REINS

Learning Inertia-Induced Geometry for Physics-Consistent Motion Representation in Clinical Gait Phenotyping

Yiran Ding[1], Yufei Zhang[2], Zijun Cui[1]

[1]Michigan State University
[2]Independent Researcher

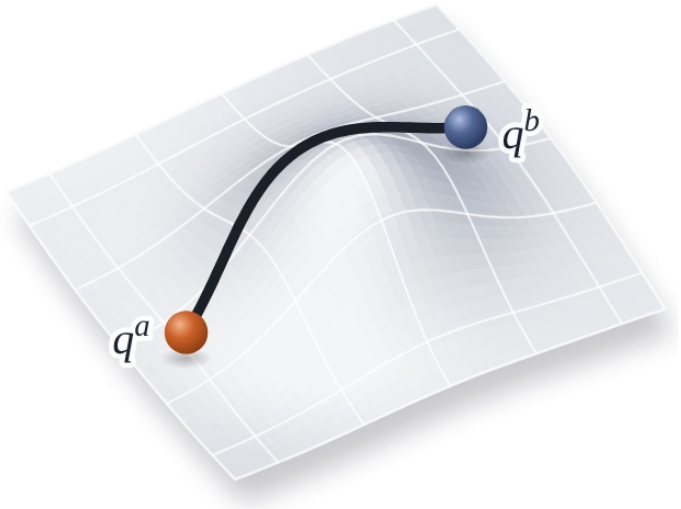

NEURAL INFORMATION PROCESSING SYSTEMS 2026

Clinical gait assessment in Parkinson's disease

More than **8.5 million** people were living with Parkinson's disease in 2019.

Its prevalence has **doubled** in the past 25 years.

How gait is rated today

MDS-UPDRS item 3.10, Gait

The patient walks at least 10 m away from the examiner and back

0	1	2	3	4
Normal	Slight	Mild	Moderate	Severe

Rated by observation: stride amplitude and speed, foot lift, heel strike, turning, arm swing

MDS-UPDRS gait scoring requires a trained examiner and is subject to inter-rater variability.

Sources: Parkinson disease fact sheet (WHO, 2023). MDS-UPDRS (International Parkinson and Movement Disorder Society, 2008). CARE-PD (Adeli et al., 2025).

Clinical gait datasets are small and site-specific

Opportunity

3D pose estimation recovers body motion from monocular video

No motion-capture facility needed

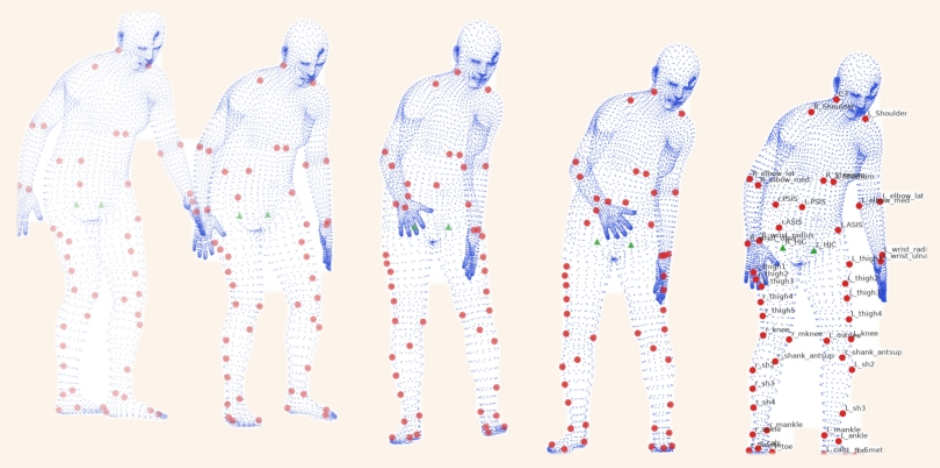

A 3D pose sequence, the input in this work

Obstacle

Clinical motion data are limited in scale, cohort diversity and label quality

- **Tens of subjects per cohort**
- **Each site has its own protocol**
- **Severity labels are sparse or missing**

Motion representations learned from small cohorts can overfit and transfer poorly across sites

NEURAL INFORMATION PROCESSING SYSTEMS 2026

How can biomechanics help us compare two walks?

