IMPERIAL

HarmonICA

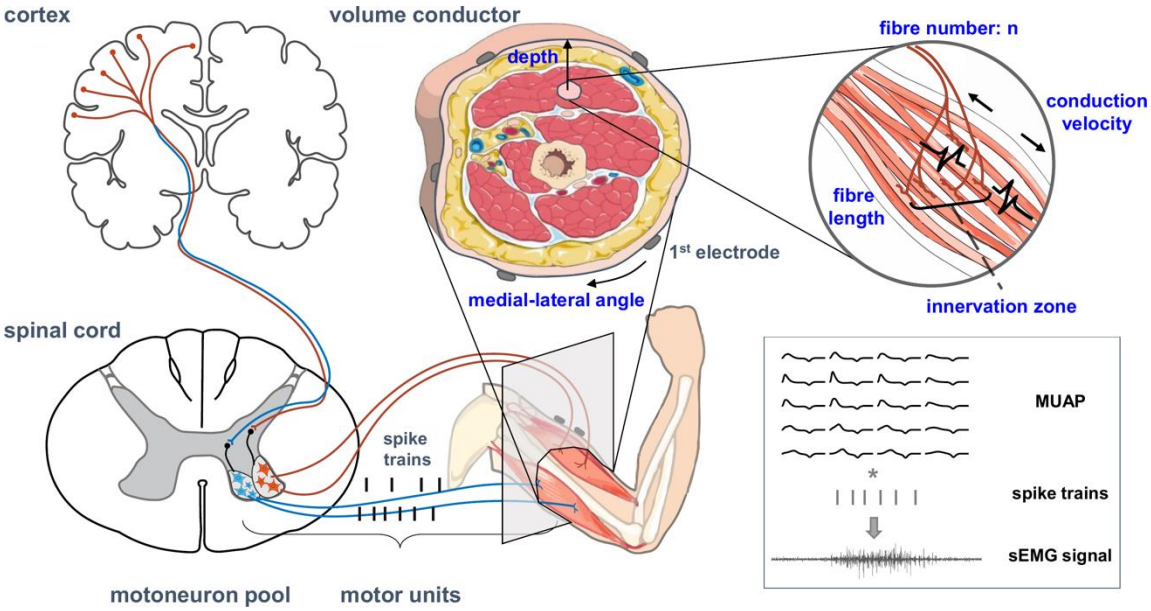
Quasi-Time-Invariant ICA for Motor Unit Decomposition during Dynamic Contractions

Alexander Kenneth Clarke, Dimitrios Halatsis, Agnese Grison, Irene Mendez Guerra, Noura Ezaz-Nikpay, Pranav Mamidanna, Shihan Ma, Silvia Muceli†, Dario Farina

Imperial College London · †Chalmers University of Technology

NeurIPS 2026

Recovering motor neuron spike trains from the skin



1 One neuron, one template
Every spike of a motor neuron stamps the same spatiotemporal template (MUAP) on the grid.

2 The grid records a mixture
Convolutive: 64–320 channels, tens of units overlapping.

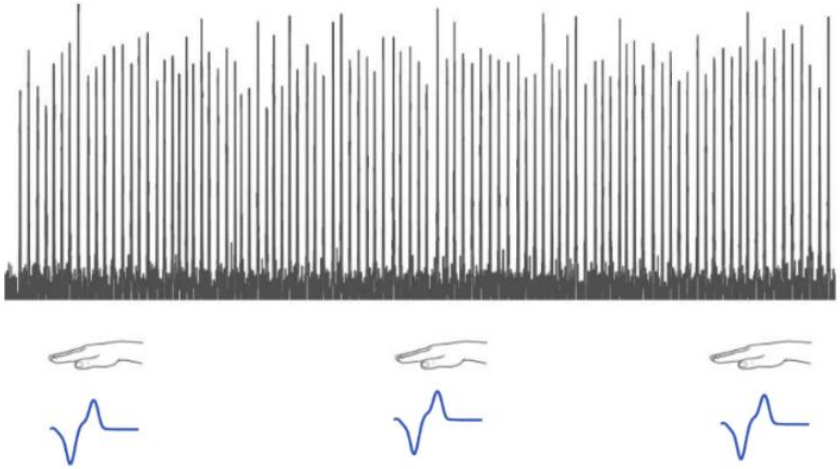
3 Decomposition inverts it
Recover every neuron's spike train with no labels: blind source separation.

$$x_i(t) = \sum_{j=1}^{M} \sum_{l=0}^{Q} h_{ij}(l)\, s_j(t-l) + \omega_i(t)$$

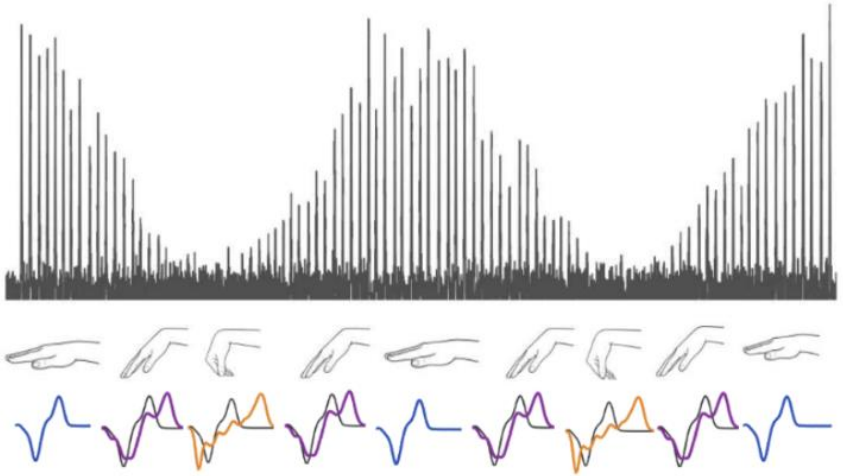
channels i = 1…N, motor units j = 1…M, MUAPs of length Q

