserted beneath the meningeal membrane at the obex to ensure unobstructed illumination during stimulation and was fixed in place with transparent dental acrylic, allowing visual confirmation of light transmission. For pharmacological interventions, a single cannula (Plastics One, C315GMN) was implanted into the cerebellar cortex (AP: -6.2 mm, ML: 0 mm, DV: -0.3 mm), or bilateral cannulae were implanted into the deep cerebellar nuclei (AP: -6.37 mm, ML: ±0.7 mm, DV: -2.12 mm). All implants, including screws, optical fibers, and cannulae, were secured with either dental cement or transparent dental acrylic as specified above. Following surgery, mice were housed individually and allowed a two-week recovery period before behavioral or imaging experiments.

Two-photon calcium imaging

Brain craniotomy. The surgical procedure was mostly similar to that described above, except for the creation of a cranial window. After removal of the overlying connective tissue and muscle, a 3 × 3 mm craniotomy was drilled over the cerebellar vermis, centered at 6 mm posterior to bregma to target lobules V-VI. A No. 1.5 glass coverslip (Deckgläser; 3 mm diameter) was placed over the opening and sealed with cyanoacrylate adhesive (DPX, Sigma-Aldrich), reinforced with dental cement (Super Bond C&B). A custom-fabricated titanium headplate was affixed to the skull for head fixation during imaging. Following surgery, mice were housed individually and allowed a 2-week recovery period before experiments.

Imaging hardware and system. Two-photon calcium imaging was performed on a custom-built microscope equipped with a mode-locked Ti: sapphire laser (MaiTai, Spectra Physics) tuned to 940 nm, delivering 150-180 mW at the sample depending on imaging depth and expression level. Scanning was controlled by custom software (JadeBio System 4.0, Southport) driving galvanometer mirrors. Images were acquired with a 16×/0.8 NA long-working-distance water-immersion objective (CFI LWD Plan Fluorite, 3.0 mm WD, Nikon) at 30.3 fps, with a field of view (FOV) of 503 x 503 μm^2. Recordings targeted the superficial molecular layer (~100 µm below the cranial window), capturing both MLIs and Purkinje cell (PC) dendrites.

Calcium signal extraction and analysis. Raw two-photon imaging data were first corrected for lateral motion using the Image Stabilizer plugin in FIJI (ImageJ, https://imagej.net/). Regions of interest (ROIs) corresponding to MLIs and nearby PC dendrites were manually delineated based on MLI soma and PC dendrite morphology and their fluorescence dynamics in the time-lapse movies. Fluorescence time-series from each ROI were extracted using custom LabVIEW pipelines and analyzed in MATLAB (R2024b, MathWorks). Calcium activity was quantified as relative fluorescence change (ΔF/F), defined as:

$$\Delta F/F = (F_t - F_0)/ F_0$$

where F_t denotes fluorescence at time t, and F_0 is the mean fluorescence across the entire recording. For consistency, all ΔF/F traces were analyzed at their native temporal resolution (frame rate: 30.3 fps) without any smoothing or filtering prior to subsequent analysis.

Correlation analysis. To determine how MLI calcium activity encoded behavioral variables, we implemented an analytical pipeline consisting of cross-correlation followed by temporal alignment and correlation measurement. The analysis proceeded through the following steps:

Preprocessing: All ΔF/F traces were z-score normalized within each trial to standardize variance across neurons and sessions. Behavioral signals, including locomotor speed, virtual reality (VR)-defined optic flow, and tail acceleration, were acquired at 30 kHz and downsampled to match the imaging frame rate (30.3 fps) for precise alignment with neural data. For population-level analysis, a population-mean activity trace was computed within each trial by averaging z-score-normalized ΔF/F traces across all recorded MLIs, such that each neuron contributed equally to the population signal.

Cross-correlation: We computed pairwise cross-correlations between each ΔF/F trace and the behavioral signal of interest using MATLAB's xcorr() function. Cross-correlations were calculated at single-frame resolution over the full temporal window to capture both leading and lagging relationships. For each neuron-behavior pair, the lag corresponding to the peak absolute correlation coefficient was identified, providing an estimate of the temporal offset at which the neuron's activity was most strongly related to the behavior. For population mean activity and weighted linear regression, the averaged or weighted activity was computed first before performing cross-correlation, to ensure all neurons shared the same temporal delay.

Delay correction: Each neuron's ΔF/F trace was then shifted in time by its optimal lag (determined above) to align neural activity with the behavioral signal. This step corrected for systematic delays due to calcium indicator kinetics and neural processing latencies, ensuring that peaks in neural activity were aligned with the corresponding behavioral events.

Correlation analysis: Using the lag-corrected fluorescence traces, we computed Pearson's correlation coefficients (r) between each neuron's activity and the behavioral signal via MATLAB's corrcoef() function. This yielded a measure of signal fidelity for each cell, reflecting how closely the delay-adjusted neural activity tracked the behavioral variable. Correlation coefficients were calculated on a trial-by-trial basis for each cell and for the population mean activity.

Shuffle control: To assess the specificity of neural-behavioral coupling, control analyses were performed by temporally shuffling the behavioral signals within each trial (running speed, optic flow, or tail movement) using a full random permutation (MATLAB randperm function), thereby destroying their temporal structure while preserving their amplitude distribution. Neural ΔF/F traces were left unchanged. The identical preprocessing, cross-correlation, delay-correction, and Pearson correlation procedures were applied to the shuffled datasets. This ensured that any structured temporal alignment or correlation observed in the real data did not arise from the analytical pipeline itself.

Behavioral paradigms

For all locomotor experiments, the mice performed self-paced running and stopped at will, making the running kinematics *unique* in each trial.

In experiments involving locomotion, mice navigated a custom-built rotary disc that converted limb-generated torque into disc rotation, enabling self-paced locomotion while head-fixed. Depending on the paradigm, locomotion was coupled to a VR environment, performed in darkness, or presented with manipulated visual flow. Locomotor speed was measured by a rotary encoder (600 ticks/rev; disc radius: 11.5 cm). Encoder ticks were digitized at 20 Hz via a National Instruments acquisition device for real-time updating of the VR environment. The instantaneous linear speed of the disc (cm/s) was derived from the encoder signal. For synchronization with neural recordings, the speed signal was simultaneously output as an analog voltage (20 rpm = 1 V calibration) and acquired at 30 kHz using the Open Ephys system. TTL pulses marking trial onsets and offsets triggered an LED visible in the calcium imaging field, enabling frame-level alignment of calcium signals acquired at 30.3 Hz. The same LED triggers were simultaneously captured by a high-speed camera with a frame rate of 120 fps. For analysis, locomotor speed traces recorded via Open Ephys (oversampling at 30 kHz) were downsampled to 30.3 Hz to match the two-photon imaging frame rate. All behavioral signals, including locomotor speed, trial event markers, 3-axis accelerometer data, and TTLs, were

acquired at 30 kHz, providing sub-millisecond synchronization precision across calcium imaging data, gait kinematics from the high-speed camera, and moving speed from the rotary disc.

Closed-loop virtual reality. In the closed-loop condition, the virtual scene was dynamically updated according to the animal's running speed. Mice could voluntarily initiate and sustain locomotion on the disc, with their movements directly coupled to optic flow in the VR display (i.e., running faster caused the visual scene to translate faster). This paradigm provided naturalistic, congruent sensorimotor feedback under head-fixed conditions, enabling direct comparison of locomotor kinematics with cerebellar neural activity during self-motion.

Pure VR playback. To dissociate visual input from locomotor output, we employed a pure playback condition in which the rotary disc was locked in place so the mice could not run freely. Predetermined VR scenes were replayed to the tested mouse, while a 3-axis accelerometer affixed to the tail root provided an independent measure of residual body/tail movements. This paradigm isolated visual optic flow stimulation from running-related locomotion and allowed us to test for visually driven responses.

Dark-field running. To eliminate visual feedback while preserving locomotor activity, mice were allowed to run on the disc in complete darkness. The VR projection was turned off, and the animal ran at will with no visual cues, using the same hardware setup.

Closed-loop VR with bidirectional gain manipulation. To test whether MLI activity encodes actual locomotor speed or visually perceived speed, we implemented a bidirectional gain manipulation in the closed-loop VR. In this paradigm, the visual scene was updated in real time by multiplying the disc-derived running speed by a gain factor. In ramp-up trials, the gain was linearly increased from 1.0× to 2.0× over 90 s within each trial, producing progressively faster optic flow despite real-time running speed. In ramp-down trials, the gain was linearly decreased from 2.0× to 1.0× over 90 s, yielding progressively slower optic flow while the animal maintained its locomotor output. Disc mechanics and encoder signals were unaffected by gain manipulation, creating a controlled mismatch between physical movement and visual feedback. This gain-manipulation paradigm directly probed the sensitivity of MLIs to mismatches between self-generated locomotion and visual optic flow.

Optogenetic Stimulation of the dorsal column sensory nuclei

Optogenetic stimulation and S1 local field potential recordings. Optogenetic stimulation was performed using a diode-pumped solid-state laser (589 nm; Cobolt) controlled by a multifunction I/O device (NI 782258-01, National Instruments), which enabled high-speed digital triggering and linear modulation of output. Light was delivered through an optical fiber ferrule (Thorlabs) implanted above the nucleus gracilis. Laser power at the fiber tip (20-40 mW) was calibrated daily with a power meter (Thorlabs) and adjusted individually for each mouse. For evoked potential (EP) analyses, stimulation consisted of 589-nm light pulses delivered at 2 Hz with a 2% duty cycle (10 ms ON, 490 ms OFF), producing 600 pulses over each 5-min stimulation block. The protocol comprised a 2-min baseline, 5-min stimulation, 2-min rest, 5-min dummy stimulation (laser powered but shutter closed), and a final 2-min rest. TTL trains and LFP signals were acquired at 30 kHz using a DAQ device (Cerebus, 349 BlackRock microsystem) for subsequent offline analysis.

Somatosensory optogenetic stimulation and MLI imaging. In a separate paradigm, the same 589-nm laser system was used to stimulate ascending proprioceptive afferents from the sensory nuclei during two-photon calcium imaging of MLIs, to test whether MLI activity reflected proprioceptive signals evoked by sensory stimulation. Stimulation consisted of trains of light pulses (13 Hz, 25% duty cycle, 4 s duration) followed by 16-s intertrial intervals, repeated for five epochs per trial. Two-photon imaging of cerebellar MLIs was performed simultaneously with a microscope operating at 23 Hz. Behavioral tracking (rotary disc speed and tail accelerometry) and TTL markers denoting trial onset and offset were recorded with Open Ephys (30 kHz) using the same procedures as in closed-

loop VR experiments. Laser pulse trains (13 Hz in 4 s) representing stimulation events were simultaneously logged on a dedicated Open Ephys channel.

Event-related analysis of MLI signals.

Raw laser TTL pulses were continuously recorded with Open Ephys (30 kHz) together with locomotor speed signals and the two-photon imaging frame clock. Each behavioral trial (~100 s) contained six 4-s stimulation trains delivered at 13 Hz. TTL rising edges were detected by thresholding at 50 % of the peak TTL amplitude; successive pulses separated by < 0.5 s were grouped into a single train, whereas inter-train intervals > 0.5 s defined distinct events. The onset of each train (first pulse) was mapped to the nearest two-photon imaging frame (frame rate: 23 fps) using the synchronized Open Ephys frame index. Peri-event epochs (−5 to +15 s) of MLI ΔF/F activity were extracted and z-scored relative to each event's own -1 to 0 s baseline. Within each event, normalized MLI traces from all identified interneurons in the field of view were averaged to yield an event-mean trace. Figures display the mean ± SD of these trial-averaged ΔF/F traces across trials. Onset latency was defined as the first frame exceeding 1.5–2.0 baseline SD for ≥ 2 consecutive frames (~87 ms at 23 Hz); trials without suprathreshold responses were assigned NaN latencies. Response amplitude was quantified as the root-mean-square (RMS) of baseline-subtracted activity within −1 to +5 s. Statistical comparisons were performed on trial-averaged values using two-sided Wilcoxon signed-rank tests.

Bicuculline infusion

Rotary disc. To examine the contribution of cerebellar inhibition to locomotor control, the $GABA_A$ receptor antagonist bicuculline methiodide (Cat. No. 2503, Tocris) was infused into either the cerebellar cortex or the deep cerebellar nuclei (DCN) during head-fixed locomotion. For cortical infusions, a cannula was inserted into lobule V of the cerebellar cortex (AP -6.2 mm; DV -0.3 mm). Bicuculline (100 μM, 1 μL) was delivered at 0.2 μL/min, and infusion was terminated at 5 minutes or immediately upon the onset of a robust behavioral response, typically a sudden acceleration in running speed. For DCN infusions, bilateral cannulae were implanted at AP -6.37 mm, ML ±0.7 mm, and DV -2.12 mm from the dura. To preserve stepping ability for gait analysis, the bicuculline concentration was reduced (20 μM, 1 μL per side) and infused at the same rate (0.2 μL/min). Infusion was terminated at 5 minutes or upon the emergence of ataxia or marked gait impairment, indicating effective disruption of cerebellar output while maintaining analyzable locomotor behavior.

Rotarod assessment. Motor endurance was assessed using a Rotarod apparatus equipped with a brushless motor controlled by myRIO hardware and LabVIEW software (National Instruments). Mice were placed on the rotating rod, which accelerated linearly from 5 to 40 rpm over 5 min. Retention time, defined as the latency to fall, was recorded. Each mouse performed three trials per day, and the longest latency among the three trials was used as the performance measure. For pharmacological manipulation, bicuculline (100 μM, 1 μL) was infused into bilateral DCN at 0.2 μL/min prior to testing. Rotarod performance was compared between baseline sessions and sessions following infusion.

Open-field assay. Mice were placed on a customized 15 × 22 cm platform equipped with cameras capturing front and top views at 30 fps. Locomotion was tracked from the top view, with body positions extracted using DeepLabCut. Distances and velocities were computed every 5 frames with a 1-frame step, yielding a 4-frame overlap for smoothing. Movement episodes longer than 10 s were selected for analysis. Bicuculline was infused using the same protocol as in the Rotarod assay.

Gait analysis

For all paradigms involving disc locomotion, head-fixed mice were placed on a freely rotating transparent acrylic plate (diameter: 11.5 cm) mounted above a 45° angled mirror, allowing a side-

mounted camera to capture the reflected view of all four limbs simultaneously. Mice walked at self-paced speeds, and high-speed video was acquired at 120 fps (UI-3040CP, IDS) for gait quantification. Paw trajectories were tracked offline using DeepLabCut, an open-source, markerless pose estimation toolbox based on deep learning.

Because mice occasionally rotated their torsos or shifted hindlimb posture on the disc, quantitative analyses were restricted to forelimb kinematics. For bicuculline infusion experiments targeting the cerebellar cortex and DCN, three parameters were extracted: stride length, step length, and step phase. Stride length was defined as the distance between toe-off and the subsequent initial contact of the same limb, while step length was defined as the distance advanced between successive contralateral paw placements (e.g., right forelimb to left forelimb). For each mouse, mean values and standard deviations (SD) were calculated across baseline and infusion conditions, and variability was further expressed as the coefficient of variation (CV = SD/mean). Step phase analysis assessed interlimb coordination by transforming forelimb trajectories with the Hilbert method to derive instantaneous phases. The phase of each step was defined as the Hilbert phase of one forelimb at the moment of contralateral contact. Phase stability was quantified by the polarity index (PI) (0 = random distribution, 1 = perfect phase-locking), and interlimb coordination variability was assessed by the SD of the phase distribution.

Gait and MLI phase analysis

Gait trajectories for the front-left (FL) and front-right (FR) limbs were extracted from synchronized video recordings and resampled to 30.3 Hz. The mean (FR+FL)/2 and difference (FR−FL)/2 traces were computed to represent the shared and alternating components of limb motion, respectively. Each trace was z-scored, and its instantaneous phase was obtained from the analytic signal via Hilbert transform. To linearize polarity and avoid discontinuities caused by phase wrapping ($-\pi$ to π), the cosine of the Hilbert phase [$\cos(\phi)$] was used for correlation analysis. For each trial, phase correlations were computed between MLI calcium traces and each gait component (FL, FR, Mean, Diff). Two reference controls were generated to benchmark rhythmic coupling: a jittered control ($ctrl^{+}$) created by adding small Gaussian noise ($\sigma = 0.5$) to the gait trace, preserving its structure but perturbing precise alignment; and a phase-shuffled control ($ctrl^{-}$) produced by randomizing phase while maintaining the original power spectrum. Welch power spectral density (PSD) estimates were used to confirm dominant step-frequency peaks (~3 Hz) in gait traces and broadband, non-periodic activity in MLI signals. To assess within- and cross-signal rhythmicity, peaks were detected in each z-scored trace, and the instantaneous Hilbert phase was sampled at those peak times. The PI, the resultant vector length of unit-phase vectors, quantified phase concentration, and the circular SD measured dispersion. Self-peak analyses evaluated phase regularity within each signal, while gait-to-MLI analysis sampled MLI phase at gait-peak times to test cross-signal phase locking. Statistical inference was performed at the trial level using two-tailed Welch's t-test with Benjamini–Hochberg false-discovery rate (FDR) correction.

Optogenetic stimulation and locomotor speed analysis

PV-Cre × Ai203 mice were implanted with a 3 × 3 mm cranial window over cerebellar lobules V/VI and head-fixed on a rotary-encoder treadmill. A 473-nm DPSS laser (CNI Laser, China) was coupled to a multimode fiber above the window to illuminate the molecular layer. Each trial consisted of a 20-s baseline, followed by 15 repeated optogenetic epochs (1-s Light-ON / 3-s Light-OFF), and a 20-s post-stimulation period. Rotary-encoder signals and laser TTL pulses were digitized at 30 kHz using Open Ephys. Running speed was extracted from encoder voltage and segmented into four analysis windows: Baseline, Light-ON, Light-OFF, and Post-stimulation. For peri-event analysis, speed traces were aligned to each TTL onset and analyzed within a −3 to +4 s window. Percent change in speed was calculated relative to the 20-s baseline period. Modulation during Light-ON and Light-OFF was quantified using a paired, sign-based effect size that captures the consistency of speed modulation

across repeated epochs. Specifically, for each stimulation epoch, speed during the ON or OFF window was compared to its own peri-epoch baseline defined as the final 3 s preceding light onset (-3 to 0 s). Positive values indicate consistent speed suppression relative to baseline, whereas negative values indicate facilitation. To assess stimulation-locked temporal entrainment, the power spectral density of the running speed was computed using Welch's method (0–2 Hz) after downsampling, focusing on power near 0.25 Hz, corresponding to the 1-s ON / 3-s OFF stimulation cycle.

Linear regression decoding and recruitment analyses

Fluorescence traces from MLIs were normalized to unit range [0,1] to equalize per-neuron dynamic range for weighted population decoding, and locomotor speed was used as the continuous target. Decoder accuracy was quantified as the Pearson correlation coefficient between predicted and actual speed on held-out data.

Population decoder analysis. Two strategies were compared. In the population mean decoder, the average activity across the recruited neurons was directly correlated with speed. In the weighted linear regression decoder, coefficients were estimated based on the selected subset of MLIs, using the ordinary least squares method without an intercept to predict locomotor speed.

Random recruitment. To test how population size influenced decoding, we implemented a random recruitment procedure. At each recruitment level (k% of the recorded MLIs), [kN] neurons were randomly selected without replacement within each repetition and used to train the linear regression decoder. This sampling-and-training procedure was repeated 100 times with independent random seeds, and the mean Pearson correlation across repetitions was plotted as a function of recruitment percentage. All computations were parallelized across CPU cores using the *joblib* library in Python. Convergence analysis showed that ~100 repeats were sufficient: increasing the number of shuffles to 200-1000 had a negligible impact on mean correlation or confidence intervals.

Ranked recruitment. To estimate a lower bound of population coding accuracy, neurons were ordered by feature importance, and subsets were constructed by progressively removing the most informative features. Feature importance scores were obtained from a Random Forest regressor (RandomForestRegressor, 50 trees, Friedman MSE criterion) trained to fit locomotor speed as a function of full MLI activity. The resulting importance scores were then used to rank neurons for the ranked recruitment analysis. At each step, the top k% most informative neurons were removed, and the remaining (100 – k)% of the population was used to retrain the linear regression decoder. Thus, a condition labeled as 10% indicates that the least informative 10% of neurons were retained, whereas "90%" indicates that 90% of the original population remained after removing only the top 10%. Linear regression was retrained on the remaining set at each step. Because the ranking is deterministic once importance scores are computed, no additional shuffling was required.

Population manifold and dimensionality analysis

Population manifold structure was quantified using principal component analysis (PCA) applied to neuronal population activity matrices. For each trial, neural activity was arranged as an N×T matrix, where N denotes the number of neurons and T the number of time points. ΔF/F traces were z-scored within each neuron across time prior to PCA. PCA was performed on the transposed activity matrix (T×N) using MATLAB's pca() function so that principal components captured covariance structure across neurons.

For visualization of population geometry, neuronal activity was projected onto the first three principal components (PC1–PC3). Each point in this three-dimensional space represents the instantaneous population activity state across neurons at a given time point. Because the number of recorded MLIs and PCs differed across experiments, neuron-number–matched subsampling was used to control for sampling bias. For each trial, neurons were randomly subsampled from each population to match the

size of the smaller population. PCA was recomputed for each subsampled dataset, and this procedure was repeated 100 times. For each subsampling iteration, the variance explained by each principal component was obtained from the eigenvalues, and cumulative variance explained was computed. The resulting cumulative variance curves and dimensionality estimates were then averaged across subsampling iterations to obtain representative population statistics while minimizing biases arising from differences in the number of recorded neurons. To enable direct comparison between populations, cumulative variance curves were aligned using a shared dimensional range defined by the minimum available dimensionality across the two populations, and curves were truncated to this shared dimensional axis prior to averaging. Population dimensionality was defined as the minimum number of principal components required to explain 80% of the total variance. Dimensionality estimates were averaged across subsampling iterations to obtain representative values for each trial.

Immunohistochemistry

Histological validation was performed after the completion of behavioral tasks. Mice were deeply anesthetized (5% isoflurane) and transcardially perfused with PBS, followed by post-fixation of 4% paraformaldehyde (PFA; Bionovas Biotechnology Co., Ltd., Toronto, Ontario, Canada) overnight in 4% at 4 °C. PFA-fixed brains were sectioned at 60 μm and processed for immunohistochemistry. Free-floating sections were blocked for 1 h at room temperature in buffer containing 5% normal donkey serum and 0.3% Triton X-100 in PBS, followed by incubation with primary antibodies overnight at 4 °C and secondary antibodies for 2 hr at room temperature. Antibodies were diluted in the same blocking solution. For MLI labeling in cerebellar cortex, mouse anti-Parvalbumin (1:1000, Millipore MAB1572) and goat anti-mouse IgG–Alexa Fluor 594 secondary antibody (1:1000, Jackson ImmunoResearch) were used. For dorsal column nuclei (nucleus gracilis and cuneatus) staining, guinea pig anti-VGLUT2 (1:500, Synaptic Systems) and donkey anti-guinea pig Alexa Fluor–conjugated secondary antibody (1:1000, Thermo Fisher) were used. DNA was counterstained with 4',6-diamidino-2-phenylindole (DAPI; Thermo Fisher Scientific, Waltham, MA, USA). After three PBS washes (3 × 10 min), sections were mounted with antifade medium. Images were acquired using a Nikon Eclipse 80i microscope equipped with 4× and 10× objectives or a Leica SP8 confocal microscope for fluorescent visualization.

Statistical analysis

For experiments involving multiple levels of data acquisition, statistical analyses were tailored to the specific physiological scale of each outcome measure, ranging from single-cell behavior correlations to trial-based decoding accuracy. Statistical tests were selected according to the level of inference.

For comparisons aimed at evaluating differences between distributions of measurements (e.g., cell-level versus trial-averaged metrics, or real versus control or shuffled datasets), statistical significance was assessed using two-tailed Welch's t-test, which does not assume equal variances or equal sample sizes between groups. This approach was applied to unpaired comparisons at the cell or trial level, enabling direct evaluation of population-level differences while preserving the full structure and variability of the underlying distributions. In contrast, for interventional analyses, where the primary objective was to characterize acute, event-related responses to manipulation (e.g., optogenetic stimulation, latency measurements, or pharmacological interventions), statistical tests were applied at the trial or mouse level, comparing experimental epochs to their corresponding baseline periods within the same trial or subject. This approach was utilized to capture high-resolution circuit dynamics and the reliability of manipulations across repeated trials. Paired comparisons were evaluated using two-tailed Wilcoxon signed-rank tests. For multiple comparisons, p-values were adjusted using the Benjamini–Hochberg false discovery rate (FDR) procedure ($q < 0.05$).

Descriptive statistics are reported as mean ± SD to reflect the natural biological variability between subjects, which is more accurately represented by the SD than the SEM in large datasets. Significance

levels were defined as: *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$, and ****$p < 0.0001$. Statistical procedures were performed using Prism (GraphPad), MATLAB (MathWorks), or custom-written scripts. Specific tests are indicated in the corresponding figure legends.

ACKNOWLEDGMENTS

We thank the support from NTU Advanced Biomedical research Center, Taiwan Centers of Excellence (NTU ABRC TCE), the Excellent Translational Medicine Research Projects of National Taiwan University College of Medicine and National Taiwan University Hospital, and the technical assistance of the First Core Labs in the National Taiwan University College of Medicine. We thank Sun Medical Co. LTD, Japan, for providing high-quality adhesives for surgical implants and the AAV Core Facility of Academia Sinica in generating recombinant AAV.

Funding:

This research was supported by National Science and Technology Council (NSTC) in Taiwan (grants 109-2326-B-002-013-MY4, 112-2628-B-002-011, 113-2628-B-002-002, and 114-2628-B-002-001 (to M.-K.P.)), National Taiwan University College of Medicine (grant no. NTUMC 110C101-011 (to M.-K.P.)), National Taiwan University Hospital (grants 113-E0001, 114-E0002 and 115-E0002 (to M.-K.P.)), Academia Sinica (grants. AS-TM-112-01-02 and AS-TM-114-01-02), National Health Research Institutes (grants NHRI-EX113-11303NI and NHRI-EX114-11303NI (to M.-K.P.)).

Author contributions:

J.-Y.H. and M.-K.P. conceived the study and designed the experimental strategy. J.-Y.H. collected the primary datasets, performed the core analyses, conducted immunostaining, and generated the figures. P.-W.C. and S.-C.W. performed the population-coding analysis. C.-A.L. conducted the control pharmacology experiments and immunostaining. J.-C.L. developed the hardware setting for data acquisition. J.-Y.H. drafted the manuscript, with editorial input from M.-K.P. All authors revised the manuscript. M.-K.P supervised the project.

Competing interests:

All authors declare that they have no competing interests.

DATA AND MATERIALS AVAILABILITY:

All data supporting the findings of this study are available in the main text or the supplementary materials. Statistical data underlying the analyses are provided in Data File S1. Raw datasets are deposited in Zenodo: https://doi.org/10.5281/zenodo.22770237. Codes for data analysis are deposited in Zenodo: https://doi.org/10.5281/zenodo.18722726.

SUPPLEMENTARY MATERIALS

Extended data Figs. 1-6

Supplementary Figs. 1-9

Supplementary Videos 1-4

Data File S1.xlsx for raw data points in main and supplementary figures

Data File S1.mat for big data arrays in main and supplementary figures.

REFERENCES and NOTES

1 Lanore, F., Cayco-Gajic, N. A., Gurnani, H., Coyle, D. & Silver, R. A. Cerebellar granule cell axons support high-dimensional representations. *Nat Neurosci* **24**, 1142-1150, doi:10.1038/s41593-021-00873-x (2021).

2 Holla, M. R., Brown, S. T. & Raman, I. M. Effects of Short-Term Synaptic Plasticity in Feedforward Inhibitory Circuits on Cerebellar Responses to Repetitive Sensory Input. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **46**, doi:10.1523/jneurosci.1850-25.2025 (2026).

3 Chabrol, F. P., Arenz, A., Wiechert, M. T., Margrie, T. W. & DiGregorio, D. A. Synaptic diversity enables temporal coding of coincident multisensory inputs in single neurons. *Nat Neurosci* **18**, 718-727, doi:10.1038/nn.3974 (2015).

4 Herzfeld, D. J., Kojima, Y., Soetedjo, R. & Shadmehr, R. Encoding of action by the Purkinje cells of the cerebellum. *Nature* **526**, 439-442, doi:10.1038/nature15693 (2015).

5 Liu, C. W. *et al.* The cerebellum shapes motions by encoding motor frequencies with precision and cross-individual uniformity. *Nat Biomed Eng* **9**, 1952-1971, doi:10.1038/s41551-025-01409-5 (2025).

6 Becker, M. I. & Person, A. L. Cerebellar Control of Reach Kinematics for Endpoint Precision. *Neuron* **103**, 335-348.e335, doi:10.1016/j.neuron.2019.05.007 (2019).

7 Calame, D. J., Becker, M. I. & Person, A. L. Cerebellar associative learning underlies skilled reach adaptation. *Nat Neurosci* **26**, 1068-1079, doi:10.1038/s41593-023-01347-y (2023).

8 Wagner, M. J. *et al.* A neural circuit state change underlying skilled movements. *Cell* **184**, 3731-3747.e3721, doi:10.1016/j.cell.2021.06.001 (2021).

9 Wagner, M. J. *et al.* Shared Cortex-Cerebellum Dynamics in the Execution and Learning of a Motor Task. *Cell* **177**, 669-682.e624, doi:10.1016/j.cell.2019.02.019 (2019).

10 Sedaghat-Nejad, E., Pi, J. S., Hage, P., Fakharian, M. A. & Shadmehr, R. Synchronous spiking of cerebellar Purkinje cells during control of movements. *Proc Natl Acad Sci U S A* **119**, e2118954119, doi:10.1073/pnas.2118954119 (2022).

11 Llinás, R. & Wolfe, J. W. Functional linkage between the electrical activity in the vermal cerebellar cortex and saccadic eye movements. *Exp Brain Res* **29**, 1-14, doi:10.1007/bf00236872 (1977).

12 Welsh, J. P., Lang, E. J., Suglhara, I. & Llinás, R. Dynamic organization of motor control within the olivocerebellar system. *Nature* **374**, 453-457, doi:10.1038/374453a0 (1995).

13 Wang, Y. M. *et al.* Neuronal population activity in the olivocerebellum encodes the frequency of essential tremor in mice and patients. *Science translational medicine*, doi:DOI:10.1126/scitranslmed.adl1408 (2024).

14 Wong, S. B. *et al.* Cerebellar Oscillations in Familial and Sporadic Essential Tremor. *Cerebellum (London, England)*, doi:10.1007/s12311-021-01309-9 (2021).

15 Pan, M. K. *et al.* Cerebellar oscillations driven by synaptic pruning deficits of cerebellar climbing fibers contribute to tremor pathophysiology. *Science translational medicine* **12**, doi:10.1126/scitranslmed.aay1769 (2020).

16 Pan, M. K. Targeting the fundamentals for tremors: the frequency and amplitude coding in essential tremor. *J Biomed Sci* **32**, 18, doi:10.1186/s12929-024-01112-8 (2025).

17 Lin, C. C. *et al.* Reduced cerebellar rhythm by climbing fiber denervation is linked to motor rhythm deficits in mice and ataxia severity in patients. *Science translational medicine* **17**, eadk3922, doi:10.1126/scitranslmed.adk3922 (2025).

18 Brown, S. T., Medina-Pizarro, M., Holla, M., Vaaga, C. E. & Raman, I. M. Simple spike patterns and synaptic mechanisms encoding sensory and motor signals in Purkinje cells and the cerebellar nuclei. *Neuron* **112**, 1848-1861.e1844, doi:10.1016/j.neuron.2024.02.014 (2024).

19 Brown, S. T. & Raman, I. M. Sensorimotor Integration and Amplification of Reflexive Whisking by Well-Timed Spiking in the Cerebellar Corticonuclear Circuit. *Neuron* **99**, 564-575.e562, doi:10.1016/j.neuron.2018.06.028 (2018).

20 Lackey, E. P. *et al.* Specialized connectivity of molecular layer interneuron subtypes leads to disinhibition and synchronous inhibition of cerebellar Purkinje cells. *Neuron* **112**, 2333-2348.e2336, doi:10.1016/j.neuron.2024.04.010 (2024).

