DISTILLING NEURAL POPULATION GEOMETRY INTO TRANSFORMERS AND GRUS FOR MOTOR CONTROL

Anonymous authors
Paper under double-blind review

ABSTRACT

We introduce neuro-distillation (Neuro-Dist), a technique designed to directly transfer computationally useful neural population structure from biological brains into artificial controllers. Using three macaque motor-cortical datasets as neuro-distillation teachers, we align controller internal states with cortical neural population geometry during training; controller deployment requires no neural data. We evaluate gated recurrent unit (GRU) networks and Transformers controlling a simulated arm. Neuro-Dist achieves the lowest median error among five methods, including distillation from movement kinematics, in 10 of 18 comparisons across trained, student-held-out, and new target directions. Seed-split aggregate statistical comparisons significantly favor Neuro-Dist over all six controls on trained and new targets, and five on student-held-out targets (Holm-corrected familywise $p < 0.05$). We furthermore find that Neuro-Dist can substantially improve controller robustness to sensory feedback noise. Mechanistically, Neuro-Dist induced low tangling of trajectories in internal controller state space—a functionally important characteristic of motor cortex, demonstrating successful transfer of a known motor-cortical population geometry feature. These findings establish neuro-distillation as a complement to behavioral distillation and point to a novel class of neuro-interfaces that tap biological neural population structure as a computational resource for improving the control of goal-directed behavior in robotics and embodied AI.

1 INTRODUCTION

Behavioral control in embodied AI agents and robotic systems is typically learned from task objectives, demonstrations, or simulated experience. Examples of large-scale efforts include collecting demonstrations with 13 physical robots over 17 months (Brohan et al., 2023) and training a manipulation policy with 131,072 parallel simulated robot instances over approximately 32 hours (Handa et al., 2024). In contrast, mammalian behavioral control builds on neural circuitry shaped by hundreds of millions of years of vertebrate evolution across successive generations of global populations (Zador, 2019). As a reference point for mammalian population scale, for example the contemporary abundance of the straw-coloured fruit bat alone—a single species with dexterous flight control—is estimated at roughly one billion individuals (Greenspoon et al., 2023; O'Mara et al., 2019). Evolution shapes inherited circuit organization, while development and lifetime learning adapt these circuits to individuals and their environments (Zador, 2019). Reproducing the breadth and adaptability of biological behavior in everyday physical environments remains an open challenge for artificial systems (Zador, 2019; Tao et al., 2025). This motivates investigating biological neural systems as a source of reusable computational structure that may be difficult to engineer *a priori* and costly to empirically rediscover through training of artificial systems alone. As an example, neural population dynamics could supply training priors that transfer useful structure to artificial agents.

We investigate this possibility in motor control, where the mammalian central nervous system combines ancient spinal circuitry with more recently evolved cortical pathways (Jung et al., 2018; Rathelot & Strick, 2009). Two converging lines of research suggest that transferring biological representations to artificial motor controllers may be feasible. First, in the visual system, supervision derived from primate cortical recordings has improved image recognition (Federer et al., 2020; Safarani et al., 2021), with representational alignment also improving visual-model robustness (Li et al., 2019; Dapello et al., 2023). Second, studies of motor-cortical population geometry describe movement preparation and execution as low-dimensional manifolds (Churchland et al.,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

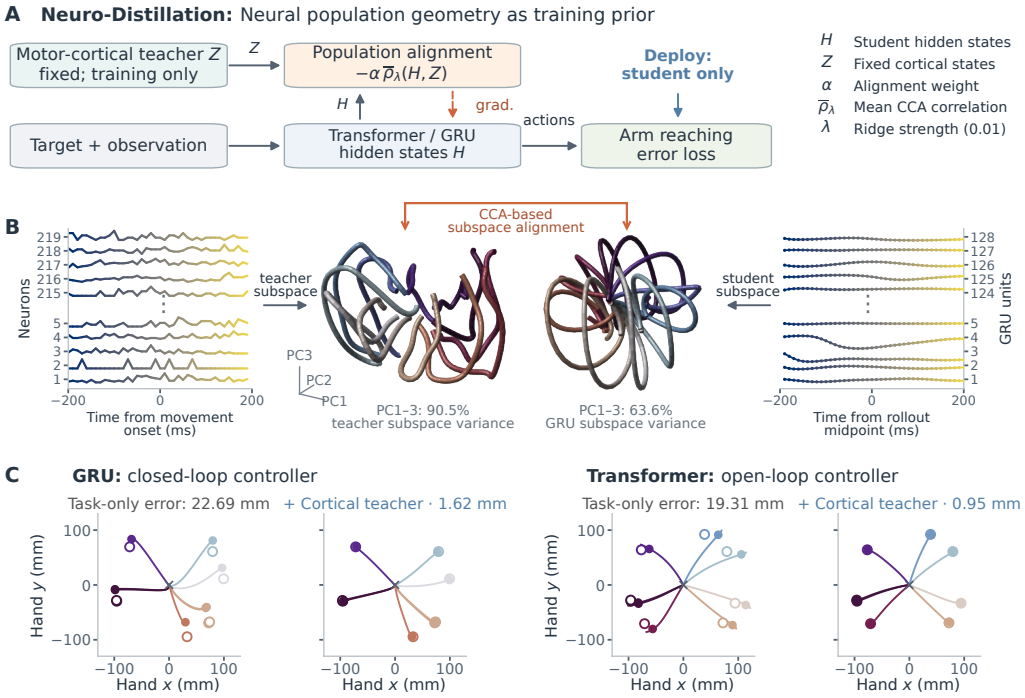

Figure 1: **Neuro-distillation of neural population geometry into artificial controllers.** (A, B) A low-dimensional subspace is constructed from the neural teacher data recorded during reaching movements. PC1–PC3 denote the first three principal components (PCs). Two controllers (GRU, Transformer) are trained in a reaching task with a simulated arm. During neuro-distillation-guided training, a subspace of the student's hidden internal state is aligned to the teacher subspace via canonical correlation analysis (CCA) metrics. The students are deployed without further supervision and without requiring neural data at this stage. (C) Neuro-distilled controllers show more accurate target reaches; they also improved robustness against sensory feedback noise (see Figure 3).

2012; Russo et al., 2018). Such neural population geometry provides a conceptual bridge between biological and artificial systems (Chung & Abbott, 2021); task-trained artificial recurrent networks can develop cortex-like dynamics (Sussillo et al., 2015). During reaching movements, a characteristic of population dynamics is *low trajectory tangling*: nearby population states tend to have similar velocity vectors. This organization distinguishes motor-cortical activity from the muscle signals it helps generate and has been linked to noise robustness: Models trained to generate the same output using less-tangled internal trajectories tolerated more state noise (Russo et al., 2018). Primate motor cortex therefore offers a biologically motivated source of neural population structure for artificial controllers.

Here, we bring these lines together and ask: *Can recorded motor-cortical population geometry serve as a training prior that improves artificial movement control?* We examine four subquestions:

(i) Does Neuro-Dist improve accuracy on trained targets relative to task-only training and alternative teachers, including kinematics-based, engineered, stochastic, and shuffled controls?

(ii) Does Neuro-Dist also improve accuracy on targets not used during controller training?

(iii) Can Neuro-Dist increase controller robustness to sensory feedback noise?

(iv) Which characteristic of neural population geometry does Neuro-Dist transfer to students?

To address these questions, we construct cortical teachers from primary motor-cortical spiking activity recorded in three macaques during reaching (Figure 1). We train GRU and Transformer controllers and compare accuracy and generalization against task-only training and alternative teachers. Our results establish a proof of concept for neuro-distillation improving artificial motor control.

2 Related work

Neural data as a teaching signal. Recorded activity may guide learning through example weighting, response prediction, representation matching, or evaluative feedback. Fong et al. (2018) weighted image-classification training errors with human functional magnetic resonance imaging (fMRI), needing no neural measurements at inference. Brain-mediated transfer improves downstream label estimation by mapping audiovisual features through fMRI-learned representations (Nishida et al., 2020). McClure & Kriegeskorte (2016) transferred representational distances between artificial networks and proposed extensions to measured brain representations. Early electroencephalography (EEG)-guided image embeddings and joint image–EEG learning (Spampinato et al., 2017; Palazzo et al., 2021) have more disputed interpretations because of temporal correlations in block-designed data (Li et al., 2020; Palazzo et al., 2020). These approaches target perceptual representations and predictions; we supervise motor controllers based on motor-cortical single-unit data.

Visual representation learning. Neural regularization may transfer stimulus relationships and individual response targets. Li et al. (2019) improved image-model robustness using similarities from a mouse-visual-cortex predictor; Federer et al. (2020) improved recognition and corrupted-label tolerance using macaque response statistics. Joint classification and macaque primary visual cortex response prediction improves image-distortion robustness (Safarani et al., 2021); inferior-temporal-cortex alignment improves human behavioral agreement and adversarial robustness (Dapello et al., 2023). A direct precursor, Pirlot et al. (2022), uses deep-CCA regularization to align convolutional neural network representations with monkey visual responses, improving classification and robustness to tested adversarial perturbations. Human visual teachers include ReAlnet's EEG-response prediction, improving neural and behavioral alignment (Lu et al., 2026); fMRI-derived predictors and category-manifold guidance, improving robustness with area-dependent benefits (Shao et al., 2025); and fMRI–image structural alignment, improving visual-category generalization (Chen et al., 2026). In our study, we extend these concepts from the visual domain to upper-extremity motor control.

Speech, language, and reasoning. For sequential inputs, Schwartz et al. (2019) fine-tuned Bidirectional Encoder Representations from Transformers (BERT) to predict fMRI and magnetoencephalography responses, preserving downstream natural language processing (NLP) performance. Brain-response supervision improves speech recognition (Li et al., 2024) and, through fMRI brain-tuning, semantic speech tasks (Moussa et al., 2025). Joint cross-participant prediction extends brain-tuning benefits to neural prediction and semantic tasks (Moussa & Toneva, 2025); fMRI-tuned auditory networks improve several downstream auditory tasks (Freteault et al., 2025); bilingual fMRI supervision improves multilingual NLP (Negi et al., 2025). Time-resolved electrocorticography can supervise speech encoders (Zhang et al., 2026); fMRI-guided parameter fine-tuning improves evaluated deductive reasoning (Xiao et al., 2026). Aquino et al. (2025) have reported improved memory-based decision learning by contrastively aligning internal network representations with human single-neuron embeddings.

Neural feedback and embodied motor learning. EEG error-related signals can teach robot control policies (Iturrate et al., 2015; Kim et al., 2017); decoded evaluative feedback accelerates robot navigation learning (Akinola et al., 2020). Here brain signals evaluate actions during deployment and provide a reward signal, e.g., for reinforcement-learning-based optimization; we distill internal controller states directly. Behaviorally trained virtual animals assess post-training neural correspondence (Aldarondo et al., 2024); neurocybernetic modeling proposes distillation from neural-data-trained models, then embodied refinement (Park et al., 2026). The closest embodied comparison, the neurally constrained nuSim variant of μSim (Almani et al., 2025), fits selected recurrent neural network (RNN) units to recorded single-unit firing rates, improving neural reconstruction at unseen cycling speeds. Its aim is to model motor cortex faithfully; controller performance or robustness benefits were not established. Our CCA-based population supervision avoids neuron-to-unit assignments and targets improved artificial control.

3 Methods

Population geometry as a training prior. Neuro-distillation combines conventional reaching supervision with an *a priori* cortical-population alignment objective. Neural teachers use publicly

Table 1: **Reaching accuracy across target sets.** Median ± bootstrap standard error (SE) of final-window coordinate mean squared error (MSE; mm^2; 100 seeds), at training-selected strengths fixed across target sets. Gaussian and Poisson columns continue in Supplementary Table 2. Bold values mark row minima. M1–M3: Monkeys 1–3; Transf.: Transformer.

Student	Monkey	Task-only	Ring	Neuro-Dist	Behavior	Shuffled
A. Trained targets (28)						
GRU	M1	58.42 ± 3.48	54.95 ± 4.54	$\mathbf{33.66} \pm 3.24$	38.57 ± 3.77	50.79 ± 2.93
GRU	M2	60.72 ± 4.79	53.94 ± 3.33	$\mathbf{26.83} \pm 3.62$	38.60 ± 2.69	50.13 ± 3.57
GRU	M3	55.59 ± 3.71	52.00 ± 2.59	48.41 ± 3.19	$\mathbf{46.56} \pm 2.55$	59.10 ± 4.08
Transf.	M1	24.29 ± 2.34	25.69 ± 1.50	12.24 ± 3.38	$\mathbf{9.77} \pm 1.37$	39.40 ± 5.08
Transf.	M2	31.40 ± 6.18	12.53 ± 2.40	$\mathbf{8.63} \pm 1.86$	15.79 ± 3.04	28.48 ± 3.57
Transf.	M3	28.16 ± 6.20	20.58 ± 2.55	$\mathbf{15.90} \pm 2.98$	16.75 ± 1.46	30.08 ± 2.19
B. Generalization to student-held-out targets (8)						
GRU	M1	62.05 ± 2.72	59.05 ± 5.90	42.98 ± 3.96	$\mathbf{42.20} \pm 5.33$	56.55 ± 5.13
GRU	M2	64.80 ± 5.31	60.09 ± 5.02	$\mathbf{31.97} \pm 3.88$	41.36 ± 2.22	55.54 ± 3.79
GRU	M3	59.62 ± 3.67	53.91 ± 3.45	$\mathbf{53.05} \pm 2.97$	53.39 ± 2.31	61.77 ± 3.73
Transf.	M1	27.66 ± 2.36	31.49 ± 2.96	18.32 ± 3.75	$\mathbf{12.60} \pm 1.78$	48.77 ± 4.34
Transf.	M2	34.13 ± 8.45	13.89 ± 2.27	$\mathbf{9.49} \pm 1.67$	16.61 ± 2.78	28.28 ± 3.65
Transf.	M3	30.38 ± 9.74	20.42 ± 2.57	17.99 ± 2.96	$\mathbf{16.60} \pm 1.67$	35.93 ± 2.57
C. Generalization to new target directions (8)						
GRU	M1	79.82 ± 3.09	76.94 ± 3.64	62.44 ± 4.84	$\mathbf{56.81} \pm 6.10$	95.63 ± 8.04
GRU	M2	58.80 ± 3.51	53.69 ± 4.20	$\mathbf{29.57} \pm 3.56$	38.90 ± 3.16	50.68 ± 3.34
GRU	M3	55.24 ± 3.03	50.61 ± 3.55	47.82 ± 3.25	$\mathbf{47.18} \pm 3.95$	53.63 ± 3.63
Transf.	M1	91.23 ± 17.11	106.44 ± 11.76	66.61 ± 15.42	$\mathbf{65.72} \pm 7.55$	125.30 ± 11.68
Transf.	M2	30.23 ± 5.54	14.86 ± 2.53	$\mathbf{10.28} \pm 1.54$	18.29 ± 2.79	31.19 ± 3.74
Transf.	M3	27.23 ± 5.38	22.44 ± 3.19	$\mathbf{16.26} \pm 2.12$	18.07 ± 1.72	35.22 ± 2.79
Sum of wins		0	0	**10**	**8**	0

available, condition-averaged motor-cortical data from three macaques performing reaching tasks: Monkey 1 from Pei et al. (2021) and Monkeys 2–3 from O'Doherty et al. (2017). We project 12 cortical PCs onto a rank-four subspace predictive of recorded hand position and velocity, retaining 12 coordinates. All 36 reaching conditions enter teacher construction, including those withheld from student training. Regularized CCA rewards correlated linear projections of student and fixed teacher activity across matched conditions and time points (Hotelling, 1936; Andrew et al., 2013). This encourages shared linear population geometry without neuron-to-unit assignments or requiring identical full-state distances or dynamics. Recordings are needed only during training; Appendix A gives implementation details and Appendix B the protocols.