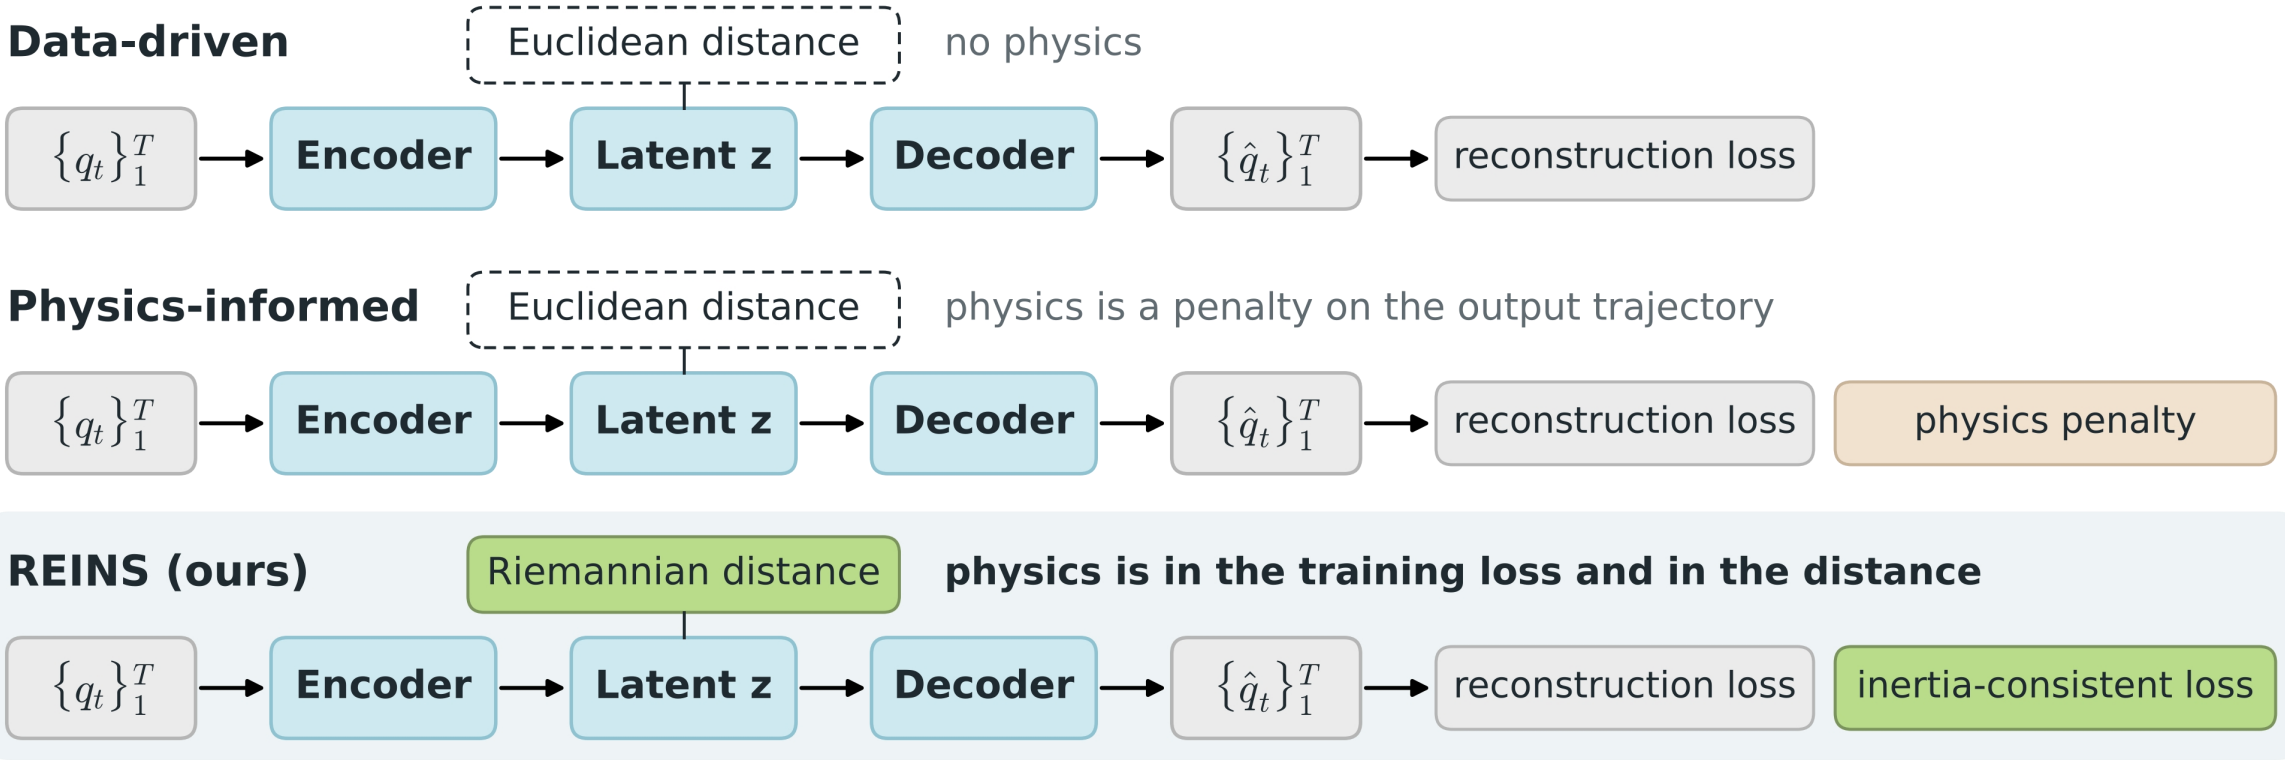

A gait distance based on body mechanics: REINS uses the body's inertia to fine-tune the motion representation and to measure how far each gait sequence is from a healthy reference.

The inertia matrix M(q) defines a metric on pose space

Euler-Lagrange equation of the body

$$M(q)\,\ddot{q} + C(q,\dot{q})\,\dot{q} + g(q) \;=\; \tau + J_C^{\top}(q)\,\lambda$$

Kinetic energy, with the same M(q)

$$\mathcal{T}(q,\dot{q}) \;=\; \tfrac{1}{2}\,\dot{q}^{\top} M(q)\,\dot{q}$$

At each pose q, M(q) defines how we measure the length of a small movement.

We compute M(q) directly from the body model and the current pose.

