# MuscleLens: A Shared Muscle-Space Pipeline for Parkinsonian Gait Analysis

Yinglei Zhu[1†] Bozhao Wang[1†] Huichan Zhao[1*]
[†]Equal contribution [*]Corresponding author

***Abstract*—Parkinson's disease (PD) gait assessment remains dominated by subjective clinical rating, while instrumented gait laboratories that could substitute for it are inaccessible to most patients. We present MuscleLens, a pipeline that lifts monocular video into a phase-locked $30 \times 80$ muscle activation code on a unified musculoskeletal model. The stack combines SMPL recovery, GMR retargeting to MyoFullBody, and a frozen MuscleMimic policy [1] that emits 80-dimensional activations at $100\,\mathrm{Hz}$. Evaluated on roughly 4 800 subject-level samples from CARE-PD, augmenting SMPL-H kinematics with muscle features (i) raises 5-fold PD balanced accuracy from $0.807$ to $0.814$ at matched AUC ($0.937$); (ii) consistently lowers cohort-normalised UPDRS MAE across PCA dimensions, reaching $0.344$ versus $0.349$ at 128 dimensions; and (iii) improves leave-one-cohort-out balanced accuracy on six of seven cohorts. Beyond accuracy, the actuator-space code exposes co-activation patterns that joint kinematics alone cannot encode, providing a biomechanically interpretable view of PD gait. The monocular branch demonstrates the end-to-end deployment route: about $1\,\mathrm{min}$ from a $15\,\mathrm{s}$ clip to a complete muscle code on a single RTX 5090, bringing muscle-level analysis within reach of routine clinical-style acquisition. Project page: https://fly-pigTH.github.io/MuscleLens.**

***Index Terms*—Parkinson's disease, gait analysis, musculoskeletal simulation, imitation learning, monocular human recovery, eldercare, MuJoCo.**

## I. Introduction

PARKINSON'S disease (PD) gait assessment remains dominated by subjective visual rating with the MDS-UPDRS scale [2], while the high-fidelity instrumented gait laboratories that could substitute for it remain inaccessible to most patients outside specialised centres. Recent progress in large-scale 3D human recovery, motion retargeting, and musculoskeletal control now makes it practical to project ordinary motion observations into a shared actuator-space representation [1], [3], [4].

Specifically, we ask whether augmenting joint-kinematic descriptors with a musculoskeletal layer yields measurable gains in PD-related supervised prediction, and whether the resulting actuator-space representation captures biomechanically interpretable structure not explicitly encoded by kinematics alone. The monocular branch is presented as a deployment demonstration that the same muscle code can be obtained directly from RGB video [3].

To address this under a fair shared input, we evaluate on the SMPL-H sequences released with the CARE-PD corpus [5] across three protocols: (a) random five-fold PD classification, (b) leave-one-cohort-out (LOCO) transfer, and (c) cohort-normalised UPDRS-gait regression.

## II. Related Work

Monocular motion-recovery systems routinely lift RGB video into SMPL-family trajectories, making markerless gait analysis feasible at scale [3]. In parallel, musculoskeletal simulators and retargeting pipelines make it practical to replay human motion on actuated biomechanical models rather than treating pose estimation as the endpoint [1], [4], [6]. MuscleLens chains these components as a fixed inference stack: video is mapped to body motion, body motion is retargeted to a muscle-actuated model, and the controller output is read as a surrogate actuator-space measurement.

The framing is motivated by PD gait analysis. Prior work demonstrates that video and 3D pose can approximate clinical gait assessment, while muscle-synergy studies show that joint kinematics alone can miss co-activation structure relevant to disease severity [7]. We therefore pose a limited, testable question: given a shared SMPL-H input on CARE-PD, do

