Does Locomotion-Intent Recognition Make a Knee Exoskeleton Better? From Wearable-Sensor Classification to Closed-Loop Simulation

Tasfia Noor Chowdhury
M.Sc. Robotics & Mechatronics Engineering, University of Dhaka
tasfiaa.chowdhury@gmail.com · November 2026

Abstract. Powered lower-limb exoskeletons need to know whether the user is walking, climbing stairs or negotiating a ramp, and they need to know it *before* the new activity starts. This work builds a complete intent-recognition pipeline on the public ENABL3S dataset (10 able-bodied subjects, 22,411 gait events) and then asks the question that classification papers usually skip: how much does the remaining classification error cost the user once the prediction drives a controller? A random forest on hand-crafted features from 300 ms windows predicts the next locomotion mode with 8.5 % error in leave-one-subject-out evaluation (24.7 % at mode transitions), on par with the dataset authors' features. Removing all EMG channels does not degrade accuracy (8.2 % vs 8.5 %), suggesting an exoskeleton can rely on its own mechanical sensors. In a MuJoCo knee-exoskeleton simulation with mode-specific impedance control, classifier-driven mode switching reduces knee tracking error from 13.5° (single fixed controller) to 4.1°, against 0.5° for perfect mode knowledge. The largest remaining errors come from a small number of "standing" predictions during walking, which lock the knee and are a safety concern.

I. INTRODUCTION

Lower-limb exoskeletons assist differently on level ground, ramps and stairs: stair ascent needs a strongly flexed knee at foot contact and large extension torque, level walking does not. A controller tuned for one mode therefore helps little, or even hinders, in another. Commercial devices often let the user select the mode with a button, which is slow and error-prone. Automatic **locomotion-mode recognition** from wearable sensors has been studied for over a decade with EMG, inertial and joint-angle signals [2], [5], and the ENABL3S benchmark [1] made bilateral recordings of these signals public.

Most of this literature reports classification error. For a wearable robot, the more relevant question is what the error does to assistance: a wrong prediction during steady walking may be harmless, while the same error rate concentrated at transitions, or a confusion that locks the knee, is not. This project therefore connects the two:

1. A reproducible pipeline from raw signals to next-mode prediction, evaluated leave-one-subject-out.
2. A sensor ablation: is EMG worth its practical cost for an exoskeleton?
3. A closed-loop MuJoCo simulation in which predictions select the knee controller, compared against perfect knowledge and against no mode recognition.

II. METHODS

A. Dataset

ENABL3S [1] contains 10 able-bodied subjects walking a circuit that includes level walking (LW), ramp ascent and descent (RA, RD), stair ascent and descent (SA, SD), standing and sitting, about 50 circuits per subject. Signals are recorded at 1 kHz from 14 surface EMG channels (7 muscles per leg), 5 six-axis IMUs (both thighs and shanks, waist) and knee and ankle goniometers on both legs. Every heel contact and toe off carries a label of the current and the next locomotion mode. A gait event is a *transition* when the two differ. This analysis uses 10 subjects.

B. Windowing and features

At each of the 22,411 gait events a 300 ms window ending at the event is extracted. Channels are ordered relative to the leg that produced the event (ipsilateral, contralateral, waist), so that left and right events share one feature space. For the 38 mechanical channels (IMU, joint angle and velocity) six statistics are computed: mean, standard deviation, minimum, maximum, first and last value. For the 14 EMG channels the time-domain set of Hudgins et al. [3] is used: mean absolute value, waveform length, zero crossings and slope sign changes, the standard set in myoelectric control [4]. This gives 284 features. The extracted events and labels match the dataset authors' feature files one-to-one, which allows a direct comparison.

C. Classifiers and evaluation

Linear discriminant analysis (LDA, shrinkage), an RBF support vector machine and a random forest (RF, 300 trees) [6] are trained on the features. A 1D convolutional network (three convolutional layers, global average pooling, ≈ 50 k parameters) is trained end-to-end on the window itself, down-sampled to 100 Hz with EMG replaced by its envelope. All models are evaluated **leave-one-subject-out**: the test person is never seen in training, and standardisation statistics come from the training subjects only. Errors are reported separately for steady-state and transition events.