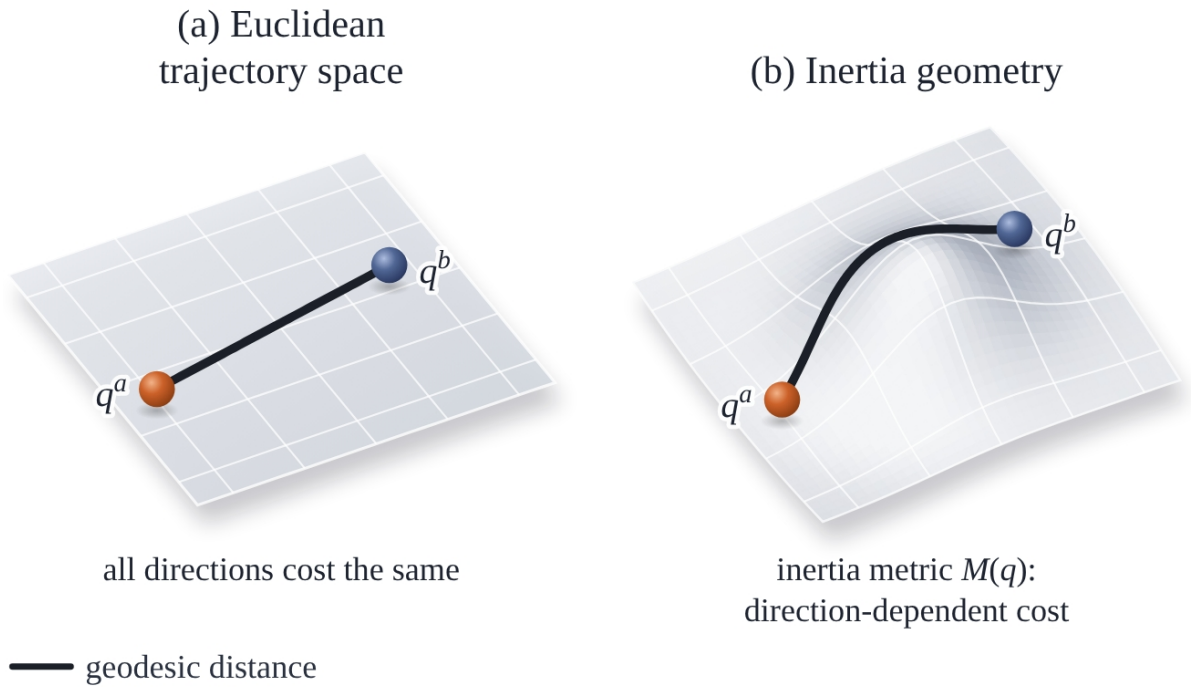

Prior work applies physical constraints to reconstructed motion

REINS uses the body's inertia to shape the motion representation and define distances within it.