21 Rieubland, S., Roth, A. & Häusser, M. Structured connectivity in cerebellar inhibitory networks. *Neuron* **81**, 913-929, doi:10.1016/j.neuron.2013.12.029 (2014).

22 Brown, S. T. *et al.* High-speed voltage imaging of action potentials in molecular layer interneurons reveals sensory-driven synchrony that augments movement. *Cell Rep* **44**, 116148, doi:10.1016/j.celrep.2025.116148 (2025).

23 Yin, X. *et al.* Delayed motor learning in a 16p11.2 deletion mouse model of autism is rescued by locus coeruleus activation. *Nat Neurosci* **24**, 646-657, doi:10.1038/s41593-021-00815-7 (2021).

24 Juczewski, K., Koussa, J. A., Kesner, A. J., Lee, J. O. & Lovinger, D. M. Stress and behavioral correlates in the head-fixed method: stress measurements, habituation dynamics, locomotion, and motor-skill learning in mice. *Scientific reports* **10**, 12245, doi:10.1038/s41598-020-69132-6 (2020).

25 Ota, T. *et al.* Evaluation of motor learning in predictable loading task using a force sense presentation device. *Exp Brain Res* **240**, 3305-3314, doi:10.1007/s00221-022-06500-w (2022).

26 Markanday, A., Hong, S., Inoue, J., De Schutter, E. & Thier, P. Multidimensional cerebellar computations for flexible kinematic control of movements. *Nat Commun* **14**, 2548, doi:10.1038/s41467-023-37981-0 (2023).

27 Moumdjian, L. *et al.* Temporal prediction and feedforward control in cerebellar ataxia during spontaneous, instructed, and adaptive auditory-motor coupling while walking. *Scientific reports* **15**, 28053, doi:10.1038/s41598-025-12316-9 (2025).

28 Nguyen, K. P. & Person, A. L. Cerebellar circuit computations for predictive motor control. *Nat Rev Neurosci* **26**, 538-553, doi:10.1038/s41583-025-00936-z (2025).

29 Gao, Z. *et al.* Excitatory Cerebellar Nucleocortical Circuit Provides Internal Amplification during Associative Conditioning. *Neuron* **89**, 645-657, doi:10.1016/j.neuron.2016.01.008 (2016).

30 Cullen, K. E. Internal models of self-motion: neural computations by the vestibular cerebellum. *Trends Neurosci* **46**, 986-1002, doi:10.1016/j.tins.2023.08.009 (2023).

31 Mardinly, A. R. *et al.* Precise multimodal optical control of neural ensemble activity. *Nat Neurosci* **21**, 881-893, doi:10.1038/s41593-018-0139-8 (2018).

32 Yu, W. *et al.* Identifying important sensory feedback for learning locomotion skills. *Nature Machine Intelligence* **5**, 919-932, doi:10.1038/s42256-023-00701-w (2023).

33

34 Floryan, D. & Graham, M. D. Data-driven discovery of intrinsic dynamics. *Nature Machine Intelligence* **4**, 1113-1120, doi:10.1038/s42256-022-00575-4 (2022).

35 Marr, D. A theory of cerebellar cortex. *J Physiol* **202**, 437-470, doi:10.1113/jphysiol.1969.sp008820 (1969).

36 Albus, J. S. A theory of cerebellar function. *Mathematical Biosciences* **10**, 25-61, doi:https://doi.org/10.1016/0025-5564(71)90051-4 (1971).

37 Raymond, J. L. & Medina, J. F. Computational Principles of Supervised Learning in the Cerebellum. *Annu Rev Neurosci* **41**, 233-253, doi:10.1146/annurev-neuro-080317-061948 (2018).

38 Farrell, M., Recanatesi, S., Moore, T., Lajoie, G. & Shea-Brown, E. Author Correction: Gradient-based learning drives robust representations in recurrent neural networks by balancing compression and expansion. *Nature Machine Intelligence* **4**, 1053-1053, doi:10.1038/s42256-022-00565-6 (2022).

39 Miall, R. C., Weir, D. J., Wolpert, D. M. & Stein, J. F. Is the cerebellum a smith predictor? *J Mot Behav* **25**, 203-216, doi:10.1080/00222895.1993.9942050 (1993).

40 Wolpert, D. M., Miall, R. C. & Kawato, M. Internal models in the cerebellum. *Trends Cogn Sci* **2**, 338-347, doi:10.1016/s1364-6613(98)01221-2 (1998).

41 Gurnani, H. & Silver, R. A. Multidimensional population activity in an electrically coupled inhibitory circuit in the cerebellar cortex. *Neuron* **109**, 1739-1753.e1738, doi:10.1016/j.neuron.2021.03.027 (2021).

42 Hu, J., Shen, L. & Sun, G. in *Proceedings of the IEEE conference on computer vision and pattern recognition.* 7132-7141.

43 Perez, E., Strub, F., De Vries, H., Dumoulin, V. & Courville, A. in *Proceedings of the AAAI conference on artificial intelligence.*

44 Becker, M. I. & Person, A. L. Cerebellar control of reach kinematics for endpoint precision. *Neuron* **103**, 335-348. e335 (2019).

45 Markanday, A., Hong, S., Inoue, J., De Schutter, E. & Thier, P. Multidimensional cerebellar computations for flexible kinematic control of movements. *Nature Communications* **14**, 2548 (2023).

46 Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51-56 (2012).

47 Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M. & Kohn, A. Cortical areas interact through a communication subspace. *Neuron* **102**, 249-259. e244 (2019).

48 Sun, P. *et al.* Algorithm–hardware co-design of neuromorphic networks with dual memory pathways. *Nature Machine Intelligence* **8**, 901-912 (2026).

49 Farrell, M., Recanatesi, S., Moore, T., Lajoie, G. & Shea-Brown, E. Gradient-based learning drives robust representations in recurrent neural networks by balancing compression and expansion. *Nature Machine Intelligence* **4**, 564-573, doi:10.1038/s42256-022-00498-0 (2022).

Figure 1

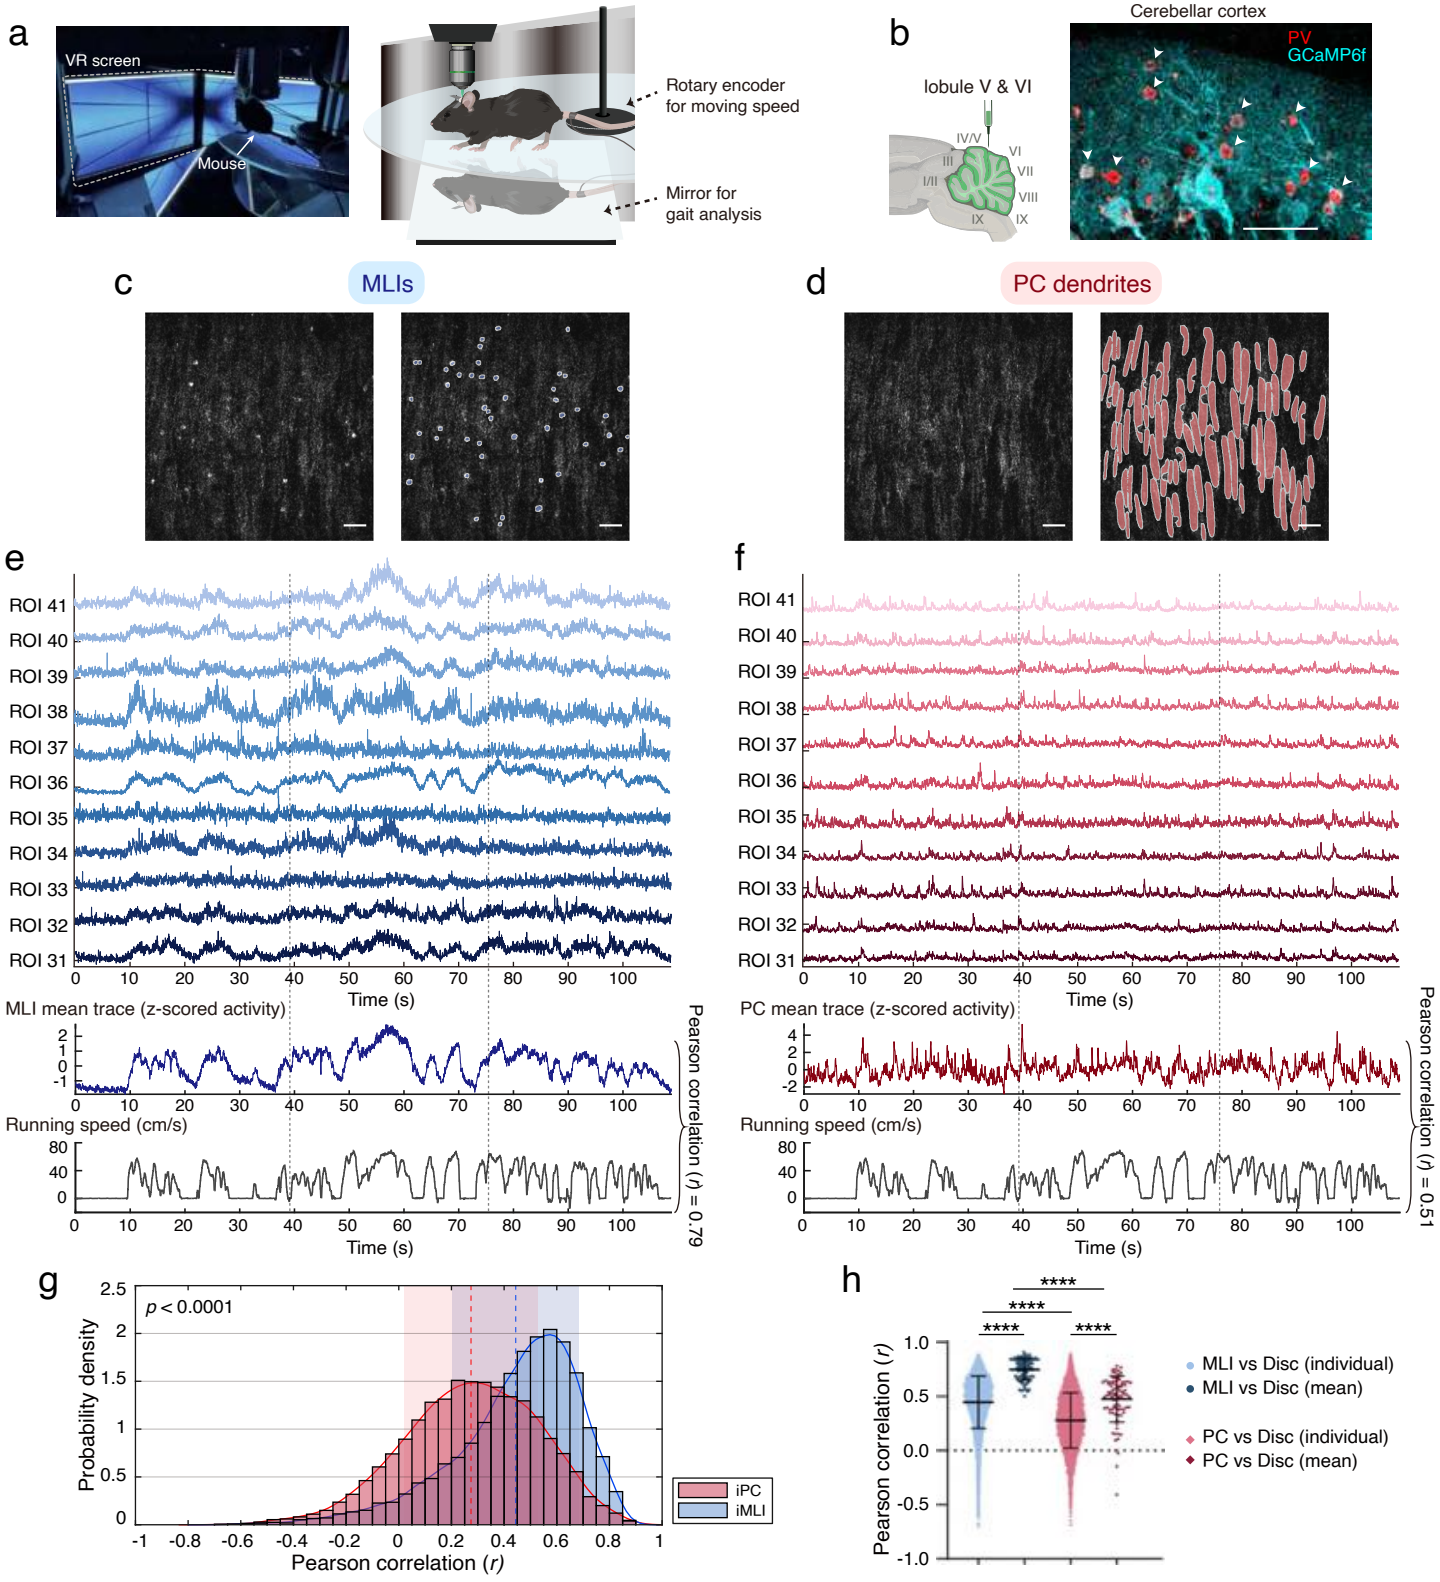


Fig. 1 | The MLI network acts as a continuous kinematic readout of locomotor speed in a closed-loop environment. a, Closed-loop virtual reality (VR) architecture. A head-fixed mouse ran on a rotary disc with real-time speed feedback coupled to visual optic flow in the VR screen. Running speed was measured via a rotary encoder, with high-dimensional gait kinematics tracked via mirror reflection. **b,** Immunofluorescence validation of the dual-labeling genetic strategy. GCaMP6f expression co-localized with parvalbumin (PV)-positive MLIs in PCP2-tTA2-GCaMP6f mice injected with AAV9-CamKII-GCaMP6f. GCaMP6f fluorescence overlaps with PV-labeled MLIs (white arrows), confirming targeted optical access to the candidate latent layer. Scale bar, 50 µm. **c,** In vivo two-photon imaging of the MLI network with identified regions of interest (ROIs) for continuous ΔF/F signal extraction. Scale bars, 50 µm. **d,** Two-photon imaging of Purkinje cell (PC) dendrites, representing the primary network output layer, with extracted ROIs. Scale bars, 50 µm. **e,** Representative MLI ΔF/F trajectories from individual MLI nodes (top), the aggregated population mean (middle), and the physical kinematic state (locomotor speed, bottom). **f,** Corresponding trajectories for the PC primary output layer. **g,** Distributions of Pearson correlations between physical locomotor speed and individual MLI versus PC ΔF/F trajectories at the single-node level (n = 6 mice; 88 trials; 8,882 individual MLI calcium traces; 5,589 individual PC calcium traces). **h,** Predictive accuracy of the kinematic representation scales with population aggregation. Scatter plots comparing Pearson correlation coefficients between locomotor speed and MLI or PC activity, evaluated for individual ΔF/F trajectories versus trial-averaged mean population responses (n = 6 mice; 88 trials; 8,882 MLI traces; 5,589 PC traces). Data are means ± SD; ****p < 0.0001; Welch's t-test.

Figure 2

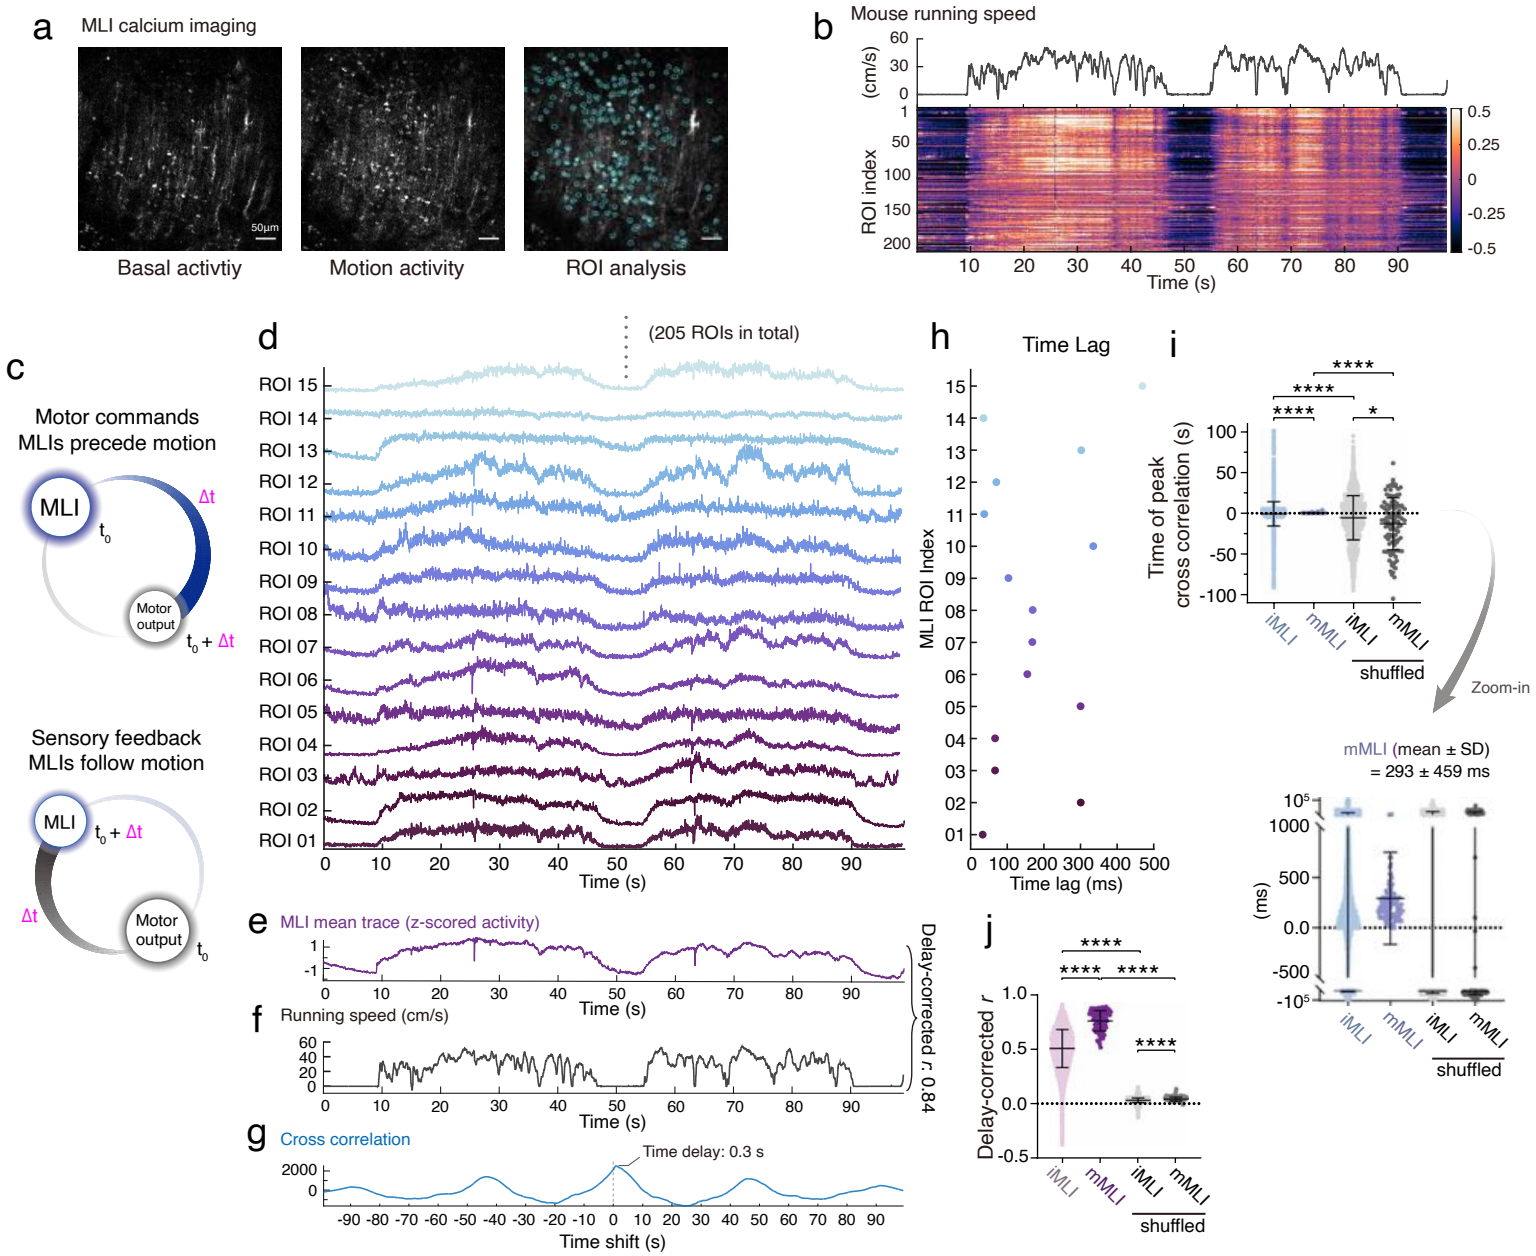


Fig. 2 | The MLI network operates as a delayed sensory state estimator. a, In vivo two-photon imaging of the MLI network showing baseline activity (left) and locomotion-evoked activity states (middle), with segmented ROIs for continuous ΔF/F signal extraction (right). Scale bars, 50 µm. **b,** Heat map of continuous ΔF/F trajectories from individual MLI nodes across a representative trial, sorted by ROI index, with the physical kinematic state (running speed) plotted above. **c,** Schematic contrasting predictive (feedforward) motor commands versus state-updating (feedback) sensory estimators. **d,** Representative ΔF/F trajectories from individual MLI nodes within a single trial, extracted from a total of 205 ROIs. **e-g,** The same representative trial as in (d), showing the aggregated population dynamics computed across all 205 MLI nodes (e), the physical kinematic state (f), and the corresponding cross-correlation function (g), which reveals a delayed processing lag of ~0.3 s. **h,** Representative time lags of individual MLIs relative to the running speed state for ROIs shown in (d), derived from cross-correlation analysis. **i,** Peak temporal lags derived from cross-correlation analysis for individual MLI trajectories (n = 9,521 traces from 6 mice), trial-averaged aggregated population dynamics (n = 107 trials, 6 mice), and shuffled controls. **j,** Delay-corrected Pearson correlations between MLI activity and the running speed state (n = 9,521 traces, 6 mice), aggregated population dynamics (n = 107 trials, 6 mice), and shuffled controls, , demonstrating enhanced numerical precision through network-level aggregation. Data are means ± SD; *$p < 0.05$, ****$p < 0.0001$; Welch's t-test.

Figure 3

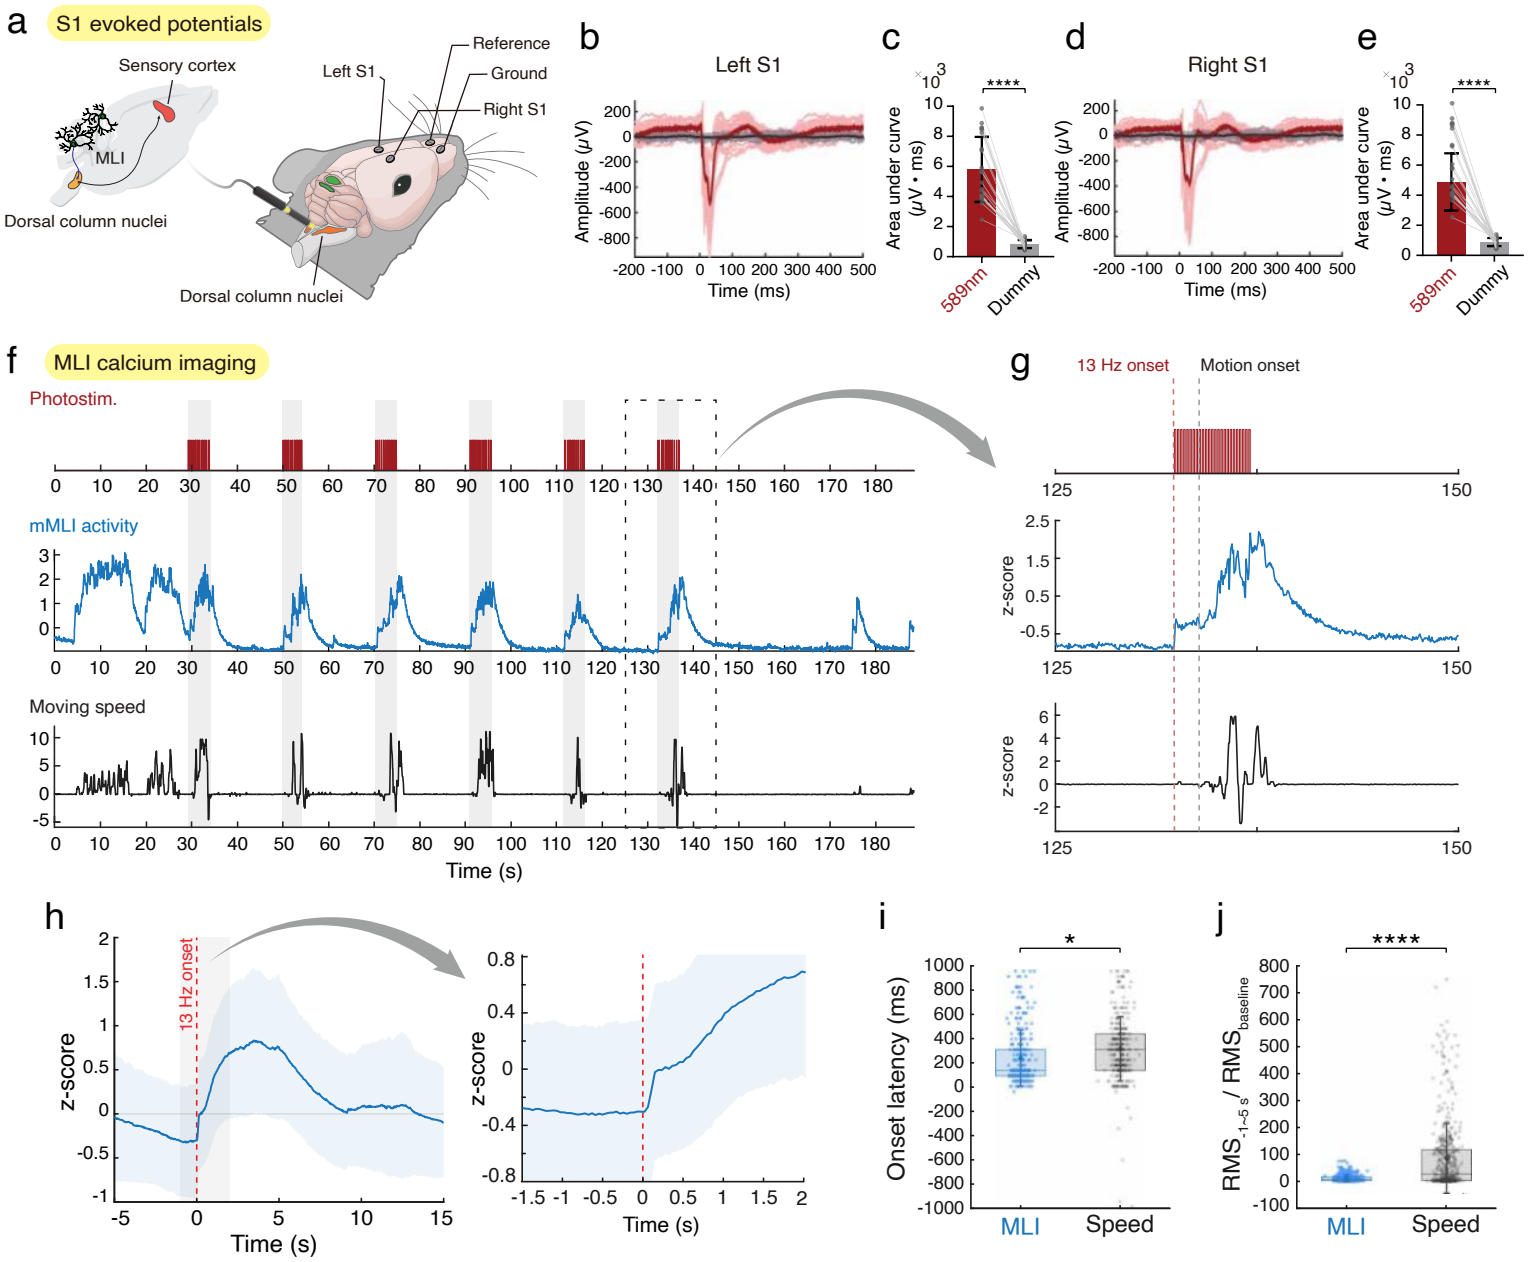


Fig. 3 | Direct actuation of ascending somatosensory pathway drives the MLI latent space. a, Schematic of the ascending sensory architecture and experimental setup. Signals from the dorsal column nuclei propagate through the dorsal column–lemniscal pathway to the cerebellar and primary somatosensory cortices (S1). The dorsal column nuclei were optogenetically actuated via an implanted optical fiber, and bilateral local field potentials (LFPs) were recorded from left and right S1. In vivo two-photon imaging of the MLI network was performed to assess whether this candidate latent layer is directly driven by ascending somatosensory inputs. **b,** Evoked potentials (EP) of left S1 during photostimulation, validating successful activation of the ascending hardware. EP traces are shown in red (photostimulation) and gray (dummy; shutter-closed), with lighter shades representing individual EP signals (n = 23 stimulating and 23 dummy trials in 5 mice). **c,** Quantification of the area under the curve (AUC) for left S1 EPs. **d-e,** EP recorded in right S1 (n = 23 stimulating and 23 dummy trials in 5 mice). **f,** MLI population dynamics during direct sensory pathway photostimulation. Example trial shows photostimulation trains (top), aggregated MLI network activity (middle), and locomotor speed (bottom). **g,** Magnified view of a single stimulation train showing that evoked MLI activity occurred between photostimulation onset and the subsequent self-evoked motor responses. **h,** Aggregated MLI network responses aligned to stimulation onset. Continuous calcium trajectories from all MLI nodes within each field of view were averaged to obtain mean network traces. Six stimulation epochs per trial were aligned to laser onset and averaged to yield trial-mean responses. The population trace shows the mean ± SD across all trial-mean traces (n = 82 trials in 6 mice). **i,** Comparison of MLI response latency versus physical locomotor onset (n = 81 trials in 6 mice; one trial without locomotion was excluded), validating that the network computation precedes overt physical execution. **j,** Peri-stimulation RMS (−1 to +5 s; baseline: -1 to 0 s) of evoked MLI activity (n = 82 trials in 6 mice) and locomotor speed (n = 81 trials in 6 mice), showing significantly smaller temporal variability in explicitly driven MLI activity. Data are means ± SD; *p < 0.05, ****p < 0.0001; Wilcoxon signed-rank test.