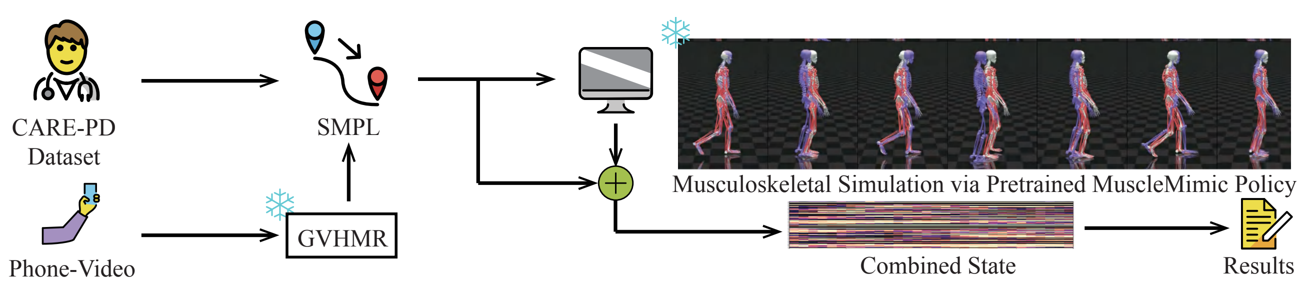


Fig. 1. **MuscleLens pipeline.** Video is brought into a common SMPL-H frame (with cross-source coordinate harmonisation applied to the CARE-PD branch only), retargeted to MyoFullBody via GMR, tracked by a frozen MuscleMimic policy, and condensed into a phase-locked muscle code for downstream analysis. Quantitative experiments validate the shared SMPL branch; the monocular branch defines the end-to-end deployment route. Qualitative video-to-muscle visualisations are available at https://fly-pigTH.github.io/MuscleLens.

muscle-derived features add predictive signal beyond joint-space descriptors?

## III. Method and Setup

*Pipeline:* MuscleLens accepts either monocular video or SMPL motion. For video we recover SMPL-H with GVHMR [3]; for CARE-PD we use the released SMPL-H sequences directly. The motion is retargeted to the MyoFullBody musculoskeletal model [6] via GMR [4] and replayed under the frozen `mm-fullbody-base` MuscleMimic policy [1], which emits 80-dimensional muscle activations at 100 Hz. A lightweight coordinate-frame harmonisation step precedes retargeting on the CARE-PD branch to keep cross-source processing consistent; we treat it as data cleaning, not a scientific contribution.

*Representation and evaluation:* Activations are segmented into gait cycles, resampled to 30 phase bins per cycle, and averaged across valid cycles to form a code $\mathbf{A} \in \mathbb{R}^{30\times 80}$ with $A_{p,m} = \mathbb{E}_k[a_m(t_p^{(k)})]$. We compare a 550-dimensional NMF summary of this code, a 5 310-dimensional kinematic descriptor, and their concatenation. Evaluation uses the public CARE-PD corpus [5] (8 476 clips, nine cohorts); the processed run analysed here covers 4 820 subject-level samples drawn from seven cohorts, with 4 669 PD labels and 2 559 UPDRS-gait labels. We report PCA+HDBSCAN clustering, random five-fold logistic models, and leave-one-cohort-out (LOCO) transfer. A third, cohort-normalised protocol z-scores each cohort separately and projects all modalities to equal-dimensional PCA spaces before regression, neutralising dimensionality and source-distribution advantages that would otherwise favour the high-dimensional kinematic vector.

*Compute and runtime:* The full inference stack runs on a single NVIDIA RTX 5090. The MuscleMimic policy emits activations at roughly 12 s per SMPL clip; on the video branch, end-to-end processing of a 15 s clip (GVHMR $\rightarrow$ SMPL $\rightarrow$ GMR $\rightarrow$ MuscleMimic) completes in about 1 min, returning both kinematics and muscle activations in a single pass. The headline number is the end-to-end budget a downstream user would experience. This places MuscleLens within reach of routine clinical-style acquisition without specialised gait laboratories.

## IV. Results and Discussion

All quantitative results below share the same SMPL-H input, isolating the effect of muscle augmentation from any upstream reconstruction error.

*Random five-fold classification:* Kinematics remain the strongest single modality (Fig. 2a). Augmenting them with muscle features preserves PD AUC ($0.937 \pm 0.029$ versus $0.938 \pm 0.030$) and raises balanced accuracy from $0.807 \pm 0.011$ to $0.814 \pm 0.025$. Paired across folds, the BAcc gain averages $+0.007$ (paired $t = 0.78$) and the AUC tie is robust (paired diff $-0.001$, $t = -0.41$). UPDRS-gait MAE in the random five-fold setting is essentially unchanged ($0.382 \pm 0.022$ versus $0.372 \pm 0.013$); the UPDRS-side improvement instead emerges once cohort structure is controlled, as shown next.