Overview of REINS

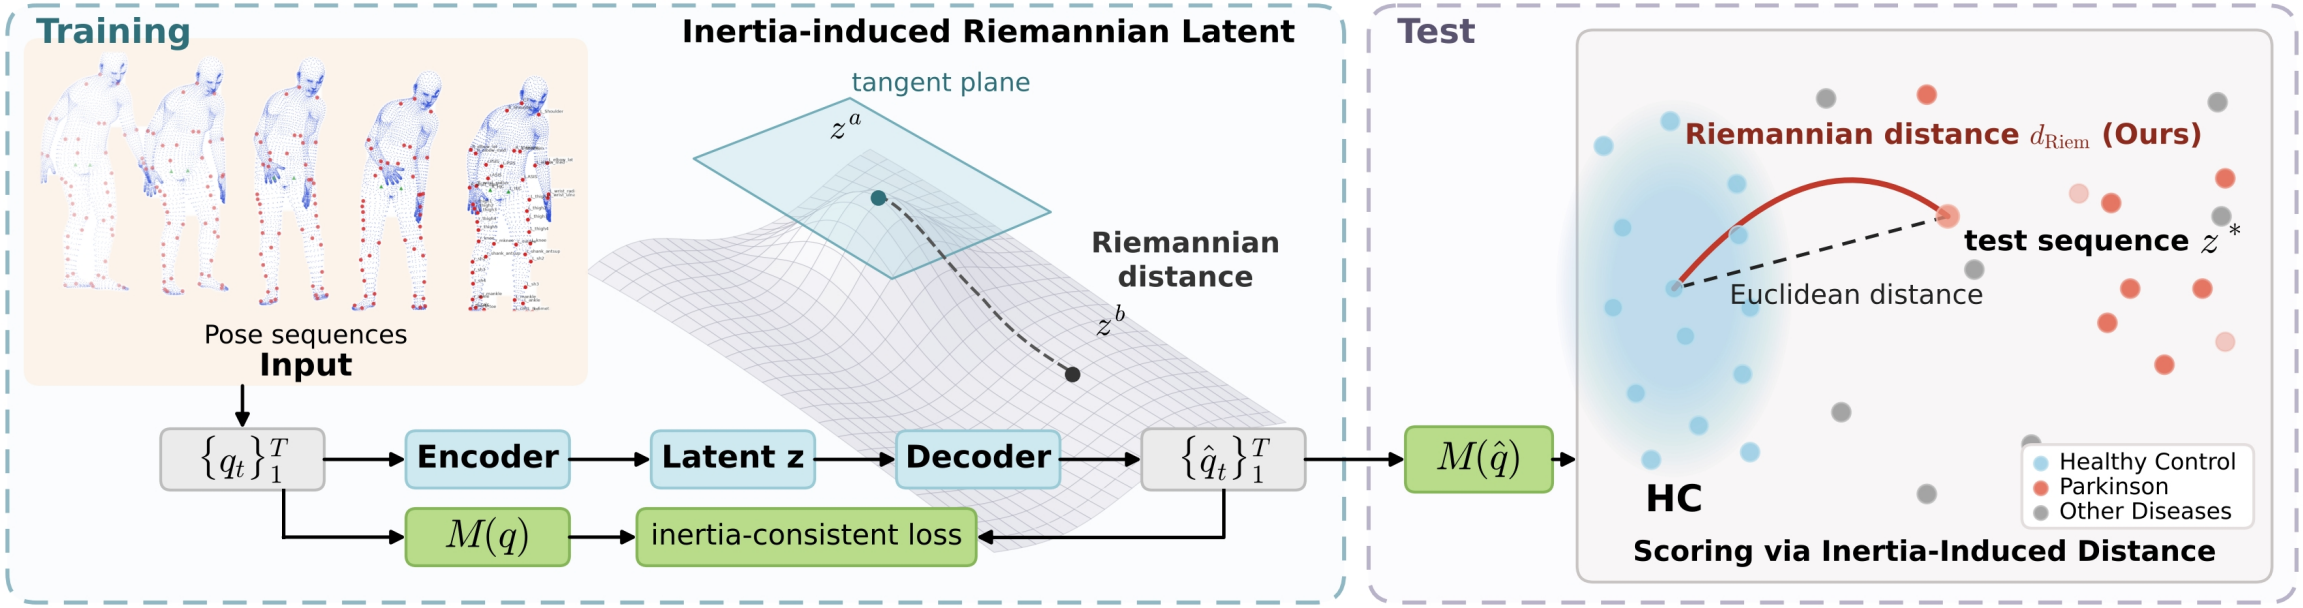

Training

The inertia-consistent loss weights errors in the reconstructed velocities by M(q)

Test

We compute the local metric using M(q̂) from the decoded pose.

The score is the Riemannian distance to the nearest healthy reference embedding

NEURAL INFORMATION PROCESSING SYSTEMS 2026

Fine-tune with an inertia-consistent loss

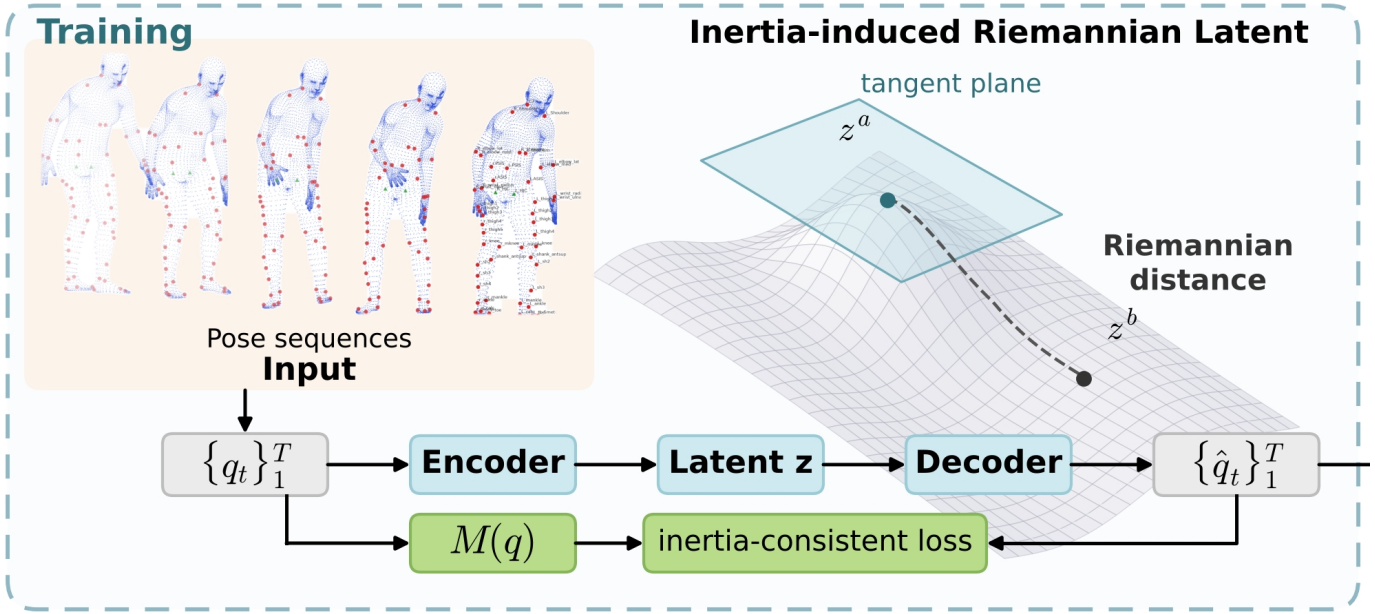

PhysPT encoder-decoder, pretrained on AMASS

Fine-tuned on healthy controls only

16-frame windows at 30 fps

Inertia-consistent loss

$$\mathcal{L}_{\text{ine}} = \sum_{t=1}^{T} (\dot{q}_t - \hat{\dot{q}}_t)^\top M(q_t) \, (\dot{q}_t - \hat{\dot{q}}_t)$$

The loss sums squared velocity errors under M(q). Each term is twice the kinetic energy of that error

Fine-tuning objective

$$\mathcal{L} = \mathcal{L}_{\text{recon}} + \mathcal{L}_{\text{force}} + \mathcal{L}_{\text{contact}} + \lambda_{\text{ine}} \mathcal{L}_{\text{ine}}$$

The PhysPT objective, with its Euler-Lagrange residual replaced by the inertia-consistent loss

How does REINS measure deviation from healthy gait?