D. Knee-exoskeleton simulation

A planar leg on a test bench is modelled in MuJoCo [7]: a fixed thigh and a 3.5 kg shank with a 1 kg foot rotate about a hinge knee driven by a 60 Nm motor. The leg is assumed passive, so the exoskeleton produces the full knee motion. For each walking mode, the reference trajectory is the knee angle averaged over all steady strides of all subjects. The controller combines inverse-dynamics feedforward with mode-specific impedance,

$$\tau = \tau_{ff}(q_r, \dot{q}_r, \ddot{q}_r) + K_m(q_r - q) + D_m(\dot{q}_r - \dot{q}),$$

with stiffer gains for stairs. Because the mode references start at different angles (stair ascent lands with the knee ≈ 60° flexed), the reference is blended smoothly from its previous

value over the first 25 % of the stride after a switch. At every right heel contact the mode for the next stride is chosen by (i) the true label (*oracle*), (ii) the leave-one-subject-out RF (*predicted*) or (iii) always level walking (*fixed*). Every circuit of every subject is simulated, and knee-angle error is measured against the reference of the **true** mode, i.e. the motion the user needed.

III. RESULTS

A. Classification

Model	Overall	Steady	Transition	Range
LDA (own features)	14.8	8.1	37.5	10.1–22.9
SVM (own features)	11.7	9.6	18.6	3.8–25.2
RF (own features)	8.5	3.8	24.7	5.6–14.6
RF (provided features)	9.0	4.2	25.1	5.8–16.8
1D-CNN (raw windows)	12.2	9.6	21.4	5.0–19.0

Table I. Leave-one-subject-out error (mean over test subjects) for next-mode prediction. Range = best and worst test subject.

The RF on the new features reaches 8.5 % overall and 24.7 % transition error, comparable to the RF on the dataset authors' 368 features (9.0 % / 25.1 %). The 1D-CNN has a higher overall error (12.2 %) but a lower transition error (21.4 %) than the RF, and it gained most from additional data: with 6 subjects it reached only 18.3 % / 28.3 %. The SVM shows the same trade-off (11.7 % overall, 18.6 % at transitions). All models make 1.9–6.5× more errors at transitions than in steady state, and every model varies strongly between test subjects (Table I, range), which points to inter-subject variability as the main remaining difficulty.

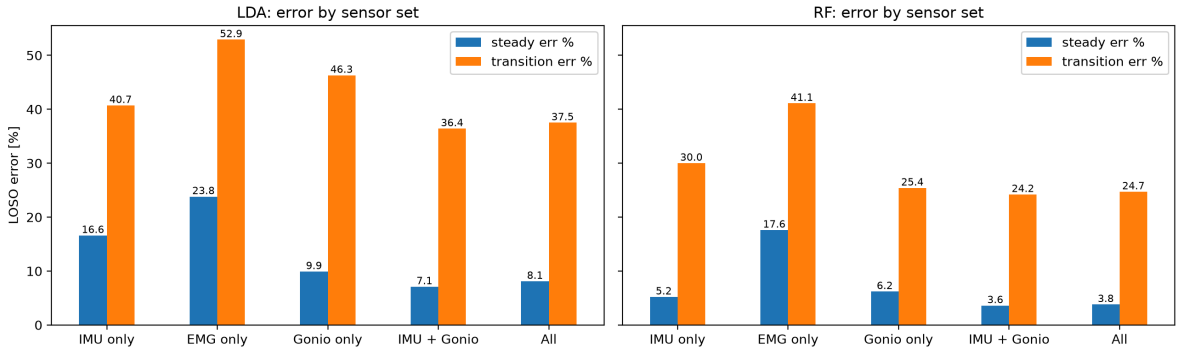


Fig. 1. Sensor ablation: steady-state and transition error for each sensor set (left LDA, right RF).

B. Which sensors matter

Sensors	LDA	RF
IMU only	22.1 / 40.7	10.9 / 30.0
EMG only	30.5 / 52.9	23.0 / 41.1
Goniometer only	18.3 / 46.3	10.5 / 25.4
IMU + Goniometer	13.8 / 36.4	8.2 / 24.2
All	14.8 / 37.5	8.5 / 24.7

Table II. Overall / transition error [%] per sensor set.