*Cohort-normalised UPDRS regression:* The clearest positive result emerges once the kinematic dimensionality advantage is removed. Across 8, 16, 32, 64, and 128 PCA dimensions, fusion lowers UPDRS MAE relative to kinematics by 0.002, 0.003, 0.008, 0.006, and 0.005 points, respectively (Fig. 2c). At 128 dimensions, fusion reaches MAE 0.344 versus 0.349 for kinematics. The gain is small in absolute terms but *consistent at every PCA dimension we tested* — the property we expect from a complementary modality.

*Leave-one-cohort-out transfer:* Table I reports per-cohort LOCO results. Pooled LOCO PD AUC looks alarmingly low for both modalities (kinematics 0.282, fusion 0.340), but the within-cohort breakdown explains why: four of the seven processed cohorts contain only PD-positive samples, so per-cohort AUC is undefined and the pooled value is dominated by between-cohort probability-calibration shift rather than by ranking failure inside any single cohort. LOCO balanced accuracy, which is invariant to that calibration shift, gives a cleaner picture: *fusion improves BAcc on six of seven cohorts*, with mean gain $+0.020$ and range $[-0.003, +0.045]$. LOCO UPDRS MAE is likewise strictly improved by fusion on all three mixed-severity cohorts. We read this as evidence that the muscle code adds a small but cohort-transferable signal on top of kinematics; we do not claim it solves cohort transfer, which remains the dominant open problem on this corpus.

TABLE I
LEAVE-ONE-COHORT-OUT PD BALANCED ACCURACY. FUSION IMPROVES BACC ON 6 OF 7 COHORTS (MEAN +0.020). COHORTS MARKED 100% PD HAVE UNDEFINED PER-COHORT AUC, SO POOLED AUC IS NOT A MEANINGFUL SUMMARY.

| Cohort | N | % PD | Kin. BAcc | Fus. BAcc | Δ |
|---|---|---|---|---|---|
| 3DGait | 88 | 72.7% | 0.466 | 0.503 | +0.036 |
| BMCLab | 779 | 100% | 0.981 | 0.985 | +0.004 |
| DNE | 303 | 38.3% | 0.480 | 0.506 | +0.026 |
| E-LC | 162 | 90.1% | 0.483 | 0.479 | −0.003 |
| KUL-DT-T | 735 | 100% | 0.899 | 0.905 | +0.005 |
| PD-GaM | 1 692 | 100% | 0.690 | 0.717 | +0.027 |
| T-LTC | 910 | 100% | 0.618 | 0.663 | +0.045 |

*Unsupervised structure:* The fused unsupervised embedding (Fig. 3) shows one dominant cluster and a small side cluster, with severity labels forming only a weak gradient. Weighted UPDRS purity is 0.472 for muscle, 0.448 for kinematics, and 0.447 for fusion; PD purity sits at $\approx 0.96$ across all modalities, essentially the labelled PD base rate of 95.1 %. The unsupervised view is therefore not diagnostic on its own; the supervised analyses above remain where the muscle-space contribution is most credible.

*Summary:* Table II condenses the supervised results. Muscle features add a small but consistent gain on top of kinematics; the gain is most defensible once cohort structure and feature dimensionality are controlled, and it appears across three independent settings: random folds, cohort-normalised PCA, and LOCO BAcc. The small absolute gains likely reflect both the already-strong kinematic baseline and the severe cohort imbalance present in CARE-PD.