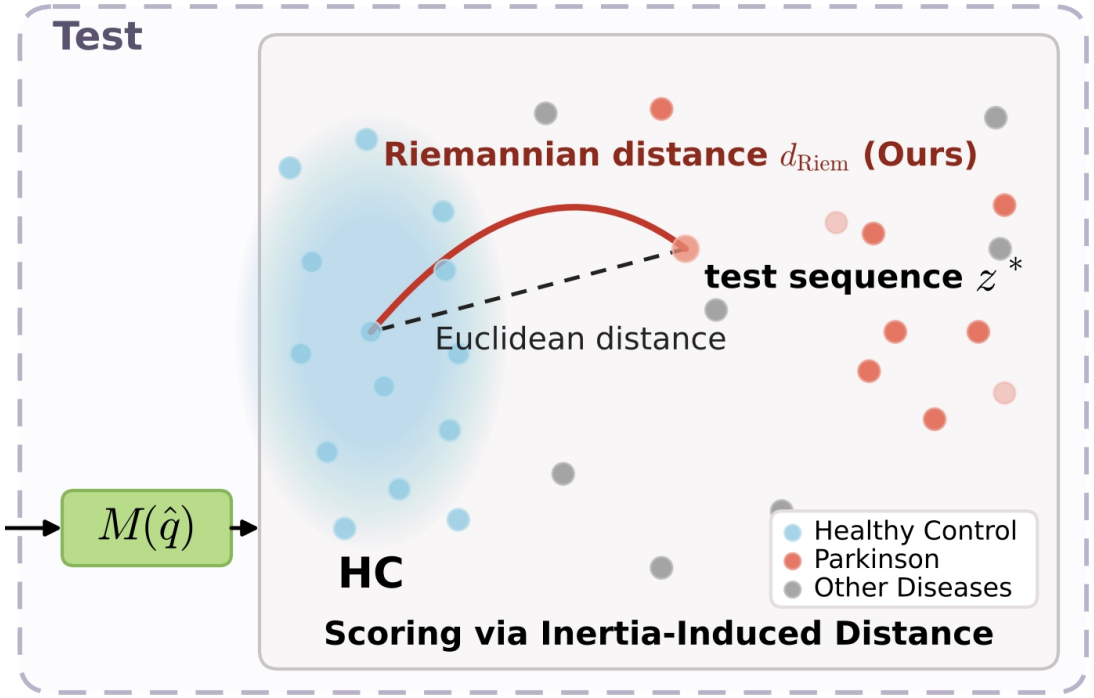

Score of a test sequence

$$s(z^*) = \min_{z \in \mathcal{Z}_{\mathcal{N}}} d_{\text{Riem}}(z^*, z)$$

Larger scores indicate greater inertia-weighted deviations from healthy motion

We compute the same score from motion for gait anomaly detection and severity assessment

Riemannian distance: the length of the shortest path between two embeddings, with each step measured by a metric g(z)

Euclidean (L2) distance is the straight-line distance between embeddings. Yet two steps of equal length in latent space can correspond to different changes in pose.

REINS therefore measures each step with (g(z)), which weights the decoded pose change by the body's inertia.

NEURAL INFORMATION PROCESSING SYSTEMS 2026

Latent metric and shortest-path approximation

1 The metric at one embedding

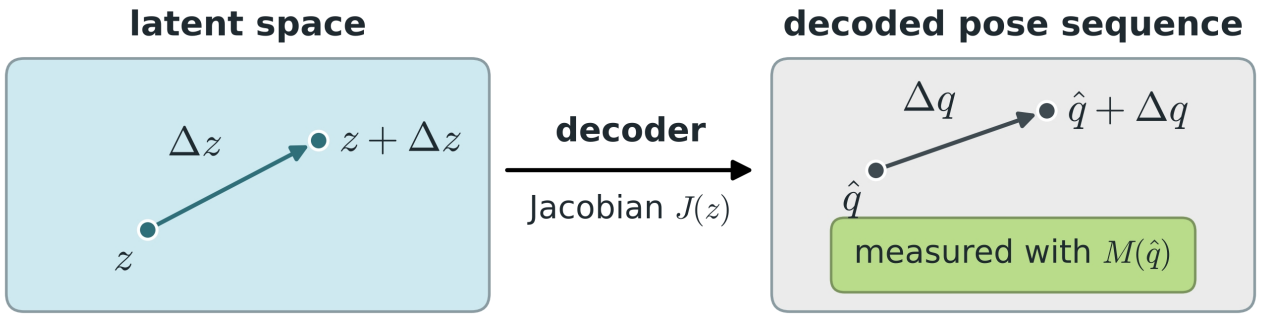

A latent step moves the decoded poses $\Delta q = J(z)\,\Delta z$

Its squared length under the inertia matrix $\Delta q^\top M(\hat{q})\,\Delta q = \Delta z^\top g(z)\,\Delta z$

Latent metric $g(z) = J(z)^\top M(\hat{q})\,J(z)$ the pullback of M through the decoder