Figure 4

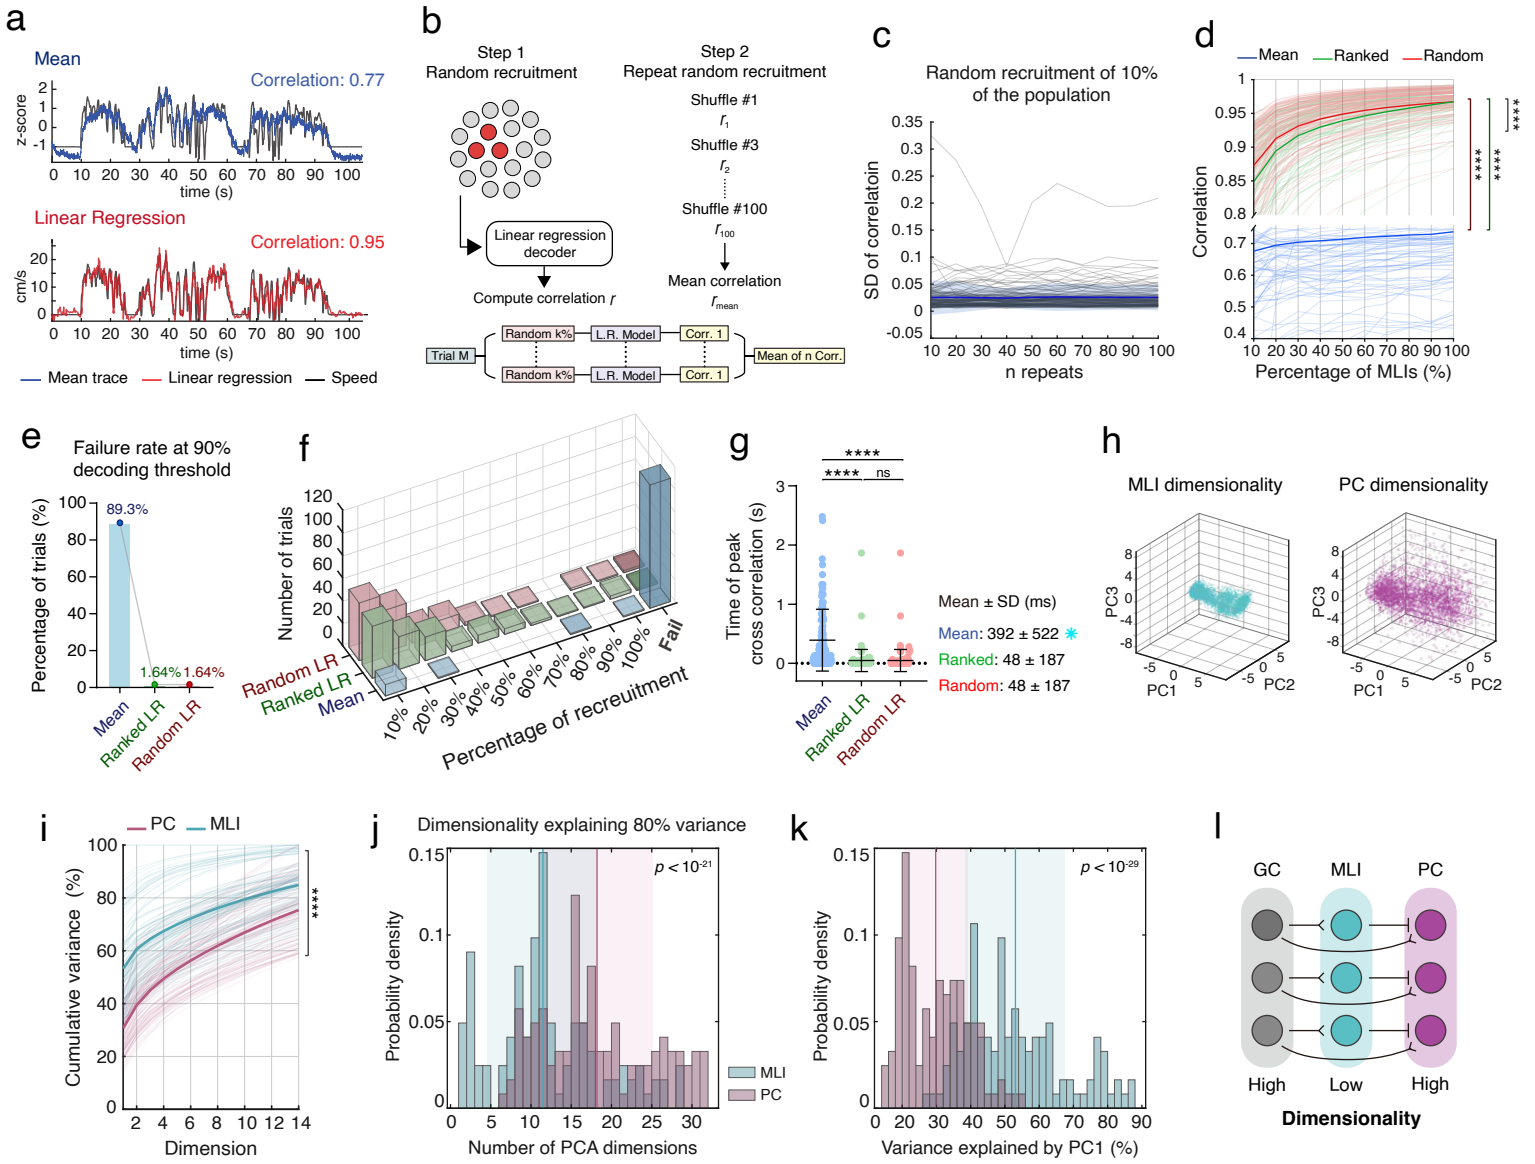


Fig. 4 | The MLI latent space extracts a low-dimensional representation of locomotor speed via natively weighted decoding. **a,** Representative trajectories from a single trial showing actual locomotor speed (gray), aggregated MLI population dynamics (blue), and the predicted scalar policy derived via weighted linear decoding (red). **b,** Schematic of the network scaling analysis. At each recruitment tier (k%), a random subset of MLI nodes was selected to train a weighted linear decoder. **c,** Example at 10% network recruitment showing decoding accuracy of the weighted linear model as a function of iteration count (10–100 iterations). Variability stabilized beyond 10 iterations, validating the use of 100 repetitions for robust performance estimation. **d,** Decoding fidelity across network scaling levels for simple uniform pooling (blue), ranked-weighted decoding (green; recruiting nodes with weaker single-cell correlations first), and random-weighted decoding (red). Thin, semi-transparent lines denote individual trajectories, and thick lines denote the mean across trials, enabling direct comparison of predictive accuracy across decoding architectures as a function of population size. Wilcoxon signed-rank tests. **e,** Proportion of continuous trajectories failing to achieve a 90% formulaic precision threshold despite full network recruitment. **f,** Distribution of trajectories exceeding the 90% precision threshold at each network scaling tier for the uniform pooling (blue), ranked-weighted (green), and random-weighted (red) architectures. Cases in which decoders with 100% recruitment failed to reach the 90% threshold are denoted as failed. **g,** Temporal processing lag (time of peak cross-correlation) between the extracted network predictions and true locomotor speed for uniform (blue), ranked (green), and random (red) weighted decoders, unmasking ultra-low latency feedback processing. *Latency estimates of uniform pooling decoder were derived using the same preprocessing pipeline as the linear regression decoder (see *Methods*), reflecting decoder-based timing rather than descriptive delays derived from z-scored activity traces as shown in Extended Data Fig. 3. **h,** Dimensionality of MLI and PC population dynamics in a representative trial, visualized via principal component analysis (PCA). Each point represents the instantaneous network state across nodes. **i,** Cumulative variance explained by the principal components for the candidate latent layer (MLIs) and the primary output layer (PCs) after node-number-matched subsampling. Thin lines represent individual trajectories; the thick line denotes the mean across trials. **j,** Distribution of the absolute number of principal components required to capture 80% of total network variance after node-number-matched subsampling. For each trial, nodes were randomly subsampled to match the size of the smaller population, and PCA was recomputed across 100 repetitions. **k,** Distribution of activity variance captured exclusively by the first principal component (PC1) across trajectories for MLI and PC populations without subsampling, confirming massive dimensionality reduction by the MLI network. **l,** Conceptual architecture of the biological information bottleneck across the GC-MLI-PC axis. The high-dimensional Granule Cell (GC) feature space is natively distilled by the parallel MLI latent space into a scalar speed policy to continuously calibrate the high-dimensional PC output. Data are from n = 122 trials in 6 mice and shown as means ± SD; ****$p < 0.0001$; d: Wilcoxon signed-rank test (Benjamini–Hochberg FDR corrected); i,k: Wilcoxon signed-rank test; j: Wilcoxon rank-sum test.

Figure 5

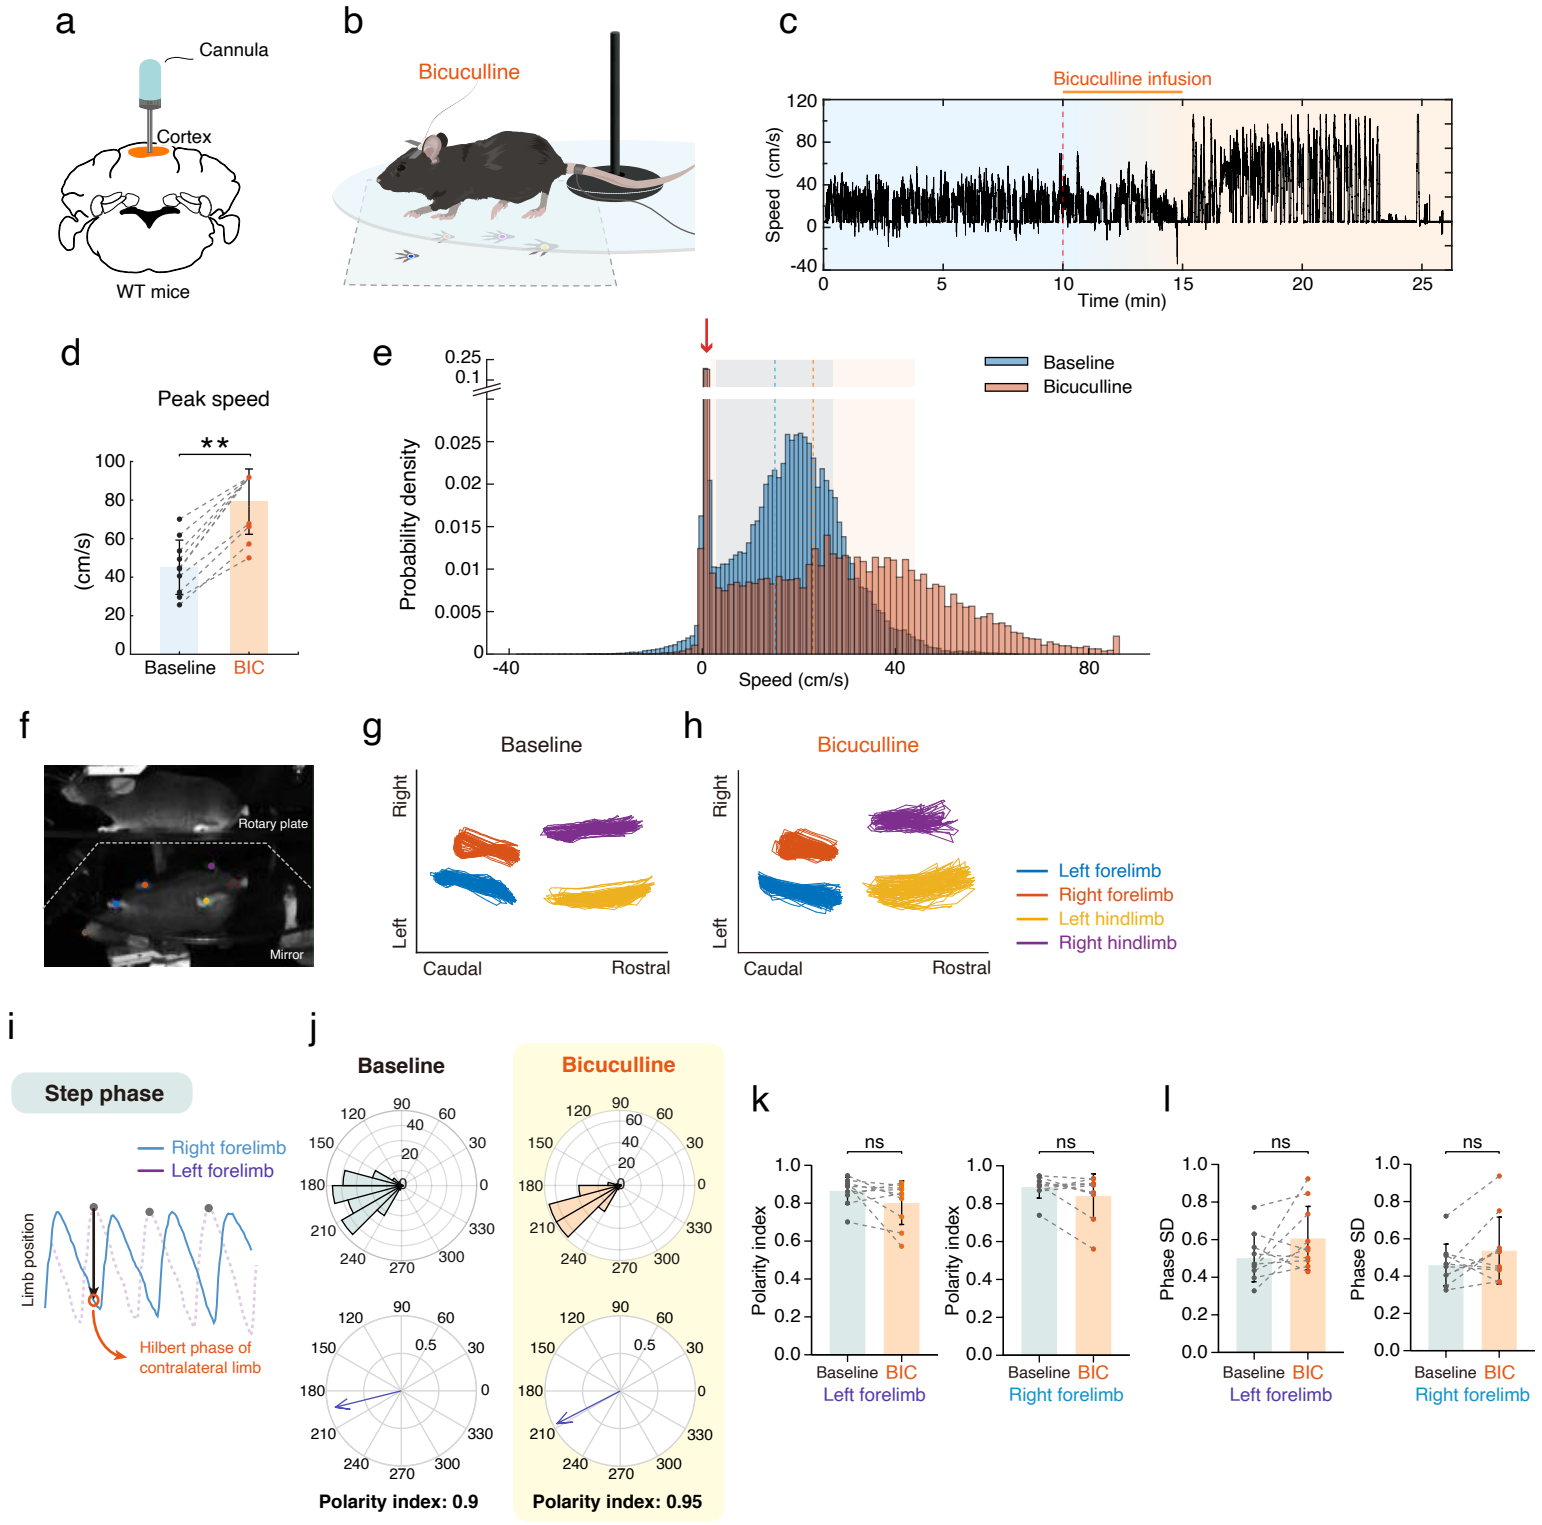


Fig. 5 | Targeted suppression of the MLI latent space recalibrates the scalar speed policy while preserving high-dimensional coordination. a, Schematic showing bicuculline infusion into the cerebellar cortex via a central cannula to transiently block the MLI-to-PC inhibitory neurotransmission. **b,** Experimental architecture showing a head-fixed mouse running on a rotary disc with a mirror underneath for high-dimensional gait tracking. **c,** Example continuous trajectory showing locomotor speed during baseline and bicuculline infusion; targeted suppression induced a compensatory acceleration in locomotor speed. **d,** Peak locomotor speed during baseline and following suppression of latent space outputs (n = 10 mice). **e,** Probability density distributions of locomotor speed pooled across subjects (n = 10 mice). **f,** High-speed videography (120 fps) enabling rigorous quantification of multidimensional limb kinematics. **g-h,** Representative forelimb and hindlimb trajectories under baseline (**g**) and bicuculline (**h**) conditions. **i,** Schematic illustrating step-phase extraction from Hilbert-transformed contralateral limb trajectories. **j,** Example polar plots of step-phase distributions confirming intact forelimb coordination during baseline and bicuculline conditions. **k,** Quantification of the polarity index (PI) for left and right forelimbs (n = 10 mice). **l,** Step-phase variability (SD) for both forelimbs (n = 10 mice), confirming the scalar policy update does not corrupt multidimensional coordination. Data are from 10 mice and presented as means ± SD; $^{**}p < 0.01$; Wilcoxon signed-rank test.

Figure 6

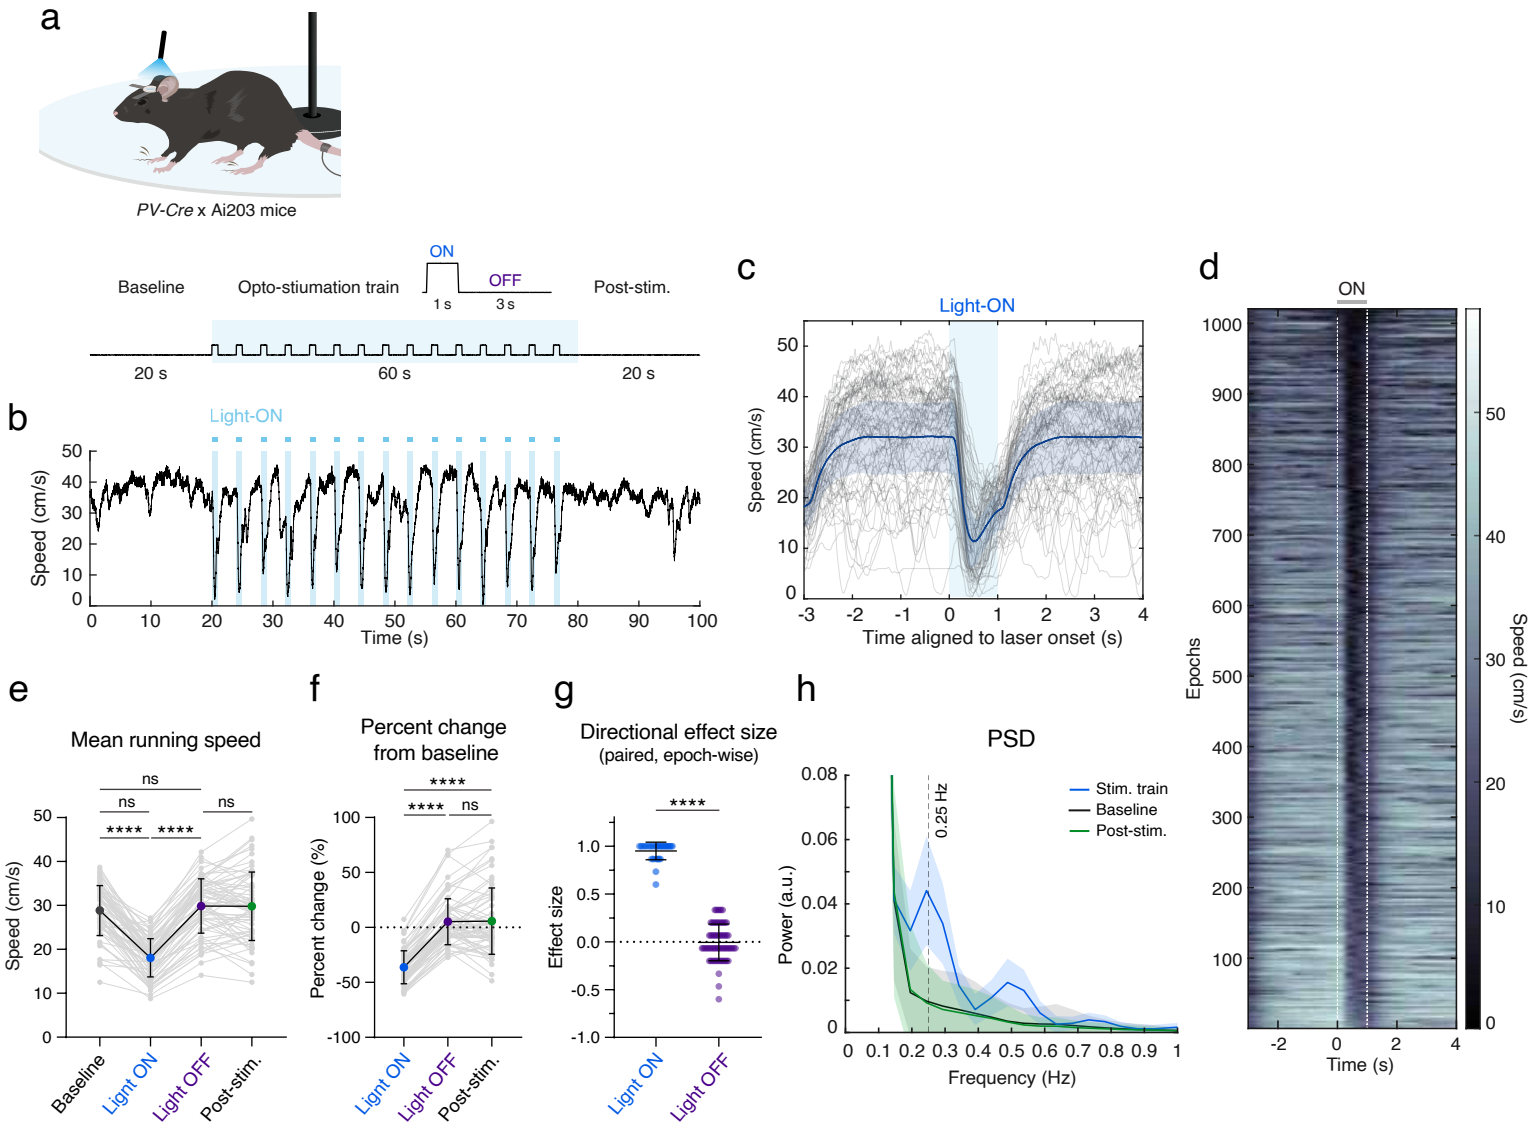
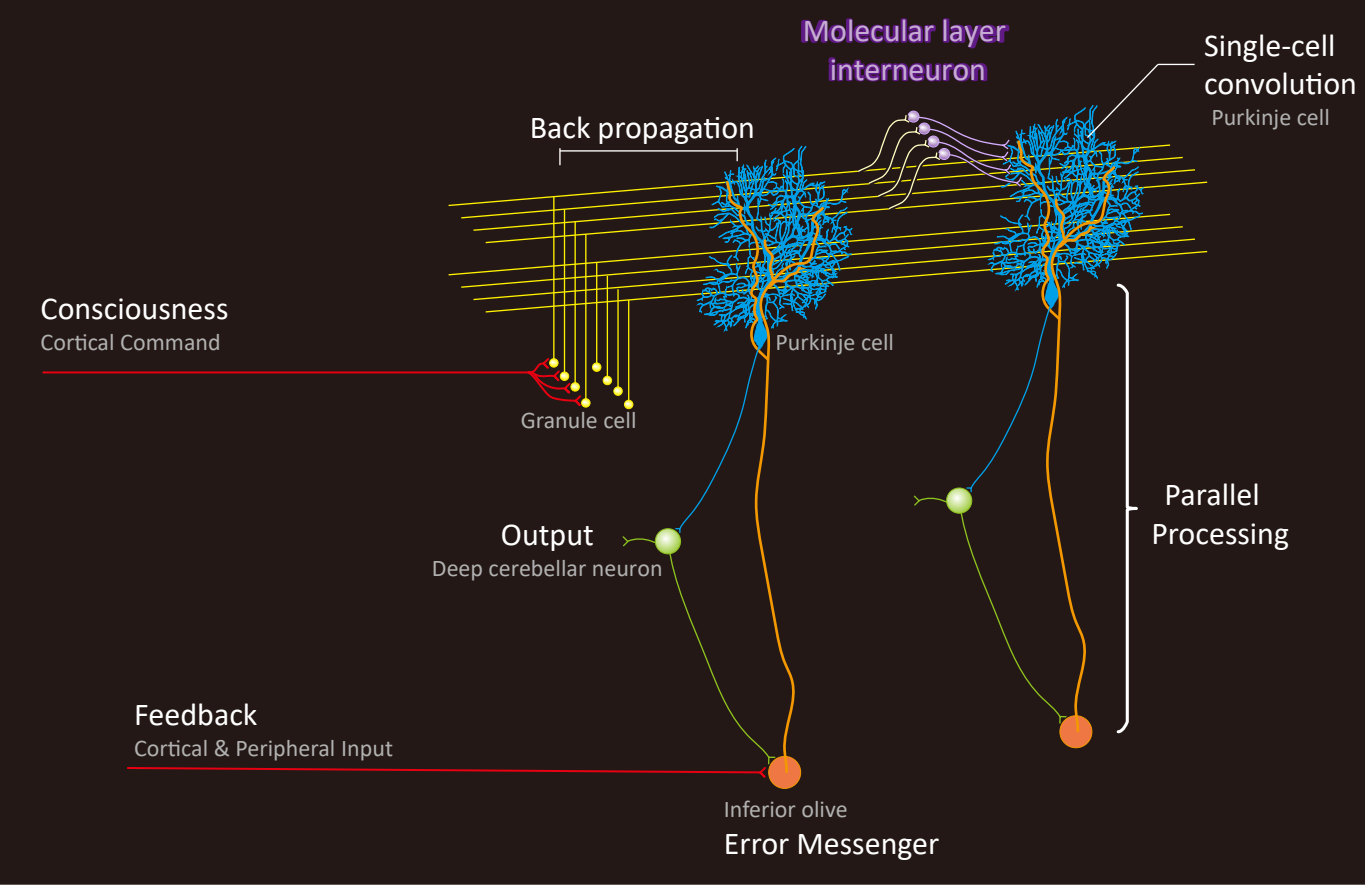


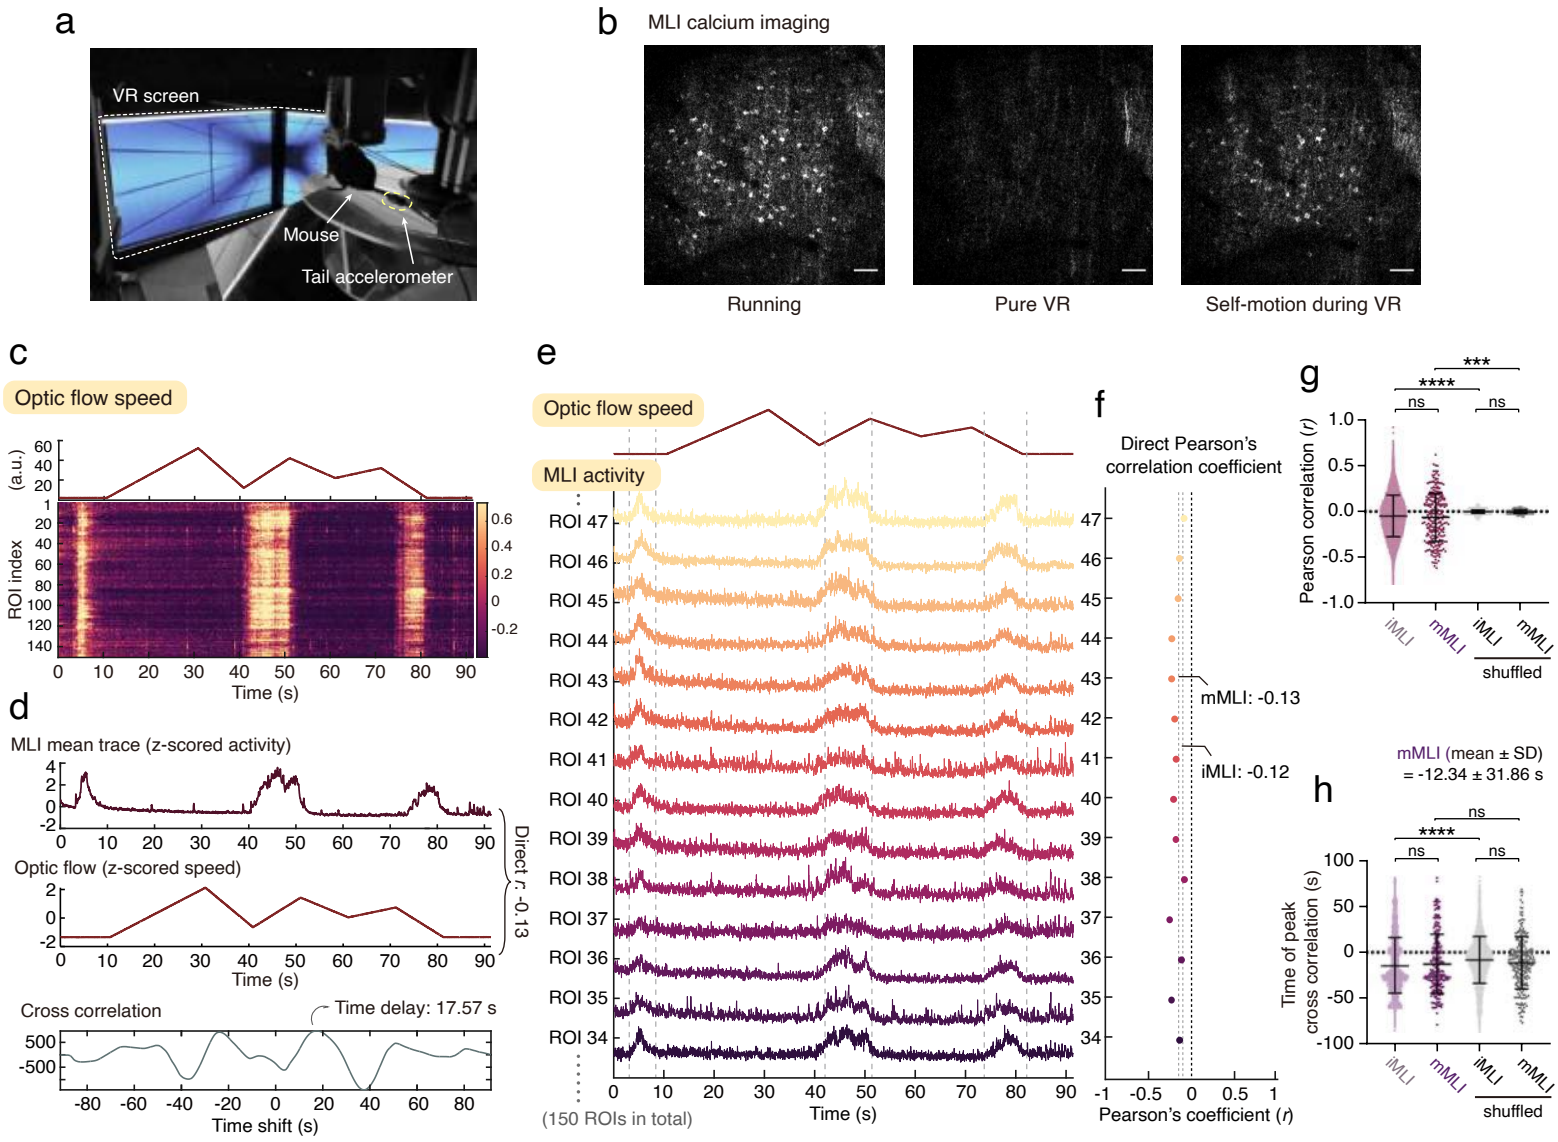
Fig. 6 | Optogenetic actuation of the MLI latent space dynamically suppresses the scalar speed policy. a, Experimental paradigm: 20-s baseline, followed by 60 s of repeated optogenetic stimulation (15 cycles of 1-s light ON / 3-s light OFF), and a 20-s post-stimulation period. **b,** An example continuous trajectory of locomotor speed. Each targeted optogenetic actuation of the MLI network produced an immediate and reproducible deceleration, visible as sharp downward deflections aligned to stimulation. **c,** Peri-event alignment of actuation epochs. For each trial, the 15 epochs were averaged to yield a single trial-level response trace (gray). The blue trace shows the pooled mean ± SD across trials. **d,** Heatmap of all actuation epochs sorted by mean locomotor speed (n = 1,020 epochs, 68 trials from 4 mice). **e,** Mean locomotor speed across baseline, light ON, light OFF, and post-actuation periods. **f,** Percent change in locomotor speed relative to baseline across the defined periods. **g,** Directional effect size reflecting the consistency of scalar policy suppression across single actuation epochs relative to the peri-epoch baseline (−3 to 0 s). A value of one indicates consistent suppression across single epochs, whereas zero indicates mixed or inconsistent modulation. **h,** PSD of locomotor speed across the three trial periods. Despite naturally heterogeneous running speeds, actuation at a fixed 1-s ON / 3-s OFF rhythm (0.25 Hz) induces a clear PSD peak at 0.25 Hz, demonstrating precise temporal entrainment of the scalar speed policy by the MLI latent space. Data are means ± SD; ****$p < 0.0001$; Wilcoxon signed-rank tests.