Task, students, and optimization. A closed-loop GRU with updated sensory feedback and an open-loop Transformer conditioned on initial observations and target information generate commands for MotorNet's differentiable two-joint, six-muscle arm (Codol et al., 2024). Students minimize

$$\mathcal{L}_i = \mathcal{L}_{\text{reach}} - \alpha \min(1, i/60)\overline{\rho}_\lambda(H, Z), \tag{1}$$

where $\mathcal{L}_{\text{reach}}$ averages squared fingertip-to-target error across coordinates, the final 15 frames, and training conditions, normalized by squared target radius. The CCA score $\overline{\rho}_\lambda$ averages 12 regularized canonical correlations between student states H and teacher states Z (Suppl. equation 2); α is the CCA alignment strength. $\alpha = 0$ gives task-only training.

Evaluation. For each animal, architecture, and aligned teacher in Table 1, we test 40 strengths ($\alpha = 0.1$–4.0), each with 100 seeds. Accuracy is the median final-15-frame coordinate MSE (mm^2) across all seeds, including high-error fits. These 100 seeds enabled robust median-based effect estimates, paired bootstrap confidence intervals, and Holm correction for multiple comparisons, with separate 50-seed groups for strength selection and evaluation (Appendix F).

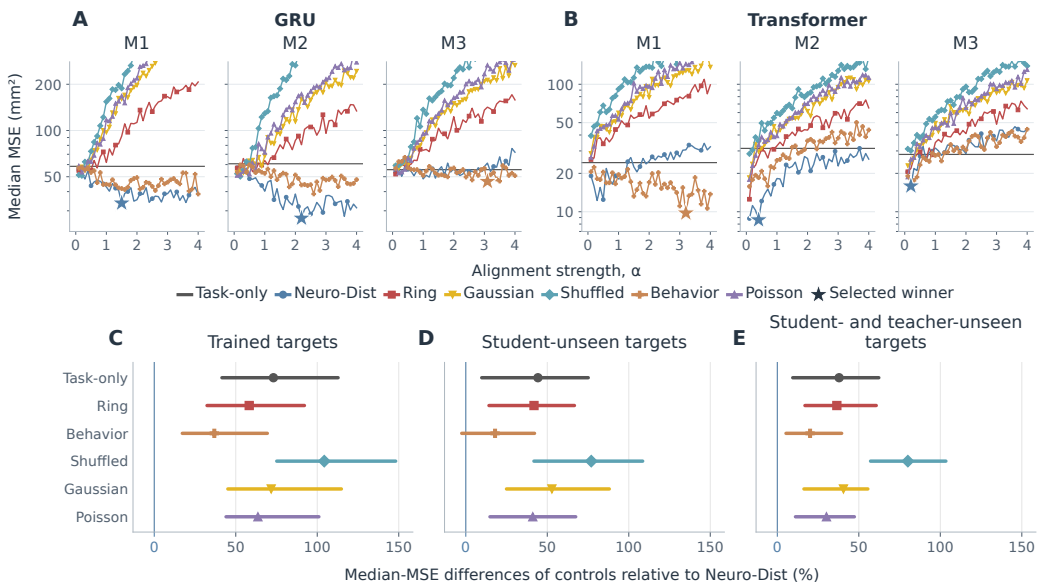

Figure 2: **Alignment response and statistical evaluation. A,B:** GRU and Transformer trained-target median MSE (100 seeds); stars mark winners across all seven methods at training-selected strengths. **C–E:** Control median-MSE differences relative to Neuro-Dist, from geometric-average ratios across both architectures and three animals (50 evaluation seeds). Positive values indicate increased control errors; 0% matches Neuro-Dist performance. Bars show pointwise 95% confidence intervals from bootstrap resampling. Statistical details: Appendix F. In aggregate comparisons, Neuro-Dist significantly outperforms all six controls on trained and new targets and five on student-held-out targets, with no control significantly outperforming Neuro-Dist (Holm-corrected $p < 0.05$).

Controls. We compare Neuro-Dist to a comprehensive set of controls, each of them designed to probe a different competing distillation strategy or possible source of benefit from teacher supervision:

- **Task-only:** establishes performance without distillation: reaching loss alone ($\alpha = 0$).
- **Behavior:** tests supervision from recorded movement: four PCs of hand position and velocity at five temporal offsets, embedded in 12 coordinates.
- **Ring:** an engineered directional prior motivated by directional organization in primary motor cortex data (see Appendix D).
- **Gaussian:** controls for population variance structure: smooth random signals matching the normalized cortical covariance eigenvalue spectrum.
- **Poisson:** tests supervision from synthetic spiking: animal-specific fitted rate distributions with smooth modulation, trial averaging, and reduction to four PCs.
- **Shuffled:** tests the importance of correct condition correspondence: whole cortical trajectories are reassigned to different reaches, preserving time order.

Sensory feedback noise and trajectory geometry. Analyses of feedback noise and internal tangling are naturally suited to the closed-loop, recurrent GRUs; our open-loop Transformer neither receives sensory feedback nor evolves a recurrent hidden state. Independent Gaussian noise was used to perturb all feedback channels, including target information; errors use final-window coordinate root mean squared error (RMSE; Appendix C). Trajectory tangling quantifies velocity disagreement between nearby full-dimensional hidden states (Russo et al., 2018), within and between reaches (Appendix E). We define $Q_i = \max_j \|\dot{x}_i - \dot{x}_j\|_2^2 / (\|x_i - x_j\|_2^2 + 0.1 s_x^2)$, where x_i is a hidden state, $\dot{x}_i$ its temporal derivative, and s_x^2 the total state variance; j ranges over reference states. Reported Q averages Q_i over evaluated states; lower values indicate less tangling. Tangling is neither directly optimized during training nor used to select distillation strengths: it tests for a population-geometric feature beyond the alignment objective.

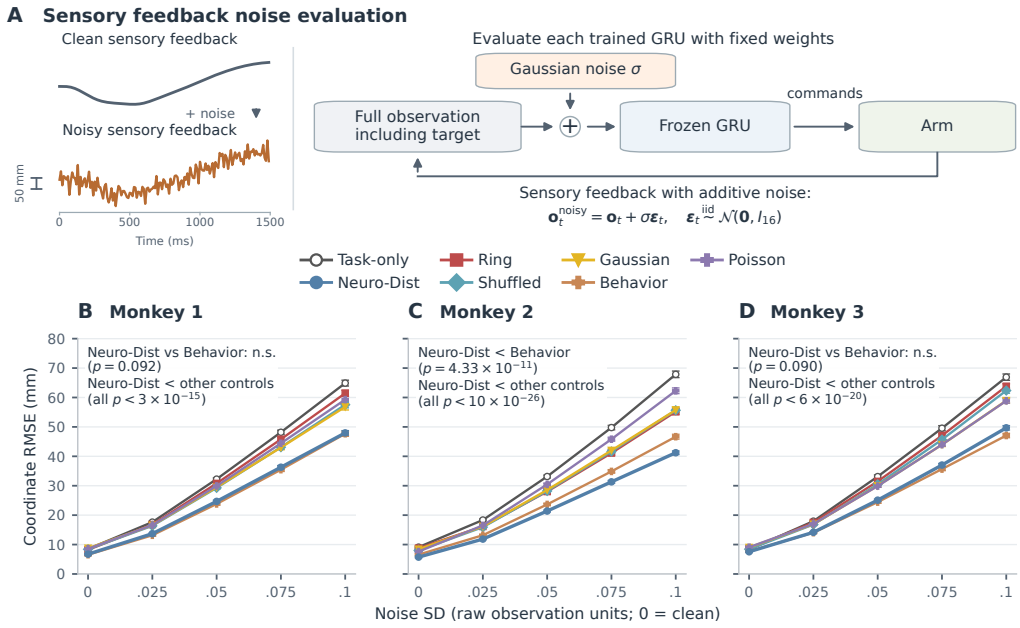

Figure 3: **Neuro-Dist improves GRU robustness to sensory feedback noise.** (A) We perturb all 16 observation channels of frozen GRUs trained without noise. Gaussian noise is independent across channels and 10-ms steps: $\mathbf{o}_t$ and $\mathbf{o}_t^{\text{noisy}}$ denote clean and noisy inputs, σ the noise standard deviation (SD) in channel units, and $\boldsymbol{\epsilon}_t$ a standard-normal vector. The separate 1.5-s illustration uses a fixed-dose controller ($\sigma = 25$ mm for positions). (B–D) Mean coordinate RMSE ± SEM (standard error of the mean) across 100 seeds, each averaging three realizations per positive amplitude. Strengths are selected on 28 clean training targets; 400-ms evaluations score eight student-held-out targets. Annotations: paired signed-rank tests of seedwise mean error ratios, Holm-corrected across 18 comparisons. "<": lower error; "n.s.": nonsignificant. Ratio definitions and confidence intervals: Supplementary Figure 8. Neuro-Dist and Behavior together have the lowest mean errors at every noise level. Their primary comparison favors Neuro-Dist in Monkey 2, with no significant difference in Monkeys 1 and 3; no control significantly outperforms Neuro-Dist in any primary comparison.

4 Results

Does Neuro-Dist improve accuracy on trained targets relative to task-only training and alternative teachers? We evaluated GRU and Transformer controllers with teachers from three monkeys and 100 seeds per setting. Table 1 reports median framewise coordinate MSE over the final 15 frames, retaining all seeds. Each aligned method uses a training-selected strength minimizing median training MSE over the same 40 positive doses; task-only by definition uses zero alignment. On the 28 trained targets, Neuro-Dist has lower median MSE than task-only training in all six animal–architecture comparisons, with reductions of 12.9–72.5% (Table 1A). Figure 2 shows the corresponding dose responses.

In a separate seed-split analysis, Neuro-Dist significantly outperforms all six controls on trained targets (tests based on paired seed-block bootstrap CIs; Holm-corrected familywise $p < 0.05$ across 18 aggregate comparisons; Appendix F). In Table 1, Neuro-Dist has the lowest median in four rows and Behavior in two; both have lower medians than Shuffled, and Neuro-Dist than Ring throughout. Gaussian and Poisson are additional controls with 40-dose coverage (Figure 2; Supplementary Table 2).

Does Neuro-Dist also improve accuracy on targets not used during controller training? We evaluated the same 3,000 selected controllers from the five methods in Table 1 on two additional target sets, without retuning alignment strengths or retraining students. *(1) Generalization to student-held-out targets.* Eight of the original 36 reaching conditions were withheld from student optimization

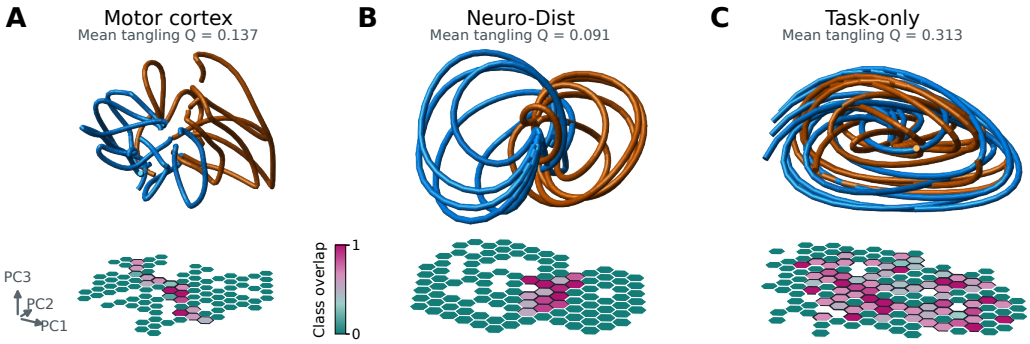

Figure 4: **State-space tangling of neural data and controllers. (A)** Example of subspace projection of Monkey 2 cortical activity from one contributing session and **(B)** of a Neuro-Dist GRU ($\alpha = 2.2$, seed 45); **(C)** its matched task-only GRU. Principal component analysis (PCA) projections show 12 representative reaching conditions; hexagons measure projected direction-class overlap. Tangling Q is low in both the motor cortical recordings and the neuro-distilled controller, but substantially higher in the task-only controller trained without any distillation.

but included in teacher construction. Neuro-Dist has lower median MSE than task-only training in all six comparisons, with reductions of 11.0–72.2% (Table 1B). Supplementary Figure 6 shows the dose responses. *(2) Generalization to new target directions.* To test generalization to conditions unseen by both teacher and student, we also evaluated eight directions per animal and seed outside the original 36-condition set used by teacher construction and student training. Neuro-Dist has lower median MSE than task-only training in all six comparisons, with reductions of 13.4–66.0% (Table 1C) (Appendix B.3).

Corrected seed-split aggregate comparisons favor Neuro-Dist over all six controls on new targets and five on student-held-out targets; no statistically significant difference is detected against Behavior on the latter. Descriptively, Neuro-Dist and Behavior each have the lowest median in three rows per regime (10 versus 8 overall; Table 1); both outperform Shuffled, and Neuro-Dist outperforms Ring throughout. No control significantly outperforms Neuro-Dist in any aggregate comparison.

Can Neuro-Dist increase controller robustness to sensory feedback noise? We tested 2,100 GRUs under white sensory feedback noise, selecting alignment strengths from clean training (Figure 3). On student-held-out targets, Neuro-Dist lowers mean error versus task-only at every amplitude in all three animals; corrected tests also favor it over Ring, Shuffled, Gaussian, and Poisson. Primary comparisons with Behavior favor Neuro-Dist in Monkey 2, with no statistically significant difference detected in Monkeys 1 and 3. Sensory feedback noise raises mean RMSE less for Neuro-Dist than task-only at every positive amplitude in all animals (Appendix C). Thus, cortical supervision can improve GRU robustness to sensory feedback noise without noise-augmented training.