2 The distance between two embeddings

g(z) is computed exactly at every node

Graph: k-nearest neighbors (κ = 20) among healthy reference windows and one test subject's windows

Edge length uses the mean of the two endpoint metrics

$$\sqrt{\Delta z^\top \tfrac{1}{2}\left(g(z_i) + g(z_j)\right)\Delta z}$$

The shortest path approximates the Riemannian distance

Subject score: median over the subject's windows

Fine-tuning on DNE healthy controls and evaluation across five cohorts

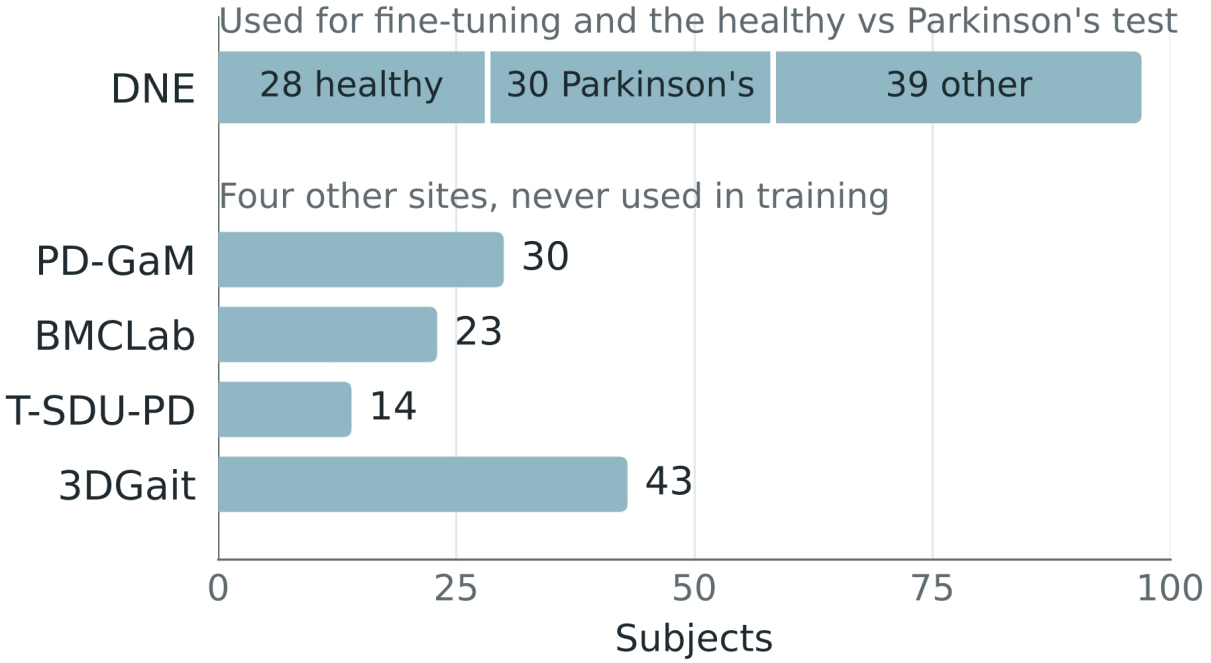

Fine-tuning and reference

Healthy gaits from DNE as a reference

We fine-tune REINS on these recordings and use them as the reference for scoring.

Test 1: healthy vs Parkinson's

Distinguishing Parkinson's gait

We compare scores for people with PD and held-out healthy controls.

Test 2: new sites

Four cohorts with UPDRS-gait ratings

We score gait against the DNE reference and compare the scores with UPDRS-gait ratings.

REINS uses the body's inertia to weight velocity errors during fine-tuning and to measure the Riemannian distance to healthy gait.

Both Euclidean baselines start from the same pretrained checkpoint as REINS.

Data: CARE-PD benchmark (Adeli et al., 2025). Backbone: PhysPT (Zhang et al., 2024).

NEURAL INFORMATION PROCESSING SYSTEMS 2026

Can a score based on healthy gait distinguish PD from controls?

Method	Main split [95% CI]	HC vs PD AUROC, main split	Ten random splits (mean ± sd)
Euclidean latent space distance: Euclidean			
Data-driven	0.5750 [0.3792, 0.7625]		0.5633 ± 0.050
Physics-supervised	0.5167 [0.3167, 0.7125]		0.5242 ± 0.049
Published motion encoders distance: Riemannian			
VPoser	0.5708 [0.3750, 0.7708]		0.5888 ± 0.046
GaitForeMer	0.7125 [0.5250, 0.8792]		0.7071 ± 0.033
MotionBERT	0.7667 [0.5458, 0.9208]		0.7621 ± 0.025
Riemannian latent space (*M*-induced, ours) distance: Riemannian			
REINS	**0.9083** [0.8042, 0.9833]		**0.8804** ± 0.027