Extended Data Fig. 1



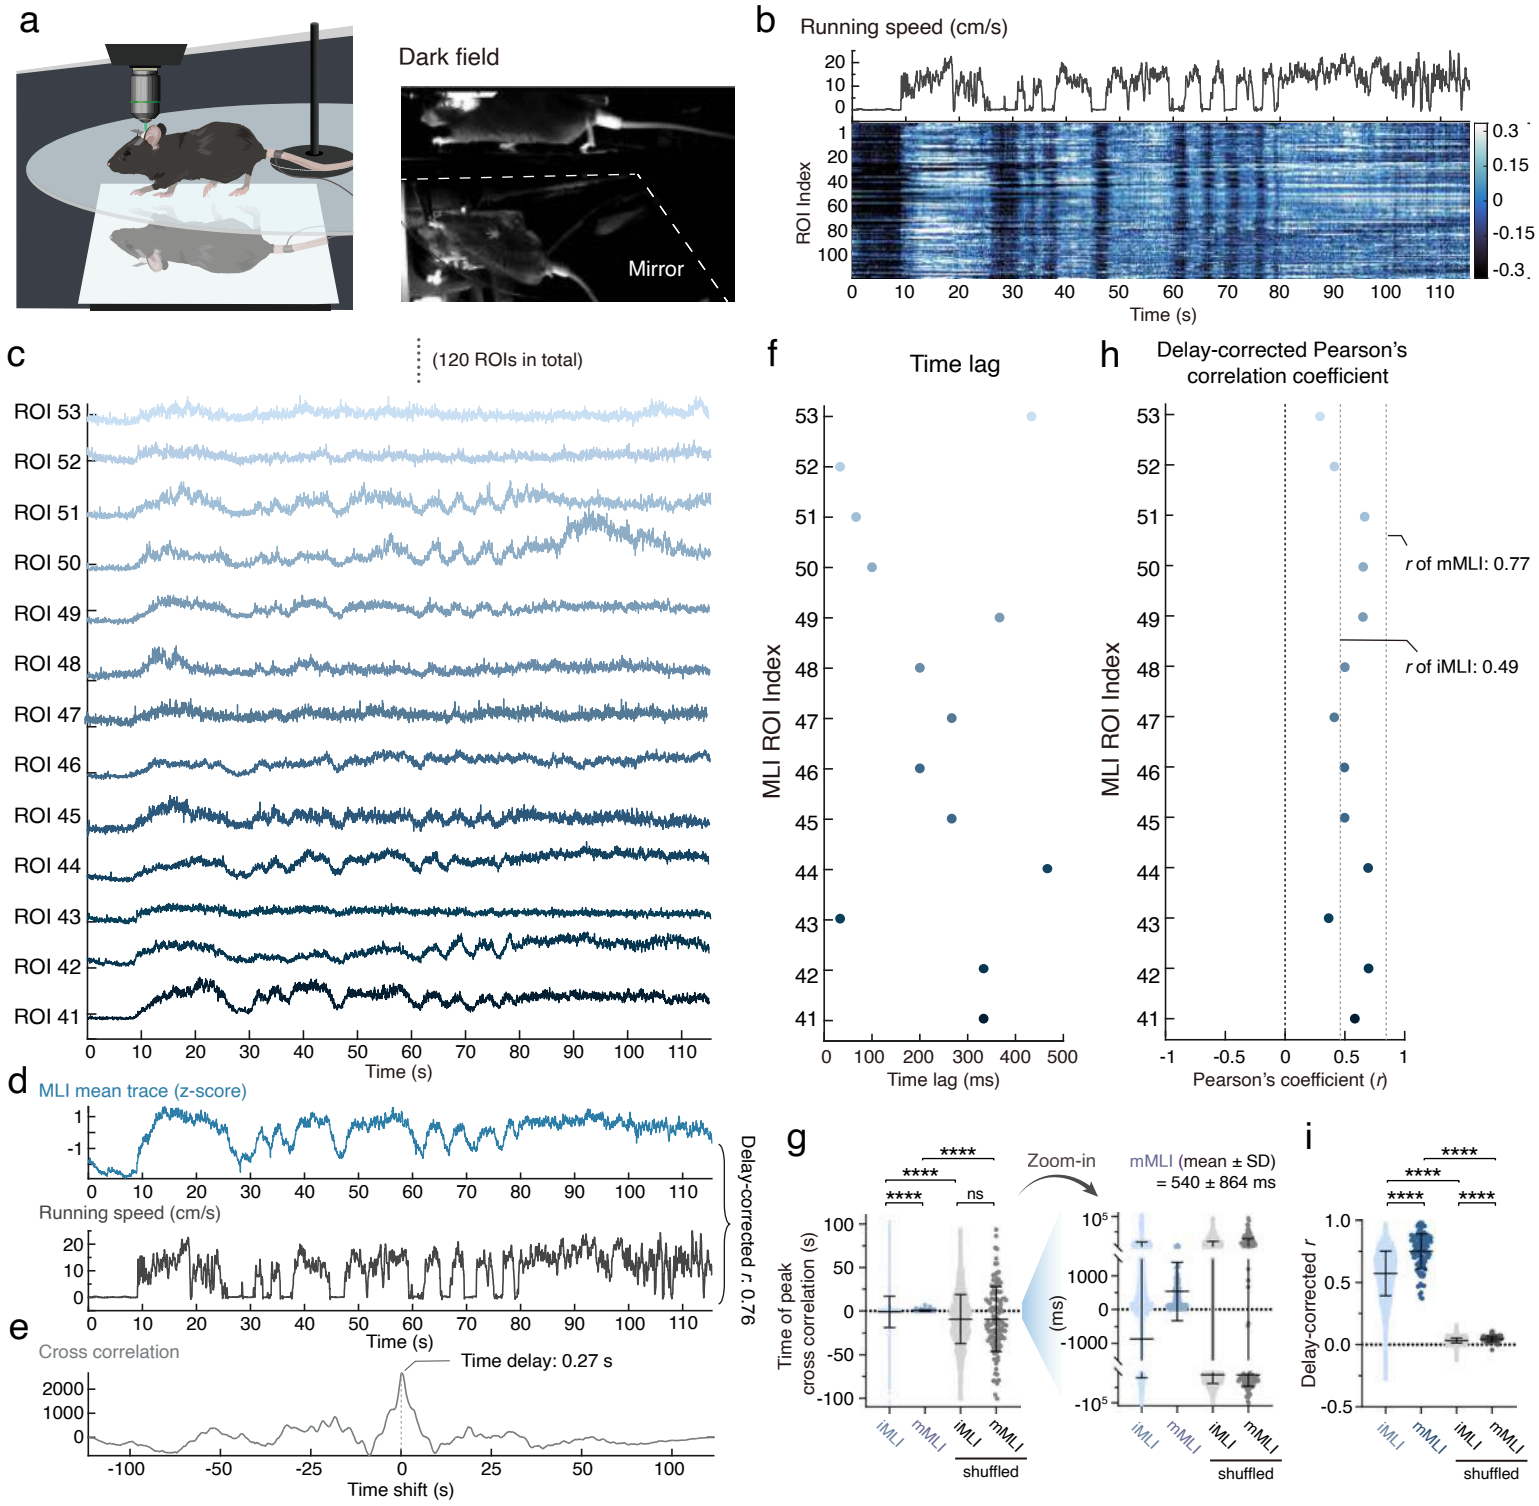
Extended Data Fig. 1 | The cerebellar microcircuit as a parallel-processing neuromorphic architecture. Cortical command signals reach the granule cell (GC) layer—the most abundant neural population accounting for ~65% of neurons in mice and ~75% in humans—establishing a massive dimensional expansion of the state space. This high-dimensional feature map is directly transmitted to Purkinje cells (PCs). PCs leverage highly enriched dendritic arbors to perform complex spatial convolutions at the single-cell level, subsequently passing this integrated output to deep cerebellar neurons to generate the final cerebellar output. For adaptive calibration, PCs act as integration nodes for teaching signals from inferior olivary neurons (error messengers) and tens of thousands of high-dimensional inputs from GC parallel fibers. Synchronous arrival of an error signal and GC input induces long-term depression (LTD) at the active parallel fiber-PC synapses, functioning as a biological equivalent to synaptic weight adjustment via backpropagation. Inferior olivary neurons, PCs, and deep cerebellar neurons operate as closed-loop modular units, facilitating massively parallel processing across millions (in mice) to tens of millions (in humans) of discrete architectures. While the primary high-dimensional feedforward and feedback learning pathways are well-defined, the specific computational role of the parallel collateral molecular layer interneuron (MLI) network remains largely uncharacterized within this neuromorphic schema.

Extended Data Fig. 2



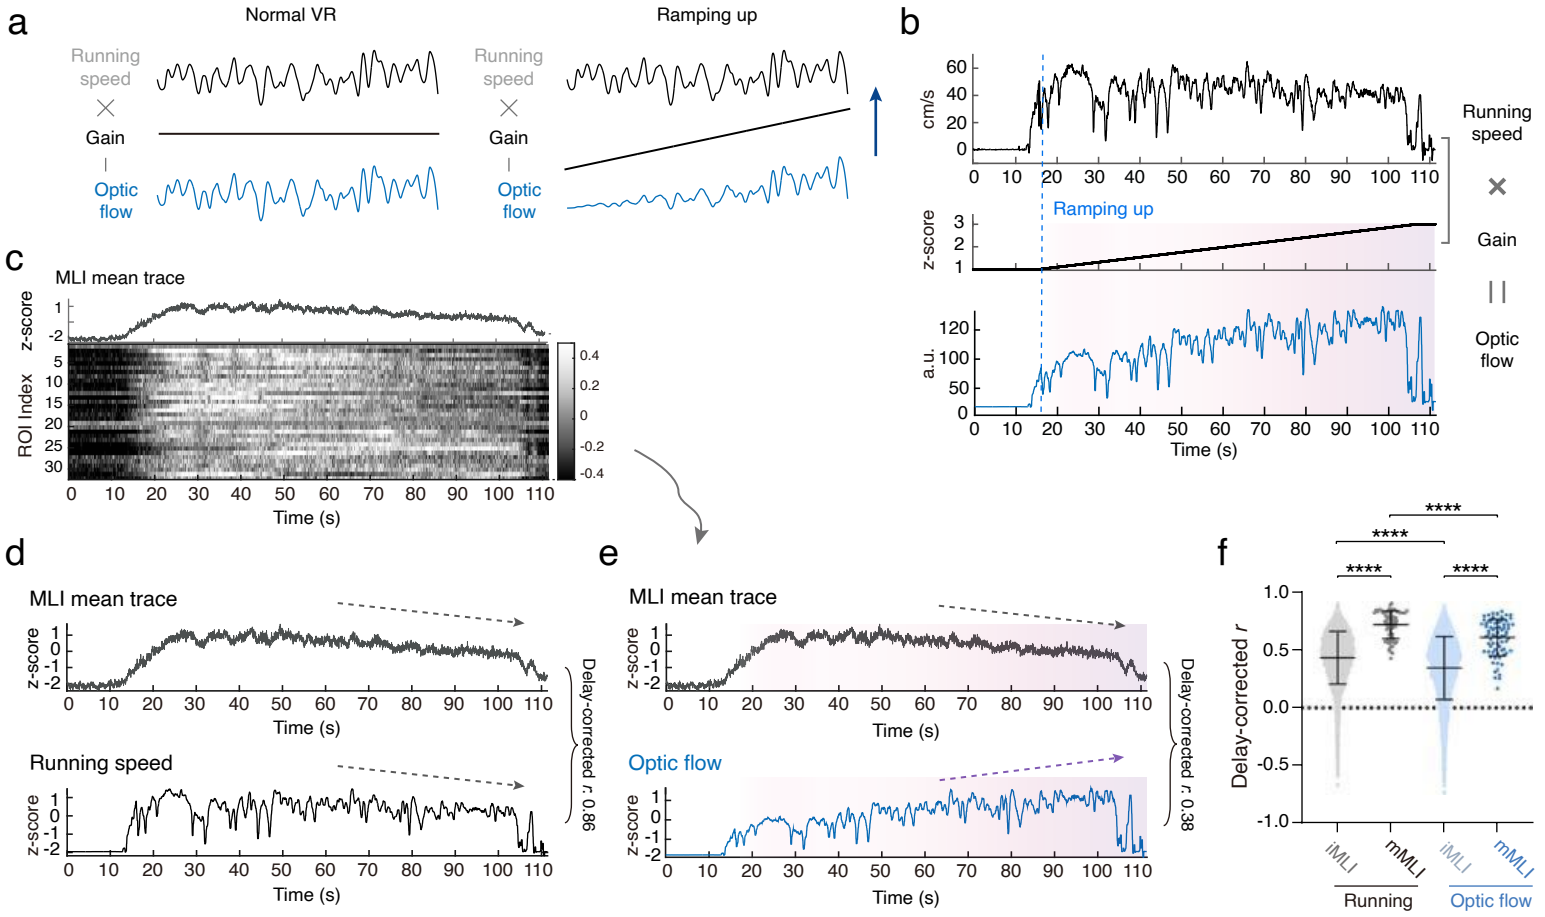
Extended Data Fig. 2 | The MLI latent space selectively filters visual optic flow during passive observation. a, Experimental architecture: a head-fixed mouse on an immobilized rotary disc viewing a preprogrammed VR movie; a tail accelerometer was attached to detect residual physical movement. **b,** In vivo two-photon images of the MLI network from the same field of view under three conditions: active self-paced locomotion (left), purely visual VR playback without physical movement (middle), and physical movement during visual playback (right). Scale bars, 50 µm. **c,** Heat map of continuous MLI activity during purely visual playback, sorted by ROI index, with optic flow speed overlaid above. **d,** Example trial showing aggregated population dynamics (top), optic flow speed (middle), and the cross-correlation function (bottom), illustrating a peak lag of ~17.6 s and a direct Pearson correlation (r) = -0.13. **e,** Example trial showing optic flow speed (top) and continuous signal trajectories from individual MLI nodes (bottom), extracted from a total of 150 ROIs. **f,** Direct Pearson correlation coefficients between MLI activity and optic flow speed for the ROIs shown in (e). Both the mean of the individual MLI–optic flow correlations across all 150 individual ROIs (iMLIs) and the correlation of the aggregated population trajectory (mMLI) with optic flow are shown. **g,** Pearson's correlations between MLI activity and optic flow speed, compared with temporally shuffled visual controls. Distributions are shown at the individual node level (iMLIs: 29,142 traces) and aggregated population level (mMLIs: 244 trial-averaged population responses) from 6 mice. **h,** Peak temporal lags of the cross-correlation between network activity and optic flow speed for the same dataset as in (g), compared with shuffled controls. Data are means ± SD; ***$p < 0.001$, ****$p < 0.0001$; Welch's t-test.

Extended Data Fig. 3



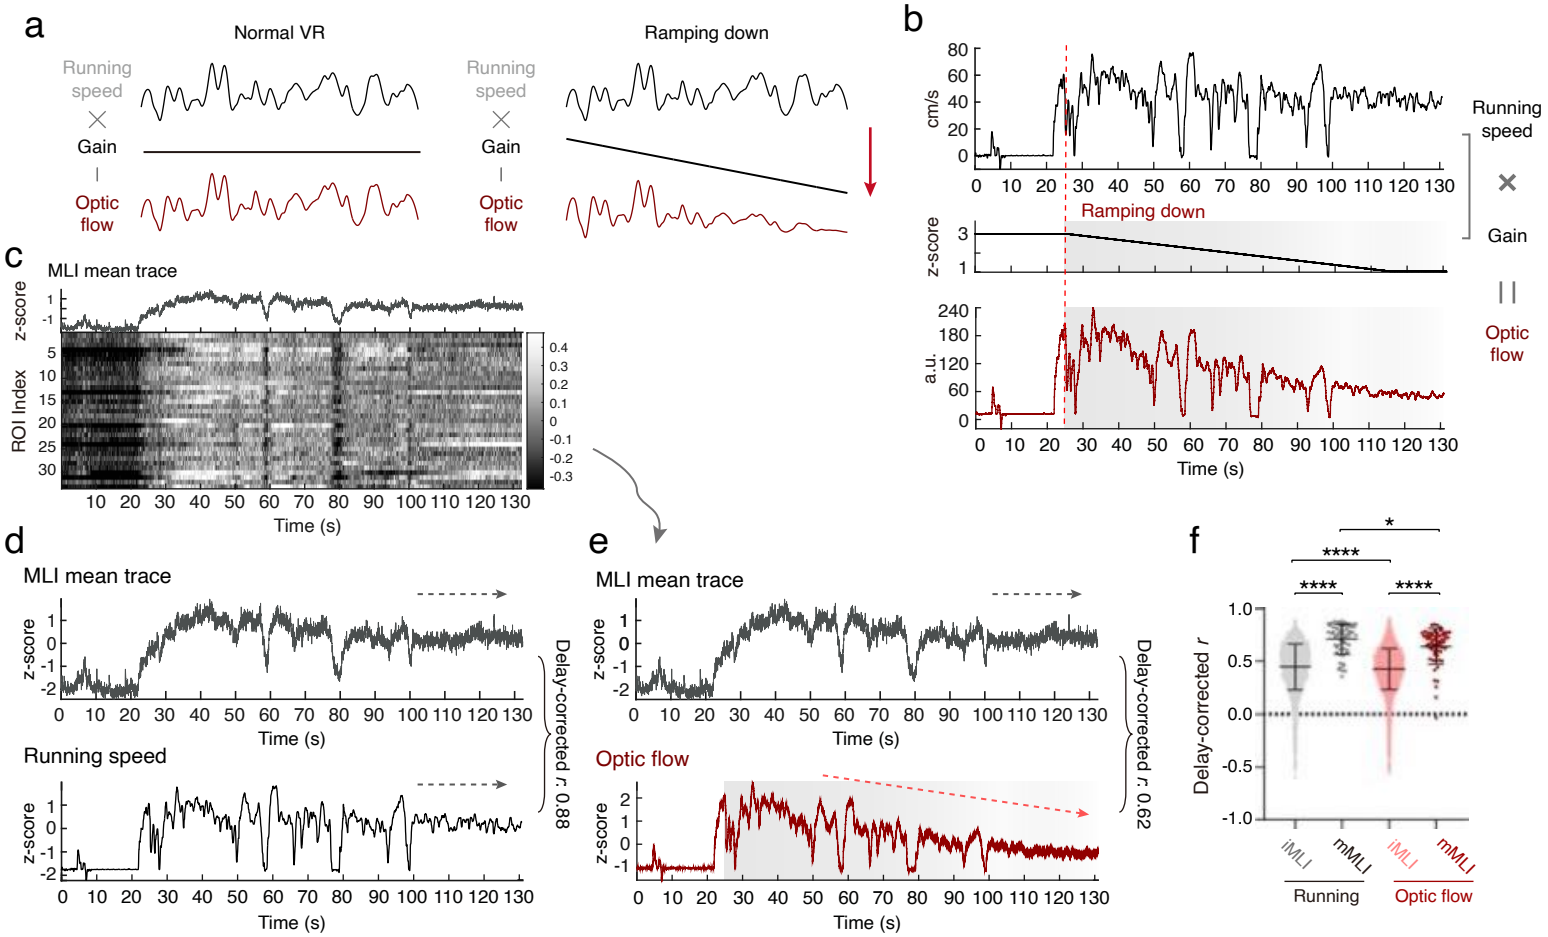
Extended Data Fig. 3 | The MLI state estimator computes locomotor speed exclusively via somatosensory channels. a, Experimental architecture: a head-fixed mouse ran volitionally on a rotary disc in complete darkness, eliminating visual optic flow. **b,** Heat map of continuous activity from individual MLI nodes during a representative trial, sorted by ROI index, with physical locomotor speed plotted above. **c,** Representative signal trajectories of individual MLI nodes from the same trial as in (b). **d,** Aggregated MLI population dynamics from the same trial plotted alongside locomotor speed (delay-corrected Pearson correlation (r) = 0.76). **e,** Cross-correlation between aggregated MLI dynamics and locomotor speed for the same trial, showing a feedback-processing delay. **f,** Temporal processing lags between individual MLI nodes and locomotor speed for ROIs in (c), derived via cross-correlation analysis. **g,** Distribution of peak processing lags between the network state and locomotor speed, compared with shuffled speed controls. Data are shown for individual MLI nodes (iMLIs; n = 16,750 traces from 6 mice) and aggregated populations (mMLIs; n = 244 trial-averaged population responses from 6 mice). An inset provides a magnified view of the distribution. **h,** Delay-corrected Pearson's correlation coefficients between individual MLI nodes and locomotor speed for the ROIs shown in (c), utilizing the latency estimated in (f). **i,** Delay-corrected Pearson's correlations between MLI network state and locomotor speed, compared with shuffled speed controls. Scatter plots illustrate coding fidelity distributions at the individual node and aggregated population levels. Data are means ± SD; ****$p < 0.0001$; Welch's t-test.

Extended Data Fig. 4



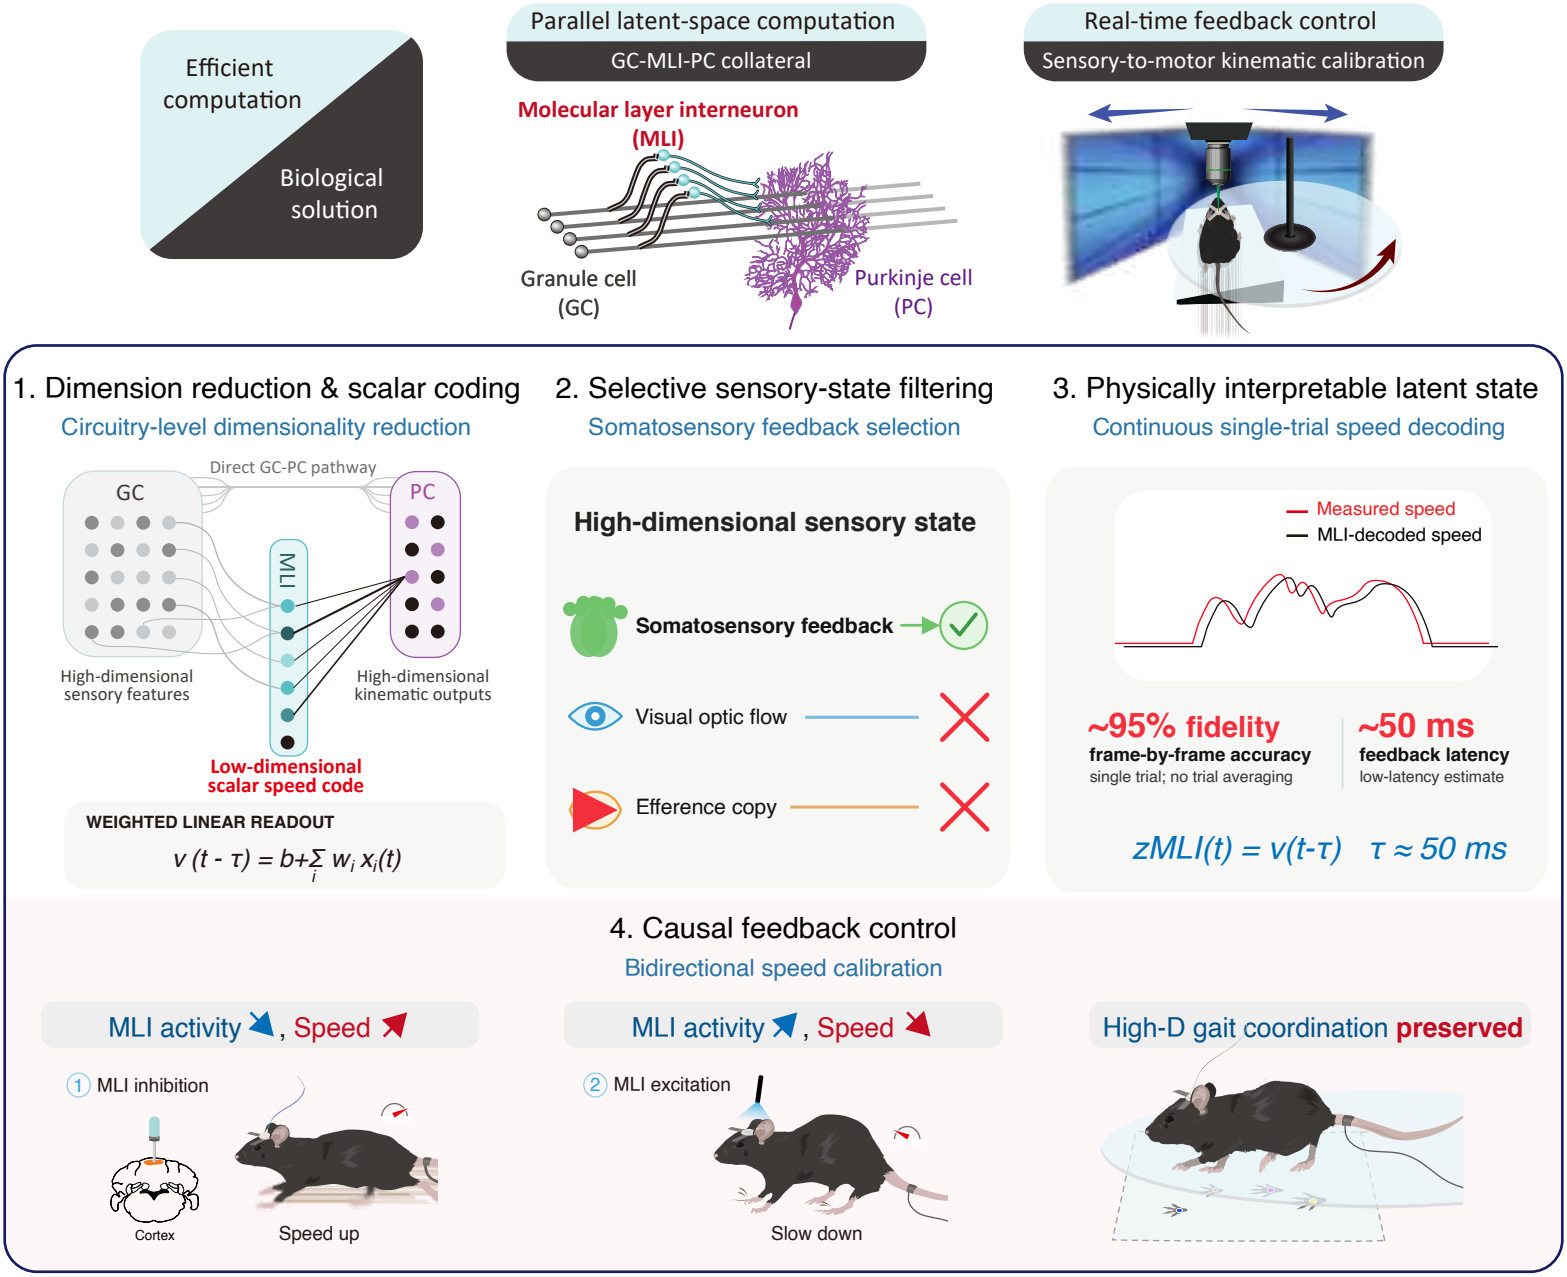
Extended Data Fig. 4 | The MLI state estimator rejects deceptive visual acceleration during dynamic gain mismatch. a, Schematic of the dynamic visual-motor gain mismatch architecture. The real-time physical locomotor speed was multiplied by a progressively increasing gain, introducing an artificial, progressively faster visual state-space within a single trial. **b,** Example trial showing actual physical locomotor speed (top), the applied ramp-up gain (middle), and the decoupled visual optic flow (bottom). **c,** MLI population dynamics during the same trial: aggregated continuous trace across all nodes (top) and heatmap of individual MLI nodes (bottom). **d,** Correlation between the aggregated MLI state and true locomotor speed within the same trial. **e,** Correlation between aggregated MLI dynamics and the manipulated visual optic flow within the same trial as (d). **f,** Correlation coefficients (delay-corrected *r*) between the MLI network state and either physical locomotor speed or the artificially manipulated visual state-space. Data are shown for individual MLI nodes (iMLIs; n = 9,050 traces from 6 mice) and aggregated population dynamics (mMLIs; n = 76 population means from 76 trials in 6 mice). Scatter plots illustrate coding fidelity at the single-node and network levels. Data are means ± SD; ****$p < 0.0001$; Welch's t-test.

Extended Data Fig. 5



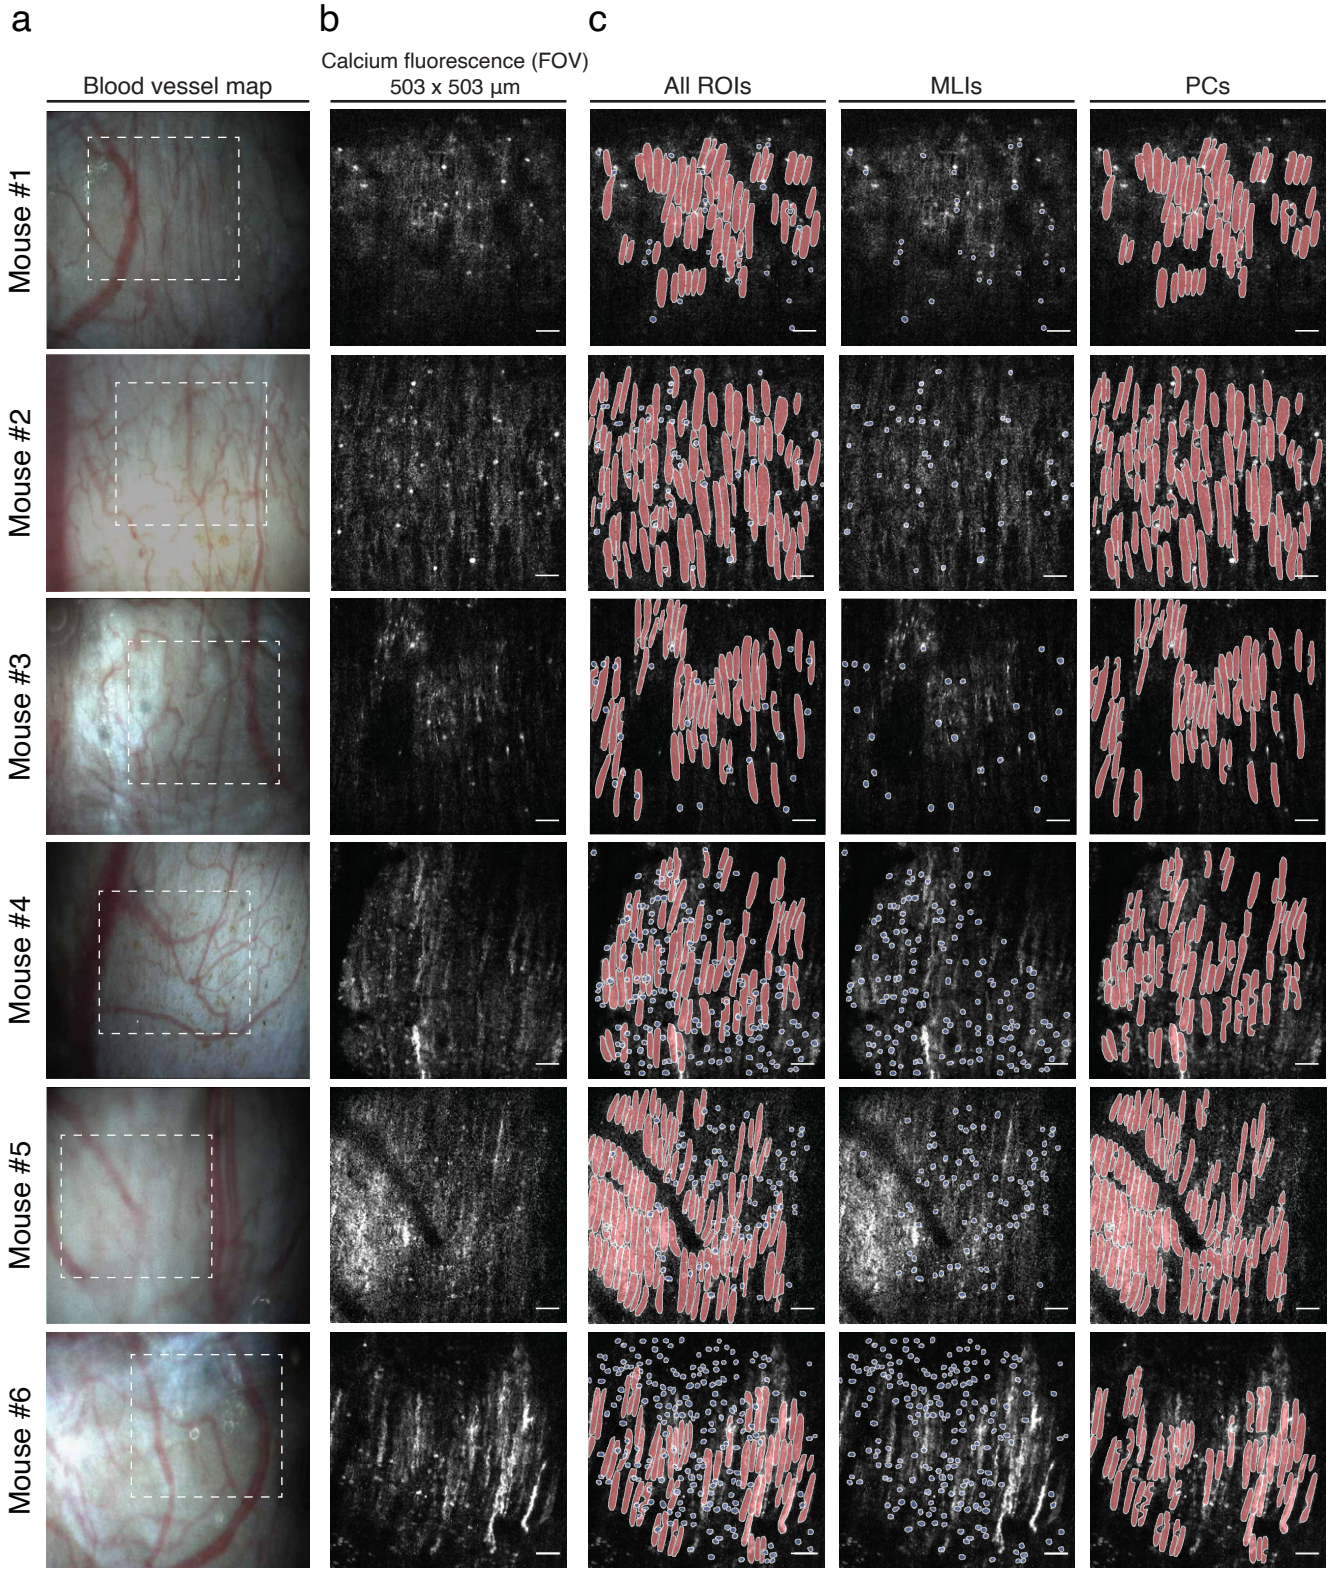
Extended Data Fig. 5 | The MLI state estimator rejects deceptive visual deceleration during dynamic gain mismatch. a, Schematic of the dynamic ramp-down VR architecture. Real-time physical locomotor speed was multiplied by a progressively decreasing gain, creating a progressively slower visual state-space mismatch within a single trial. **b,** Example trial showing true physical locomotor speed (top), the dynamic ramp-down gain (middle), and the manipulated visual optic flow (bottom). **c,** MLI population dynamics during the same trial: aggregated continuous trace (top) and heatmap of individual nodes (bottom). **d,** Correlation between the aggregated MLI state and true locomotor speed within the same trial. **e,** Correlation between the aggregated MLI state and the manipulated ramp-down visual state-space within the same trial. **f,** Delay-corrected correlation coefficients between the MLI network state and either physical locomotor speed or the manipulated visual variables. Data are shown for individual MLI nodes (iMLIs; $n = 6{,}974$ traces from 6 mice) and aggregated population dynamics (mMLIs; $n = 61$ population means from 61 trials in 6 mice). Scatter plots illustrate coding fidelity at the single-node and network levels. Data are means ± SD; $^{*}p < 0.05$, $^{****}p < 0.0001$; Welch's t-test.