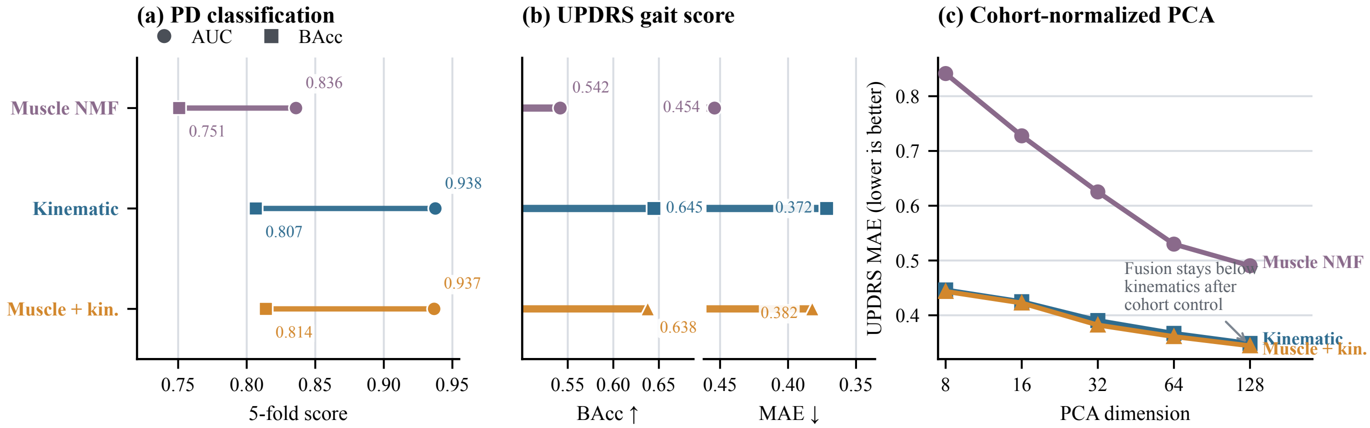


Fig. 2. **Modality comparison.** (a) Random five-fold PD classification. (b) Random five-fold UPDRS-gait prediction. (c) Under cohort-wise z-scoring and equal-dimensional PCA, fusion yields consistently lower UPDRS MAE than kinematics across all tested dimensions.

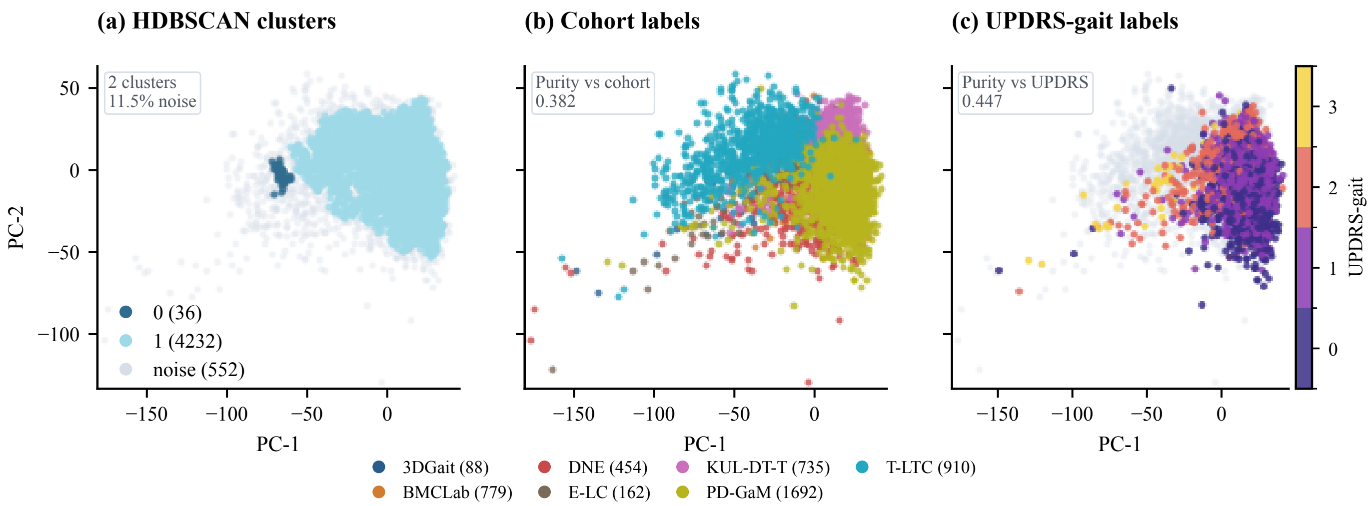


Fig. 3. **Fusion embedding overview.** (a) HDBSCAN identifies one dominant cluster and a small side cluster. (b) Cohort labels retain source structure. (c) UPDRS-gait labels form a weak ordering rather than clean separation.