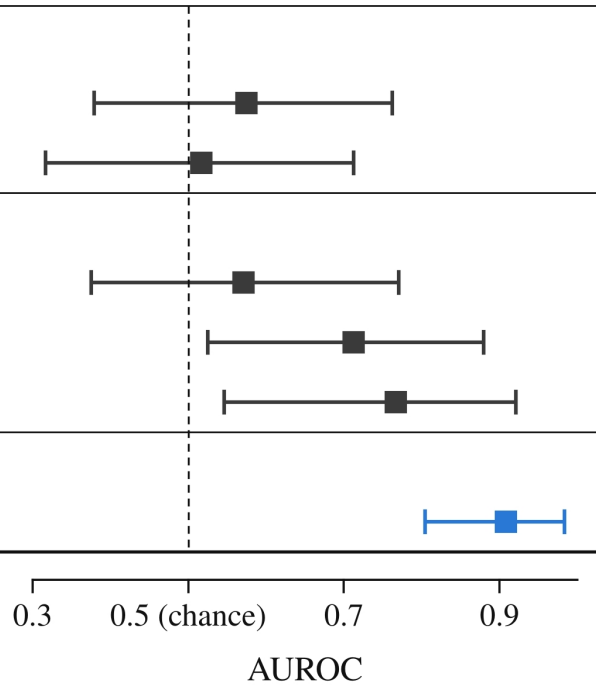

With REINS, people with PD tended to have a greater distance from the healthy gait reference than held-out controls.

On DNE (CARE-PD), we compare 8 held-out healthy controls with 30 participants with PD. Error bars show 95% bootstrap confidence intervals.

How much of that gain comes from the inertia matrix?

Scoring distance	HC vs PD AUROC [95% CI]
Fine-tuning: Physics-supervised (w/o $\mathcal{L}_{\text{ine}}$)	
Euclidean	0.5167 [0.3167, 0.7125]
$J^\top J$ Riemannian	0.5542 [0.3583, 0.7750]
$J^\top MJ$ Riemannian	0.5917 [0.3833, 0.7708]
Fine-tuning: REINS (w/ $\mathcal{L}_{\text{ine}}$)	
Euclidean	0.6208 [0.4250, 0.8000]
$J^\top J$ Riemannian	0.7333 [0.5375, 0.8667]
$J^\top MJ$ **Riemannian**	**0.9083** [0.8042, 0.9833]

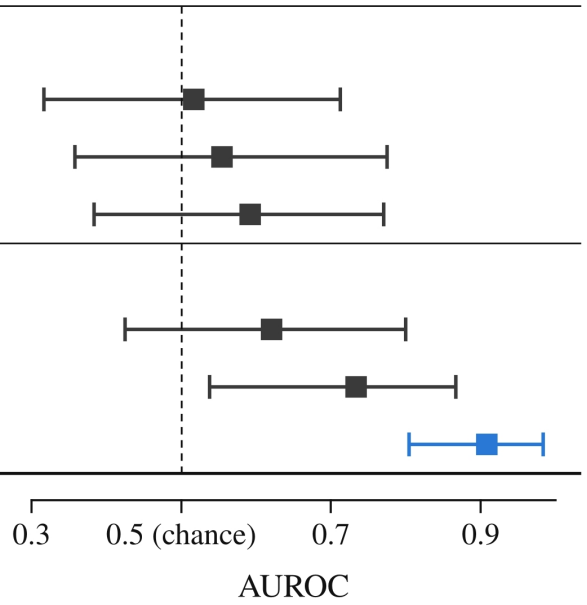

HC vs PD AUROC, 20 healthy training subjects. Intervals: 95% bootstrap.

With the REINS model fixed, adding the body's inertia to the decoder's metric improves PD discrimination.

The same scoring rule separates PD from healthy controls better when fine-tuning also uses the inertia-consistent loss.

Can REINS still distinguish PD with fewer healthy training subjects?

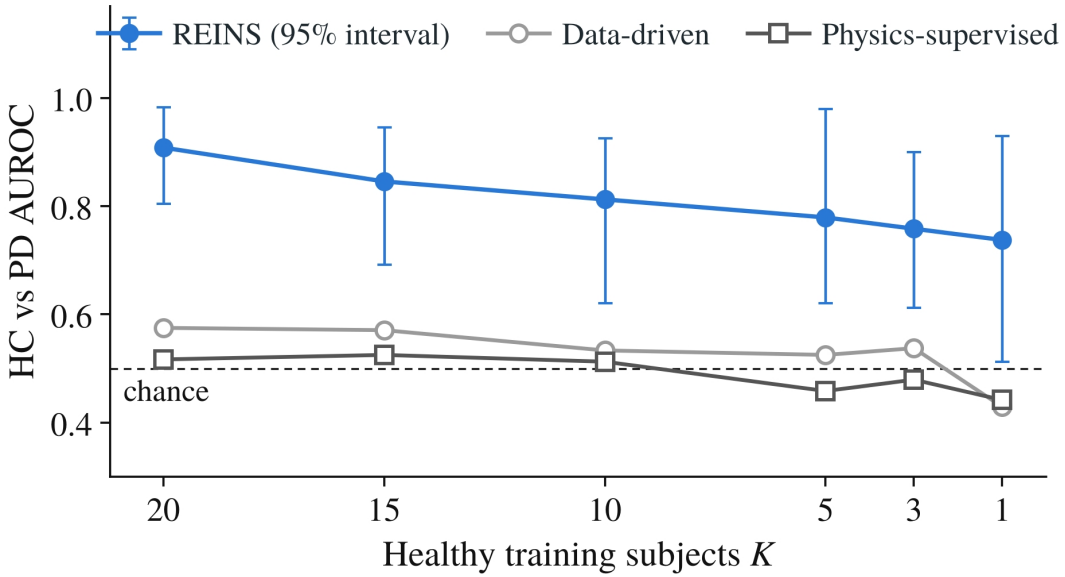

Method	K=20	K=15	K=10	K=5	K=3	K=1
Euclidean latent space distance: Euclidean						
Data-driven	0.5750	0.5708	0.5333	0.5250	0.5375	0.4292
Physics-supervised	0.5167	0.5250	0.5125	0.4583	0.4792	0.4417
Riemannian latent space (M-induced, ours) distance: Riemannian						
REINS	**0.9083**	**0.8458**	**0.8125**	**0.7792**	**0.7583**	**0.7375**

We used progressively fewer healthy participants to fine-tune REINS and build the reference. When we used gait data from one healthy participant for both, the score still distinguished people with PD from held-out healthy controls.