Extended Data Fig. 6



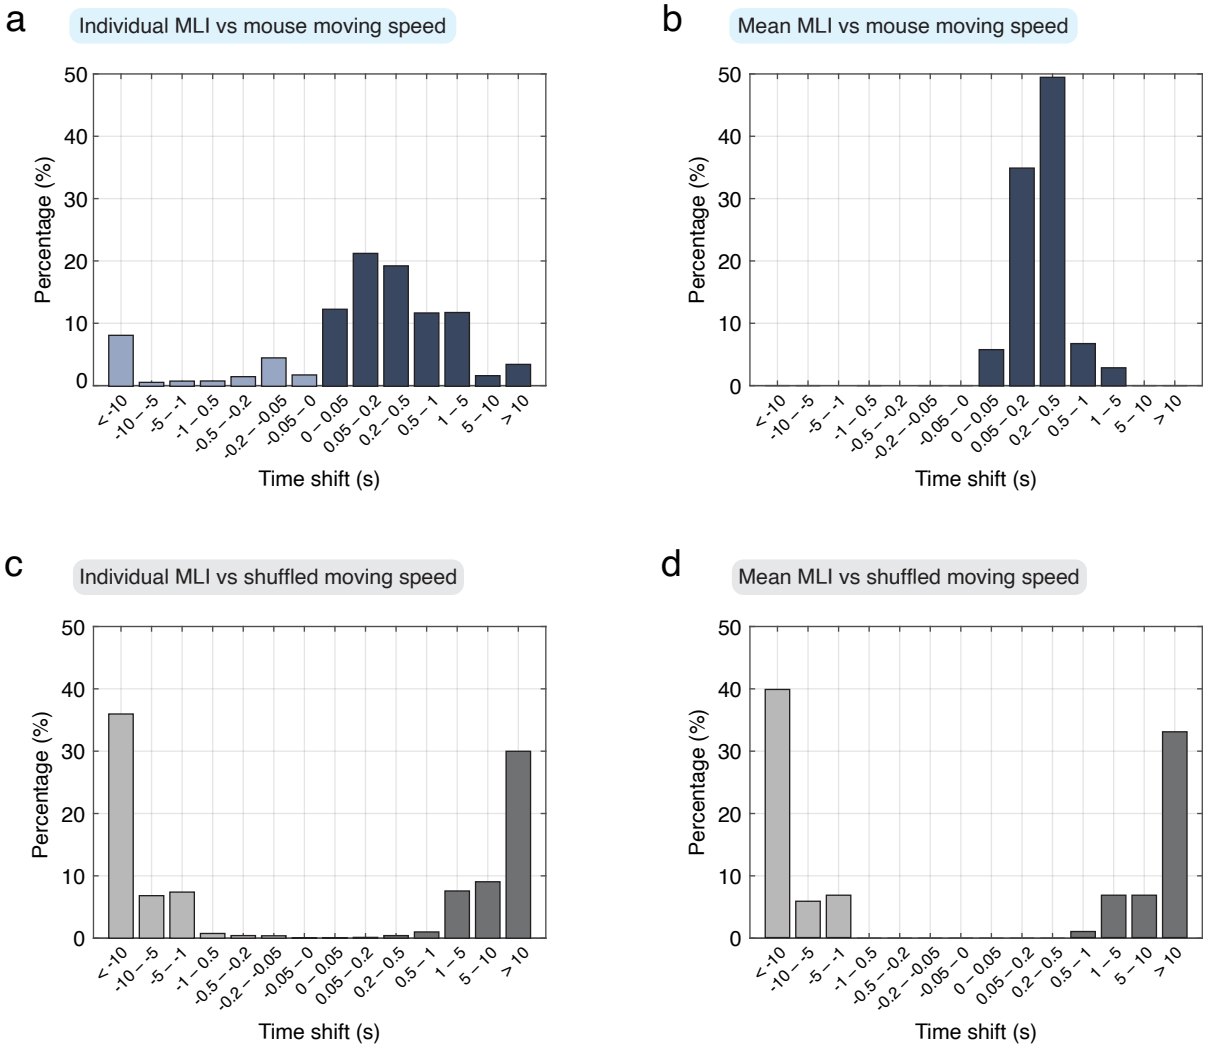
Extended Data Fig. 6 | Graphical summary: the cerebellar MLI network implements a low-dimensional control policy for real-time kinematic calibration. Mammalian brains achieve accurate and reliable sensorimotor control under stringent computational and energetic constraints. Here, we identify a biological circuit solution that compresses high-dimensional sensory information into a low-dimensional state for rapid kinematic calibration. High-dimensional sensory features represented in the granule cell (GC) layer are transmitted directly to Purkinje cells (PCs) to support multidimensional kinematic outputs, whereas the parallel molecular layer interneuron (MLI) collateral performs circuitry-level dimensionality reduction and extracts a scalar representation of locomotor speed through weighted population decoding. The MLI network selectively processes somatosensory feedback while rejecting visual optic flow and an internally generated efference copy of motion. The resulting population code forms a physically interpretable latent state and achieves approximately 95% frame-by-frame decoding fidelity in individual self-paced trials without trial averaging. MLI population activity lags locomotor speed by approximately 50 ms, summarized as $zMLI(t) \approx v(t-\tau)$, where $\tau \approx 50$ ms. Pharmacological suppression of MLI-to-PC inhibitory transmission accelerates locomotion while preserving high-dimensional gait coordination, whereas optogenetic excitation of MLIs decelerates locomotion. Together, these findings identify the MLI latent space as a selective, low-latency feedback controller that bidirectionally calibrates the scalar speed policy without disrupting multidimensional motor coordination.

Supplementary Fig. 1



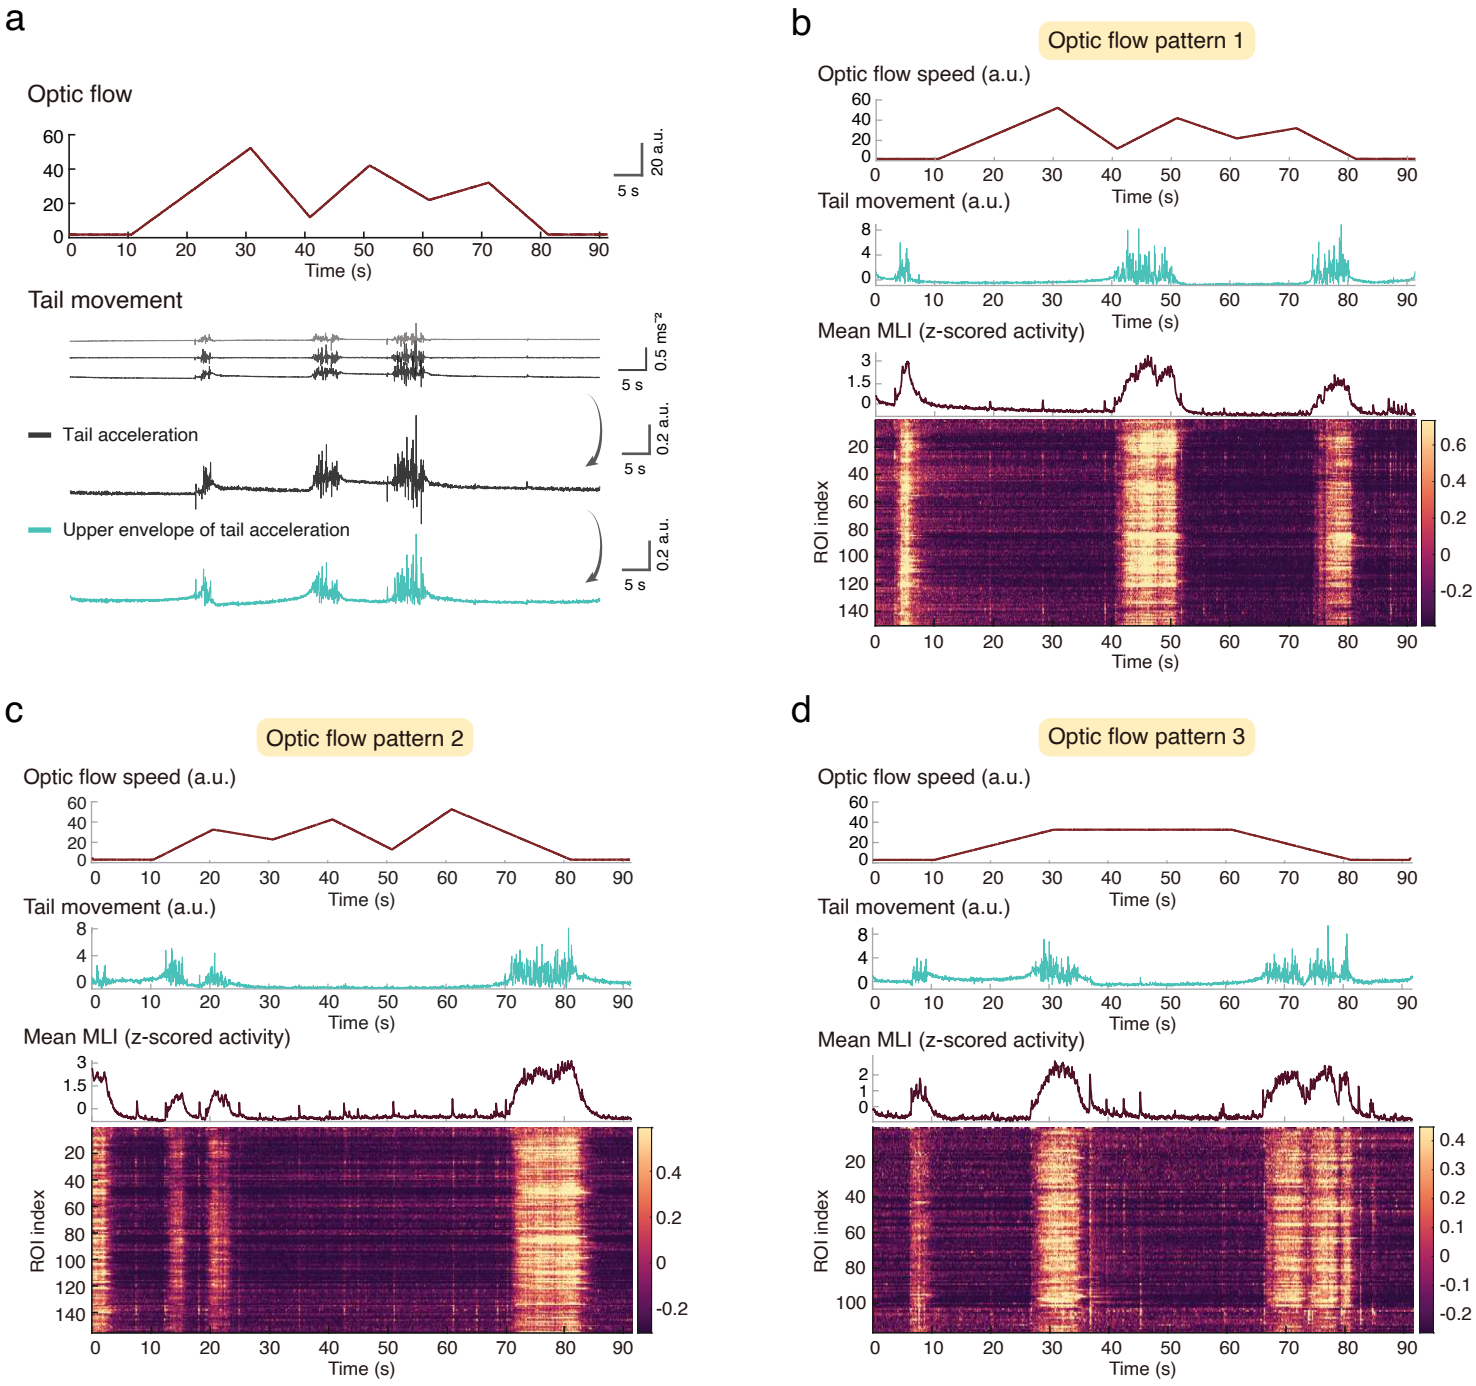
Supplementary Fig. 1 | Vascular landmarks, two-photon imaging fields, and node segmentation across subjects. a, Wide-field vascular maps showing surface vasculature landmarks used to guide two-photon imaging and to align imaging fields across sessions (n = 107 in MLI-containing trials; n = 88 in PC-containing trials; n = 6 mice). **b,** Representative two-photon calcium images from each mouse within lobules V and VI, capturing MLI somata (the candidate latent layer) and PC dendrites (the primary network output layer). **c,** ROI segmentation maps for extracting continuous signal trajectories from individual MLI and PC nodes across all 6 mice. Scale bars, 50 µm for (**b**) and (**c**).

Supplementary Fig. 2



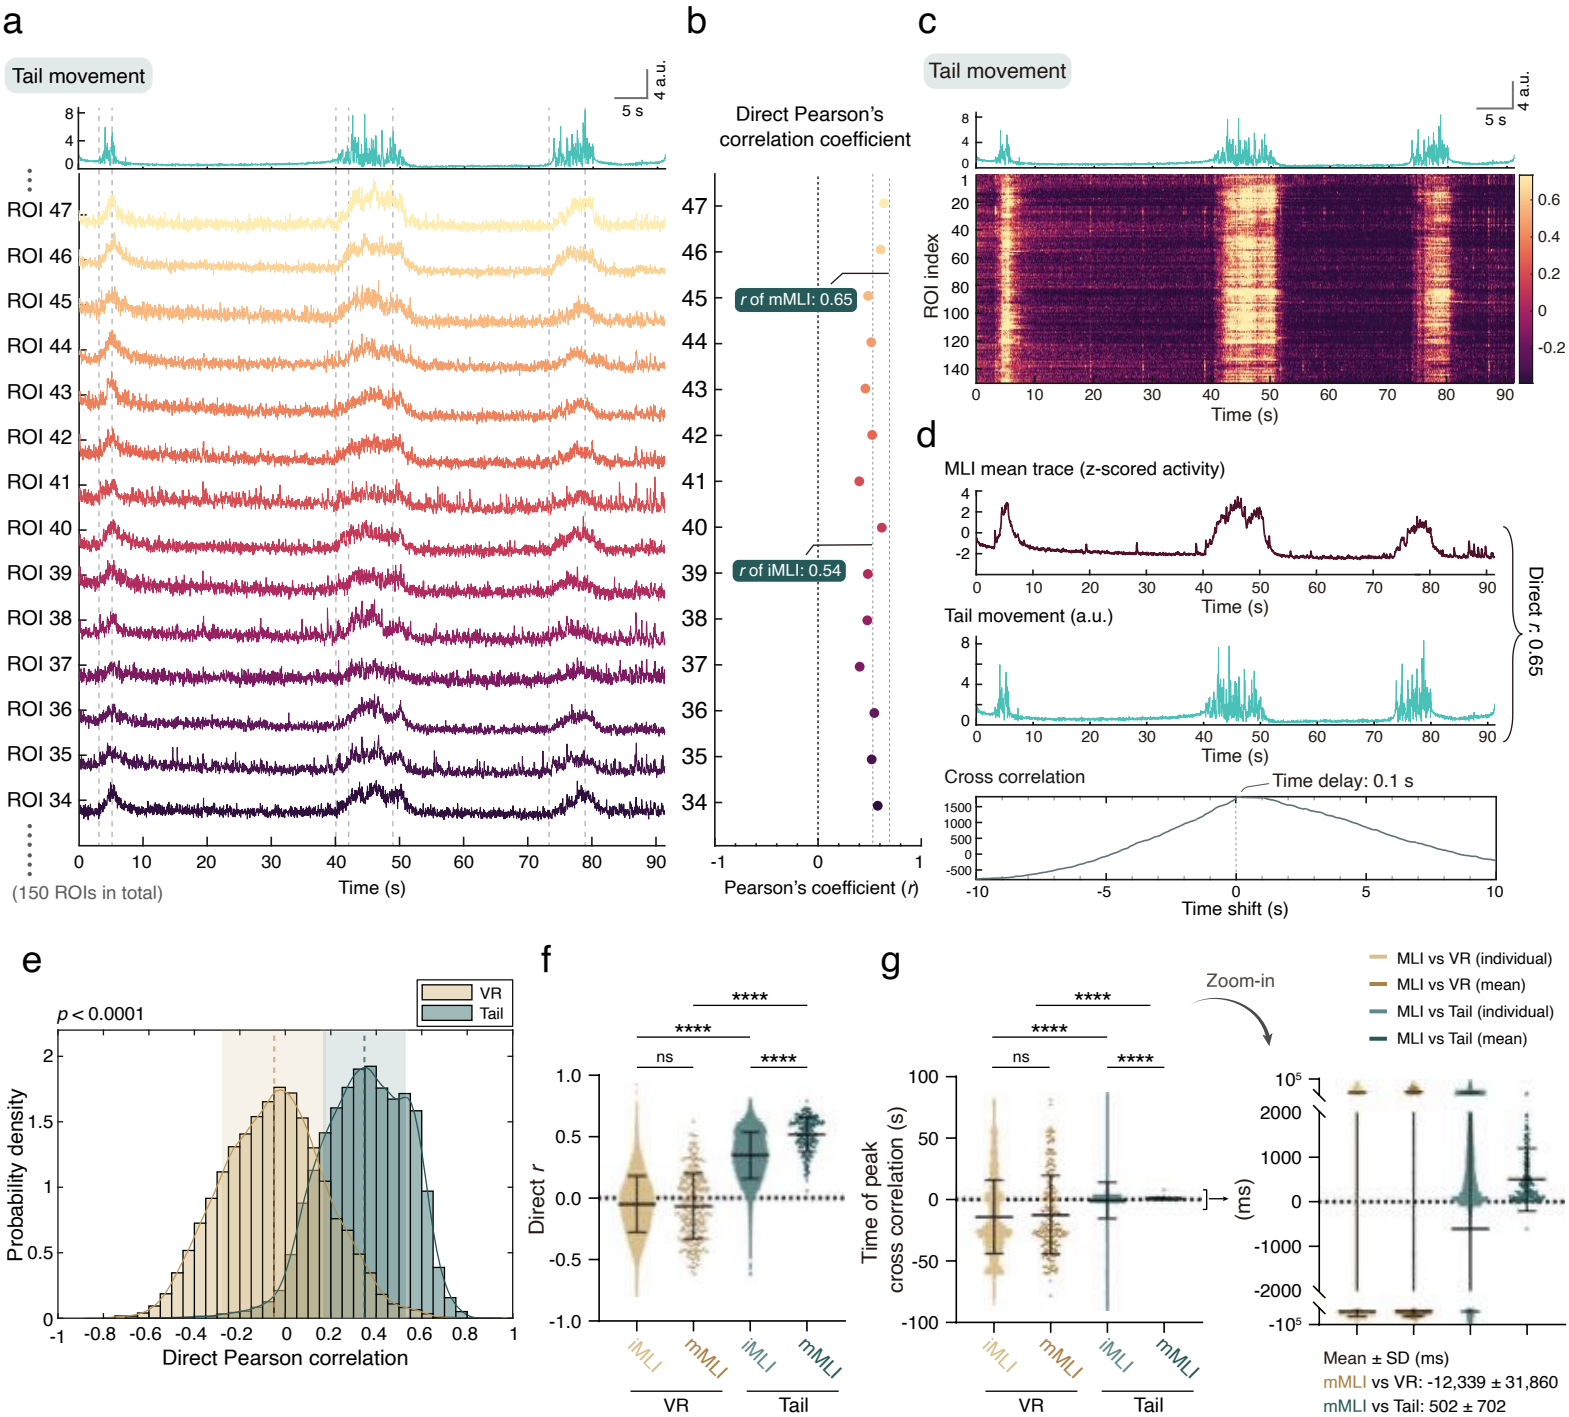
Supplementary Fig. 2 | Temporal relationship quantification of MLI activity as a delayed feedback estimator. a, Distribution of peak cross-correlation lags between individual MLI activity and the running speed state in the closed-loop VR dataset (n = 107 trials, n = 6 mice). The y-axis represents the proportion of individual MLI nodes exhibiting each peak-lag value. **b,** Distribution of temporal lags for the aggregated population dynamics relative to the running speed. The y-axis shows the proportion of trials. In contrast to individual nodes, the aggregated MLI population uniformly lagged the physical state of running speed across all continuous trajectories, validating a robust network-level feedback processing delay. **c,** Shuffled control for (a). Internal negative control comparing individual MLI trajectories against temporally shuffled speed states from the same trials. **d,** Shuffled control for (b). Internal negative control for the aggregated population dynamics. Temporal shuffling abolished the uniform processing lag observed in (b), confirming that the network-level delay represents a genuine biological state computation rather than a statistical artifact.

Supplementary Fig. 3



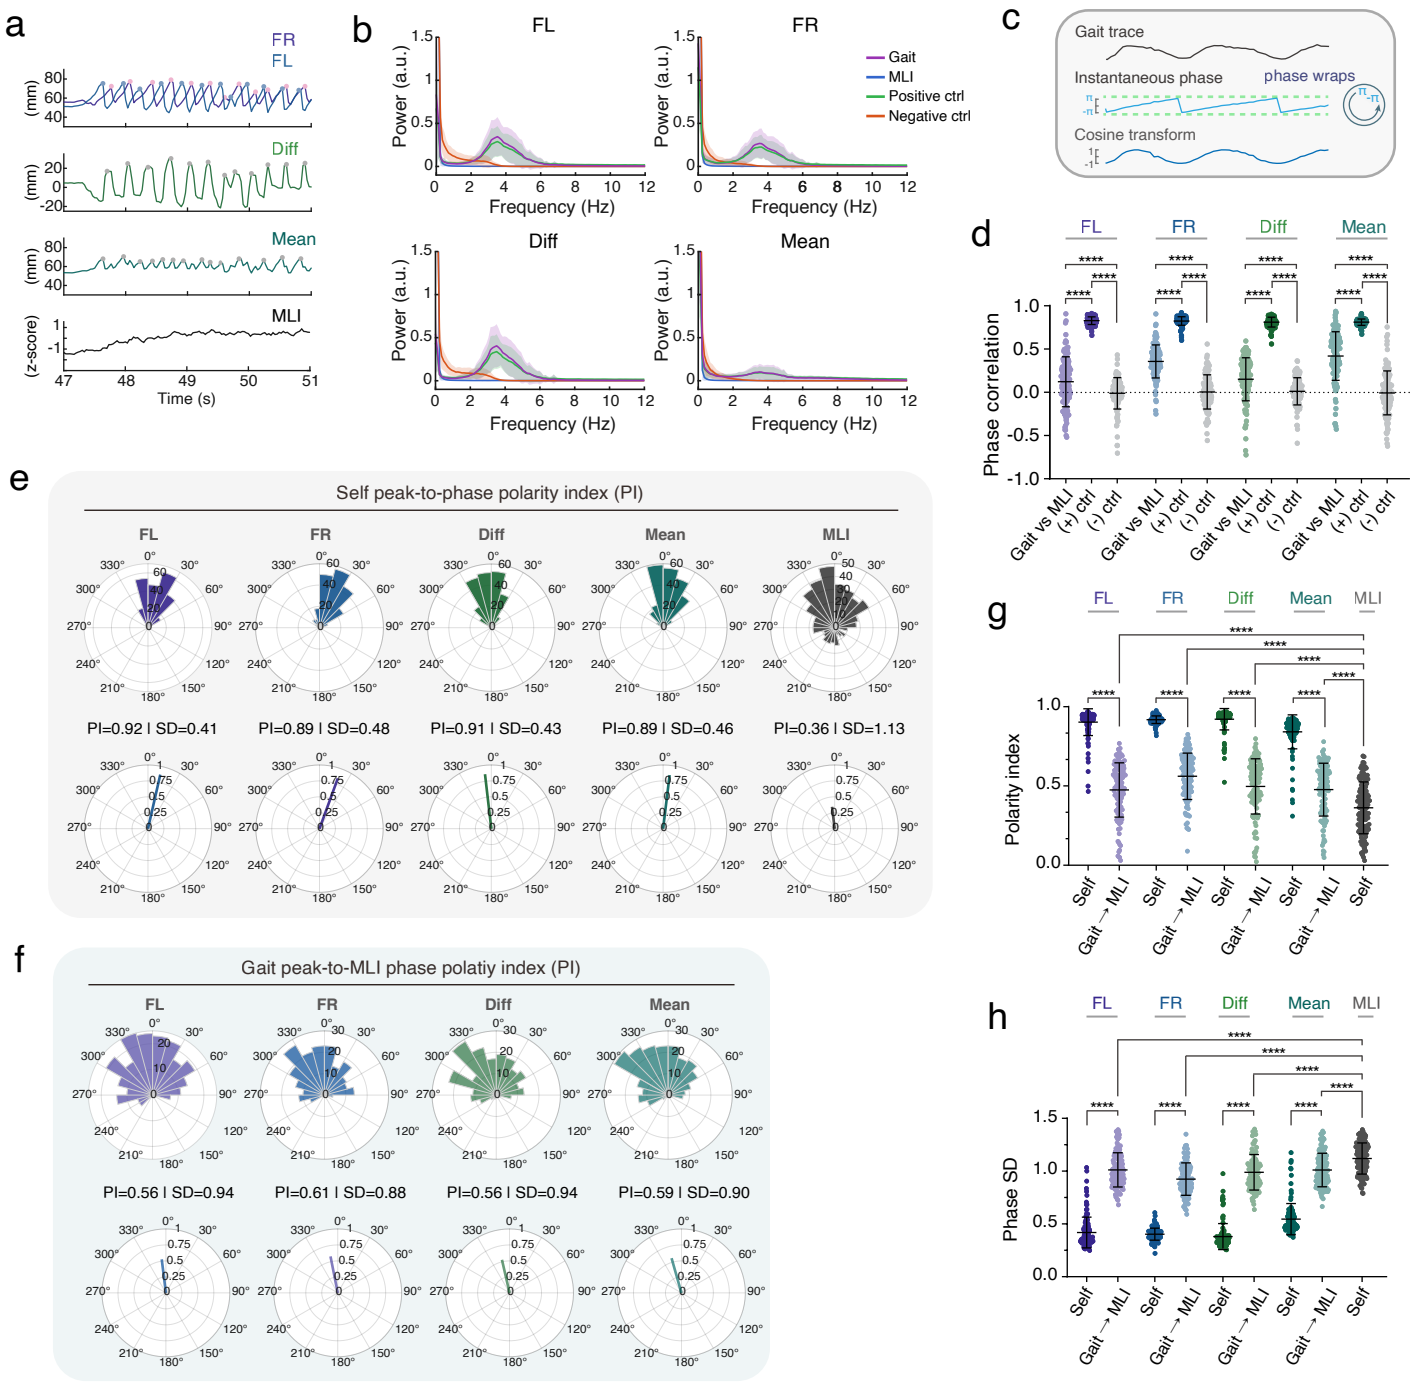
Supplementary Fig. 3 | Monitoring MLI network activity and tail kinematics across various visual playback patterns. a, Schematic of physical movement extraction using a three-axis accelerometer attached to the mouse tail. Acceleration signals along three orthogonal axes were integrated to compute the total acceleration magnitude, and the upper envelope was extracted to represent physical movement amplitude. This enabled the dissociation of visual state-space variables from actual somatosensory-driven physical motion. **b-d,** Example trials under open-loop conditions showing three distinct preprogrammed optic flow patterns. For each pattern, the top traces show the visual optic flow speed, the middle traces show the physical movement envelope derived from the accelerometer, and the bottom panels show the aggregated MLI population dynamics (z-scored) alongside heatmaps of individual MLI nodes from a representative trial. These examples demonstrate how MLIs were recorded across variable visual state-spaces while physical movement is simultaneously monitored.