TABLE II
COMPACT SUMMARY OF THE MAIN SUPERVISED RESULTS. PD AUC, PD BACC, AND UPDRS MAE ARE 5-FOLD MEAN ± STD; "128-D MAE" IS THE COHORT-NORMALISED PCA SETTING AT 128 DIMENSIONS.

| Modality | PD AUC | PD BAcc | UPDRS MAE | 128-D MAE |
|---|---|---|---|---|
| Muscle NMF | .836±.023 | .751±.027 | .454±.037 | 0.490 |
| Kinematic | .938±.030 | .807±.011 | .372±.013 | 0.349 |
| Muscle + kin. | .937±.029 | .814±.025 | .382±.022 | **0.344** |

## V. CONCLUSION AND LIMITATIONS

MuscleLens establishes a shared route from monocular video or SMPL-H motion to phase-locked muscle activation codes, and provides initial evidence that the resulting actuator-space representation carries information complementary to joint kinematics: *five-fold PD balanced accuracy improves, cohort-normalised UPDRS regression improves at every PCA dimension we tested, and LOCO balanced accuracy improves on six of seven cohorts.* The activation code also surfaces co-activation structure that joint kinematics alone cannot encode, which we see as a step toward more interpretable PD gait analysis.

Three limitations bound the present claim. First, the monocular branch is a deployment demonstration rather than a clinically validated front end; large-scale patient-video evaluation remains future work, with qualitative video → SMPL → muscle results documented at the project page https://fly-pigTH.github.io/MuscleLens. Second, the musculoskeletal model and tracking policy are not yet calibrated to older adults with PD-specific movement strategies, which likely caps co-activation fidelity. Third, the residual cohort-transfer gap on pooled LOCO PD AUC reflects a between-cohort probability-calibration problem more than a within-cohort ranking failure; domain-invariant training and cohort-calibrated scoring are natural next steps and the most direct route to strengthening the present claim.

# REFERENCES


[1] C. Li, C. Wang, B. Ziliotto, M. Simos, J. Kovecses, G. Durandau, and A. Mathis, "Towards embodied AI with MuscleMimic: Unlocking full-body musculoskeletal motor learning at scale," *arXiv preprint arXiv:2603.25544*, 2026. [Online]. Available: https://arxiv.org/abs/2603.25544

[2] C. G. Goetz, B. C. Tilley, S. R. Shaftman, G. T. Stebbins, S. Fahn, P. Martinez-Martin, W. Poewe, C. Sampaio, M. B. Stern, R. Dodel *et al.*, "Movement Disorder Society-sponsored revision of the Unified Parkinson's Disease Rating Scale (MDS-UPDRS): scale presentation and clinimetric testing results," *Movement Disorders*, vol. 23, no. 15, pp. 2129–2170, 2008.

[3] Z. Shen, H. Pi, Y. Xia, Z. Cen, S. Peng, Z. Hu, H. Bao, R. Hu, and X. Zhou, "World-grounded human motion recovery via gravity-view coordinates," in *SIGGRAPH Asia 2024 Conference Papers*, 2024, arXiv:2409.06662. [Online]. Available: https://arxiv.org/abs/2409.06662

[4] J. P. Araujo, Y. Ze, P. Xu, J. Wu, and C. K. Liu, "Retargeting matters: General motion retargeting for humanoid motion tracking," *arXiv preprint arXiv:2510.02252*, 2025, to appear, ICRA 2026. [Online]. Available: https://arxiv.org/abs/2510.02252

[5] V. Adeli, I. Klabucar, J. Rajabi, B. Filtjens, S. Mehraban, D. Wang, H. Seo, T.-H. Hoang, M. N. Do, C. Muller, C. Oliveira, D. B. Coelho, P. Ginis, M. Gilat, A. Nieuwboer, J. Spildooren, L. McKay, H. Kwon, G. Clifford, C. Esper, S. Factor, I. Genias, A. Dadashzadeh, L. Shum, A. Whone, M. Mirmehdi, A. Iaboni, and B. Taati, "CARE-PD: A multi-site anonymized clinical dataset for Parkinson's disease gait assessment," in *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*, 2025. [Online]. Available: https://arxiv.org/abs/2510.04312

[6] V. Caggiano, H. Wang, G. Durandau, M. Sartori, and V. Kumar, "MyoSuite – a contact-rich simulation suite for musculoskeletal motor control," in *Proceedings of the 4th Conference on Learning for Dynamics and Control (L4DC)*, 2022, arXiv:2205.13600. [Online]. Available: https://arxiv.org/abs/2205.13600

[7] F. Hulzinga, N. D'Cruz, C. McCrum, P. Ginis, M. Gilat, and A. Nieuwboer, "The impact of Parkinson's disease on adaptive split-belt walking: A muscle synergy analysis," *Clinical Neurophysiology*, 2025.