Which feature of neural population geometry does Neuro-Dist transfer to the students? To address this question, we focused on low trajectory tangling, an established hallmark of motor-cortical population dynamics (Russo et al., 2018). We tested whether cortical alignment induces this feature in the student by measuring velocity disagreement between nearby hidden states. Tangling is neither a training penalty nor a dose-selection criterion. We analyzed all 128 hidden coordinates in 2,100 GRUs, with seven methods, three datasets, and 100 seeds; strengths minimize median training MSE over each aligned method's complete 40-dose grid. Figure 5 summarizes each model's mean Q over all 40 frames of its 28 trained reaches, using all 36 conditions as references (Appendix E).

Neuro-Dist reduces median total Q relative to task-only training by 55.5%, 68.1%, and 48.1% in Monkeys 1–3, respectively. All five alternative teachers also reduce median tangling relative to task-only. Notably, Neuro-Dist has the lowest median total Q among the seven methods in Monkeys 1 and 2—the same monkeys where Neuro-Dist substantially reduced reaching errors: tangling Q was 0.138 and 0.105, versus 0.201 and 0.138 for Behavior. In Monkey 3, Behavior's Q (and reaching error) is lower (0.134 versus 0.177), while Neuro-Dist remains below the other five controls. These

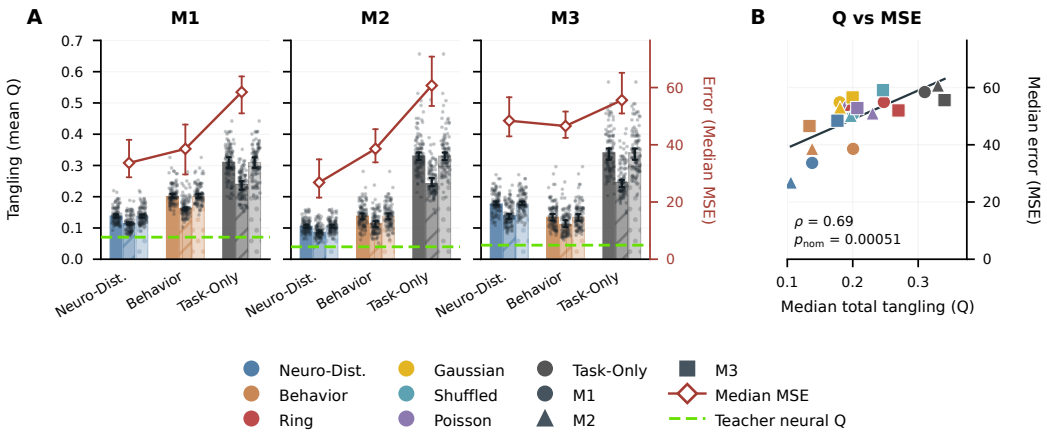

Figure 5: **Neuro-distillation reduces GRU state-space trajectory tangling. A,** Bars: median total, within-reach and between-reach Q (left to right); dots: individual models. Red diamonds: median training MSE (mm^2); whiskers: pointwise 95% bootstrap intervals across 100 seeds. Green dashed lines: cortical references. Neuro-Dist has the lowest median tangling and training error in M1–M2; Behavior has both minima in M3 (Table 1A, GRU rows). **B,** Lower tangling correlates with lower reaching error across 21 dataset–method conditions. See Appendix E for analysis and display details.

Neuro-Dist–Behavior relationships also hold for within-reach and between-reach tangling. Crucially, and confirming the observation in these examples, across all 21 animal–method summaries, median tangling is positively and significantly associated with median training MSE (Spearman $\rho = 0.692$, nominal $p = 5.07 \times 10^{-4}$). Together with the biological importance of low tangling for robust motor control, these findings suggest transfer of a low-tangling state as a mechanism of the neuro-distillation success in achieving robust movement control also in artificial controllers.

The Neuro-Dist–Behavior ordering persists on student-held-out reaches and new-target evaluations, with all seven methods included in both analyses. Supplementary Figure 11 reports these comparisons and decompositions of velocity disagreement and state separation. Figure 4 illustrates the geometry using Monkey 2 as an example; the population comparisons above use all seeds and full hidden states.

5 Discussion, Limitations and Conclusions

In the present study we show a proof of concept for neuro-distillation, a technique designed to directly transfer computationally useful neural population structure from biological brains into artificial controllers. Using three macaque intracortical datasets as neuro-distillation teachers, we successfully align the internal states of GRU and Transformer controllers with primary motor cortex population geometry extracted from single-unit spiking data.

We selected CCA-based subspace alignment for neuro-distillation *a priori*, at the start of the project. This choice was motivated by evidence that CCA can reveal shared, behaviorally relevant structure across neuronal populations, and is supported by successful applications of CCA-based alignment in brain-guided learning and model distillation (Wang et al., 2015; Pirlot et al., 2022; Fan et al., 2025). During reaching, heterogeneous single-neuron responses collectively form organized, low-dimensional population trajectories (Churchland et al., 2012). Using CCA, Gallego et al. (2020) revealed stable motor-cortical latent dynamics across recording sessions despite changes in the recorded neuronal sample, supporting reliable decoding of movement over time. Safaie et al. (2023) subsequently used CCA to demonstrate preserved latent dynamics across animals performing similar behaviors, enabling movement decoding across individuals. Together, these findings provided a physiological rationale for targeting shared linear population structure when aligning cortical activity with artificial controller states. CCA accommodates different population sizes without requiring correspondence between individual neurons and artificial units. The current study shows that CCA-

based neural supervision can be extended to the motor control domain, by aligning controller states directly with behavior-predictive neural teacher subspaces. Our findings demonstrate that this alignment concept improved reaching accuracy (Table 1) and increased robustness to sensory feedback noise in the recurrent controllers (Figure 3B–D). Interstingly, neuro-distillation based on trial-shuffled data was among the least accurate teachers (Table 1). This highlights that cortical statistics alone cannot explain the beneficial effects of neuro-distillation.

Neuro-distillation reduced trajectory tangling in GRU population activity (Figure 5A), despite tangling being neither an explicit training objective nor a criterion for model selection. Low tangling characterizes motor-cortical activity, and recurrent-network experiments have demonstrated its contribution to noise robustness (Russo et al., 2018). Thus, our findings replicate this neurobiologically documented association between low tangling and improved reaching movements in our distilled controllers, across animal–method groups (Figure 5B; Appendix E). These observations support a functional role of reduced tangling in the observed distillation control benefits. However, our analyses are correlative and do not yet establish exactly to what extent other neural population features may have contributed to the observed control benefits. An intriguing possibility therefore is that reduced tangling reflects only a facet of a broader transfer of computationally useful population structure, including features that our analyses may not yet have identified. Candidates for further investigation include the rotational organization of movement-related population trajectories (Churchland et al., 2012) and their movement-dependent arrangement in neural state space (Sabatini & Kaufman, 2024), which may provide a structured—and possibly distillable—basis for generating diverse motor commands. Together, these considerations also raise the question of whether cortical supervision may convey useful structure beyond that provided by any alternative, neuro-agnostic teacher such as the stochastic, behavior-based, and engineered teachers tested here. Future studies could probe these possibilities by selectively preserving or disrupting candidate features in neural teachers and, alongside matched controls, testing their causal contribution to controller accuracy and robustness. Establishing which features account for the functional value of neuro-distillation would guide its development and could also offer a new paradigm to generate testable hypotheses to advance our mechanistic understanding of biological motor control.

Beyond the aspects discussed so far, the present proof of concept of neuro-distillation from biological neural systems to artificial controllers opens a broad research agenda: key questions are which aspects of biological population computation can improve motor control and goal-directed behavior in robotic systems and embodied AI agents, and what roles the teacher and student tasks and the kinds of effectors used (e.g., their biomechanical realism) play. Another crucial direction is translation from simulations, as studied here, to real-world settings, to investigate whether neuro-distillation can bring significant, relevant and lasting advantages in real-world applications.

Such applications range from dexterous manipulation and articulate speech production to whole-body coordination or expressive facial movements in humanoids, and the control of walking, flying, or swimming robots. Benefits could encompass accuracy and robustness, but also learning efficiency, adaptation, and generalization to unfamiliar bodies or environments. Different cortical areas, combinations of areas, species, and behavioral repertoires may offer useful structure for different computational demands. Comparing these sources could distinguish broadly reusable principles from specializations associated with particular behaviors or embodiments. At single-neuron resolution, technologies such as Neuropixels probes (Trautmann et al., 2025), wireless microelectrode arrays (Schwarz et al., 2014), and miniaturized calcium-imaging microscopes (Zhang et al., 2024) provide complementary ways to sample population activity during behavior. These approaches could help establish how population coverage, temporal resolution, and access to natural behavior influence the value of neural teachers. Equally central is the design of potentially novel, optimized distillation objectives that preserve and allow transfer of the relevant aspects of population geometry, temporal organization, or coordination between neural processes. Together, these topics and questions outline a research field on neuro-distillation as a new form of neuro-interfacing, linking the neurobiology of neural population computations and machine learning, with the aim to transfer useful computational structure to guide the computations learned by artificial controllers across tasks and embodiments.

In this study, we used publicly available macaque datasets; human neural data can be expected to offer important perspectives beyond animal recordings relevant to neuro-distillation in our sense, for example for dexterous hand and finger control or speech production. Such opportunities, however, would also come with new ethical questions, such as those regarding data privacy, which we advocate addressing as a crucial, integral research component right from the beginning (UNESCO, 2025).

AI USE STATEMENT

In this work, generative AI tools assisted research execution through coding support, including implementation of analysis and figure-generation scripts. They also supported literature retrieval and manuscript editing. We did not use generative AI tools for research ideation, methodological critique, or interpretation of results. The authors reviewed and tested AI-assisted code, checked numerical results, and verified every aspect of the scientific work and its presentation. We take responsibility for the final content of this work, including text, claims, code, and other artifacts produced with AI assistance.

ETHICS STATEMENT

This study reuses publicly available macaque motor-cortical recordings (Pei et al., 2021; O'Doherty et al., 2017) and evaluates artificial controllers in simulation. No new animal procedures or human-participant studies were conducted. The Neural Latents Benchmark publication reports institutional animal-care approvals for its source recordings (Pei et al., 2021). Our findings concern three macaque datasets and a simulated reaching task. Any future translation to physical robots, clinical devices, or human neural interfaces would require separate validation and appropriate oversight. Future extensions using human neural data should address informed consent, protection of neural-data privacy, and safeguards against unauthorized reuse from the outset (UNESCO, 2025).

REPRODUCIBILITY STATEMENT

We take a comprehensive and systematic approach to supporting the reproducibility of our methods and results. We use a configuration-driven workflow with separate stages for training, evaluation, statistical analysis, and figure generation. Experiment configurations specify model and optimization settings, teacher conditions, parameter sweeps, and random seeds. Run records preserve configurations, software versions, and hashes of source code and inputs alongside saved checkpoints, trajectories, and numerical summaries. Downstream analyses reuse these fixed outputs, while provenance records connect manuscript tables and figures to their numerical sources. Validation includes replaying saved models and independently recomputing evaluation metrics and statistical summaries. Section 3 and Appendices A–F describe the scientific protocols. The anonymous code supplement provides implementations, configurations, processed inputs, selected-model results, and analysis and figure-generation scripts, with dependency and usage documentation. Together, these practices provide a transparent and traceable basis for independently verifying, reproducing, and extending the work.

REFERENCES

Iretiayo Akinola, Zizhao Wang, Junyao Shi, Xiaomin He, Pawan Lapborisuth, Jingxi Xu, David Watkins-Valls, Paul Sajda, and Peter Allen. Accelerated robot learning via human brain signals. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3799–3805, 2020. doi: 10.1109/ICRA40945.2020.9196566.

Diego Aldarondo, Josh Merel, Jesse D. Marshall, Leonard Hasenclever, Ugne Klibaite, Amanda Gellis, Yuval Tassa, Greg Wayne, Matthew Botvinick, and Bence P. Ölveczky. A virtual rodent predicts the structure of neural activity across behaviours. *Nature*, 632(8025):594–602, August 2024. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-024-07633-4. URL https://www.nature.com/articles/s41586-024-07633-4.

Muhammad Noman Almani, John Lazzari, and Shreya Saxena. muSim: A goal-driven framework for elucidating the neural control of movement through musculoskeletal modeling. *bioRxiv*, 2025. doi: 10.1101/2024.02.02.578628. URL https://www.biorxiv.org/content/early/2025/11/07/2024.02.02.578628.

Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In Sanjoy Dasgupta and David McAllester (eds.), *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 1247–1255, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL https://proceedings.mlr.press/v28/andrew13.html.

Tomas Aquino, Victoria Qinwanxian Liu, Habiba Azab, Raissa Mathura, Andrew J Watrous, Eleonora Bartoli, Benjamin Y Hayden, Paul Sajda, Sameer A. Sheth, and Nuttida Rungratsameetaweemana. Brain2Model learning: Training sensory and decision models with human neural activity as a teacher. In *NeurIPS 2025 Workshop on Foundation Models for the Brain and Body*, 2025. URL https://openreview.net/forum?id=iV6UJyFrjH.

Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil J Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deeksha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez, Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Kevin Sayed, Jaspiar Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong Tran, Vincent Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. RT-1: Robotics transformer for real-world control at scale. arXiv:2212.06817, 2023. URL https://arxiv.org/abs/2212.06817.

Gavin C. Cawley and Nicola L. C. Talbot. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11(70):2079–2107, 2010. URL http://jmlr.org/papers/v11/cawley10a.html.

Jiaxuan Chen, Yu Qi, Yueming Wang, and Gang Pan. Human-like cognitive generalization for large models via mental representation-guided supervision. *Nature Communications*, 17(1), April 2026. ISSN 2041-1723. doi: 10.1038/s41467-026-71267-5. URL http://dx.doi.org/10.1038/s41467-026-71267-5.

Kyunghyun Cho, Bart van Merrienboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734. Association for Computational Linguistics, 2014. doi: 10.3115/v1/d14-1179. URL http://dx.doi.org/10.3115/v1/D14-1179.

SueYeon Chung and L.F. Abbott. Neural population geometry: An approach for understanding biological and artificial neural networks. *Current Opinion in Neurobiology*, 70:137–144, October 2021. ISSN 0959-4388. doi: 10.1016/j.conb.2021.10.010. URL http://dx.doi.org/10.1016/j.conb.2021.10.010.

Mark M. Churchland, John P. Cunningham, Matthew T. Kaufman, Justin D. Foster, Paul Nuyujukian, Stephen I. Ryu, and Krishna V. Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, June 2012. ISSN 1476-4687. doi: 10.1038/nature11129. URL http://dx.doi.org/10.1038/nature11129.