Supplementary Fig. 4



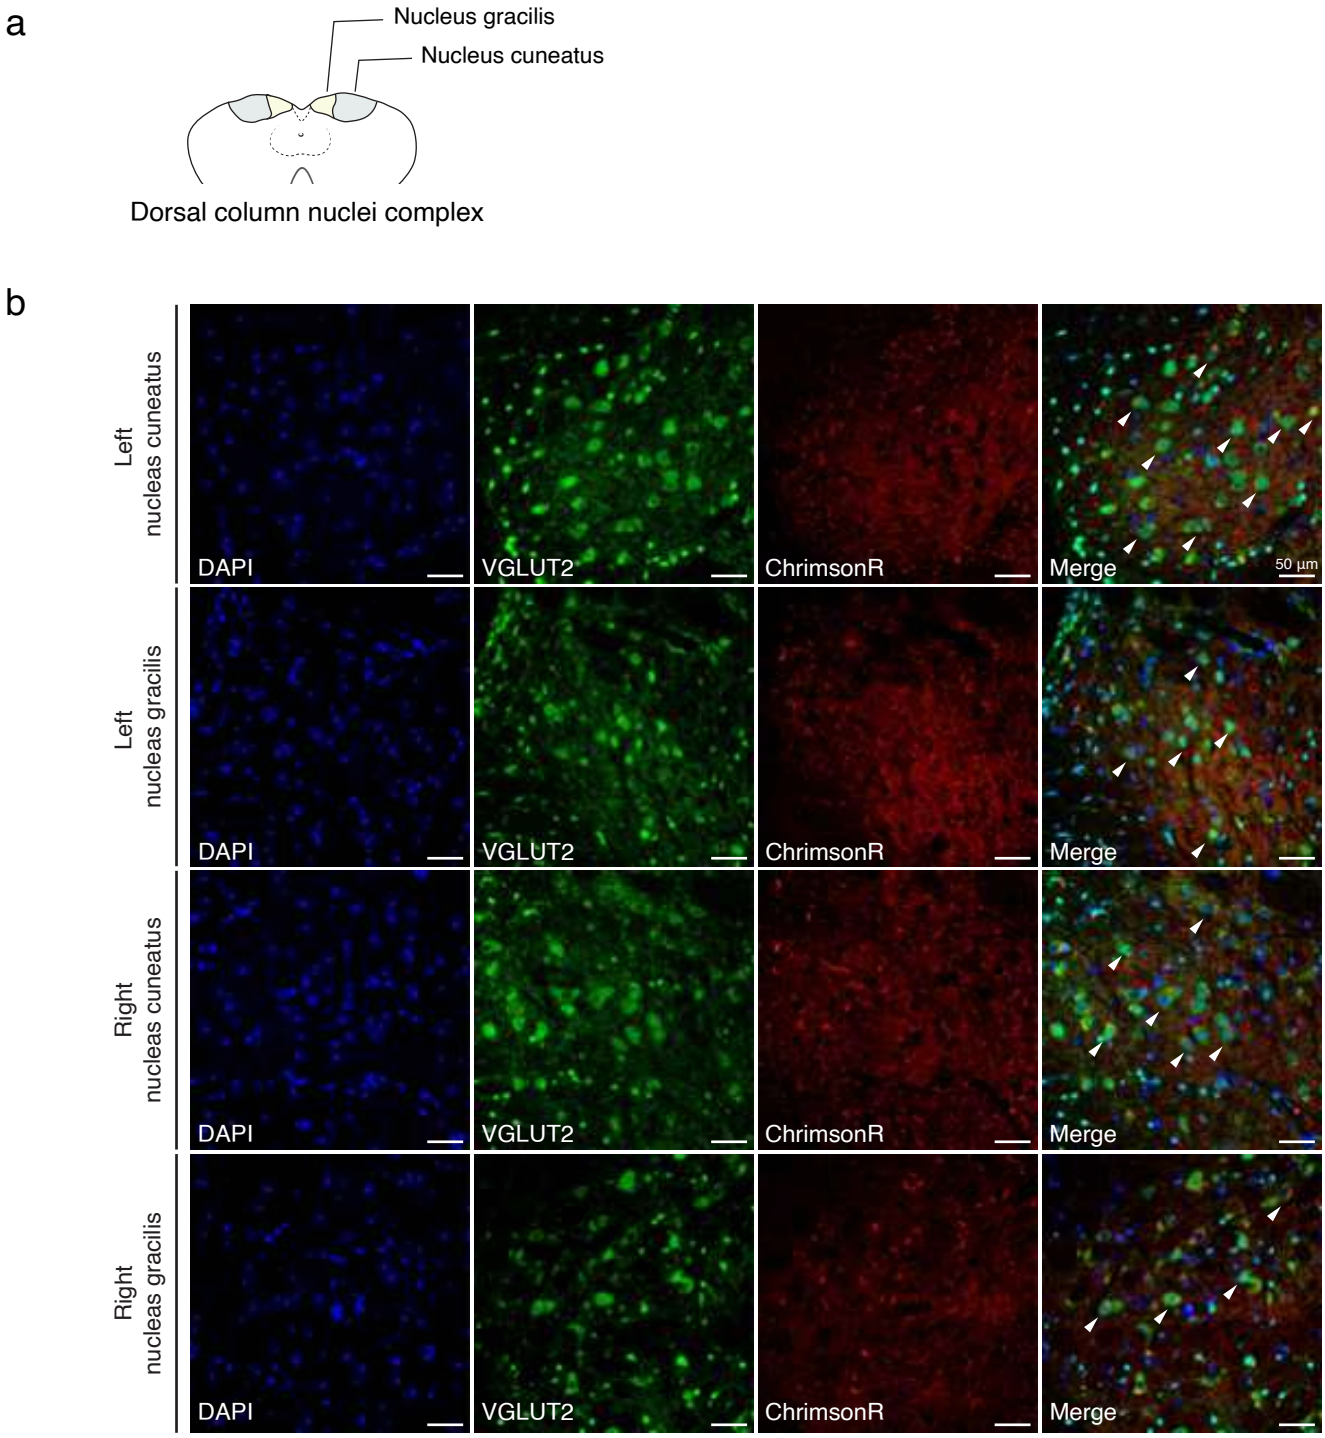
Supplementary Fig. 4 | MLI activity and its correlation with optic flow versus tail movement. a, Example trial showing tail-movement amplitude (top, cyan) and MLI activity (bottom). **b,** Pearson's correlation coefficients between tail movement and MLI activity for the ROIs shown in (**a**). Correlations are shown for individual MLI nodes (iMLIs) and aggregated population dynamics (mMLI). **c,** Heatmap of MLI population activity aligned to tail movement during the example trial. **d,** An example trace of aggregated MLI dynamics (top), tail-movement amplitudes (middle), and their cross-correlation function (bottom). **e,** Probability density distributions of Pearson's correlation coefficients between MLI activity and either tail-movement amplitude or VR optic flow speed, shown at the level of individual MLI calcium traces (n = 29,142 traces). **f,** Scatter plots comparing correlation coefficients for individual MLI traces (iMLIs; n = 29,142 traces from 6 mice) and aggregated population dynamics (mMLIs; n = 244 population means from 244 trials in 6 mice) with tail movement and VR optic flow. **g,** temporal processing lags derived from cross-correlation analyses between MLI activity and either VR optic flow or tail movement. Data are shown for individual MLI traces (iMLIs; n = 29,142 traces, 6 mice) and mMLIs (mMLIs; n = 244 population means from 244 trials in 6 mice). Insets show magnified distributions of delay values. Data are means ± SD; ****$p < 0.0001$; Welch's t-test.

Supplementary Fig. 5



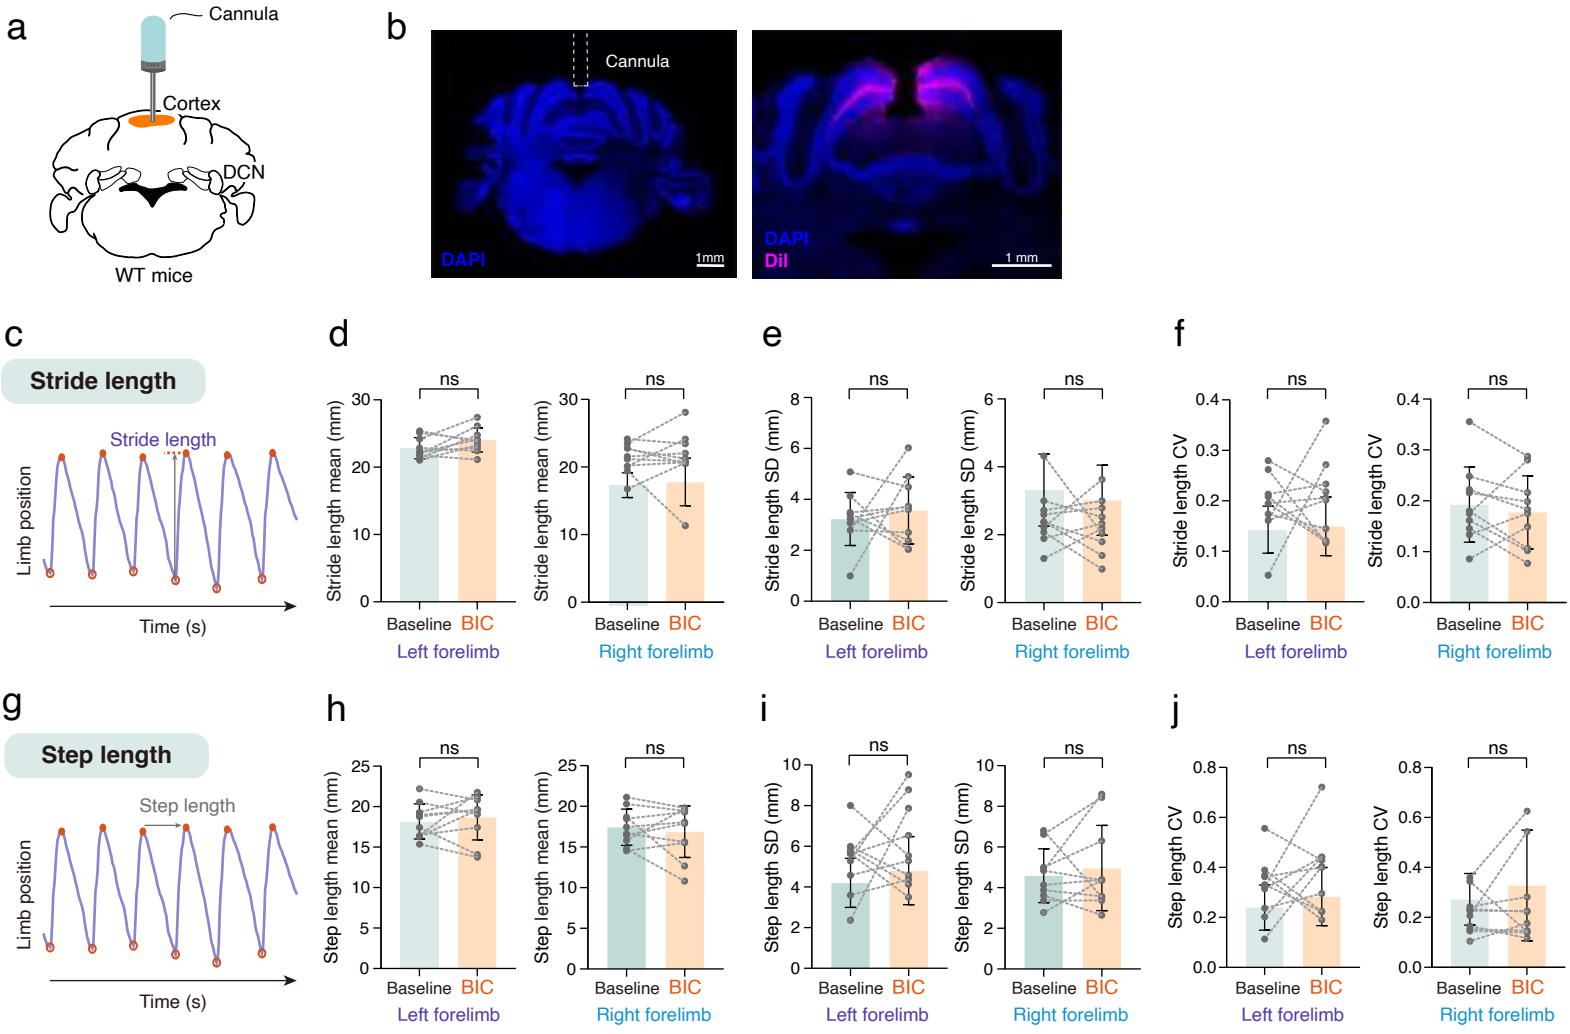
Supplementary Fig. 5 | Phase relationship analysis confirms the MLI latent space extracts a continuous scalar speed representation rather than cyclic gait variables. Analyses were performed on visual-deprivation trials to isolate somatosensory-driven locomotion and eliminate optic-flow interference. **a,** Example segment of high-dimensional gait traces alongside aggregated MLI dynamics for phase analysis. Front-left (FL) and front-right (FR) limb trajectories were extracted from video-based tracking. A difference component (Diff = (FR − FL)/2) and a mean component (Mean = (FR + FL)/2) were computed to represent alternating and shared components of limb kinematics, respectively. **b,** Spectral analysis of gait kinematics and MLI signals. FL and FR traces exhibited clear cyclic peaks near the step frequency (~3 Hz), while the MLI latent space displayed a broadband, non-periodic spectrum consistent with continuous scalar speed encoding rather than cyclic phase modulation. Shaded areas denote ± SD across trials. **c,** Schematic of phase extraction and transformation. Each trace was converted to its analytic signal using the Hilbert transform to obtain the instantaneous phase (–π to π). The cosine of this phase was computed to linearize polarity and remove discontinuities from phase wrapping, providing a continuous measure of oscillatory state for correlation analysis. **d,** Phase correlations between gait components and MLI activity under original, jittered, and shuffled conditions. To test whether MLI activity shared rhythmic phase structure with gait, both signals were Hilbert-transformed, cosine-projected, and correlated in phase space. Gait-MLI correlations remained exceptionally low across all conditions, indicating the MLI network does not share the step-locked phase structure of the multidimensional motor output. **e,** Example self peak-phase polarity index (PI) of gait and MLI signals from a single trial. Rhythmic gait components exhibited tightly clustered peak phases (PI ≈ 0.9), whereas the MLI latent space showed dispersed phase distributions with low PI (≈ 0.36), reflecting an absence of rhythmic phase locking. **f,** Example gait-to-MLI phase PI distributions from the same trial, confirming the network computes a continuous scalar magnitude rather than a discrete step-related cyclic variable. **g,** Population summary of PI across trials (n = 122 trials, from 6 mice). **h,** Gait signals displayed low SDs, indicating stable step-locked timing, whereas MLI self-phase SDs were large, reflecting continuous, non-periodic feedback extraction. Scatter plots represent individual trials (n = 122). Data are means ± SD; ****$p < 0.0001$; Welch's t-test.

Supplementary Fig. 6



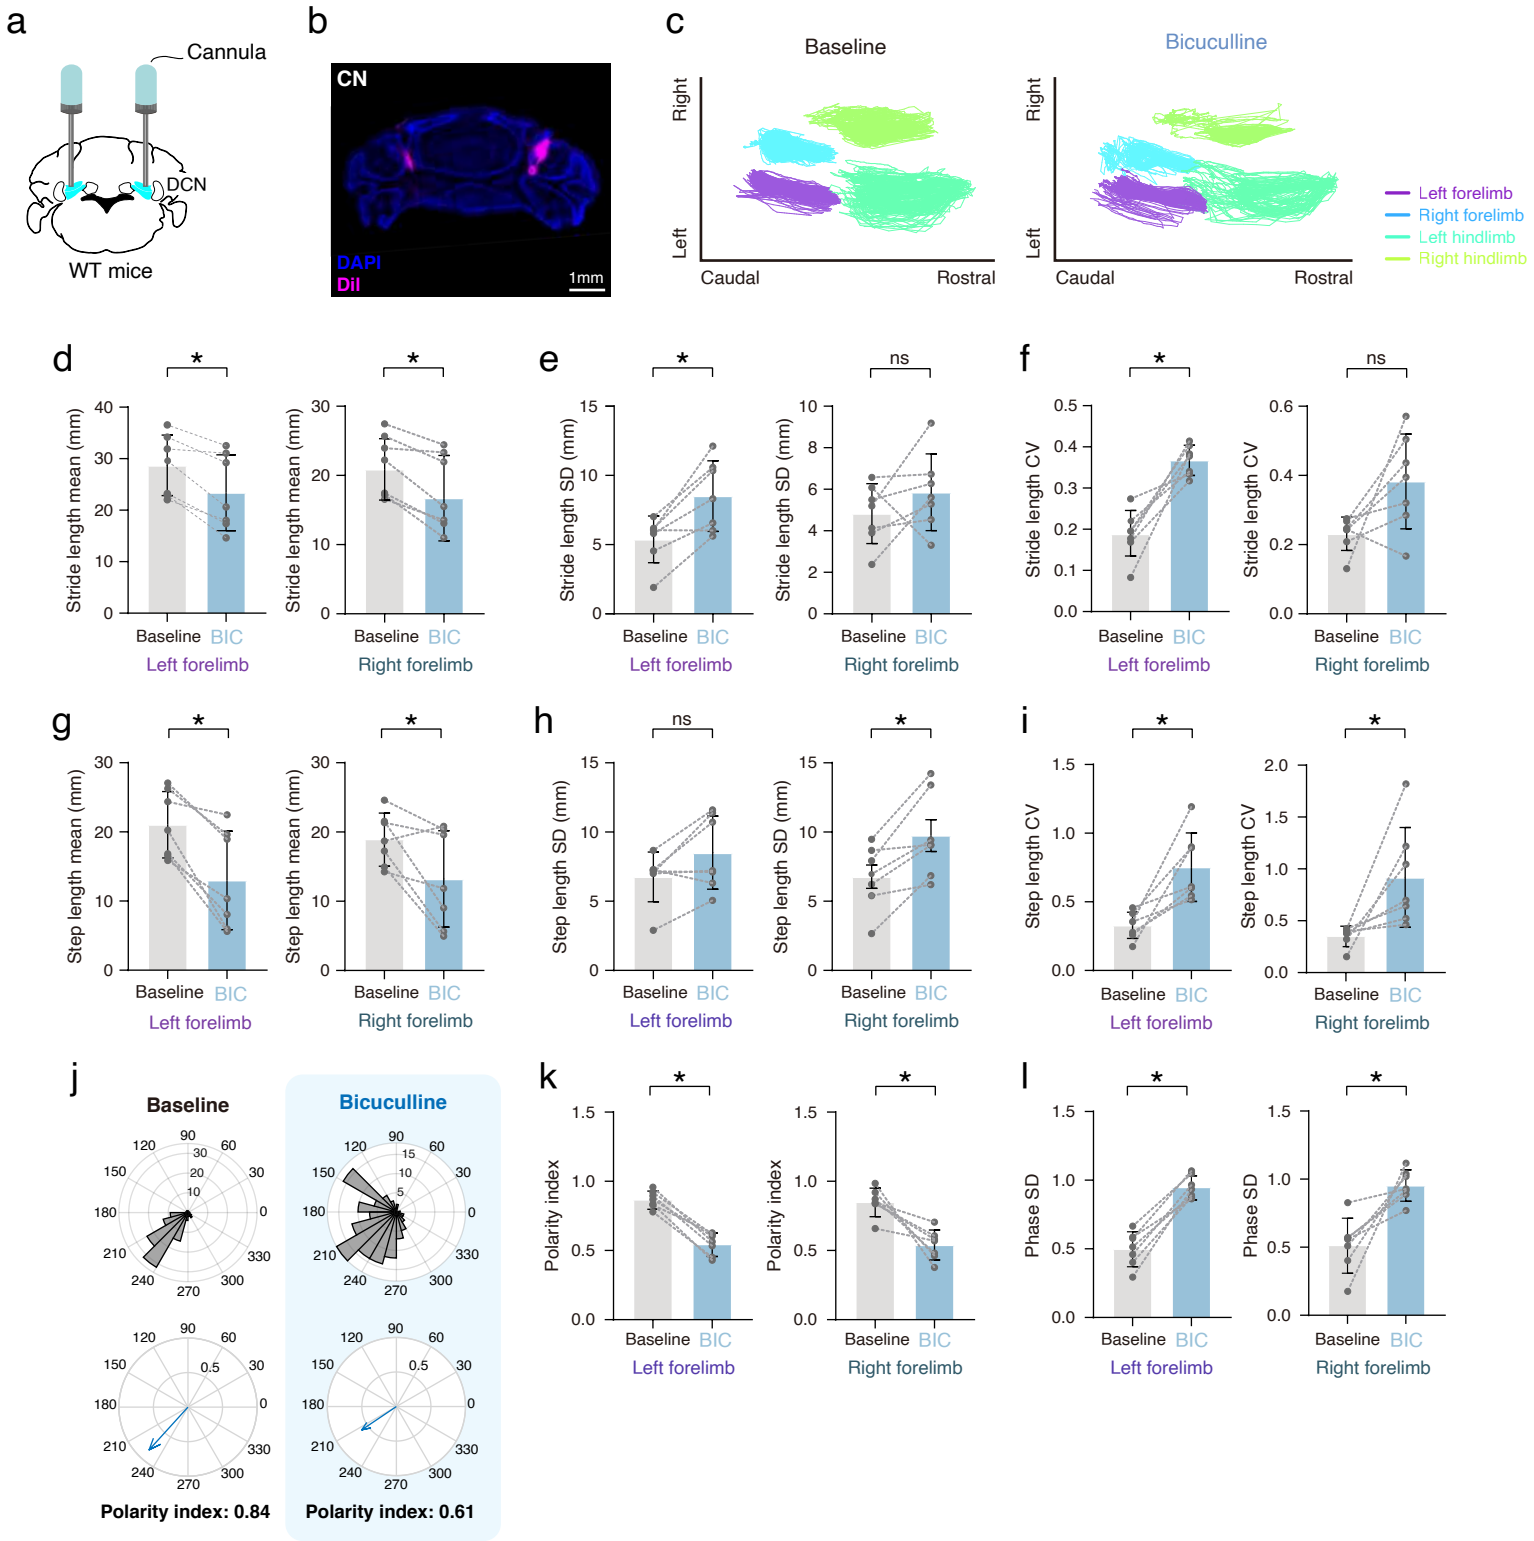
Supplementary Fig. 6 | Anatomical validation of targeted actuation in dorsal column nuclei, the ascending somatosensory-relay hardware. a, Schematic of the dorsal column nuclei complex showing the positions of the nucleus gracilis (NG) and nucleus cuneatus (CU) within the dorsal medulla. **b,** High-magnification (40x) images of the left and right nuclei demonstrating colocalization of tdTomato and VGLUT2 (white arrows) in large-soma neurons. These larger excitatory neurons correspond to the principal relay cells of the dorsal column nuclei that transmit low-threshold mechanosensory input to higher centers. The observed ChrimsonR-tdTomato/VGLUT2 colocalization confirms viral targeting of glutamatergic sensory-relay neurons in the nucleus gracilis and cuneatus. Scale bars, 50 µm.

Supplementary Fig. 7



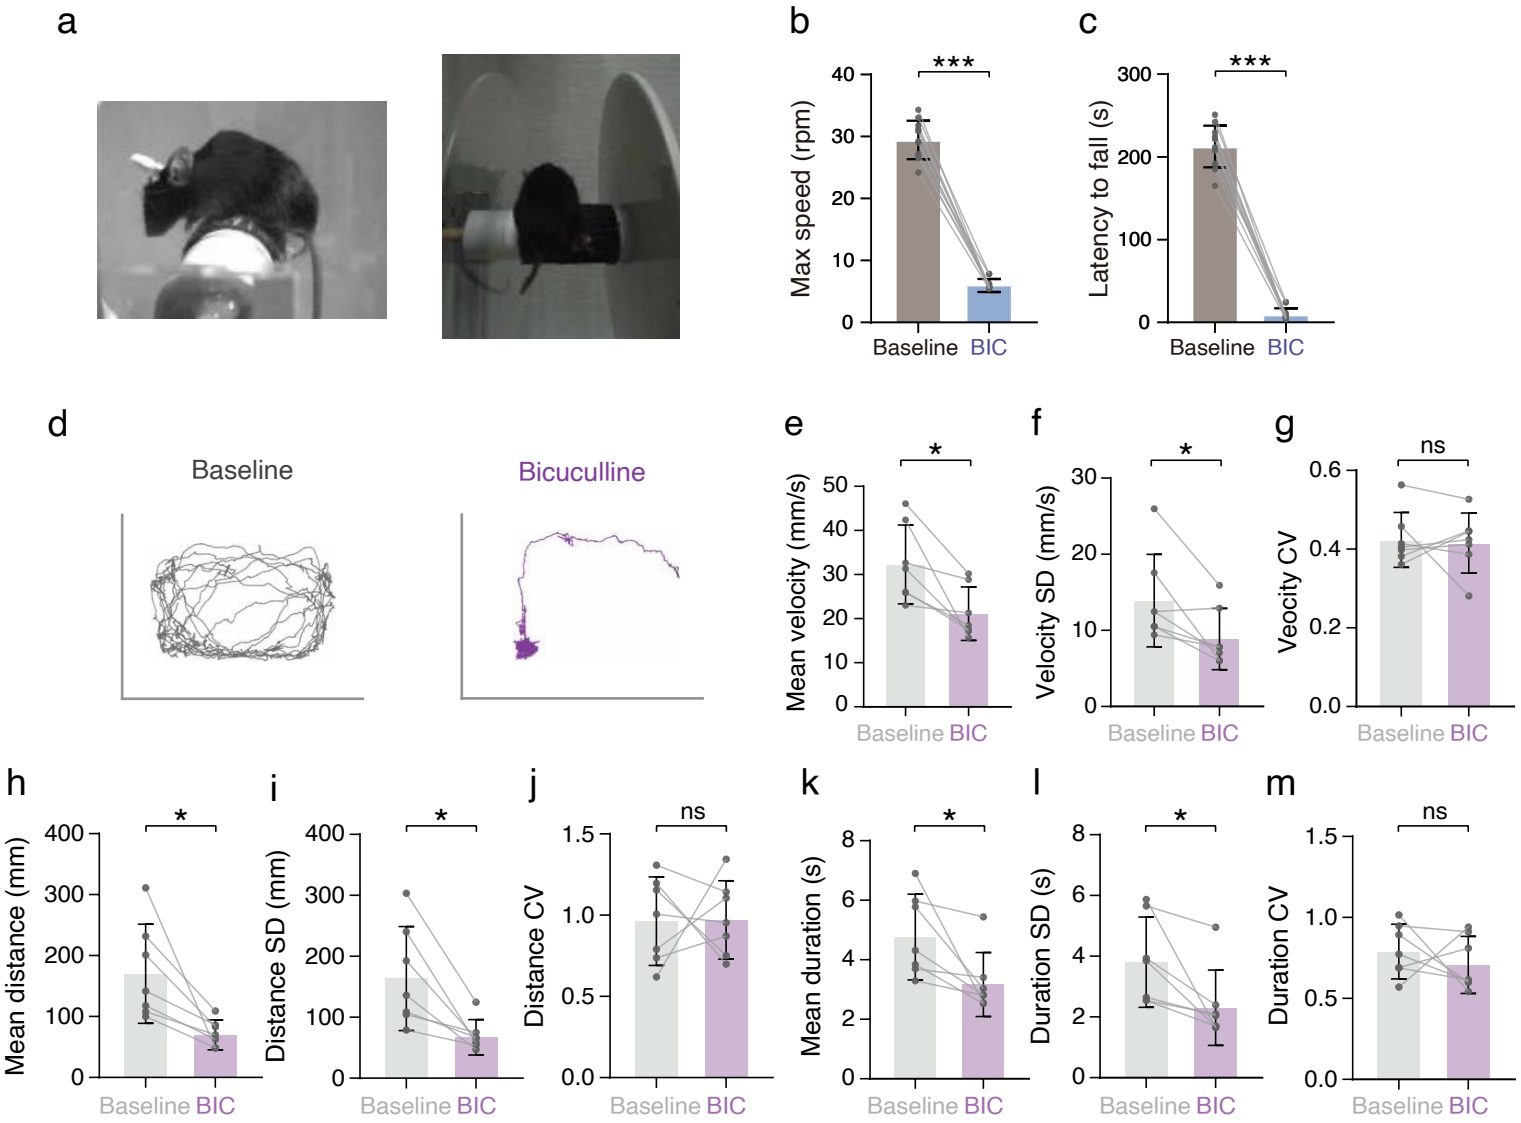
Supplementary Fig. 7 | Targeted suppression of the MLI latent space strictly preserves high-dimensional limb kinematics. **a,** Schematic of surgical targeting for bicuculline infusion into the cerebellar cortex via a centrally positioned cannula. **b,** Histological verification of cannula placement within the cerebellar cortex. Coronal section showing a central cannula track targeting lobule V, visualized by DiI tissue labeling (magenta) and DAPI nuclear staining (blue). Scale bars, 1mm. **c,** Schematic illustrating the definition of stride length, a fundamental component of the multidimensional motor output, derived from paw trajectories captured by high-speed videography during locomotion. **d-f,** Stride length metrics for the left and right forelimbs: mean stride length (**d**), stride length variability (SD; **e**), and coefficient of variation (CV; **f**) across baseline and bicuculline conditions. **g,** Schematic illustrating the definition of step length measured from inter-paw distances across consecutive gait cycles. **h-j,** Step length metrics for the left and right forelimbs: mean step length (**h**), step length variability (SD; **i**), and step length CV (**j**) across baseline and bicuculline conditions. Data are from 10 mice and shown as means ± SD; Wilcoxon signed-rank test.

Supplementary Fig. 8



Supplementary Fig. 8 | Disruption of the primary output node of DCN compromises global gait coordination. a, Schematic of bilateral cannula placement targeting bilateral DCN for primary output disruption by bicuculline. **b,** Histological verification of bilateral cannula placement within the DCN. Coronal cerebellar section showing bilateral cannula tracks in the DCN. Cannula locations were visualized by DiI tissue labeling (magenta) and counterstained with DAPI (blue). Scale bar, 1mm. **c,** Example multidimensional paw tracking during locomotion showing forelimb and hindlimb trajectories under baseline and bicuculline conditions. **d-f,** Paired stride length metrics for the left and right forelimbs: mean stride length (**d**), stride length variability (SD; **e**), and stride length CV (**f**) across baseline and bicuculline conditions ($n = 7$ mice). **g-i,** Paired step length metrics for the left and right forelimbs: mean step length (**g**), step length variability (SD; **h**), and step length CV (**i**) across baseline and bicuculline conditions. **j,** PI quantifying inter-forelimb coordination before and after bicuculline infusion. **k,** Quantification of the PI for left and right forelimbs. **l,** Step-phase variability (SD) for left and right forelimbs in baseline and under bicuculline. Data are from 7 mice and shown as means ± SD; *$p < 0.05$; Wilcoxon signed-rank test.

Supplementary Fig. 9



Supplementary Fig. 9 | Disruption of the primary output node of DCN compromises forced and spontaneous locomotor activity. a, Photographs of the rotarod apparatus used for motor endurance testing. **b-c,** Rotarod performance before and after bicuculline infusion: maximum rotation speed (**b**) and latency to fall (**c**) (n = 11 mice). **d,** Representative open-field trajectories from a single mouse under baseline and bicuculline conditions. **e-g,** Quantification of open-field velocity: mean velocity (**e**), variability (SD; **f**), and coefficient of variation (CV; **g**) (n = 7 mice). **h-j,** Quantification of open-field distance traveled: mean distance (**h**), variability (SD; **i**), and CV (**j**). **k-m,** Movement bout duration metrics: mean duration (**k**), variability (SD; **l**), and CV (**m**) (n = 7 mice). Data are means ± SD; *$p < 0.05$, ***$p < 0.001$; Wilcoxon signed-rank test.