Intervals for REINS: 95% bootstrap.

Do REINS scores correlate with clinical gait ratings in new cohorts?

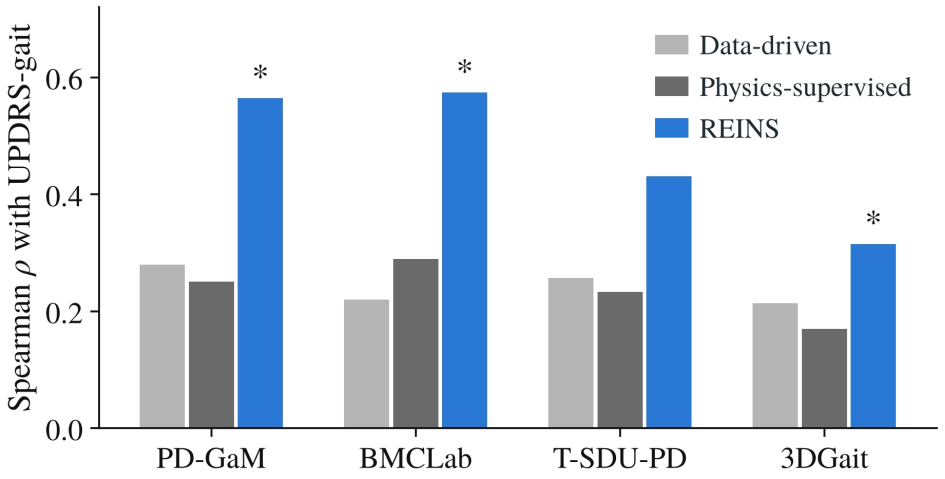

Method	PD-GaM (n=30)		BMCLab (n=23)		T-SDU-PD (n=14)		3DGait (n=43)	
	ρ	τ	ρ	τ	ρ	τ	ρ	τ
Euclidean latent space distance: Euclidean								
Data-driven	0.2802	0.1648	0.2200	0.1972	0.2569	0.2405	0.2138	0.1891
Physics-supervised	0.2510	0.1405	0.2895	0.2648	0.2337	0.1243	0.1700	0.1567
Riemannian latent space (M-induced, ours) distance: Riemannian								
REINS	**0.5649***	**0.4582***	**0.5742***	**0.4400***	**0.4308**	**0.3078**	**0.3154***	**0.2132**

We scored patients at other sites using the healthy gait reference from the main experiment.

Within each cohort, patients with more severe gait impairment on the UPDRS tended to have higher REINS scores.

Subject-level correlation between the score and the UPDRS-gait rating. * $p < 0.05$

SUMMARY

NEURAL INFORMATION PROCESSING SYSTEMS 2026

A biomechanical distance distinguishes Parkinson's gait and correlates with clinical gait ratings

1 A Riemannian distance based on body mechanics

REINS uses inertia in fine-tuning and scoring. How much a change in motion counts depends on which joints move and on the body's posture.

2 Performance with limited clinical data

On DNE, REINS could distinguish PD from healthy controls using the same small group of healthy participants for fine-tuning and as the reference.

3 Correlation with clinical gait ratings across cohorts

At other sites, participants whose gait was farther from the DNE healthy reference tended to have higher UPDRS-gait ratings.

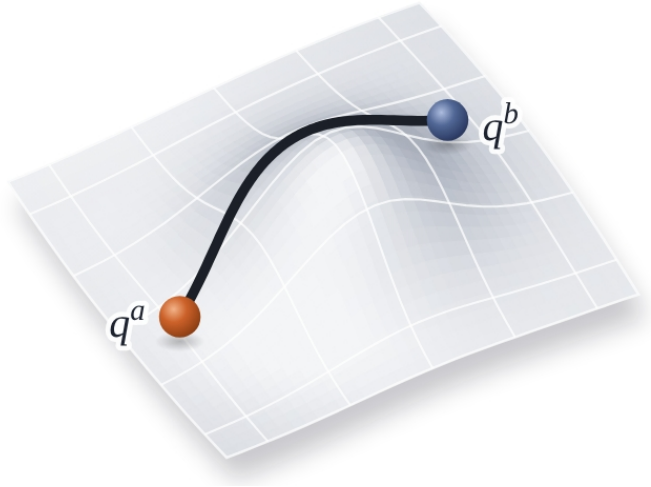

Code github.com/DominoAI-Lab/REINS-NeurIPS-2026

Paper openreview.net/forum?id=6Ku1AhEjKI

Thank you

Acknowledgments

This work was supported by Michigan State University.

Yiran Ding yiran_ding@alumni.brown.edu

Poster Session 4.
Room: Hall C1 Thu 16:30 EST in Atlanta

Code: github.com/DominoAI-Lab/REINS-NeurIPS-2026

Paper: openreview.net/forum?id=6Ku1AhEjKl