Olivier Codol, Jonathan A Michaels, Mehrdad Kashefi, J Andrew Pruszynski, and Paul L Gribble. MotorNet, a Python toolbox for controlling differentiable biomechanical effectors with artificial neural networks. *eLife*, 12, July 2024. ISSN 2050-084X. doi: 10.7554/elife.88591.4. URL http://dx.doi.org/10.7554/eLife.88591.4.

Joel Dapello, Kohitij Kar, Martin Schrimpf, Robert Baldwin Geary, Michael Ferguson, David Daniel Cox, and James J. DiCarlo. Aligning model and macaque inferior temporal cortex representations improves model-to-human behavioral alignment and adversarial robustness. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=SMYdcXjJh1q.

Bradley Efron. Better bootstrap confidence intervals. *Journal of the American Statistical Association*, 82(397):171–185, March 1987. ISSN 1537-274X. doi: 10.1080/01621459.1987.10478410. URL http://dx.doi.org/10.1080/01621459.1987.10478410.

Wen-Shu Fan, Su Lu, Shangyu Xing, Xin-Chun Li, and De-Chuan Zhan. Maximizing the effectiveness of larger BERT models for compression. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 21975–21990, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.1067. URL https://aclanthology.org/2025.acl-long.1067/.

Callie Federer, Haoyan Xu, Alona Fyshe, and Joel Zylberberg. Improved object recognition using neural networks trained to mimic the brain's statistical properties. *Neural Networks*, 131:103–114, November 2020. ISSN 0893-6080. doi: 10.1016/j.neunet.2020.07.013. URL http://dx.doi.org/10.1016/j.neunet.2020.07.013.

Ruth C. Fong, Walter J. Scheirer, and David D. Cox. Using human brain activity to guide machine learning. *Scientific Reports*, 8(1), March 2018. ISSN 2045-2322. doi: 10.1038/s41598-018-23618-6. URL http://dx.doi.org/10.1038/s41598-018-23618-6.

Maëlle Freteault, Maximilien Le Clei, Loic Tetrel, Lune Bellec, and Nicolas Farrugia. Alignment of auditory artificial networks with massive individual fMRI brain data leads to generalisable improvements in brain encoding and downstream tasks. *Imaging Neuroscience*, 3, 2025. ISSN 2837-6056. doi: 10.1162/imag_a_00525. URL http://dx.doi.org/10.1162/imag_a_00525.

Juan A. Gallego, Matthew G. Perich, Raeed H. Chowdhury, Sara A. Solla, and Lee E. Miller. Long-term stability of cortical population dynamics underlying consistent behavior. *Nature Neuroscience*, 23(2):260–270, January 2020. ISSN 1546-1726. doi: 10.1038/s41593-019-0555-4. URL http://dx.doi.org/10.1038/s41593-019-0555-4.

Lior Greenspoon, Eyal Krieger, Ron Sender, Yuval Rosenberg, Yinon M. Bar-On, Uri Moran, Tomer Antman, Shai Meiri, Uri Roll, Elad Noor, and Ron Milo. The global biomass of wild mammals. *Proceedings of the National Academy of Sciences*, 120(10), February 2023. ISSN 1091-6490. doi: 10.1073/pnas.2204892120. URL http://dx.doi.org/10.1073/pnas.2204892120.

Ankur Handa, Arthur Allshire, Viktor Makoviychuk, Aleksei Petrenko, Ritvik Singh, Jingzhou Liu, Denys Makoviichuk, Karl Van Wyk, Alexander Zhurkevich, Balakumar Sundaralingam, Yashraj Narang, Jean-Francois Lafleche, Dieter Fox, and Gavriel State. DeXtreme: Transfer of agile in-hand manipulation from simulation to reality. arXiv:2210.13702, 2024. URL https://arxiv.org/abs/2210.13702.

Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2):65–70, 1979. ISSN 03036898, 14679469. URL http://www.jstor.org/stable/4615733.

Harold Hotelling. Relations between two sets of variates. *Biometrika*, 28(3-4):321–377, December 1936. ISSN 1464-3510. doi: 10.1093/biomet/28.3-4.321. URL http://dx.doi.org/10.1093/biomet/28.3-4.321.

Iñaki Iturrate, Ricardo Chavarriaga, Luis Montesano, Javier Minguez, and José del R. Millán. Teaching brain-machine interfaces as an alternative paradigm to neuroprosthetics control. *Scientific Reports*, 5(1), September 2015. ISSN 2045-2322. doi: 10.1038/srep13893. URL http://dx.doi.org/10.1038/srep13893.

Heekyung Jung, Myungin Baek, Kristen P. D'Elia, Catherine Boisvert, Peter D. Currie, Boon-Hui Tay, Byrappa Venkatesh, Stuart M. Brown, Adriana Heguy, David Schoppik, and Jeremy S. Dasen. The ancient origins of neural substrates for land walking. *Cell*, 172(4):667–682.e15, February 2018. ISSN 0092-8674. doi: 10.1016/j.cell.2018.01.013. URL http://dx.doi.org/10.1016/j.cell.2018.01.013.

Maryna Kapitonova and Tonio Ball. Neuro-inspired inverse learning for planning and control. arXiv:2605.24152, 2026. URL https://arxiv.org/abs/2605.24152.

Su Kyoung Kim, Elsa Andrea Kirchner, Arne Stefes, and Frank Kirchner. Intrinsic interactive reinforcement learning – using error-related potentials for real world human-robot interaction. *Scientific Reports*, 7(1), December 2017. ISSN 2045-2322. doi: 10.1038/s41598-017-17682-7. URL http://dx.doi.org/10.1038/s41598-017-17682-7.

Roger Koenker and Kevin F Hallock. Quantile regression. *Journal of Economic Perspectives*, 15 (4):143–156, November 2001. ISSN 0895-3309. doi: 10.1257/jep.15.4.143. URL http://dx.doi.org/10.1257/jep.15.4.143.

Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3519–3529. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/kornblith19a.html.

HengYu Li, Kangdi Mei, Zhaoci Liu, Yang Ai, Liping Chen, Jie Zhang, and Zhenhua Ling. Refining self-supervised learnt speech representation using brain activations. In *Interspeech 2024*, interspeech_2024, pp. 1480–1484. ISCA, September 2024. doi: 10.21437/interspeech.2024-604. URL http://dx.doi.org/10.21437/Interspeech.2024-604.

Ren Li, Jared S. Johansen, Hamad Ahmed, Thomas V. Ilyevsky, Ronnie B. Wilbur, Hari M. Bharadwaj, and Jeffrey Mark Siskind. The perils and pitfalls of block design for EEG classification experiments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2020. ISSN 1939-3539. doi: 10.1109/tpami.2020.2973153. URL http://dx.doi.org/10.1109/TPAMI.2020.2973153.

Zhe Li, Wieland Brendel, Edgar Walker, Erick Cobos, Taliah Muhammad, Jacob Reimer, Matthias Bethge, Fabian Sinz, Zachary Pitkow, and Andreas Tolias. Learning from brains how to regularize machines. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/70117ee3c0b15a2950f1e82a215e812b-Paper.pdf.

Zitong Lu, Yile Wang, and Julie D. Golomb. Achieving more human brain-like vision via human EEG representational alignment. *Communications Biology*, 9(1), February 2026. ISSN 2399-3642. doi: 10.1038/s42003-026-09685-w. URL http://dx.doi.org/10.1038/s42003-026-09685-w.

Patrick McClure and Nikolaus Kriegeskorte. Representational distance learning for deep neural networks. *Frontiers in Computational Neuroscience*, 10, December 2016. ISSN 1662-5188. doi: 10.3389/fncom.2016.00131. URL http://dx.doi.org/10.3389/fncom.2016.00131.

Omer Moussa and Mariya Toneva. Brain-tuning improves generalizability and efficiency of brain alignment in speech models. In *Advances in Neural Information Processing Systems 38*, NeurIPS 2025, pp. 135575–135600. Neural Information Processing Systems Foundation, Inc. (NeurIPS), 2025. doi: 10.52202/085713-4079. URL http://dx.doi.org/10.52202/085713-4079.

Omer Moussa, Dietrich Klakow, and Mariya Toneva. Improving semantic understanding in speech language models via brain-tuning. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 59823–59849, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/95b7a93e60fdfd10cc202f44fd6adf5f-Paper-Conference.pdf.

Anuja Negi, SUBBAREDDY OOTA, Anwar Nunez-Elizalde, Manish Gupta, and Fatma Deniz. Brain-informed fine-tuning for improved multilingual understanding in language models. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 152057–152086. Curran Associates, Inc., 2025. doi: 10.52202/085713-5084. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/df311d51f4f1a3d0adc9d3f1725d0850-Paper-Conference.pdf.

Satoshi Nishida, Yusuke Nakano, Antoine Blanc, Naoya Maeda, Masataka Kado, and Shinji Nishimoto. Brain-mediated transfer learning of convolutional neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):5281–5288, April 2020. ISSN 2159-5399. doi: 10.1609/aaai.v34i04.5974. URL http://dx.doi.org/10.1609/aaai.v34i04.5974.

Joseph E. O'Doherty, Mariana M. B. Cardoso, Joseph G. Makin, and Philip N. Sabes. Nonhuman primate reaching with multichannel sensorimotor cortex electrophysiology, 2017. URL https://zenodo.org/record/583331. Dataset.

M. Teague O'Mara, Anne K. Scharf, Jakob Fahr, Michael Abedi-Lartey, Martin Wikelski, Dina K. N. Dechmann, and Kamran Safi. Overall dynamic body acceleration in straw-colored fruit bats increases in headwinds but not with airspeed. *Frontiers in Ecology and Evolution*, 7, 2019. ISSN 2296-701X. doi: 10.3389/fevo.2019.00200. URL https://www.frontiersin.org/journals/ecology-and-evolution/articles/10.3389/fevo.2019.00200.

Simone Palazzo, Concetto Spampinato, Joseph Schmidt, Isaak Kavasidis, Daniela Giordano, and Mubarak Shah. Correct block-design experiments mitigate temporal correlation bias in EEG classification. arXiv:2012.03849, 2020. URL https://arxiv.org/abs/2012.03849.

Simone Palazzo, Concetto Spampinato, Isaak Kavasidis, Daniela Giordano, Joseph Schmidt, and Mubarak Shah. Decoding brain representations by multimodal learning of neural activity and visual features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3833–3849, November 2021. ISSN 1939-3539. doi: 10.1109/tpami.2020.2995909. URL http://dx.doi.org/10.1109/TPAMI.2020.2995909.

Il Memming Park, Ayesha Vermani, Gonzalo G. de Polavieja, Juan Álvaro Gallego, Kathleen Esfahany, Shreya Saxena, Michael Orger, Auke Ijspeert, Matthew Dowling, Daniel McNamee, Srinivas C. Turaga, Zachary Mainen, Joseph J. Paton, and Alfonso Renart. Integrative neurocybernetic modeling in the era of large-scale neuroscience. arXiv:2604.23903, 2026. URL https://arxiv.org/abs/2604.23903.

Felix Pei, Joel Ye, David Zoltowski, Anqi Wu, Raeed Chowdhury, Hansem Sohn, Joseph O'Doherty, Krishna V Shenoy, Matthew Kaufman, Mark Churchland, Mehrdad Jazayeri, Lee Miller, Jonathan Pillow, Il Memming Park, Eva Dyer, and Chethan Pandarinath. Neural Latents Benchmark '21: Evaluating latent variable models of neural population activity. In J. Vanschoren and S. Yeung (eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/979d472a84804b9f647bc185a877a8b5-Paper-round2.pdf.

Cassidy Pirlot, Richard C. Gerum, Cory Efird, Joel Zylberberg, and Alona Fyshe. Improving the accuracy and robustness of CNNs using a deep CCA neural data regularizer. arXiv:2209.02582, 2022. URL https://arxiv.org/abs/2209.02582.

Jean-Alban Rathelot and Peter L. Strick. Subdivisions of primary motor cortex based on cortico-motoneuronal cells. *Proceedings of the National Academy of Sciences*, 106(3):918–923, January 2009. ISSN 1091-6490. doi: 10.1073/pnas.0808362106. URL http://dx.doi.org/10.1073/pnas.0808362106.

Abigail A. Russo, Sean R. Bittner, Sean M. Perkins, Jeffrey S. Seely, Brian M. London, Antonio H. Lara, Andrew Miri, Najja J. Marshall, Adam Kohn, Thomas M. Jessell, Laurence F. Abbott, John P. Cunningham, and Mark M. Churchland. Motor cortex embeds muscle-like commands in an untangled population response. *Neuron*, 97(4):953–966.e8, February 2018. ISSN 0896-6273. doi: 10.1016/j.neuron.2018.01.004. URL http://dx.doi.org/10.1016/j.neuron.2018.01.004.

David A. Sabatini and Matthew T. Kaufman. Reach-dependent reorientation of rotational dynamics in motor cortex. *Nature Communications*, 15(1), August 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-51308-7. URL http://dx.doi.org/10.1038/s41467-024-51308-7.

Mostafa Safaie, Joanna C. Chang, Junchol Park, Lee E. Miller, Joshua T. Dudman, Matthew G. Perich, and Juan A. Gallego. Preserved neural dynamics across animals performing similar behaviour. *Nature*, 623(7988):765–771, November 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06714-0. URL http://dx.doi.org/10.1038/s41586-023-06714-0.

Shahd Safarani, Arne Nix, Konstantin Willeke, Santiago Cadena, Kelli Restivo, George Denfield, Andreas Tolias, and Fabian Sinz. Towards robust vision by multi-task learning on monkey visual cortex. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 739–751. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/06a9d51e04213572ef0720dd27a84792-Paper.pdf.

Dan Schwartz, Mariya Toneva, and Leila Wehbe. Inducing brain-relevant bias in natural language processing models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/2b8501af7b64d1aaae7dd832805f0709-Paper.pdf.

David A Schwarz, Mikhail A Lebedev, Timothy L Hanson, Dragan F Dimitrov, Gary Lehew, Jim Meloy, Sankaranarayani Rajangam, Vivek Subramanian, Peter J Ifft, Zheng Li, Arjun Ramakrishnan, Andrew Tate, Katie Z Zhuang, and Miguel A L Nicolelis. Chronic, wireless recordings of large-scale brain activity in freely moving rhesus monkeys. *Nature Methods*, 11(6):670–676, April 2014. ISSN 1548-7105. doi: 10.1038/nmeth.2936. URL http://dx.doi.org/10.1038/nmeth.2936.