Supplementary Materials for

Cerebellar interneurons execute a low-dimensional control policy for real-time kinematic calibration

Jou-Yu Ho, Po-Wei Chen, Chih-An Lin, Jye-Chang Lee, Shun-Chi Wu, Ming-Kai Pan*

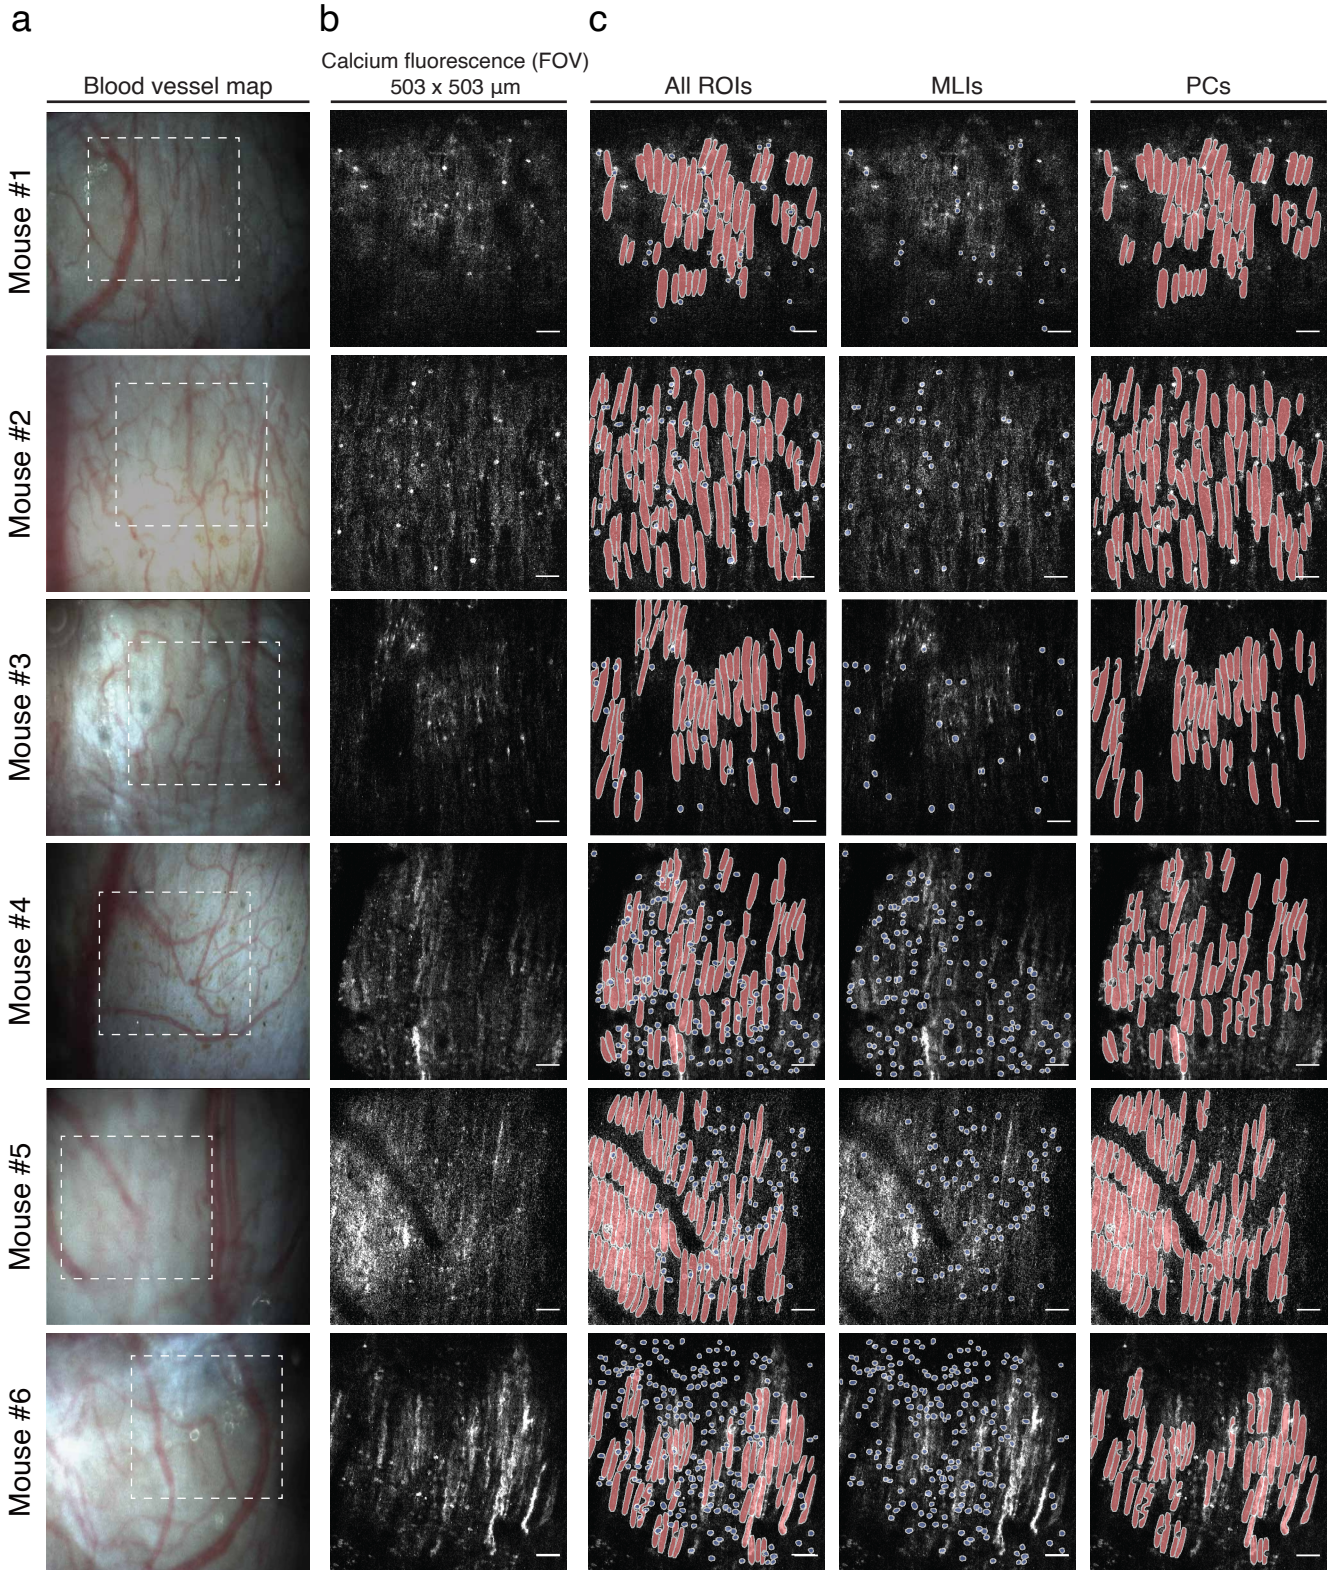
Corresponding author: Ming-Kai Pan, emorymkpan@ntu.edu.tw

The PDF file includes:

Supplementary Figs. 1 to 9

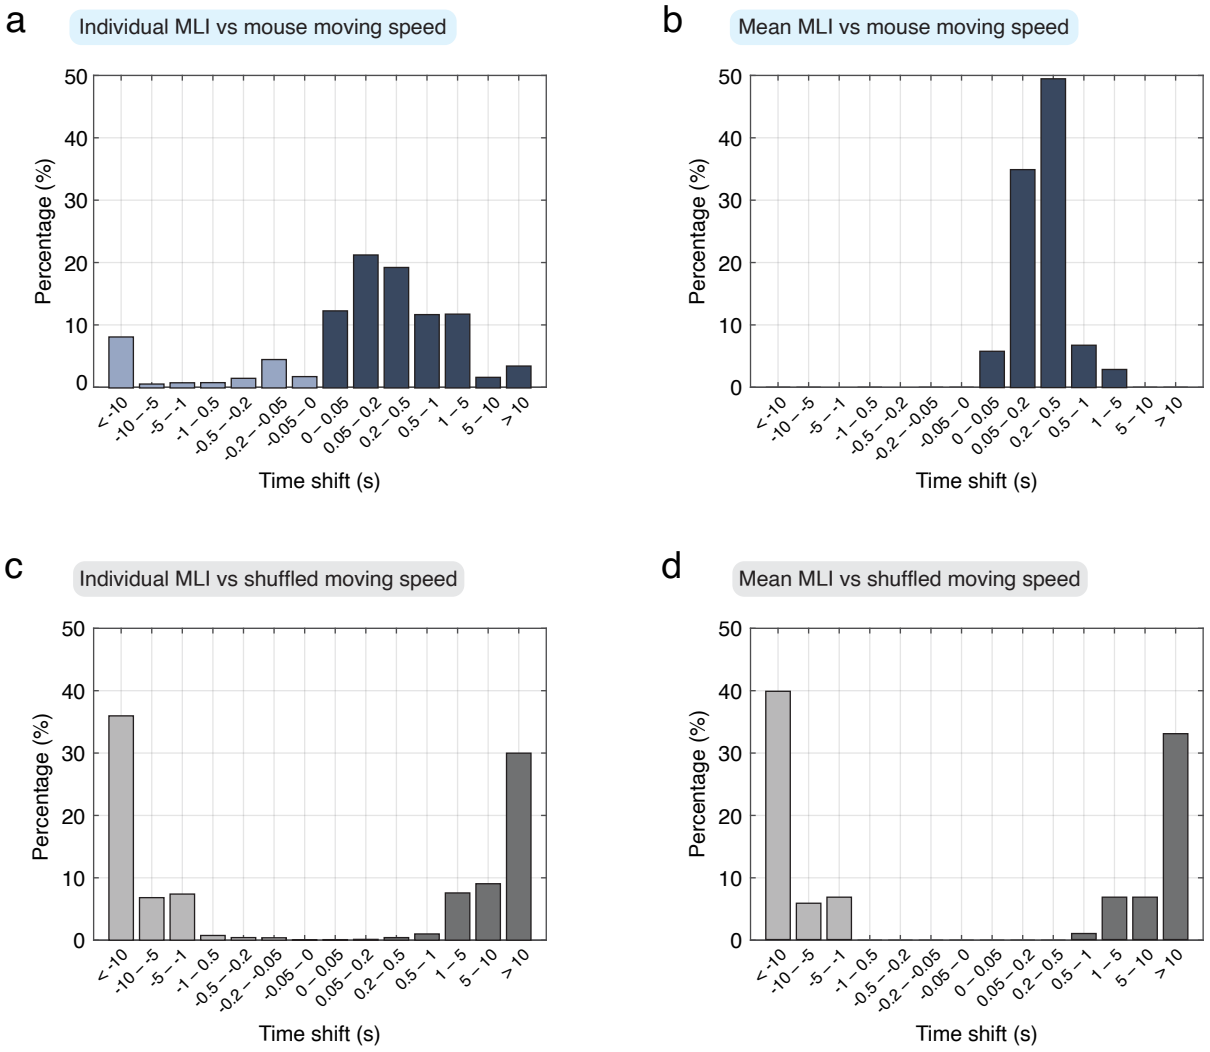
Supplementary Videos S1 to S4 (legend only)

Supplementary Fig. 1



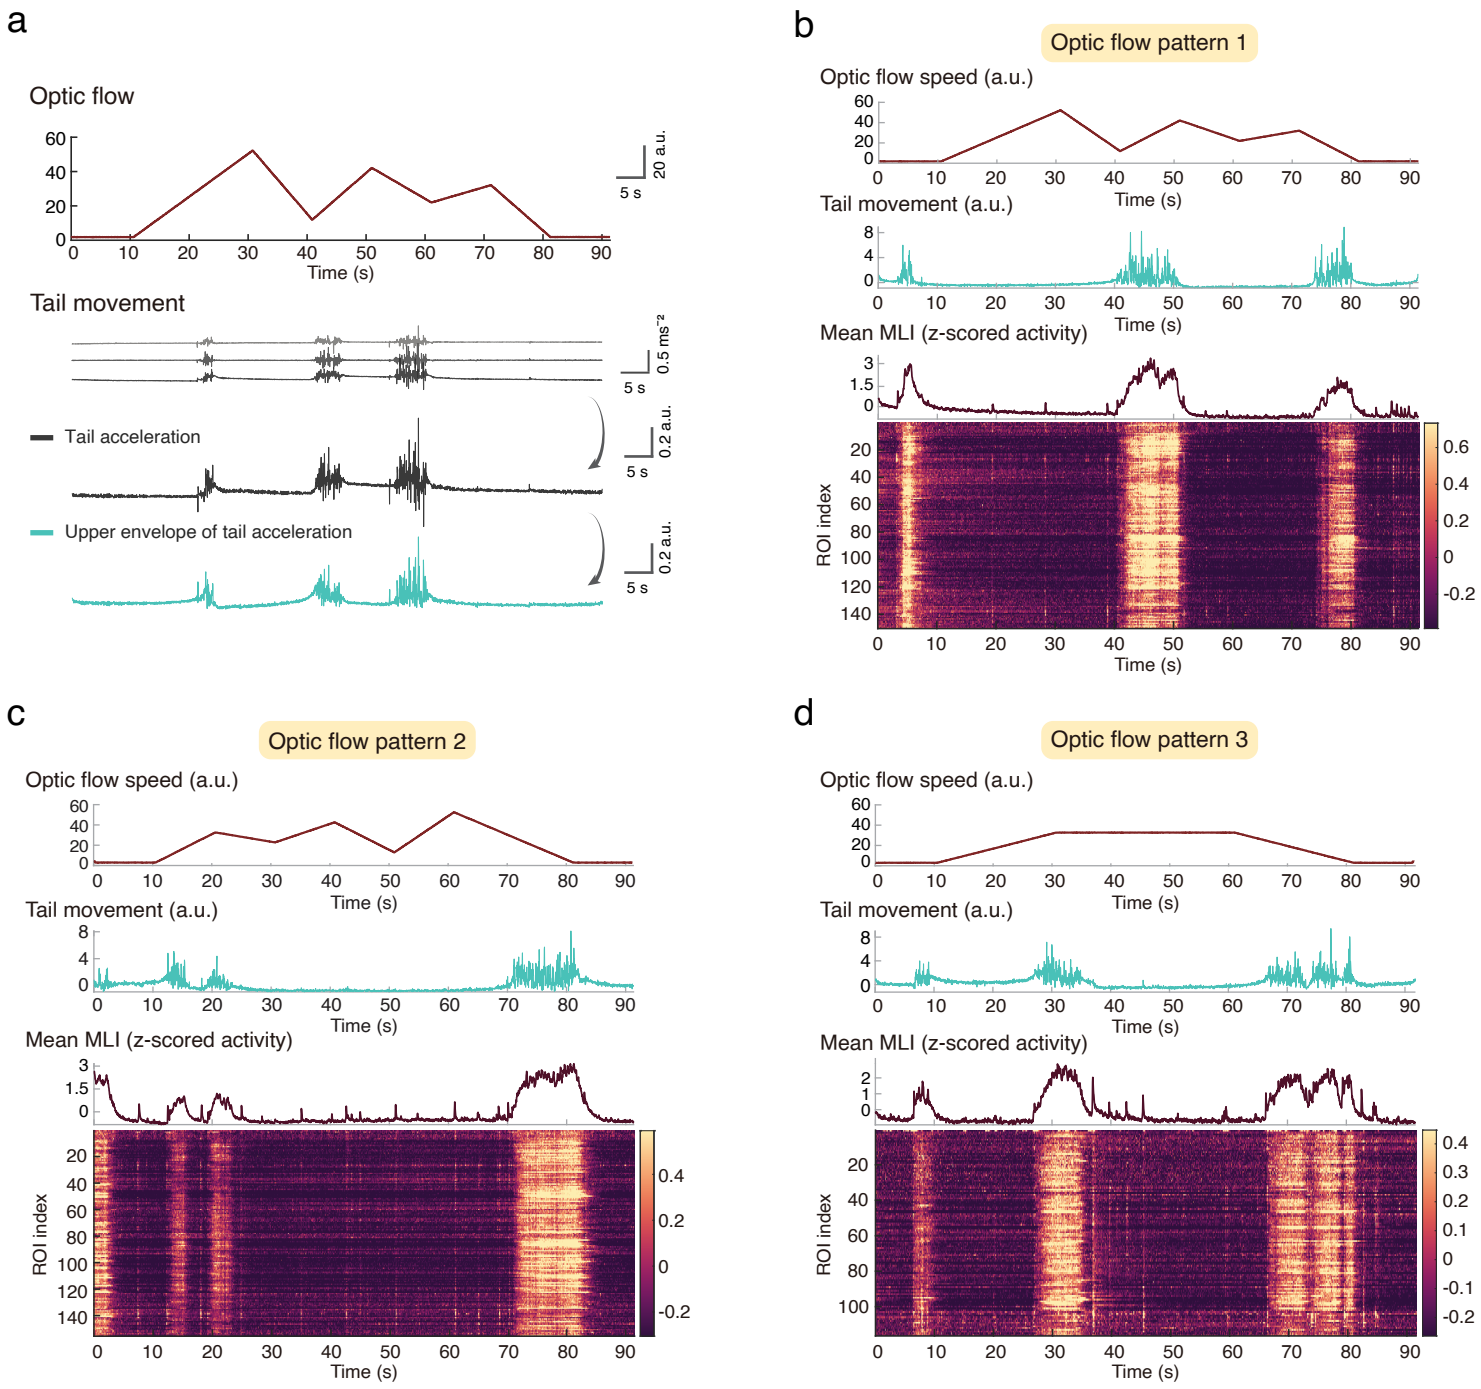
Supplementary Fig. 1 | Vascular landmarks, two-photon imaging fields, and node segmentation across subjects. a, Wide-field vascular maps showing surface vasculature landmarks used to guide two-photon imaging and to align imaging fields across sessions (n = 107 in MLI-containing trials; n = 88 in PC-containing trials; n = 6 mice). **b,** Representative two-photon calcium images from each mouse within lobules V and VI, capturing MLI somata (the candidate latent layer) and PC dendrites (the primary network output layer). **c,** ROI segmentation maps for extracting continuous signal trajectories from individual MLI and PC nodes across all 6 mice. Scale bars, 50 µm for (**b**) and (**c**).

Supplementary Fig. 2



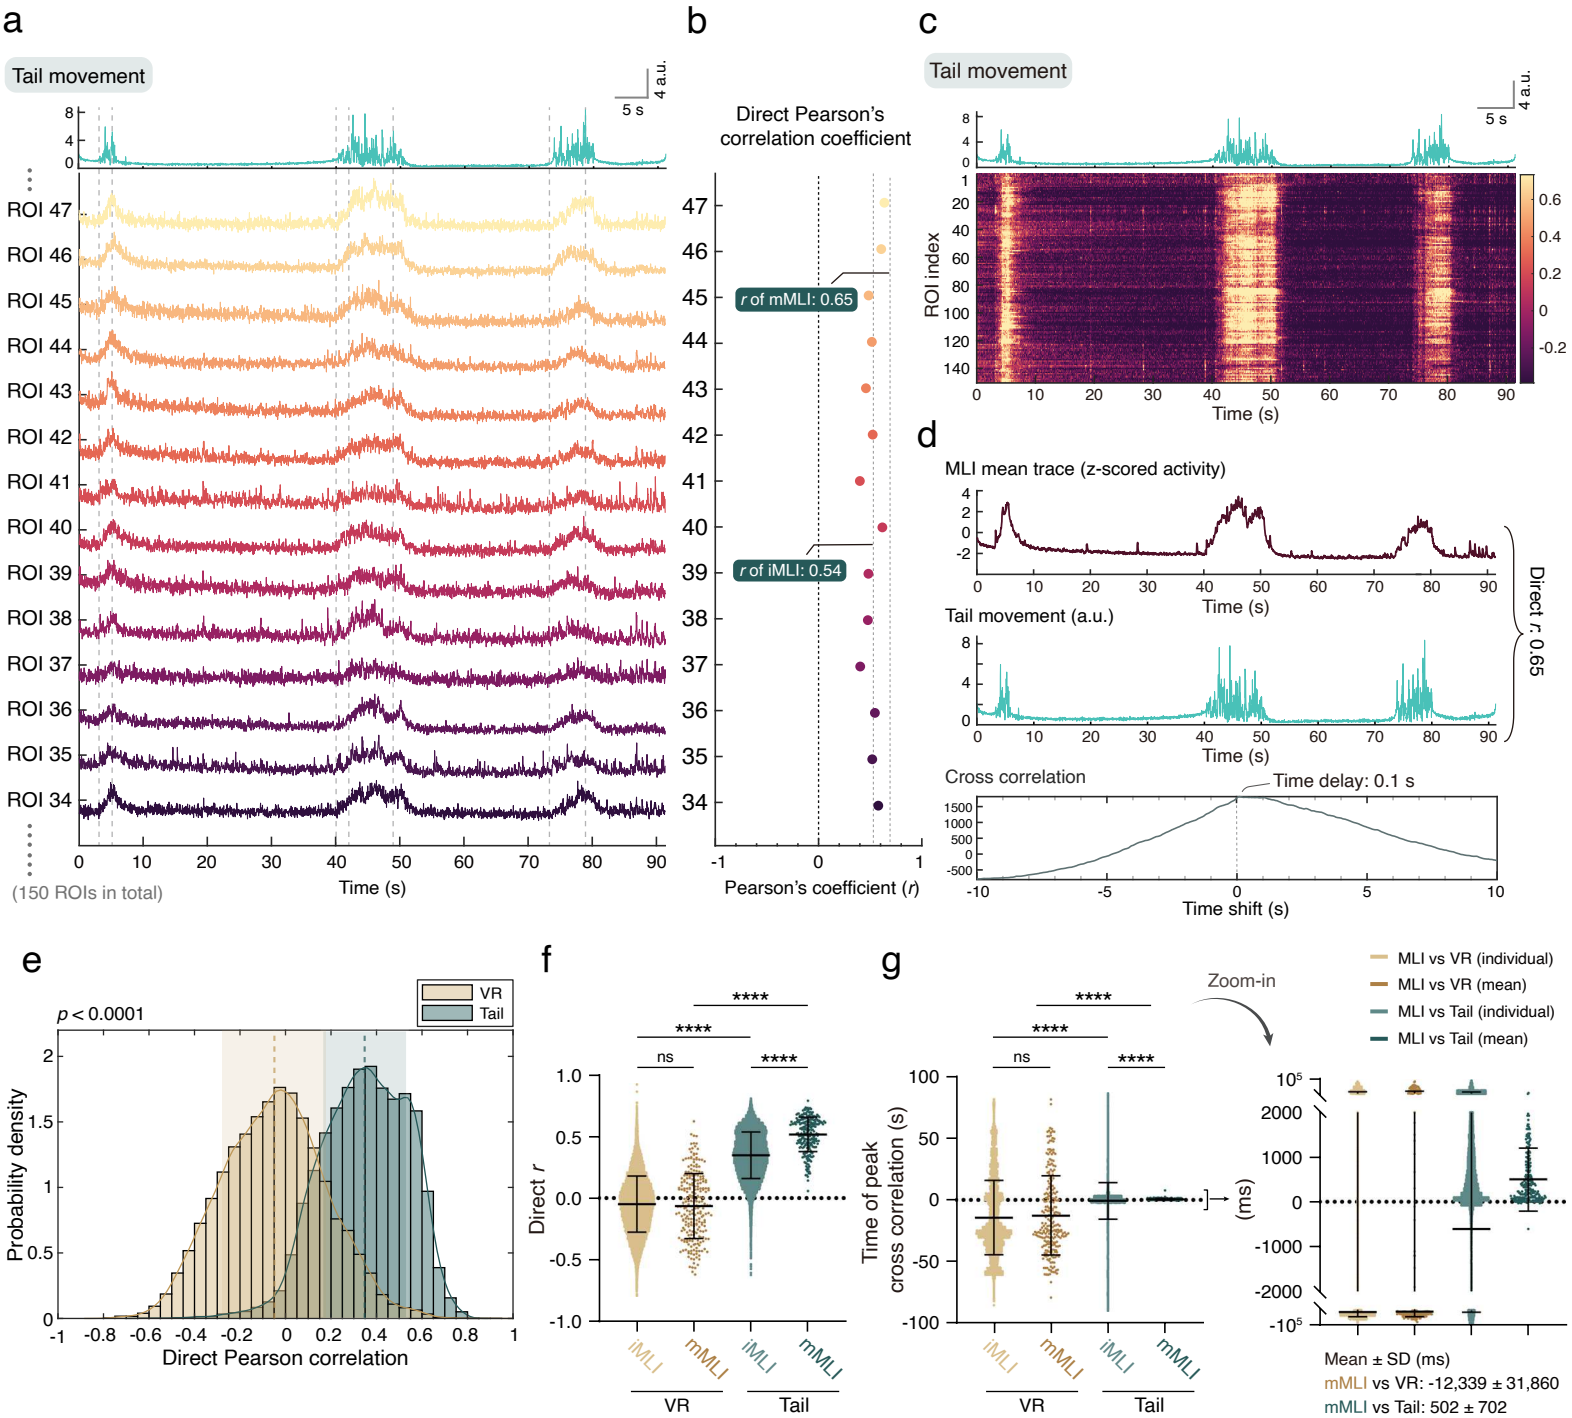
Supplementary Fig. 2 | Temporal relationship quantification of MLI activity as a delayed feedback estimator. a, Distribution of peak cross-correlation lags between individual MLI activity and the running speed state in the closed-loop VR dataset (n = 107 trials, n = 6 mice). The y-axis represents the proportion of individual MLI nodes exhibiting each peak-lag value. **b,** Distribution of temporal lags for the aggregated population dynamics relative to the running speed. The y-axis shows the proportion of trials. In contrast to individual nodes, the aggregated MLI population uniformly lagged the physical state of running speed across all continuous trajectories, validating a robust network-level feedback processing delay. **c,** Shuffled control for (a). Internal negative control comparing individual MLI trajectories against temporally shuffled speed states from the same trials. **d,** Shuffled control for (b). Internal negative control for the aggregated population dynamics. Temporal shuffling abolished the uniform processing lag observed in (b), confirming that the network-level delay represents a genuine biological state computation rather than a statistical artifact.

Supplementary Fig. 3



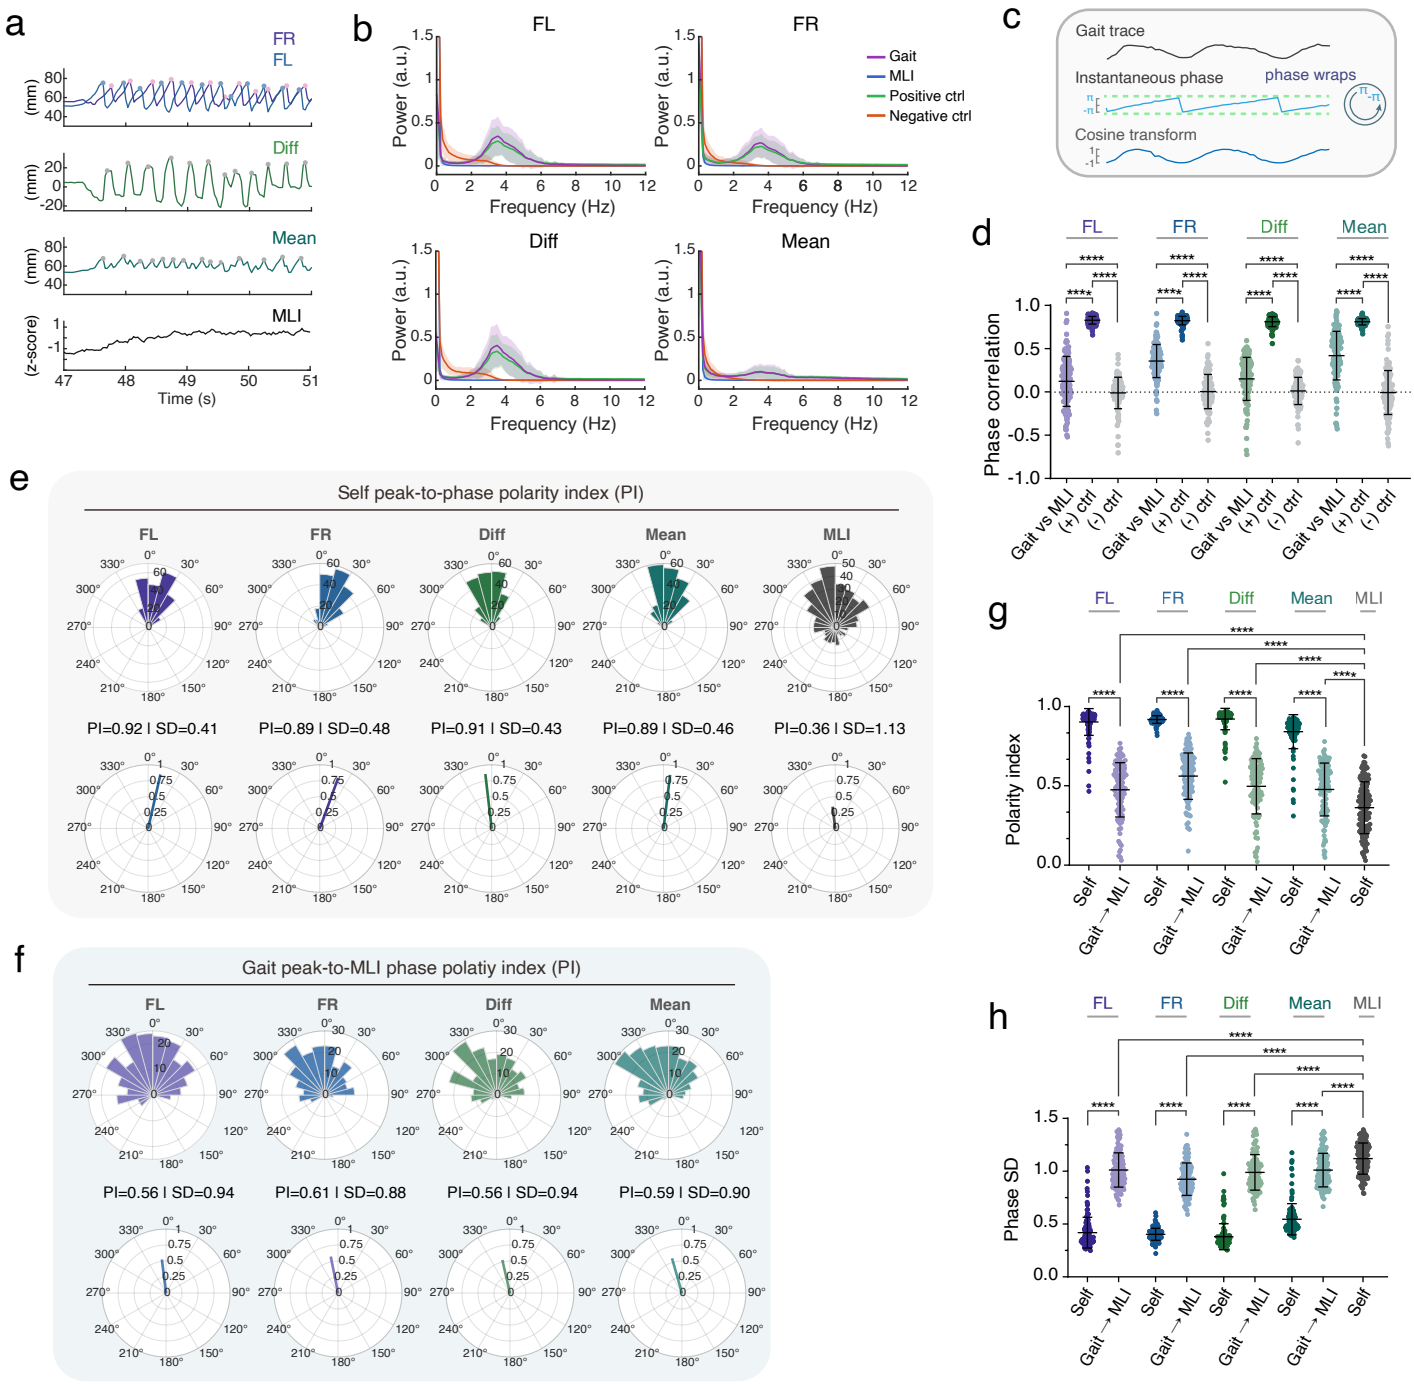
Supplementary Fig. 3 | Monitoring MLI network activity and tail kinematics across various visual playback patterns. a, Schematic of physical movement extraction using a three-axis accelerometer attached to the mouse tail. Acceleration signals along three orthogonal axes were integrated to compute the total acceleration magnitude, and the upper envelope was extracted to represent physical movement amplitude. This enabled the dissociation of visual state-space variables from actual somatosensory-driven physical motion. **b-d,** Example trials under open-loop conditions showing three distinct preprogrammed optic flow patterns. For each pattern, the top traces show the visual optic flow speed, the middle traces show the physical movement envelope derived from the accelerometer, and the bottom panels show the aggregated MLI population dynamics (z-scored) alongside heatmaps of individual MLI nodes from a representative trial. These examples demonstrate how MLIs were recorded across variable visual state-spaces while physical movement is simultaneously monitored.