With the RF, mechanical sensors alone (IMU + goniometer, 8.2 %) perform as well as all sensors (8.5 %), and EMG alone is the weakest set (23.0 %). The random-forest importance agrees: 18 % of the total importance falls on EMG features. For an exoskeleton, which has joint encoders and can carry IMUs at little cost, EMG electrodes add set-up time, sweat and fatigue sensitivity without a measurable gain here.

C. How early can the decision be made?

Delay [ms]	All	IMU + Gonio
0	24.7	24.2
30	23.7	23.8
60	23.9	23.3
90	22.8	23.1
120	22.3	22.0

Table III. RF transition error [%] when the window ends 0–120 ms after the gait event.

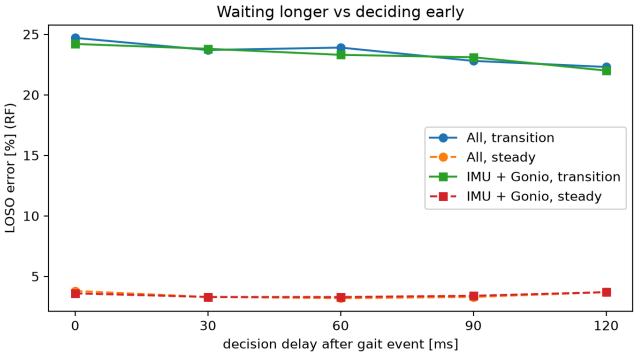


Fig. 2. Error vs decision delay (RF).

Waiting after the gait event lets the window include the start of the new movement, but the gain is small: transition error falls only from 24.7 % to 22.3 % at 120 ms, so deciding at the gait event costs little. The hardest transitions are LW → RD (65 % error, mostly predicted as LW) and LW → Stand (60 % error, mostly predicted as LW): the classifier tends to keep the old mode one step too long, i.e. it switches late rather than wrongly.

D. Closed-loop exoskeleton simulation

Mode selection	RMSE [°]	Max error [°]	Mean torque [Nm]
Oracle (true mode)	0.5	4.1	8.2
Predicted (RF)	4.1	24.3	8.2
Fixed (level walking)	13.5	55.4	7.7

Table IV. Knee tracking error relative to the true mode's reference, averaged over all simulated circuits. Max error = largest error within a

circuit, averaged over circuits.