Zhenan Shao, Linjian Ma, Yiqing Zhou, Yibo Jacky Zhang, Sanmi Koyejo, Bo Li, and Diane M. Beck. Probing human visual robustness with neurally-guided deep neural networks. arXiv:2405.02564, 2025. URL https://arxiv.org/abs/2405.02564.

Concetto Spampinato, Simone Palazzo, Isaak Kavasidis, Daniela Giordano, Nasim Souly, and Mubarak Shah. Deep learning human mind for automated visual classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

David Sussillo, Mark M Churchland, Matthew T Kaufman, and Krishna V Shenoy. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience*, 18(7): 1025–1033, June 2015. ISSN 1546-1726. doi: 10.1038/nn.4042. URL http://dx.doi.org/10.1038/nn.4042.

Tony Tao, Mohan Srirama, Jason Liu, Kenneth Shaw, and Deepak Pathak. DexWild: Dexterous human interactions for in-the-wild robot policies. In *Robotics: Science and Systems XXI*, RSS2025. Robotics: Science and Systems Foundation, June 2025. doi: 10.15607/rss.2025.xxi.075. URL http://dx.doi.org/10.15607/RSS.2025.XXI.075.

Eric M. Trautmann, Janis K. Hesse, Gabriel M. Stine, Ruobing Xia, Shude Zhu, Daniel J. O'Shea, Bill Karsh, Jennifer Colonell, Frank F. Lanfranchi, Saurabh Vyas, Andrew Zimnik, Elom Amematsro, Natalie A. Steinemann, Daniel A. Wagenaar, Marius Pachitariu, Alexandru Andrei, Carolina Mora Lopez, John O'Callaghan, Jan Putzeys, Bogdan C. Raducanu, Marleen Welkenhuysen, Mark Churchland, Tirin Moore, Michael Shadlen, Krishna Shenoy, Doris Tsao, Barundeb Dutta, and Timothy Harris. Large-scale high-density brain-wide neural recording in nonhuman primates. *Nature Neuroscience*, 28(7):1562–1575, June 2025. ISSN 1546-1726. doi: 10.1038/s41593-025-01976-5. URL http://dx.doi.org/10.1038/s41593-025-01976-5.

UNESCO. Recommendation on the ethics of neurotechnology. Adopted by the General Conference of UNESCO at its 43rd session, Samarkand, November 2025. URL https://www.unesco.org/en/legal-affairs/recommendation-ethics-neurotechnology.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

Shangfei Wang, Yachen Zhu, Lihua Yue, and Qiang Ji. Emotion recognition with the help of privileged information. *IEEE Transactions on Autonomous Mental Development*, 7(3):189–200, 2015. doi: 10.1109/TAMD.2015.2463113.

Mingqing Xiao, Kai Du, and Zhouchen Lin. Beyond representational alignment with brain-guided language models for robust reasoning. *Nature Machine Intelligence*, 8(8):1275–1289, August 2026. ISSN 2522-5839. doi: 10.1038/s42256-026-01278-w. URL http://dx.doi.org/10.1038/s42256-026-01278-w.

Anthony M. Zador. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 10(1), August 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-11786-6. URL http://dx.doi.org/10.1038/s41467-019-11786-6.

Yuanlong Zhang, Lekang Yuan, Qiyu Zhu, Jiamin Wu, Tobias Nöbauer, Rujin Zhang, Guihua Xiao, Mingrui Wang, Hao Xie, Zengcai Guo, Qionghai Dai, and Alipasha Vaziri. A miniaturized mesoscope for the large-scale single-neuron-resolved imaging of neuronal activity in freely behaving mice. *Nature Biomedical Engineering*, 8(6):754–774, June 2024. ISSN 2157-846X. doi: 10.1038/s41551-024-01226-2. URL http://dx.doi.org/10.1038/s41551-024-01226-2.

Zhejun Zhang, Wenqing Zhou, Haozhe Xu, Lin Zhang, and Lei Li. Temporal precision matters: Brain-tuning speech language models with millisecond-resolution neural signals. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 41208–41226. Association for Computational Linguistics, 2026. doi: 10.18653/v1/2026.acl-long.1911. URL http://dx.doi.org/10.18653/v1/2026.acl-long.1911.

SUPPLEMENTARY MATERIAL

A DETAILED METHODS

Neural teachers. We use frozen condition-averaged motor-cortical teachers from three macaques: Monkey 1 (MC_Maze, distributed through the Neural Latents Benchmark; Pei et al., 2021) and Monkeys 2 and 3 (O'Doherty et al., 2017). Each teacher contains 36 conditions and 40 time bins. Preprocessing uses 10-ms bins, Gaussian smoothing with a 30-ms SD, soft normalization, and 12 PCs. O'Doherty conditions combine 12 direction bins and three reach-magnitude bins; session-specific latent trajectories are aligned by orthogonal Procrustes and averaged. Neural windows cover execution for Monkey 1 (approximately 0–400 ms from movement onset) and peri-movement activity for Monkeys 2–3 (-200 to 200 ms). These different windows reflect the task paradigms: Monkey 1 had an explicitly defined movement-execution phase in an instructed-delay task, whereas Monkeys 2–3 performed continuous, self-paced reaching without instructed epochs, motivating peri-movement windows centered on kinematically detected movement onset.

For centered cortical scores Y_c and recorded hand-position and velocity variables K_c, we fit $B = \arg\min_B \|Y_c B - K_c\|_F^2$ and set $Z = Y_c U_B U_B^\top$, where U_B spans the columns of B. The teacher has rank four, stored in 12 coordinates; it retains cortical variation within these directions, rather than replacing activity with kinematics. Construction uses all 36 conditions before student splits. Thus, eight conditions are held out only from student optimization. Because conditions include magnitude bins, withholding a condition need not remove its entire direction class. New evaluation directions lie outside the original condition set (Appendix B.3).

Control teachers. The task-only baseline and five alternative teachers probe different possible sources of benefit from alignment. Each teacher replaces the cortical representation in the same alignment objective; task-only omits this term.

- **Task-only:** the controller minimizes reaching loss alone ($\alpha = 0$), without a teacher, using the same 28 training conditions and seed-specific split as the aligned controllers.
- **Behavior:** this control supplies geometry derived directly from recorded movement rather than neural activity. We selected the *Behavior* teacher as the strongest overall control among several candidates in preliminary tests, based on achieving the lowest error in the largest number of comparisons. It uses the same recorded x, y, v_x, v_y variables that select the cortical subspace. We concatenate offsets of $-100, -50, 0, 50, 100$ ms, clip out-of-window indices to boundary samples, standardize each of the 20 features across all 36×40 teacher-fit condition–time observations, and retain four unwhitened PCs with deterministic signs. Positive offsets use future samples only in this offline teacher; recordings are unnecessary at inference. A fixed orthonormal map embeds the representation into 12 coordinates, preserving rank four without adding noise; it does not match the cortical covariance spectrum or effective alignment gradients. One frozen teacher per animal is shared across seeds, architectures, and doses. Monkey 1 uses the verified straight-reach execution trajectories underlying cortical-subspace selection. The behavioral representation preserves the original condition and sample ordering. For Monkeys 2–3, execution-window kinematics select a peri-movement neural subspace, so physical time windows are not matched.

- **Ring:** this engineered teacher tests a simple directional prior, motivated by the directional organization observed in cortical activity and trained controllers. See Appendix D for the full analysis of directional organization, the control's motivation, and its construction and visualization. Its first two coordinates multiply the cosine and sine of the target angle by a shared temporal envelope, increasing linearly from 0.3 to 1.0 before Gaussian smoothing with a 30-ms SD. Each coordinate is scaled using the corresponding leading cortical covariance eigenvalue. Direction stays fixed within a reach while amplitude increases, so trajectories extend radially rather than rotate around the ring. Ten additional coordinates are generated from temporally smoothed Gaussian signals, centered, rotated to their principal axes, and scaled using the remaining cortical eigenvalues, independently for each model seed. Concatenating the directional and random components gives 12 coordinates; their cross-covariances prevent exact matching of the complete cortical covariance spectrum.

- **Gaussian:** this control preserves population variance structure without imposing reach-direction geometry. Independent Gaussian signals are smoothed along time with a 30-ms SD, centered over all 36×40 samples, rotated to their principal axes, and rescaled using the cortical covariance eigenvalues. The resulting 12-coordinate teachers match the normalized cortical covariance spectrum and are generated independently per model seed. Spectrum matching applies to the pooled 36 conditions, not necessarily to each 28-condition training subset.

- **Poisson:** this control supplies synthetic population activity using animal-specific fitted lognormal unit rates with Gaussian log-rate modulation (amplitude 0.4, 30-ms smoothing). Twenty Poisson trials are averaged in 10-ms bins, retaining mean counts per bin rather than converting to rates in Hz. After 30-ms temporal smoothing, soft normalization by unit range plus five, centering, and unwhitened PCA over all 36 conditions, four PC scores are retained and eight exact-zero coordinates appended. One frozen teacher per animal and seed is shared across architectures and doses. This matches cortical teacher rank and stored dimension, not its covariance spectrum. Fitted rate summaries use execution windows, which differ from the peri-movement cortical windows for Monkeys 2–3.

- **Shuffled:** this control disrupts condition correspondence while preserving complete cortical trajectories. An independent condition derangement is drawn for each animal and model seed, rejecting fixed points among all 36 conditions so that no trajectory remains assigned to its original condition. Complete trajectories are reassigned with their time order unchanged, and the permutation remains fixed throughout training and is shared across architectures and doses. Because different conditions can share a direction class, derangement of condition indices does not ensure that every assigned reach direction changes.

Table 2 continues Table 1 with Gaussian and Poisson columns, using the same metric, target sets, and training-based strength selection. Each control has complete 40-dose, 100-seed sweeps in both architectures and all three animals. For each control, 600 training-selected controllers were evaluated on the same 300 frozen new-target sets, yielding 4,800 new reaches. No retraining or strength retuning was performed.

Task and student architectures. Students control MotorNet's differentiable planar arm with two joints, six rigid-tendon Hill-type muscles, and 10-ms simulation steps (Codol et al., 2024). Rollouts last 40 steps and begin at the midpoint of joint ranges with zero joint velocity. Targets lie 10 cm from the initial fingertip in each teacher's stored directions, irrespective of the recorded magnitude bin. Main-benchmark observations and rollouts are deterministic.

The closed-loop GRU (Cho et al., 2014) has 128 hidden units and updated observations at every step. The open-loop Transformer (Vaswani et al., 2017) has three pre-normalized layers, width 64, four attention heads, feedforward width 128, and zero dropout. Initial observations, absolute and relative target coordinates, and target distance condition all time tokens, with learned positional embeddings and causal attention. It emits all 40 muscle commands in one forward pass without subsequent feedback; output biases initialize to -2. The Transformer is trained as an Inverter, following Kapitonova & Ball (2026), by backpropagating the reaching objective through the differentiable arm model over the full action sequence. The architectures are not parameter-count matched.

Table 2: **Continuation of Table 1: Gaussian and Poisson controls.** Median ± bootstrap SE of final-window coordinate MSE (mm^2; 100 seeds, 100,000 seed resamples), using the same target sets and strength-selection protocol as Table 1. Each strength minimizes median training MSE over 40 doses and remains fixed across target sets; SEs condition on that selection. The sum of wins counts lowest medians across all seven methods in this table and Table 1, over the same 18 comparisons. Wins are descriptive, not significance tests.

Student	Monkey	Gaussian	Poisson
A. Trained targets (28)			
GRU	M1	54.85 ± 3.46	53.52 ± 2.86
GRU	M2	52.97 ± 2.79	50.97 ± 4.20
GRU	M3	56.55 ± 3.64	52.83 ± 3.67
Transf.	M1	28.64 ± 2.29	26.70 ± 3.59
Transf.	M2	18.52 ± 2.17	17.89 ± 1.43
Transf.	M3	23.00 ± 2.16	18.73 ± 2.22
B. Generalization to student-held-out targets (8)			
GRU	M1	59.86 ± 4.83	60.23 ± 4.35
GRU	M2	59.39 ± 3.91	52.54 ± 2.57
GRU	M3	62.57 ± 3.67	59.59 ± 4.12
Transf.	M1	35.12 ± 4.33	32.89 ± 4.55
Transf.	M2	20.60 ± 2.31	19.61 ± 1.78
Transf.	M3	25.43 ± 3.44	20.30 ± 2.59
C. Generalization to new target directions (8)			
GRU	M1	97.99 ± 9.05	88.68 ± 6.23
GRU	M2	52.83 ± 2.56	49.32 ± 3.49
GRU	M3	54.53 ± 3.64	49.53 ± 2.19
Transf.	M1	119.49 ± 14.44	121.12 ± 14.56
Transf.	M2	19.77 ± 2.86	19.88 ± 1.65
Transf.	M3	22.16 ± 2.94	20.52 ± 2.02
Sum of wins		0	0

Population alignment objective. CCA rewards correlated linear projections across corresponding condition–time samples without neuron-to-unit assignments (Hotelling, 1936; Andrew et al., 2013). We flatten training-condition and time axes, center each feature, and divide each activity matrix by its overall root-mean-square (RMS) magnitude. For normalized student states H and teacher states Z, define $C_{HH} = H^\top H/(n-1) + \lambda I$, $C_{ZZ} = Z^\top Z/(n-1) + \lambda I$, and $C_{HZ} = H^\top Z/(n-1)$, with $\lambda = 0.01$. The score is

$$\overline{\rho}_\lambda(H, Z) = \frac{1}{12} \sum_{j=1}^{12} \mathrm{clip}_{[0,1]} \left[\sigma_j \left(C_{HH}^{-1/2} C_{HZ} C_{ZZ}^{-1/2} \right) \right]. \tag{2}$$

Cholesky whitening and differentiable singular values implement the score; alignment gradients update only the student. Orthogonal coordinate changes preserve the score, whereas isotropic ridge regularization prevents invariance to arbitrary invertible transforms. Twelve singular values are averaged despite teacher rank four, giving an approximate ceiling of $4/12$. We retain this fixed normalization across teacher constructions rather than renormalizing by effective rank. For an exactly rank-four teacher, with its representation and regularization fixed, averaging all 12 rather than the leading four singular values divides the alignment term by three, a constant rescaling that can be absorbed into α. Alignment rewards shared linear structure, not identical full-state distances or dynamical laws; functional benefit is evaluated through control performance (Kornblith et al., 2019).