Supplementary Fig. 4



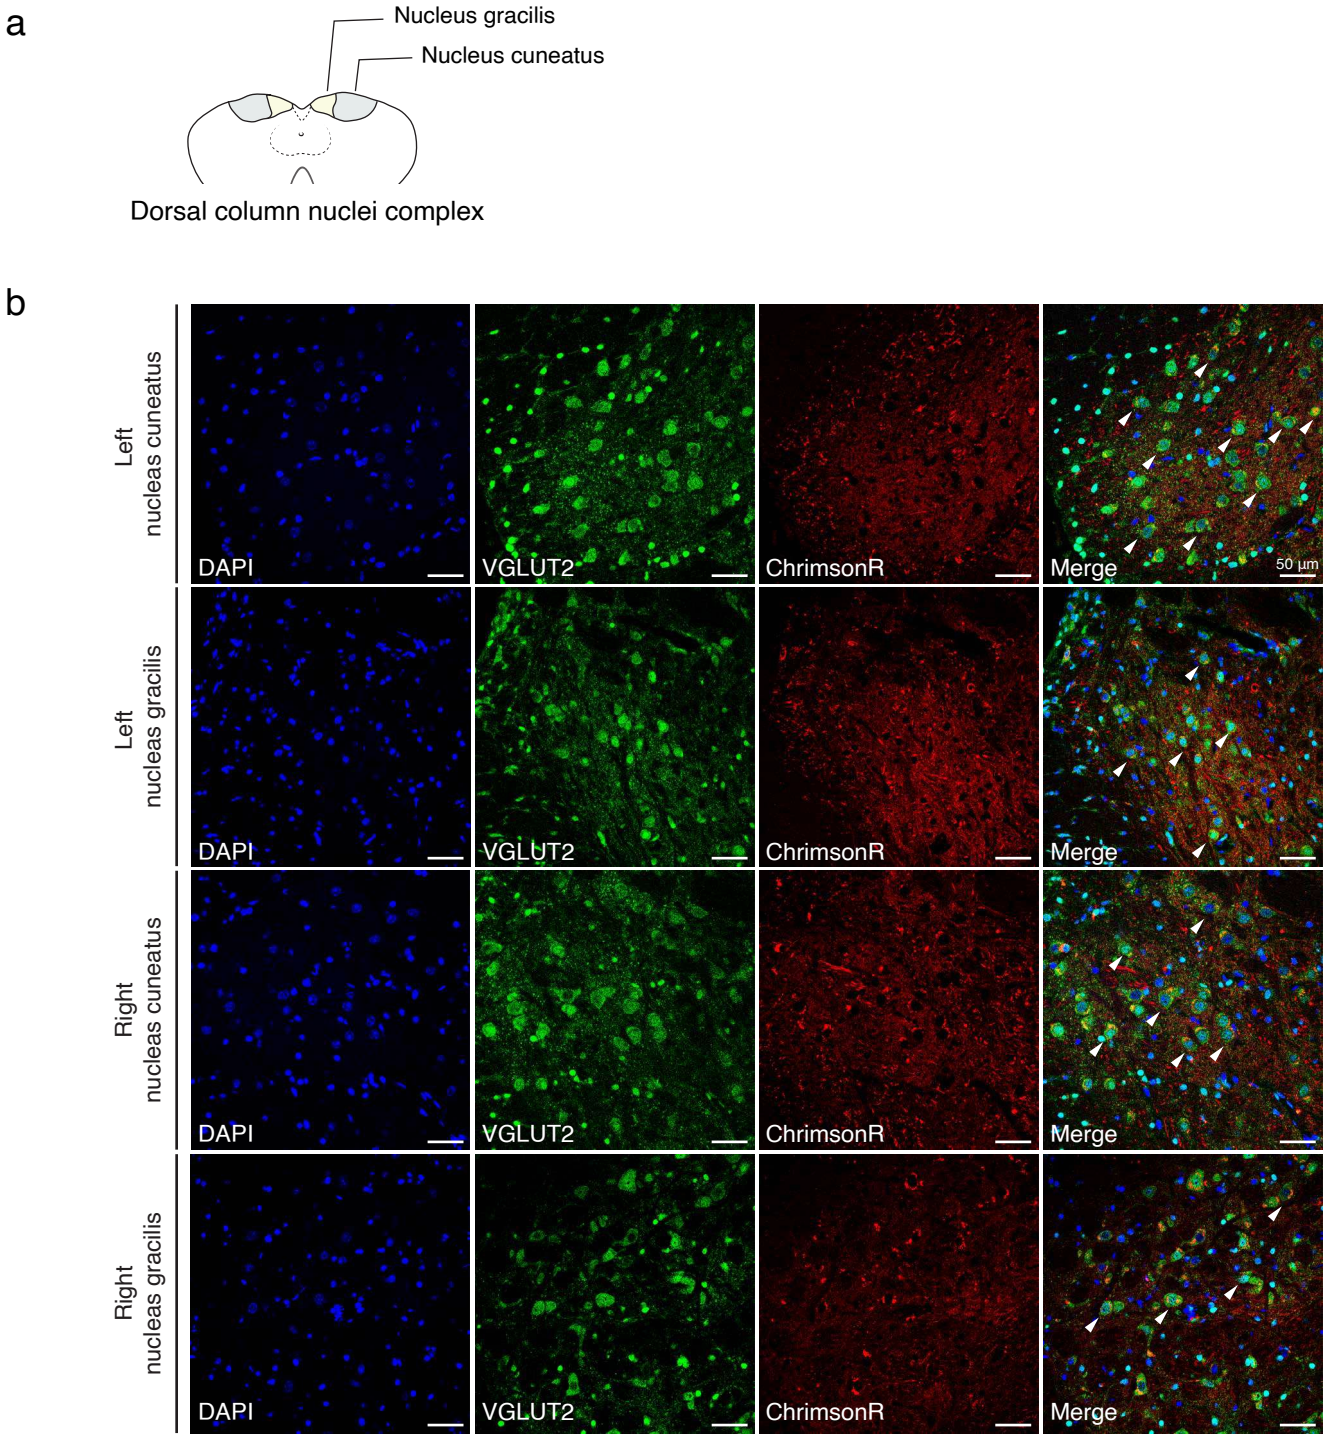
Supplementary Fig. 4 | MLI activity and its correlation with optic flow versus tail movement. a, Example trial showing tail-movement amplitude (top, cyan) and MLI activity (bottom). **b,** Pearson's correlation coefficients between tail movement and MLI activity for the ROIs shown in (**a**). Correlations are shown for individual MLI nodes (iMLIs) and aggregated population dynamics (mMLI). **c,** Heatmap of MLI population activity aligned to tail movement during the example trial. **d,** An example trace of aggregated MLI dynamics (top), tail-movement amplitudes (middle), and their cross-correlation function (bottom). **e,** Probability density distributions of Pearson's correlation coefficients between MLI activity and either tail-movement amplitude or VR optic flow speed, shown at the level of individual MLI calcium traces (n = 29,142 traces). **f,** Scatter plots comparing correlation coefficients for individual MLI traces (iMLIs; n = 29,142 traces from 6 mice) and aggregated population dynamics (mMLIs; n = 244 population means from 244 trials in 6 mice) with tail movement and VR optic flow. **g,** temporal processing lags derived from cross-correlation analyses between MLI activity and either VR optic flow or tail movement. Data are shown for individual MLI traces (iMLIs; n = 29,142 traces, 6 mice) and mMLIs (mMLIs; n = 244 population means from 244 trials in 6 mice). Insets show magnified distributions of delay values. Data are means ± SD; ****$p < 0.0001$; Welch's t-test.

Supplementary Fig. 5



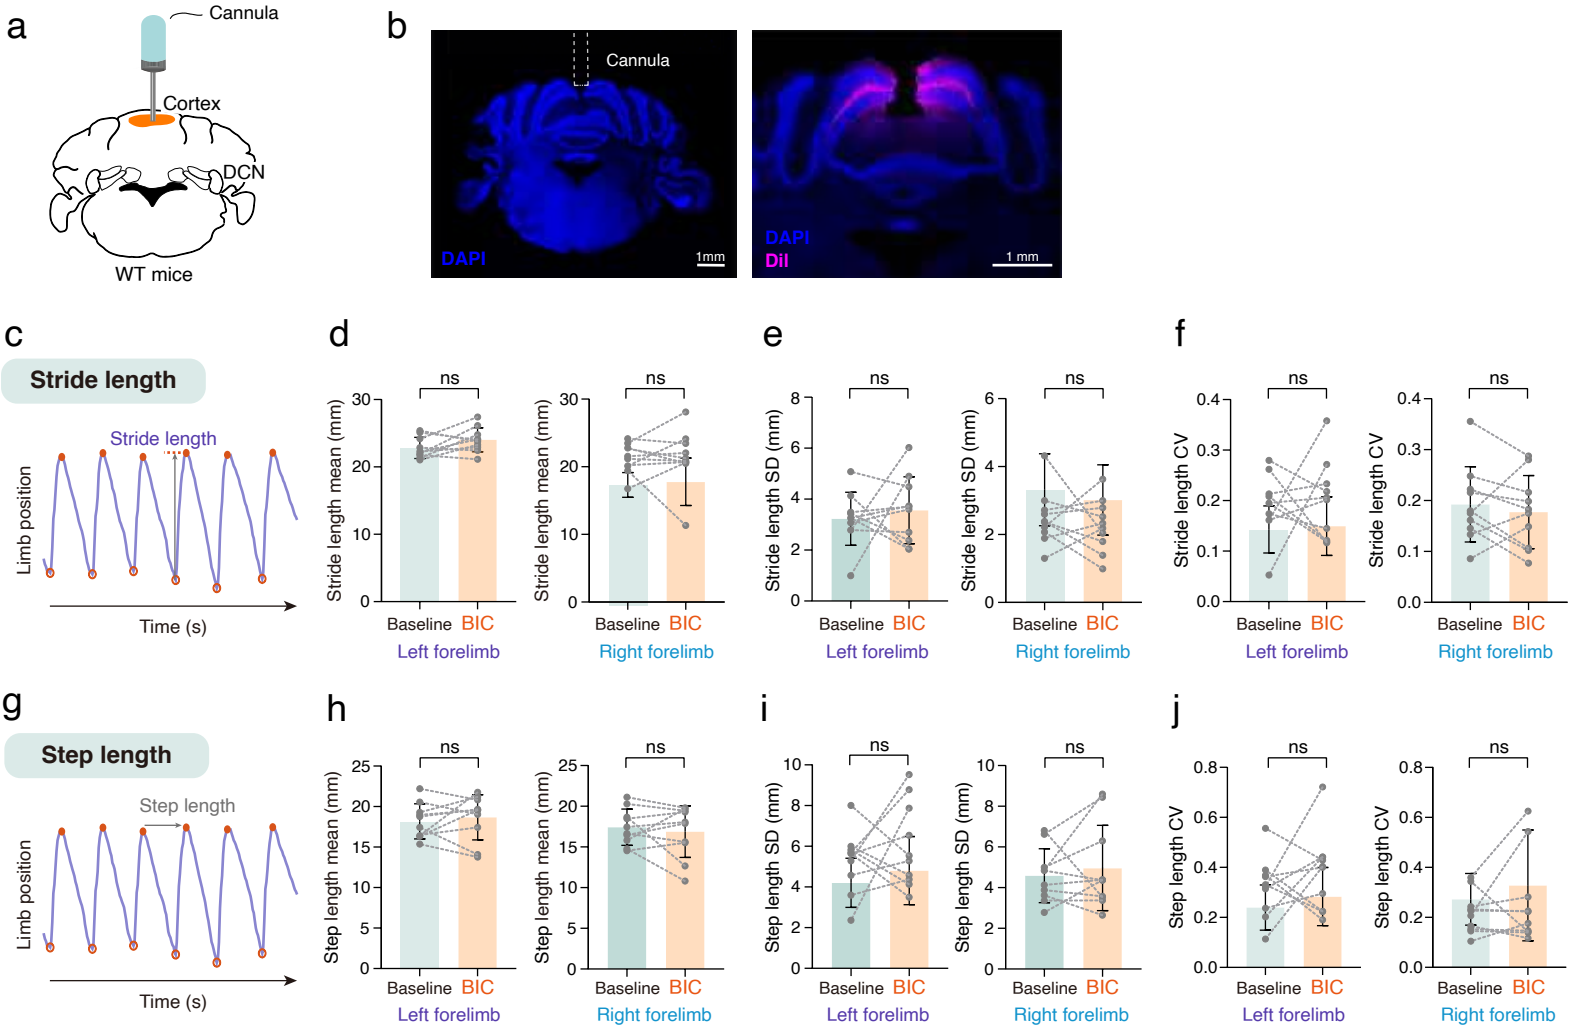
Supplementary Fig. 5 | Phase relationship analysis confirms the MLI latent space extracts a continuous scalar speed representation rather than cyclic gait variables. Analyses were performed on visual-deprivation trials to isolate somatosensory-driven locomotion and eliminate optic-flow interference. **a,** Example segment of high-dimensional gait traces alongside aggregated MLI dynamics for phase analysis. Front-left (FL) and front-right (FR) limb trajectories were extracted from video-based tracking. A difference component (Diff = (FR − FL)/2) and a mean component (Mean = (FR + FL)/2) were computed to represent alternating and shared components of limb kinematics, respectively. **b,** Spectral analysis of gait kinematics and MLI signals. FL and FR traces exhibited clear cyclic peaks near the step frequency (~3 Hz), while the MLI latent space displayed a broadband, non-periodic spectrum consistent with continuous scalar speed encoding rather than cyclic phase modulation. Shaded areas denote ± SD across trials. **c,** Schematic of phase extraction and transformation. Each trace was converted to its analytic signal using the Hilbert transform to obtain the instantaneous phase (–π to π). The cosine of this phase was computed to linearize polarity and remove discontinuities from phase wrapping, providing a continuous measure of oscillatory state for correlation analysis. **d,** Phase correlations between gait components and MLI activity under original, jittered, and shuffled conditions. To test whether MLI activity shared rhythmic phase structure with gait, both signals were Hilbert-transformed, cosine-projected, and correlated in phase space. Gait-MLI correlations remained exceptionally low across all conditions, indicating the MLI network does not share the step-locked phase structure of the multidimensional motor output. **e,** Example self peak-phase polarity index (PI) of gait and MLI signals from a single trial. Rhythmic gait components exhibited tightly clustered peak phases (PI ≈ 0.9), whereas the MLI latent space showed dispersed phase distributions with low PI (≈ 0.36), reflecting an absence of rhythmic phase locking. **f,** Example gait-to-MLI phase PI distributions from the same trial, confirming the network computes a continuous scalar magnitude rather than a discrete step-related cyclic variable. **g,** Population summary of PI across trials (n = 122 trials, from 6 mice). **h,** Gait signals displayed low SDs, indicating stable step-locked timing, whereas MLI self-phase SDs were large, reflecting continuous, non-periodic feedback extraction. Scatter plots represent individual trials (n = 122). Data are means ± SD; ****$p < 0.0001$; Welch's t-test.

Supplementary Fig. 6



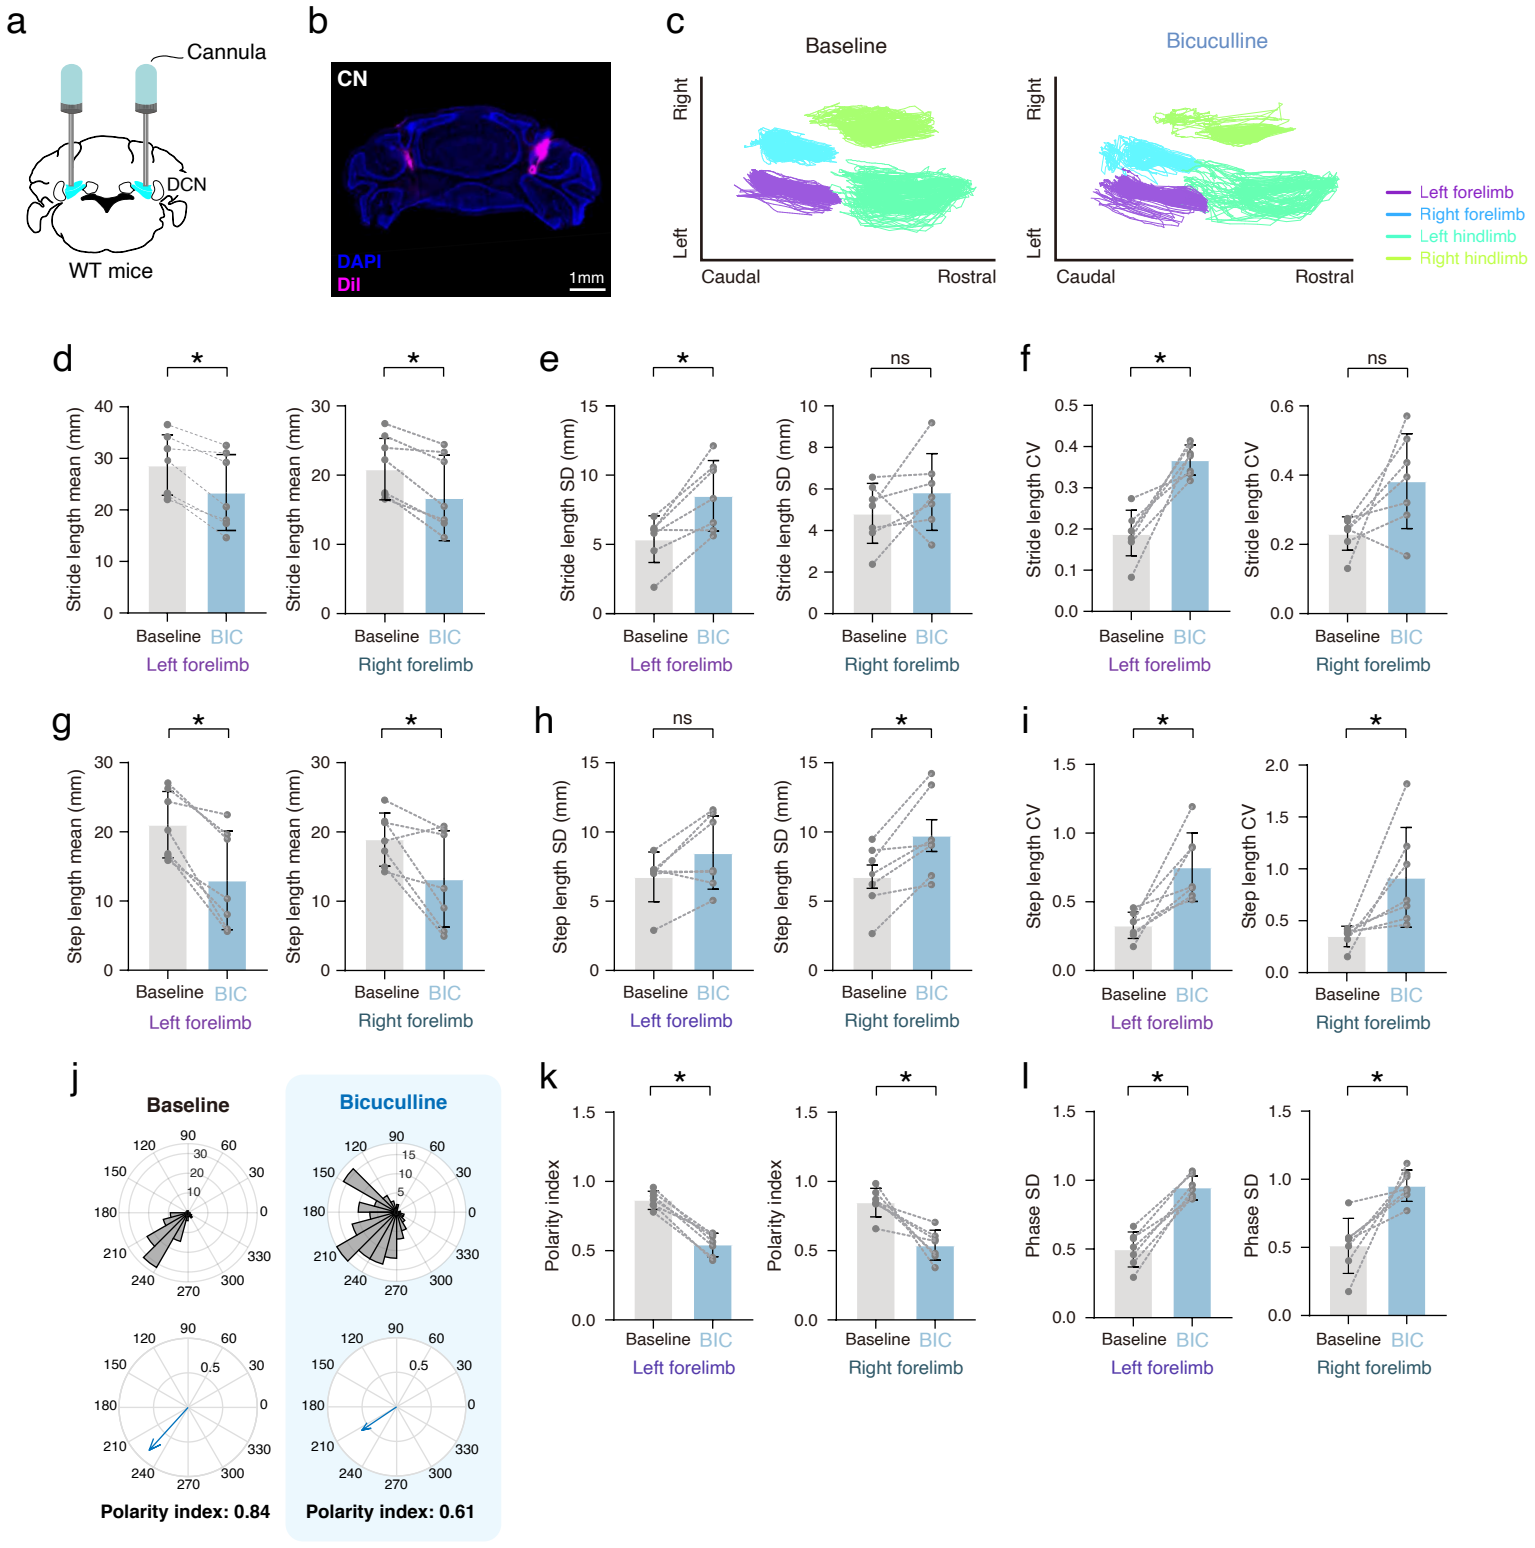
Supplementary Fig. 6 | Anatomical validation of targeted actuation in dorsal column nuclei, the ascending somatosensory-relay hardware. a, Schematic of the dorsal column nuclei complex showing the positions of the nucleus gracilis (NG) and nucleus cuneatus (CU) within the dorsal medulla. **b,** High-magnification (40x) images of the left and right nuclei demonstrating colocalization of tdTomato and VGLUT2 (white arrows) in large-soma neurons. These larger excitatory neurons correspond to the principal relay cells of the dorsal column nuclei that transmit low-threshold mechanosensory input to higher centers. The observed ChrimsonR-tdTomato/VGLUT2 colocalization confirms viral targeting of glutamatergic sensory-relay neurons in the nucleus gracilis and cuneatus. Scale bars, 50 µm.

Supplementary Fig. 7



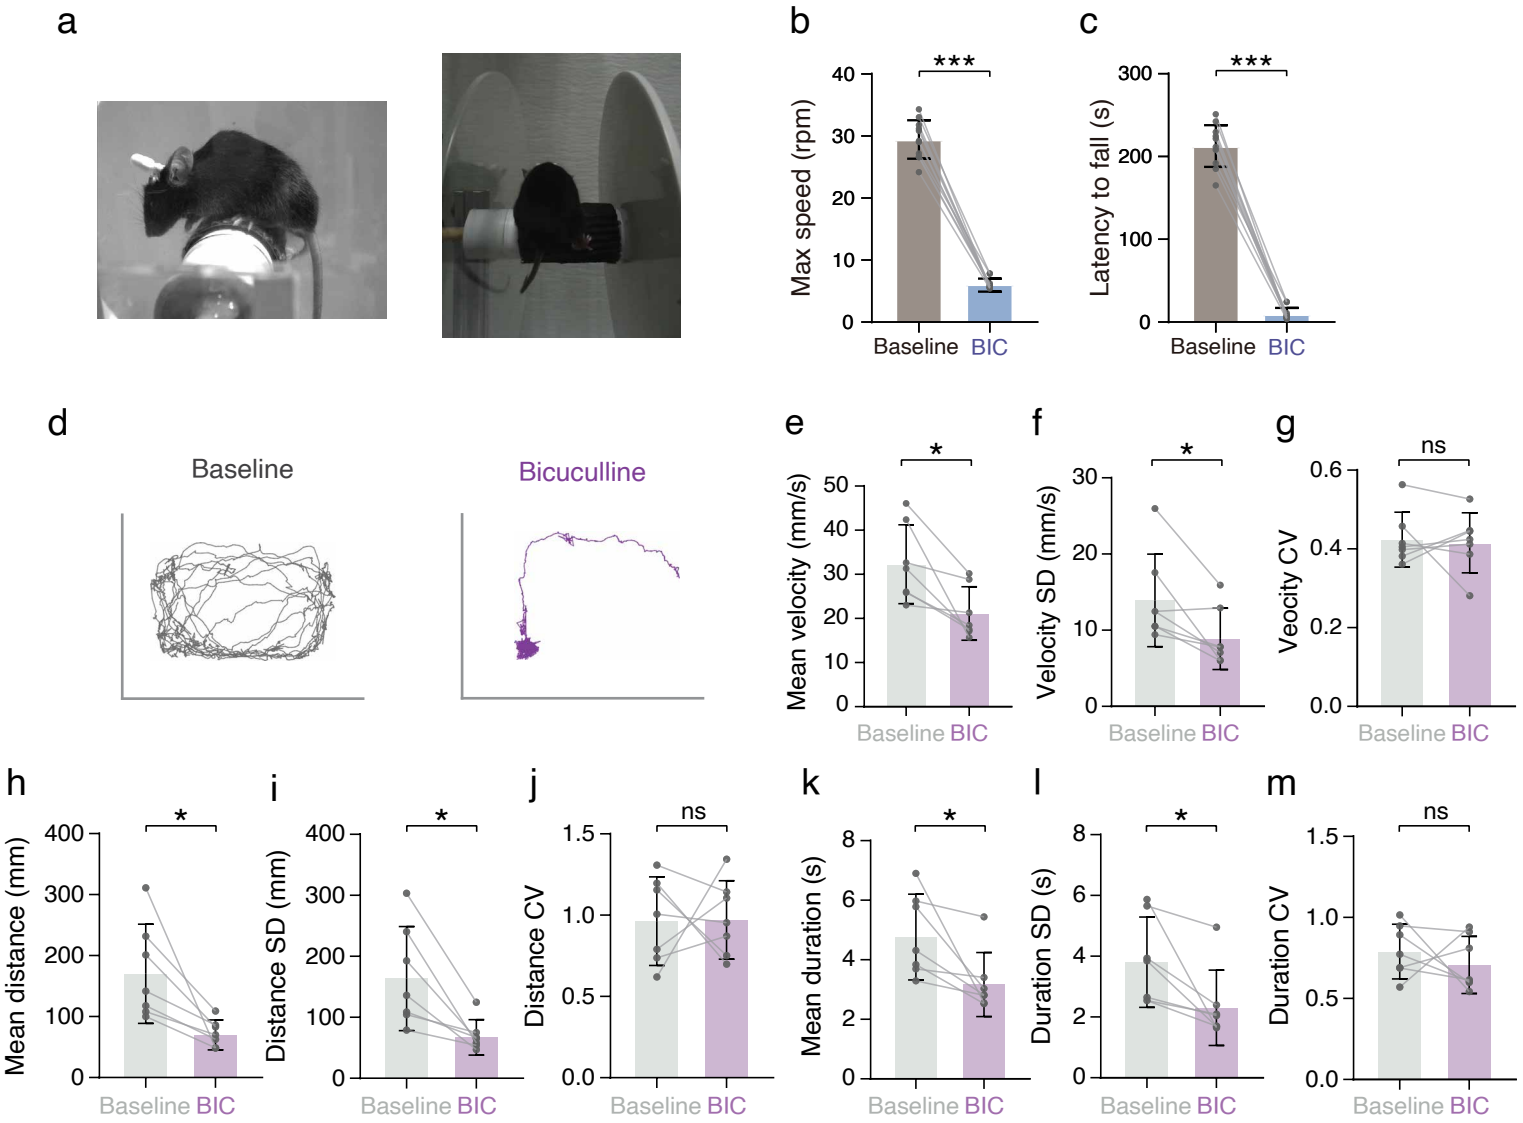
Supplementary Fig. 7 | Targeted suppression of the MLI latent space strictly preserves high-dimensional limb kinematics. **a,** Schematic of surgical targeting for bicuculline infusion into the cerebellar cortex via a centrally positioned cannula. **b,** Histological verification of cannula placement within the cerebellar cortex. Coronal section showing a central cannula track targeting lobule V, visualized by Dil tissue labeling (magenta) and DAPI nuclear staining (blue). Scale bars, 1mm. **c,** Schematic illustrating the definition of stride length, a fundamental component of the multidimensional motor output, derived from paw trajectories captured by high-speed videography during locomotion. **d-f,** Stride length metrics for the left and right forelimbs: mean stride length (**d**), stride length variability (SD; **e**), and coefficient of variation (CV; **f**) across baseline and bicuculline conditions. **g,** Schematic illustrating the definition of step length measured from inter-paw distances across consecutive gait cycles. **h-j,** Step length metrics for the left and right forelimbs: mean step length (**h**), step length variability (SD; **i**), and step length CV (**j**) across baseline and bicuculline conditions. Data are from 10 mice and shown as means ± SD; Wilcoxon signed-rank test.

Supplementary Fig. 8



Supplementary Fig. 8 | Disruption of the primary output node of DCN compromises global gait coordination. a, Schematic of bilateral cannula placement targeting bilateral DCN for primary output disruption by bicuculline. **b,** Histological verification of bilateral cannula placement within the DCN. Coronal cerebellar section showing bilateral cannula tracks in the DCN. Cannula locations were visualized by DiI tissue labeling (magenta) and counterstained with DAPI (blue). Scale bar, 1mm. **c,** Example multidimensional paw tracking during locomotion showing forelimb and hindlimb trajectories under baseline and bicuculline conditions. **d-f,** Paired stride length metrics for the left and right forelimbs: mean stride length (**d**), stride length variability (SD; **e**), and stride length CV (**f**) across baseline and bicuculline conditions ($n = 7$ mice). **g-i,** Paired step length metrics for the left and right forelimbs: mean step length (**g**), step length variability (SD; **h**), and step length CV (**i**) across baseline and bicuculline conditions. **j,** PI quantifying inter-forelimb coordination before and after bicuculline infusion. **k,** Quantification of the PI for left and right forelimbs. **l,** Step-phase variability (SD) for left and right forelimbs in baseline and under bicuculline. Data are from 7 mice and shown as means ± SD; *$p < 0.05$; Wilcoxon signed-rank test.

Supplementary Fig. 9



Supplementary Fig. 9 | Disruption of the primary output node of DCN compromises forced and spontaneous locomotor activity. a, Photographs of the rotarod apparatus used for motor endurance testing. **b-c,** Rotarod performance before and after bicuculline infusion: maximum rotation speed (**b**) and latency to fall (**c**) (n = 11 mice). **d,** Representative open-field trajectories from a single mouse under baseline and bicuculline conditions. **e-g,** Quantification of open-field velocity: mean velocity (**e**), variability (SD; **f**), and coefficient of variation (CV; **g**) (n = 7 mice). **h-j,** Quantification of open-field distance traveled: mean distance (**h**), variability (SD; **i**), and CV (**j**). **k-m,** Movement bout duration metrics: mean duration (**k**), variability (SD; **l**), and CV (**m**) (n = 7 mice). Data are means ± SD; *$p < 0.05$, ***$p < 0.001$; Wilcoxon signed-rank test.

Video S1. Closed-loop VR locomotion with two-photon imaging. A head-fixed mouse ran on a rotary disc in a closed-loop virtual reality environment, where visual flow was coupled to locomotor speed. Simultaneous two-photon calcium imaging in the cerebellar cortex monitored calcium activity of MLIs and PC dendrites.

Video S2. Bicuculline infusion into the cerebellar cortex. Head-fixed mice performed self-paced locomotion on a rotary disc during local infusion of the $GABA_A$ receptor antagonist bicuculline. Infusion into the cerebellar cortex (lobule V) produced a marked acceleration in running speed while preserving coordinated stepping.

Video S3. Rotarod and open-field behavior after bicuculline infusion into the DCN. Freely moving mice performed on an accelerating rotarod and explored an open-field arena before and after bilateral bicuculline infusion via DCN cannulas. Infusion impaired rotarod performance and altered locomotor behavior in the open field.

Video S4. Optogenetic activation of MLIs suppresses locomotion. A *PV-Cre* × Ai203 mouse ran on a rotary-encoder treadmill while MLIs in lobules V/VI were stimulated with 473-nm light (1 s ON / 3 s OFF). Each light pulse produced a rapid and reversible decrease in running speed.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- DatafileS1.mat
- VideoS2.mp4
- DatafileS1.xlsx
- VideoS3.mp4
- VideoS1.mp4
- VideoS4.mp4