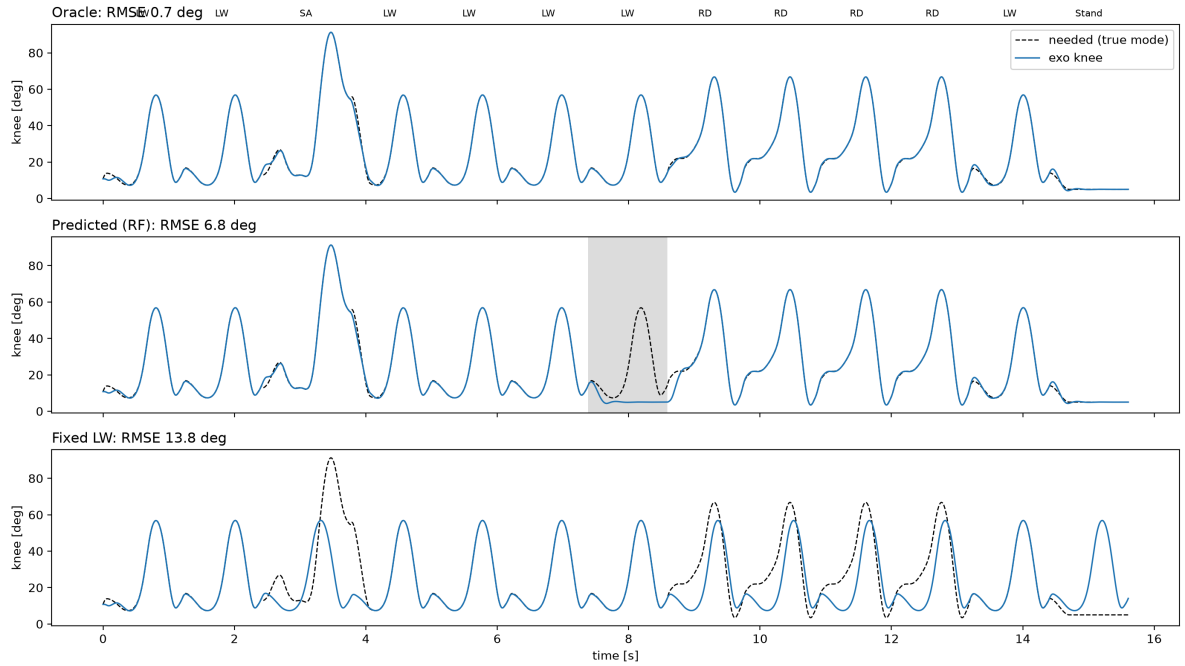


Fig. 3. One circuit of a test subject. Dashed: knee motion the user needed; solid: simulated exoskeleton knee. Grey: strides where the classifier chose the wrong mode.

Classifier-driven switching reduces the tracking error by a factor of 3.3 compared with a single level-walking controller (13.5° → 4.1° RMSE), recovering most of the benefit of perfect mode knowledge (0.5°). The fixed controller fails exactly where assistance matters most: it under-flexes the knee by tens of degrees on stairs (Fig. 3, bottom). The remaining error is not spread evenly. Fig. 3 (middle) shows the most harmful failure: when a walking stride is classified as *standing*, the controller holds the knee straight while the user needs to swing it; such an error would trip a real user.

IV. DISCUSSION

What the classifier error costs. The 8.5 % classification error translated into a 4.1° average tracking error, but the largest error within a circuit averages 24°, because errors cluster at transitions and in a few confusions. Reporting only accuracy hides this; the closed-loop simulation exposes which errors are dangerous. This motivates confidence-gated switching (keep the current mode unless the classifier is confident) and mode-transition rules that forbid physically implausible jumps such as walking → standing within one stride.

EMG. In this dataset mechanical signals carry the information needed for mode recognition. EMG may still matter for estimating *how much* torque to give, or for users with altered gait, but not for *which* mode to select.

Limitations. The simulation is a swing-only bench model without ground contact or body weight, the leg is passive, the mode is selected once per stride, and references are population averages. Results are for 10 able-bodied subjects; users with impairments, the target group of exoskeletons, typically show larger inter-subject variability.

Future work. Mid-stride switching at toe off, confidence-gated switching, personalised references learned from a few demonstrations, and validation on a wearable sensor prototype (Project A of this portfolio).

V. CONCLUSION

A leave-one-subject-out locomotion-intent classifier built from wearable signals reaches 8.5 % error and, when it selects the controller of a simulated knee exoskeleton, reduces knee tracking error 3.3-fold compared with no mode recognition. EMG was not needed. The simulation also shows that a few specific misclassifications dominate the remaining error, which points to where intent recognition for exoskeletons should improve next.

Code and data. Project page with the demonstration video: noortasfia.github.io/projects/knee-exo. Code is available on request. The ENABL3S data are public [1].

REFERENCES

[1] B. Hu, E. Rouse, and L. Hargrove, "Benchmark datasets for bilateral lower-limb neuromechanical signals from wearable sensors during unassisted locomotion in able-bodied individuals," *Front. Robot. AI*, vol. 5, p. 14, 2018, doi:10.3389/frobt.2018.00014.

[2] H. Huang, T. A. Kuiken, and R. D. Lipschutz, "A strategy for identifying locomotion modes using surface electromyography," *IEEE Trans. Biomed. Eng.*, vol. 56, no. 1, pp. 65–73, 2009.

[3] B. Hudgins, P. Parker, and R. N. Scott, "A new strategy for multifunction myoelectric control," *IEEE Trans. Biomed. Eng.*, vol. 40, no. 1, pp. 82–94, 1993, doi:10.1109/10.204774.

[4] K. Englehart and B. Hudgins, "A robust, real-time control scheme for multifunction myoelectric control," *IEEE Trans. Biomed. Eng.*, vol. 50, no. 7, pp. 848–854, 2003, doi:10.1109/TBME.2003.813539.

[5] A. J. Young and D. P. Ferris, "State of the art and future directions for lower limb robotic exoskeletons," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 25, no. 2, pp. 171–182, 2017.

[6] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[7] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *Proc. IEEE/RSJ IROS*, 2012, pp. 5026–5033.