Optimization and selection. For each seed $s = 0, \ldots, 99$, we randomly partition the 36 conditions into 28 training and eight evaluation conditions, using the same seed and split across teachers. All training conditions are optimized together for 700 Adam updates at learning rate 2×10^{-3}. Gradient

norms are clipped to one; updates with nonfinite norms are skipped. For fingertip position p_{ct}, target g_c, and radius $R = 0.10$ m,

$$\mathcal{L}_i = \frac{1}{2 \cdot 15\,|\mathcal{C}_{\mathrm{tr}}|R^2} \sum_{c \in \mathcal{C}_{\mathrm{tr}}} \sum_{t=26}^{40} \|p_{ct} - g_c\|_2^2 - \alpha \min(1, i/60)\overline{\rho}_\lambda(H, Z). \tag{3}$$

Neuro-Dist, Ring, Behavior, Gaussian, independently Shuffled, and Poisson teachers each have 40 positive strengths 0.1–4.0 in steps of 0.1, with 100 seeds per cell in both architectures and all three animals. Task-only uses $\alpha = 0$. For Table 1, each teacher's strength minimizes median training coordinate MSE, with ties favoring the smaller strength. Models and strengths stay fixed across trained, student-held-out, and new targets. Gaussian and Poisson are additional controls outside the five-method main table; Supplementary Table 2 uses the same strength-selection rule. The selected Shuffled models are shared by the noise and tangling analyses (Appendices C and E). All 100 seeds are retained in every cell, including high-error fits; no cells or runs are excluded.

Using 100 seeds per experimental setting allows us to characterize variability across model initializations and training/held-out condition splits. In the retrospective accuracy analysis, this also enabled separate groups of 50 seeds for alignment-strength selection and evaluation, with the evaluation seeds providing paired replication units for bootstrap confidence intervals and Holm-corrected comparisons (Appendix F).

Evaluation and metrics. Clean-error tables and dose-response figures report final-window coordinate MSE,

$$E_{\mathrm{MSE}}(\mathcal{C}) = \frac{10^6}{2 \cdot 15\,|\mathcal{C}|} \sum_{c \in \mathcal{C}} \sum_{t=26}^{40} \|p_{ct} - g_c\|_2^2 \quad (\mathrm{mm}^2). \tag{4}$$

Here, $\mathcal{C}$ is the set of evaluated reaching conditions and $|\mathcal{C}|$ is its size. This matches the reaching-loss reductions across the final 15 frames, conditions, and coordinates, in physical units and excluding alignment. Each model is scored separately on its 28 trained targets, eight student-held-out targets, or eight new directions; summaries are medians over all 100 seeds. New directions are constructed from target geometry alone, shared across methods, and at least 1 degree from every original direction up to numerical rounding (Appendix B.3). Sensory feedback noise evaluation uses coordinate RMSE, $\sqrt{E_{\mathrm{MSE}}}$, for each run and condition subset.

Tables 1 and 2 report the sample SD of $B = 100{,}000$ bootstrap medians (denominator $B - 1$). Each resample draws 100 whole-seed records with replacement, sharing the sampled seed indices across methods and target sets within each animal–controller group. Strengths remain fixed at their training-selected values. These SEs quantify uncertainty in the median across seeds, conditional on the recorded datasets and selected strengths; they exclude dose-selection uncertainty and are not paired tests of between-method differences.

The illustrative reaching examples report settled-endpoint error,

$$E(\mathcal{C}) = 1000 \sqrt{\frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \left\| \frac{1}{15} \sum_{t=26}^{40} p_{ct} - g_c \right\|_2^2} \quad (\mathrm{mm}). \tag{5}$$

Averaging positions before measuring error omits terminal temporal variability; this endpoint metric also lacks the coordinate factor used by MSE and RMSE. Seeds vary initialization and condition split, not animal or recording session. Interquartile-range bands describe run dispersion, not uncertainty; differences of marginal medians are not medians of paired differences. Figure 2 scores trained targets and Figure 6 student-held-out targets. Their stars mark table-row winners at training-selected doses, not significance or held-out curve minima. Logarithmic axes and clipping affect display only; outcomes from all 100 seeds enter summaries.

Trajectory tangling. Following Russo et al. (2018), nearby states with dissimilar velocities have high tangling. For condition–time states x_i, finite-difference velocities $\dot{x}_i$, and total variance $s_x^2 = N^{-1} \sum_i \|x_i - \bar{x}\|_2^2$,

$$Q_i = \max_j \frac{\|\dot{x}_i - \dot{x}_j\|_2^2}{\|x_i - x_j\|_2^2 + 0.1 s_x^2}, \qquad Q_S = \frac{1}{|S|} \sum_{i \in S} Q_i. \tag{6}$$

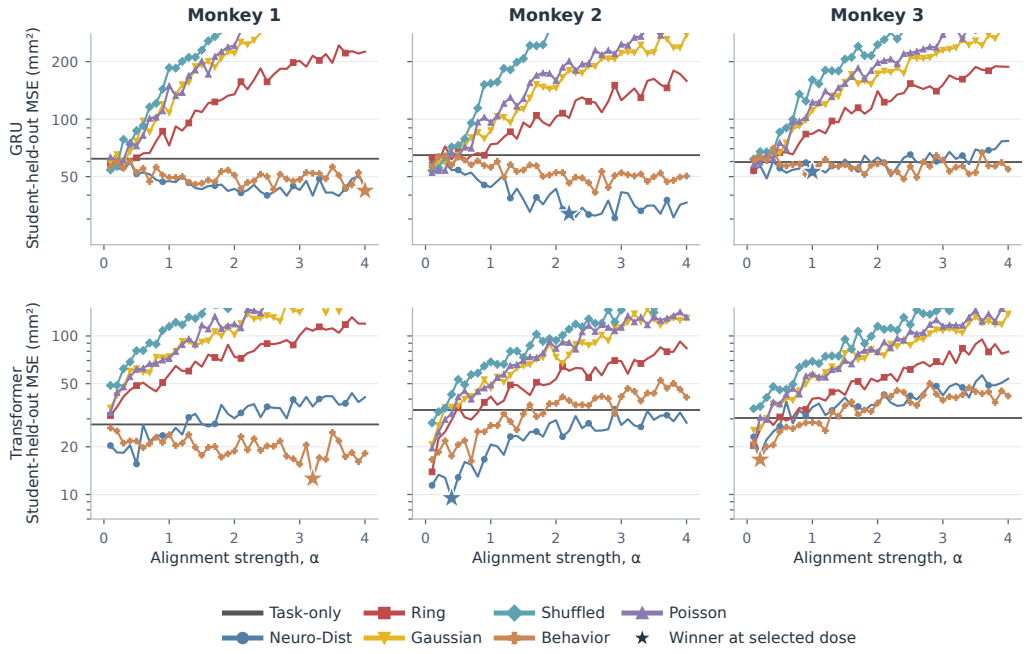

Figure 6: **Neuro-Dist gains extend to student-held-out targets.** Median coordinate MSE on eight targets excluded from student training but included in teacher construction; 100 seeds per dose, without uncertainty bands. Stars compare the five methods in Table 1 at training-selected strengths. Gaussian and Poisson are additional full-grid controls. The zoom clips off-scale values.

The native 128-dimensional GRU states are used without PCA. This measures the full recurrent-state geometry without choosing a variance-retention threshold: PCA truncation can alter state and velocity distances and discard low-variance directions relevant to tangling. Training and held-out query subsets S share all 36 conditions as references and for variance normalization. Restricting partners to the same or other reaches defines within- and between-reach tangling; total pointwise Q is their maximum. Appendix E specifies derivatives, uncertainty, cortical references, new-target reference pools, and velocity/distance and speed/angular decompositions.

Sensory feedback noise evaluation. The main evaluation uses 2,100 frozen GRUs: seven methods, three animals, and 100 seeds, selected using clean training MSE. Independent Gaussian noise perturbs all 16 observations, including target information, at every 10-ms step. Four positive amplitudes $\sigma \in \{0.025, 0.05, 0.075, 0.1\}$ each use three matched realizations. Coordinate RMSE scores only the eight student-held-out targets. Appendix C gives the dose selection, sampling, paired ratio endpoint, multiple-comparison correction, and uncertainty protocol.

B Supplementary Experimental Protocols

B.1 Student-Held-Out Dose Responses

Figure 6 shows the full dose sweeps on the eight original conditions withheld from student optimization, with stars at the training-selected strengths used for Table 1. These conditions participated in teacher construction. Every aligned teacher has 40 positive strengths and all 100 seeds are retained at each strength, with no exclusions. Stars compare the five table methods at their selected doses, rather than minima chosen on held-out error. Gaussian and Poisson are additional plotted controls. The logarithmic display clips off-scale values; full numerical values are retained in the source artifacts.

B.2 NEW-TARGET DOSE RESPONSES

Figure 7 evaluates all 144,600 saved GRU and Transformer controllers on the eight new directions defined in Appendix B.3. Each of the six aligned teachers has 40 strengths and 100 seeds per strength, alongside the shared task-only baseline. These descriptive curves retain the original training-selected operating points; no new-target outcome selects a strength or changes the statistical comparisons.



Figure 7: **Dose responses on new target directions.** Median coordinate MSE on eight directions absent from both student training and teacher construction; 100 seeds per dose, without uncertainty bands. Stars compare the five methods in Table 1 at their unchanged training-selected strengths. Gaussian and Poisson are additional full-grid controls.

B.3 GENERALIZATION TO NEW TARGET DIRECTIONS

We construct eight additional directions for each animal and seed, yielding 300 target sets shared across methods and architectures. Construction uses only original target directions and condition splits, not neural activity or controller outcomes. With the eight original held-out directions fixed, we maximize minimum circular separation over old–new and new–new pairs by enumerating the 6,435 allocations of eight new targets among eight cyclic gaps. Fixed old–old pairs do not enter the objective. New directions then receive the smallest displacement needed to remain at least 1 degree from every original direction, with deterministic tie-breaking and a 10^{-4}-degree numerical tolerance. The arm retains the same start posture and 10-cm target radius.

The 3,000 controllers and their strengths in Table 1 are frozen before evaluation, producing 24,000 new reaches. Each new rollout resets controller and environment state and includes the eight new targets alongside the original 28 training targets; only the new targets contribute to the reported error. No teacher refitting, retraining, dose retuning, or model exclusions occur.

These are directions outside the discrete teacher-construction and student-training sets, sometimes only 1 degree from an original direction. This evaluates within-task target transfer, not a teacher's predictions or transfer across sectors, sessions, tasks, or postures. Figure 7 extends the same evaluation protocol to the full dose grid; the selected-strength comparisons remain unchanged.

B.4 Overview figure and example selection

Figure 1B uses Monkey 2 session `indy_20160627_01`, a contributing session rather than the 37-session pooled training teacher. The cortical trace display uses unsmoothed, condition-averaged 10-ms spike counts from 219 nonempty sorted primary motor cortex units, without additional isolation-quality filtering. Ten displayed units are individually normalized by their maxima. The population paths instead retain the teacher's 30-ms Gaussian smoothing and this session's frozen projection. Twelve middle-magnitude direction conditions are displayed; three PCs explain 90.5% of rank-four subspace variance across all 36 conditions, not all firing-rate variance.

The paired population example is a Neuro-Dist GRU (seed 71, $\alpha = 2.2$). Ten traces use one common $2\times$ display gain and vertical offsets; all 128 units enter the independent PCA, whose first three PCs explain 63.6% of hidden-state variance. Cortical bins span -200 to 200 ms relative to movement onset; GRU post-update states span 10–400 ms and are displayed relative to the rollout midpoint. Matching display times or colors therefore do not imply physical-time correspondence. The connector illustrates the CCA training objective; the independent display projections are not a fitted CCA alignment.

Panel C shows examples selected to represent the observed reaching behavior: Monkey 2 GRU seed 53 ($\alpha = 3.6$) and Transformer seed 61 ($\alpha = 0.4$). Strengths were selected by minimizing population median training settled-endpoint error. All eight student-held-out reaches are shown at a common spatial scale; the two seeds have different splits. Filled dots denote final-window mean positions, and printed errors use Equation 5, rather than framewise MSE.

C GRU sensory feedback noise on student-held-out targets

Cohort and strength selection. Figure 3 and Supplementary Figure 8 evaluate 2,100 frozen GRUs: seven methods, three animals, and 100 seeds. For each teacher and animal, one strength minimizes median clean final-15-frame coordinate MSE on the 28 training conditions, with ties favoring the smaller strength. No noisy, held-out, or per-seed outcome selects the dose. All candidate cells contain 100 seeds, and all 2,100 selected models have finite clean and noisy errors; no cells or models are excluded. Neuro-Dist, Ring, Gaussian, Behavior, Poisson, and Shuffled have 40 strengths (0.1–4.0, step 0.1).

For Monkeys 1/2/3, selected strengths are $1.5/2.2/1.0$ for Neuro-Dist, $0.1/0.6/0.1$ for Ring, $0.2/0.2/0.1$ for Shuffled, $0.3/0.3/0.2$ for Gaussian, $0.2/0.1/0.2$ for Poisson, and $4.0/2.6/3.1$ for Behavior; task-only uses zero. The original seed permutation defines the 28/8 split, matched across teachers. All 36 conditions enter teacher construction. No models were retrained or doses reselected for this decomposition.

Perturbation and scoring. Independent Gaussian noise perturbs all 16 observation channels at every 10-ms feedback step: two target coordinates, two visual fingertip coordinates, and 12 muscle-length and muscle-velocity channels, normalized by optimal fiber length and maximum shortening velocity, respectively. The four positive amplitudes are $\sigma \in \{0.025, 0.05, 0.075, 0.1\}$, with three realizations per model seed, matched across methods and amplitudes. Positional coordinates are in meters, so their SDs are 2.5–10 cm for a 10-cm reach; muscle channels have different normalized units. Noise is not demeaned or rescaled by realized RMS.

The rollout includes all 36 original targets and lasts 40 simulation steps. Only the eight student-held-out targets contribute to the final-15-frame coordinate RMSE, averaged across realizations. Clean error is evaluated separately. Figure 3A is a separate fixed-dose, 1.5-s illustration; quantitative comparisons use these 400-ms selected-strength rollouts.

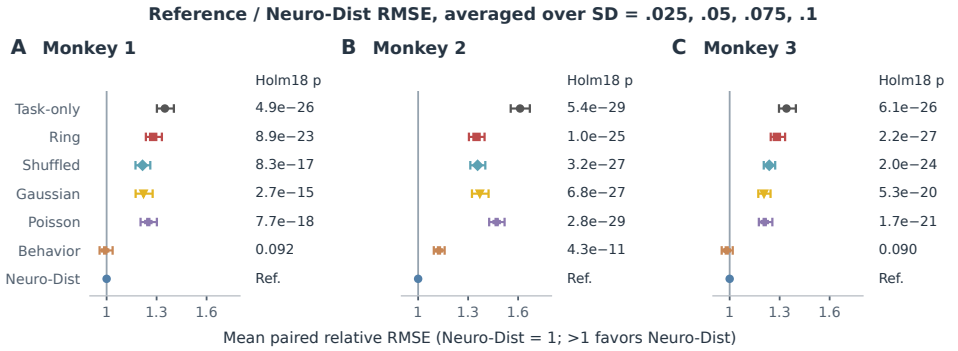

Figure 8: **Primary comparisons of noisy student-held-out accuracy.** (A–C) Monkeys 1–3 for the same 2,100 selected-strength GRUs as in Figure 3. Strengths minimize clean error on 28 student-training targets; evaluation scores the eight student-held-out targets, included in all-36-condition teacher construction. Three noise realizations are averaged per amplitude. Reference errors are divided by the Neuro-Dist error for the same seed separately at each positive amplitude ($\sigma = 0.025, 0.05, 0.075, 0.1$), then averaged within seed. Markers show the mean ratio across 100 seeds; values above one favor Neuro-Dist. Whiskers are pointwise 95% BCa intervals from 50,000 seed resamples. Neuro-Dist's self-reference is exactly one and has no independent interval or test. Primary paired signed-rank p-values use one Holm correction across all six references and three animals (18 tests).

Paired inference. For each reference j, animal a, and model seed s, we compute

$$q_{asj} = \frac{1}{4} \sum_{\sigma \in \{0.025, 0.05, 0.075, 0.1\}} \frac{E_{asj}(\sigma)}{E_{as,\text{Neuro-Dist}}(\sigma)},$$

where each E first averages realization-level RMSE. Thus, normalization precedes averaging and positive levels receive equal weight. Values above one favor Neuro-Dist; this score measures relative noisy accuracy rather than degradation from clean performance.

Eighteen two-sided signed-rank tests (six references by three animals) share one Holm correction. Pointwise 95% bias-corrected and accelerated (BCa) intervals for mean ratios use 50,000 paired-seed resamples, preserving all amplitudes within each seed. The statistical unit is the model seed ($n = 100$), not a target, frame, amplitude, or realization. Signed-rank inference assumes symmetry and does not directly test mean ratios.

Results and scope. The signed-rank tests favor Neuro-Dist over task-only, Ring, Shuffled, Gaussian, and Poisson in every animal; all 15 comparisons have $p \leq 2.68 \times 10^{-15}$ after Holm18 correction including Behavior. Mean ratios of task-only error to Neuro-Dist error are 1.350/1.612/1.341 in Monkeys 1/2/3. Neuro-Dist has lower mean error than task-only at every positive amplitude; averaging amplitudes gives 30.65/26.45/31.49 mm versus 40.74/42.28/41.88 mm.

Behavior has mean ratios 0.991/1.125/0.983. The signed-rank tests favor Neuro-Dist in Monkey 2 (corrected $p = 4.33 \times 10^{-11}$) and detect no statistically significant difference in Monkeys 1 and 3 ($p = 0.0924$ and 0.0897, respectively). Behavior mean-ratio intervals include one in Monkeys 1 and 3: $[0.9578, 1.0360]$ and $[0.9522, 1.0185]$.

Neuro-Dist also has smaller mean absolute increases from clean error than task-only in all 12 animal–amplitude combinations. Mean increases across amplitudes are 23.92/20.72/23.94 mm versus 32.26/33.16/33.58 mm. These increases are descriptive; the paired tests concern noisy accuracy, not uniformly smaller relative degradation. Comparison with Behavior remains dataset dependent. Inference conditions on frozen strengths: the same seeds enter selection and estimation, so selection uncertainty is not propagated.

D DIRECTIONAL ORGANIZATION AND THE ENGINEERED RING CONTROL

D.1 DIRECTIONAL ORGANIZATION MOTIVATES A SIMPLE CONTROL

Figure 9 illustrates reach-direction organization in the Monkey 3 primary motor cortex teaching representation and Transformer states. Neighboring regions in the displayed state space preferentially represent different reach directions. This organization is visible in the cortical reference, task-only controllers, and controllers trained with Neuro-Dist or Behavior supervision. It therefore motivates a control that supplies a simple, explicitly constructed directional geometry. The Ring teacher asks whether this accessible feature of task-related population organization is sufficient to account for benefits of cortical supervision; its functional consequences are evaluated by the accuracy and robustness experiments, rather than inferred from the appearance of these maps.

The example also shows directionally organized Ring-trained states. Thus, directional organization is shared across several training regimes and is not a signature unique to Neuro-Dist. The displayed cortical reference is the behavior-selected rank-four primary motor cortex representation used for teaching, not the complete recorded population.

D.2 CONSTRUCTION OF THE DIRECTIONAL DENSITY MAPS

For each animal we fit an unwhitened four-dimensional PCA basis to the fixed cortical teacher. Within each model seed, student PCA and a four-dimensional orthogonal Procrustes alignment are fitted using only the 28 training conditions. A single scalar matches the centered four-dimensional training-state RMS magnitude to the cortical reference; translation matches the cortical training mean. The fitted transformation is then applied to all 36 conditions, without per-axis scaling or nonlinear warping. Figure 9 displays the first two cortical coordinates, pooling all 40 frames of every condition. Each student panel includes all 100 seeds with equal weight.

For occupancy D and complex directional mass R, each condition–time location contributes equal probability weight $1/(36 \times 40)$ to D and this weight times $\exp(\mathrm{i}\theta_c)$ to R. We average these arrays across seeds and smooth both arrays on a 128×128 grid with a Gaussian kernel of standard deviation 2 pixels. Hue is the circular mean direction $\arg R$, saturation is the local concentration $|R|/D$, and color fades toward white according to relative occupancy $(D/\max D)^{0.55}$. The maximum is taken separately within each panel, so brightness encodes occupancy relative to that panel's peak.

All displayed maps use one common square viewport derived from the pooled 99.5th-percentile maximum absolute coordinate, $\max(|\mathrm{PC1}|, |\mathrm{PC2}|)$, across both architectures, all teachers and strengths, enlarged by 10% and expanded to include every cortical reference point. In this figure, at most 0.496% of pooled state samples lie outside any one panel; the largest individual-seed fraction outside is 10.14%. No seed is excluded and cropped probability is not renormalized. The cortical teacher was constructed from all 36 conditions, including those withheld from student training.

D.3 FROM AN IDEALIZED CIRCLE TO THE ACTUAL ENGINEERED TEACHER

Figure 10 separates the schematic directional construction from an actual teacher generated by the training implementation. For condition angle θ_c and frame t, the directional pair is

$$z_{ct,1} = a_1 g_t \cos\theta_c, \qquad z_{ct,2} = a_2 g_t \sin\theta_c, \tag{7}$$

where g is a linear envelope from 0.3 to 1.0, smoothed with a Gaussian kernel of standard deviation 3 frames. Each a_d scales the variance of its directional coordinate using the corresponding leading cortical covariance eigenvalue. The angle is constant within each reach while the amplitude increases; trajectories extend radially rather than rotate around the ring. Because the two scale factors differ, the direction-indexed curve is generally an ellipse in construction coordinates.

Ten additional coordinates are obtained from temporally smoothed Gaussian signals, centered, rotated to their principal axes, and scaled using the remaining cortical eigenvalues. These components are generated independently for each model seed and concatenated with the directional pair to give the 12-coordinate Ring teacher. Cross-covariances with the directional coordinates mean that this construction does not exactly match the complete cortical covariance spectrum. In the displayed Monkey 3 seed-0 teacher, the first pair accounts for approximately 79.0% of total coordinate variance

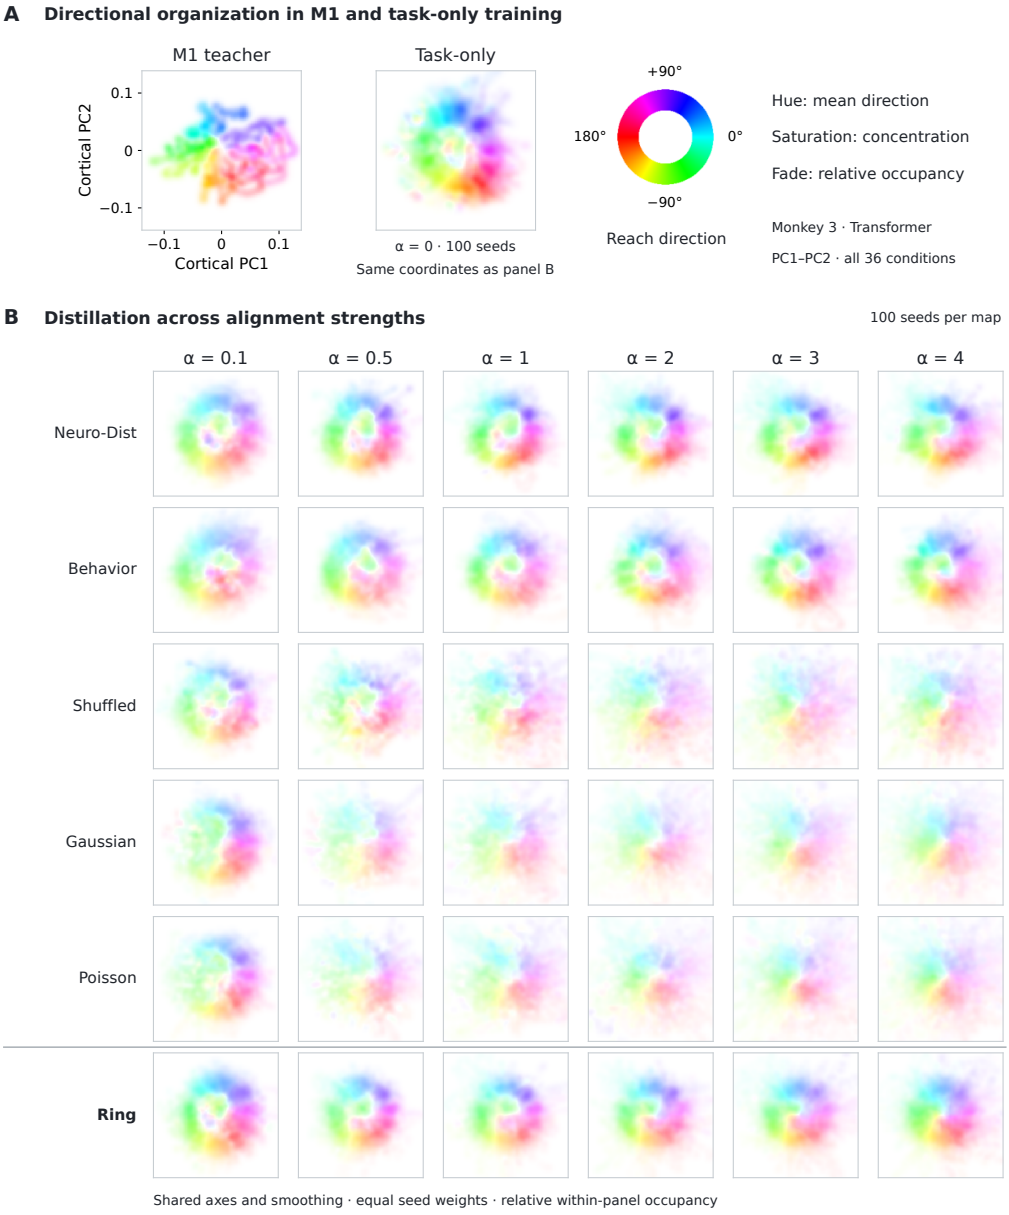

Figure 9: **Directional organization motivating the engineered Ring control.** In this figure only, M1 denotes primary motor cortex. **A:** Monkey 3 M1 teaching representation and task-only Transformer states in common cortical PC1–PC2 coordinates. **B:** Transformer states across six illustrative alignment strengths. Every student map pools 100 seeds and all 36 conditions; alignment is fitted on the 28 training conditions per seed. Hue indicates reach direction, saturation directional concentration, and fading relative occupancy within each panel. The same cortical reference, axes, and smoothing apply throughout. Viewport cropping and projection details are specified in Appendix D. Directional organization is visible under multiple training regimes and is not specific to cortical supervision.

and coordinates 3–4 for 21.0%; coordinates 5–12 contain only numerical residual variance. The plotted axes are construction coordinates, not PCs of the complete concatenated teacher.

The displayed teacher uses the first seed, without performance-based selection, and is generated from the cortical reference using the same construction as in training. The schematic circle explains the design principle; the numerical panels show the actual scaled directional paths and stochastic

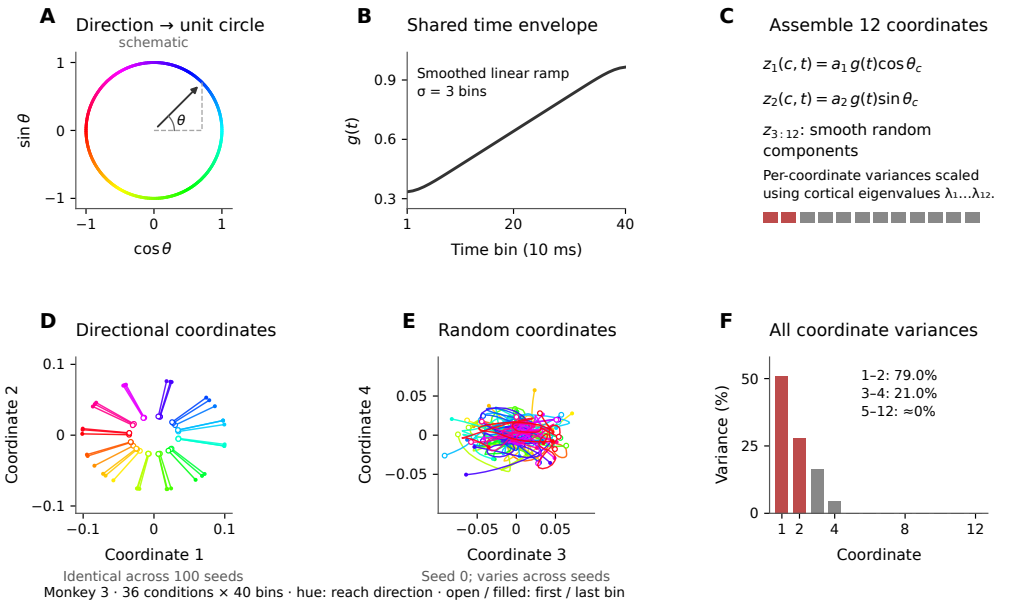

Figure 10: **The engineered Ring teacher: schematic and numerical construction. A:** Idealized direction-to-circle mapping. **B:** Actual shared temporal envelope. **C:** Twelve-coordinate construction. **D:** Actual directional coordinates for all 36 Monkey 3 conditions, identical across seeds. **E:** Smooth random coordinates 3–4 for seed 0. **F:** Variance of all 12 coordinates. Hue denotes reach direction as in Figure 9; open and filled markers identify the first and last frames. Panels D–F show the numerical teacher from the training implementation. Its first pair is anisotropically scaled and evolves radially, while the remaining components provide smooth random structure. These are construction coordinates, not PCs of the full teacher.

remainder. This construction provides a transparent test of a simple directional prior, without assuming that cortical population activity follows the engineered trajectories or their temporal envelope.

E TRAJECTORY-TANGLING ANALYSIS

We analyzed all 2,100 GRUs spanning seven methods, three animals, and 100 seeds per cell (Figure 5). Aligned strengths minimize median training MSE across the 40-dose grid. Neuro-Dist reduces median training tangling by 48–68% relative to task-only. The Neuro-Dist–Behavior comparison favors Neuro-Dist in Monkeys 1–2 and Behavior in Monkey 3 across trained, student-held-out, and new targets.

We compute Equation 6 in all 128 hidden coordinates, using centered interior and one-sided endpoint differences with $\Delta t = 1$ frame. All 36×40 original states define the reference pool and variance normalization. Within-model means summarize total, within-reach, and between-reach tangling separately for trained and held-out queries; pointwise total Q is the maximum of the latter two components. We separate velocity disagreement from state separation at shared condition–time pairs selected using task-only models, and resolve velocity disagreement into speed and angular contributions at each model's maximizing pairs (Figure 11A–C).

For new-target analysis, the same controllers reach eight additional directions shared across methods for each animal and seed. We recompute Q using a common pool of 44 original and new conditions and its variance, then summarize trained, held-out, and new queries separately (Figure 11D). A sensitivity analysis compares held-out and new queries against the 28 training conditions with fixed original-condition variance.

We report medians across seeds and medians of paired seed-wise differences. Pointwise 95% percentile intervals use 10,000 bootstrap resamples shared across methods, preserving model–target

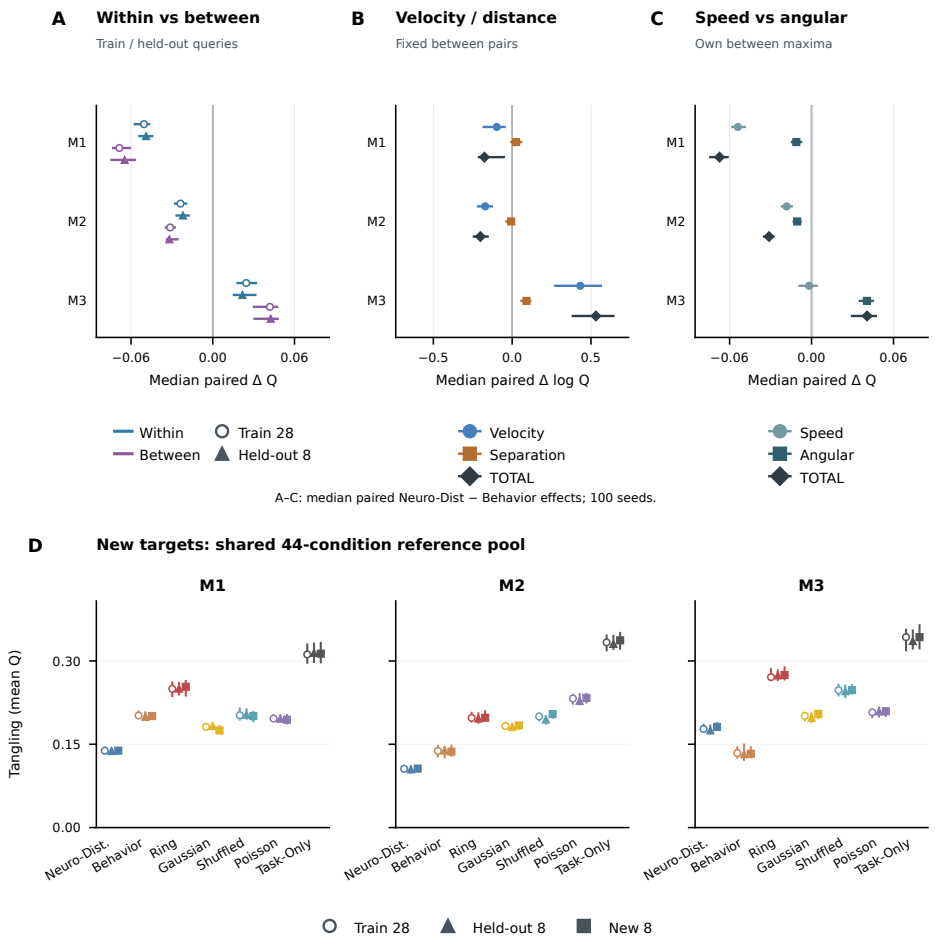

Figure 11: **Tangling differences persist at new target directions.** (A–C) Median seed-paired Neuro-Dist-minus-Behavior differences: within- and between-reach Q (A), velocity and separation contributions at shared task-only-selected pairs (B), and speed and angular contributions at each model's maximizing pairs (C). Negative separation effects indicate larger normalized denominators; component medians need not sum to the total. A uses trained and held-out queries; B–C use all 36. (D) Medians across seeds of model-level mean total Q for all seven methods at trained, held-out, and new targets, using a common 44-condition reference pool and variance instead of the original 36. Whiskers are pointwise 95% bootstrap intervals for the medians across 100 seeds, conditional on selected doses and fixed datasets.

pairing and conditioning on selected strengths and fixed datasets. Across the 21 animal–method summaries, we assess the descriptive tangling–error association with Spearman correlation and show an ordinary least-squares fit.

Cortical Q provides a descriptive reference from the parent 12-PC representation, with session pooling for Monkeys 2–3 and distinct temporal preprocessing. Figure 4 illustrates a representative Monkey 2 session and matched GRU pair (seed 45) using independent PCA projections of 12 representative reaches; its printed Q uses all 36 conditions in the native representations.

F STATISTICAL COMPARISON OF REACHING ACCURACY

F.1 METHODS

Question and analysis scope. We compare Neuro-Dist with each of Task-only, Ring, Behavior (delayed position–velocity), Shuffled, Gaussian, and Poisson in GRUs and Transformers across the three animal datasets. The outcome is each controller's final-window coordinate MSE (Equation 4): squared position residuals averaged over the final 15 post-action frames, both coordinates, and evaluated targets, in mm^2. All evaluated MSEs are finite, and all runs are retained, including high-error fits. Tables 1 and 2 report descriptive medians and bootstrap SEs across all 100 seeds at training-selected strengths. These SEs condition on the selected strengths and do not correct for reuse of the same seeds for selection and estimation (Cawley & Talbot, 2010); they supply no selection-adjusted significance claims. For statistical comparisons, we use a *seed-split validation*, separating strength selection from evaluation.

Independent strength-selection seeds. Before computing validation contrasts, we fixed a deterministic partition of the 100 seeds into 50 selection and 50 evaluation seeds, shared across methods, animals, and architectures. For each teacher, animal, and architecture, the selected α minimizes median MSE on the 28 trained targets among 40 strengths $0.1{:}0.1{:}4.0$, using only selection seeds and resolving ties toward smaller strengths. Task-only retains the original 28-target baseline at $\alpha = 0$. The resulting dose map differs from the table's map in 16 of 42 cells (Table 3). Each strength stays fixed across evaluation seeds, target regimes, and bootstrap draws; it is not optimized per seed or for held-out/new-target error.

Table 3: **Frozen strengths for seed-split validation.** Each α minimizes trained-target median MSE on the 50 selection seeds; results use the other 50. Parentheses give the original 100-seed table strength where it differs. Each selected strength applies unchanged to all three target regimes. M1–M3 denote Monkeys 1–3.

Cohort	Task-only	Ring	Neuro-Dist	Behavior	Shuffled	Gaussian	Poisson
GRU/M1	0.0	0.3 (0.1)	2.4 (1.5)	2.3 (4.0)	0.2	0.3	0.2
GRU/M2	0.0	0.6	2.2	1.8 (2.6)	0.2	0.4 (0.3)	0.1
GRU/M3	0.0	0.1	1.4 (1.0)	2.8 (3.1)	0.3 (0.1)	0.1 (0.2)	0.2
Transf./M1	0.0	0.2 (0.1)	0.3	3.7 (3.2)	0.1	0.1	0.1
Transf./M2	0.0	0.1	0.1 (0.4)	0.7 (0.1)	0.2 (0.1)	0.1	0.1
Transf./M3	0.0	0.2 (0.1)	0.2	0.2	0.1 (0.2)	0.1	0.1

Three evaluation regimes. The other 50 seeds supply results on the 28 trained targets, eight targets withheld from student training but included in teacher construction, and eight directions unseen by both teacher and student. The last regime uses the fixed animal/seed target bundles of Appendix B.3. The 2,100 evaluation-half controllers and their selected strengths remain fixed across all three regimes, with no retraining, teacher refitting, or exclusions. The new directions remain within the original task, radius, and start posture.

Factorial median model and effects. We estimate a saturated raw-scale median model with one parameter for every method j, architecture a, animal m, and target regime r: $7 \times 2 \times 3 \times 3 = 126$ cells. This is a fully interacted median-regression model (Koenker & Hallock, 2001), fitted directly by cell medians, without imposing homogeneous teacher effects. For even sample sizes, the median is the arithmetic mean of the two central observations. With $\widehat{\mu}_{jamr}$ denoting a raw MSE median, we define

$$d_{jamr} = \log \frac{\widehat{\mu}_{\mathrm{ND},amr}}{\widehat{\mu}_{jamr}}, \qquad D_{jr} = \frac{1}{6} \sum_{a,m} d_{jamr}, \qquad G_{jr} = 100\big(1 - e^{D_{jr}}\big). \tag{8}$$

Positive G favors Neuro-Dist. This is an equally weighted geometric-average median-MSE reduction across the six animal–architecture cells. We take logs *after* raw-scale medians. For percentages

referenced to Neuro-Dist, the equivalent contrast is $H_{jr} = 100(e^{-D_{jr}} - 1)$: positive values indicate increased control MSE, and zero indicates equal performance. The combined dose/comparison figure uses this scale, transforming the log-interval endpoints in reverse order; the tables retain G_{jr}.

Table 4: **Aggregate seed-split validation.** Geometric-average median-MSE reduction (%) and pointwise 95% BCa interval versus each comparator. All three target regimes form one 18-comparison family; † marks Holm-adjusted $p < 0.05$, not simultaneous interval coverage. Results use 50 evaluation seeds with independently selected, frozen strengths.

Comparator	Trained targets	Student-held-out	New targets
Task-only	42.2 [29.3, 53.0]†	30.7 [8.9, 42.9]†	27.5 [8.6, 38.4]†
Ring	36.8 [24.5, 48.0]†	29.5 [12.5, 40.0]†	26.8 [14.5, 37.8]†
Behavior	26.9 [14.7, 41.0]†	15.2 [−2.6, 29.7]	16.7 [5.0, 28.3]†
Shuffled	51.1 [42.9, 59.7]†	43.5 [29.5, 52.0]†	44.5 [36.4, 50.8]†
Gaussian	41.8 [31.1, 53.4]†	34.6 [20.0, 46.8]†	28.9 [14.0, 35.7]†
Poisson	38.9 [30.6, 50.3]†	29.1 [12.9, 40.3]†	23.2 [10.0, 32.1]†

Paired bootstrap and multiplicity. One whole-seed index vector is resampled with replacement and shared across every method, animal, architecture, and target regime. This preserves dependencies from initialization, condition splits, and seed-specific new-target sets; neither frames nor reaches are independent replications. The validation uses 4,000,000 resamples of 50 evaluation blocks; the selection half remains fixed. We compute pointwise 95% bias-corrected and accelerated (BCa) intervals (Efron, 1987), using tie-aware bias correction (the average strict/weak bootstrap rank), paired leave-one-seed-out acceleration, and linear quantile interpolation. Two-sided marginal p-values invert these confidence curves: the smallest α for which the $1 - \alpha$ interval excludes zero, preserving strict exclusion at ties. Values below 10^{-5} are conservatively reported as upper bounds. Holm's step-down correction (Holm, 1979) uses $K = 18$ for the six-comparator-by-three-regime family. For ordered marginal values $p_{(1)} \le \cdots \le p_{(K)}$, the adjusted value is $\widetilde{p}_{(k)} = \min\{1, \max_{\ell \le k}(K - \ell + 1)p_{(\ell)}\}$. At $\widetilde{p} < 0.05$, negative/positive estimated log ratios support lower Neuro-Dist/control error; no statistically significant difference is detected in the other comparisons, which does not establish equivalence. Holm requires no additional dependence assumption; familywise control remains approximate because BCa inference uses a finite sample of independent seed blocks.

Table 5: **Adjusted aggregate significance.** Two-sided BCa confidence-curve p-values with Holm correction across the same 18 aggregate comparisons as Table 4. Values preceded by $\le$ are conservative upper bounds resulting from the marginal numerical floor.

Comparator	Trained targets	Student-held-out	New targets
Task-only	$\le 1.8 \times 10^{-4}$	0.0284	0.0284
Ring	$\le 1.8 \times 10^{-4}$	0.0040	9×10^{-4}
Behavior	0.0046	0.0931	0.0284
Shuffled	$\le 1.8 \times 10^{-4}$	$\le 1.8 \times 10^{-4}$	$\le 1.8 \times 10^{-4}$
Gaussian	$\le 1.8 \times 10^{-4}$	$\le 1.8 \times 10^{-4}$	$\le 1.8 \times 10^{-4}$
Poisson	$\le 1.8 \times 10^{-4}$	0.0046	0.0046

Numerical checks and inferential scope. The four-million-draw budget was fixed before evaluating the new-dose new-target outcomes. Numerical checks assess median ties and gaps, bootstrap-tail resolution, interval endpoints, and Holm-corrected decisions in the full bank and its two equal halves, without fallback, filtering, or significance-dependent resampling. Inference concerns seed variability in three *fixed* animal datasets, conditional on the frozen selection half and dose map.

F.2 RESULTS

Aggregate comparisons. Figure 2 and Tables 4–5 summarize the seed-split validation. Holm-corrected comparisons support lower Neuro-Dist median MSE against all six comparators on trained targets (geometric-average reductions 26.9%–51.1%), five on student-held-out targets (29.1%–43.5%), and all six on new directions (16.7%–44.5%). Task-only reductions are 30.7% on student-held-out targets and 27.5% on new directions (both adjusted $p = 0.0284$). The new-direction reduction against Behavior is 16.7% (adjusted $p = 0.0284$); no statistically significant difference is detected against Behavior on student-held-out targets (15.2%, adjusted $p = 0.0931$). Thus, 17 of 18 aggregate comparisons favor Neuro-Dist, with none favoring a control.

Numerical resolution. All 18 aggregate Holm decisions agree across the full bootstrap bank and both halves. No cell has a central median tie.