Figure: Ma, Clarke et al., IEEE TNNLS 2025 (BioMime)

The problem: movement breaks a fixed separator



Wrist still
Same filter, clean spikes for the whole recording

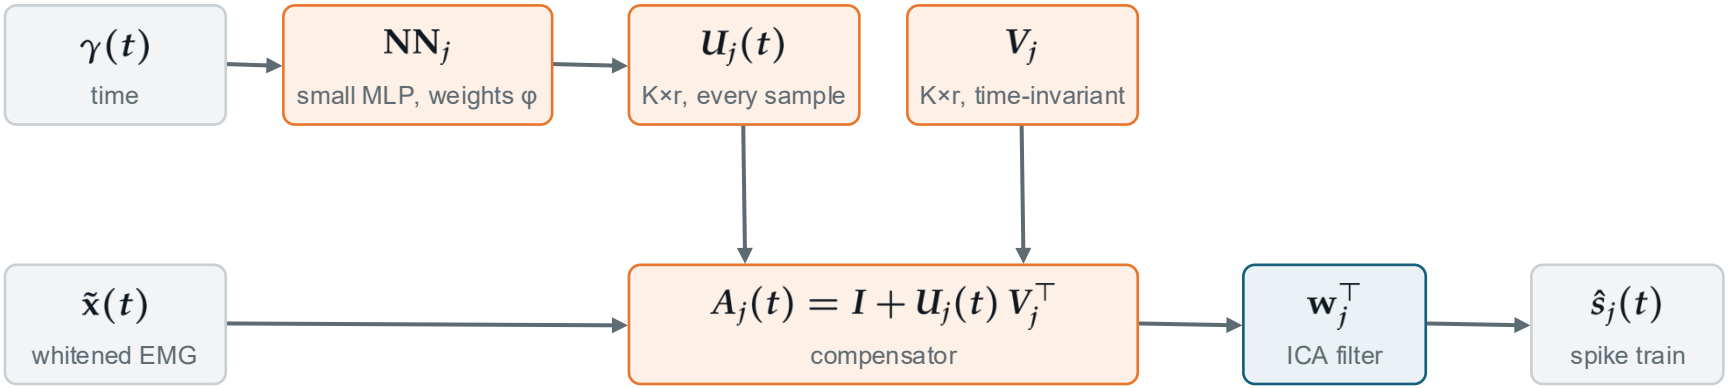


Wrist moving
The waveform drifts from the filter's template: spikes become modulated and unfindable

$$\tilde{\mathbf{x}}(t) = H(t)\,\tilde{\mathbf{s}}(t) + \tilde{\omega}(t)$$

the mixing moves with the joint; a fixed w cannot follow it

The model: a learned compensator in front of ICA



$$\gamma(t) = \left(\sin(2\pi t/T_p),\ \cos(2\pi t/T_p)\right)_{p=1}^{P}$$

$$A_j(t)^{-1} = I - U_j(t)\left(I_r + V_j^\top U_j(t)\right)^{-1} V_j^\top$$

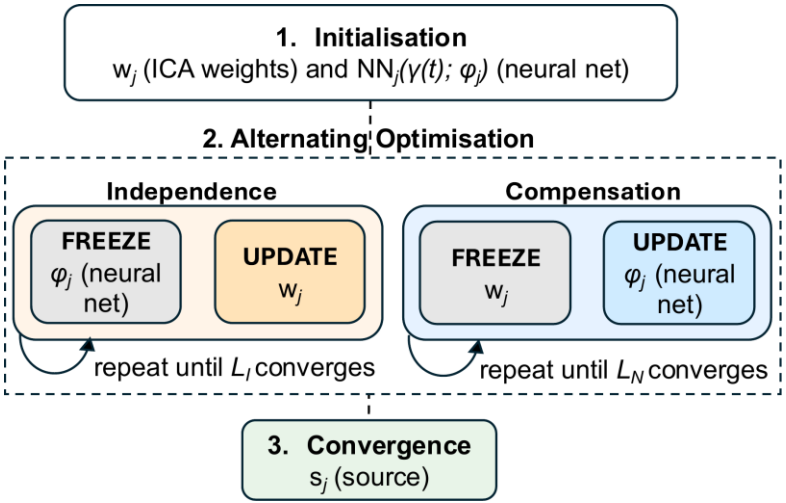
Identity + rank r: reduces to ICA, cheap, and invertible in closed form: Woodbury, then Cramer's rule for the r × r inverse.

Conditioned on time, trained on the EMG alone: unsupervised, like ICA.

ABLATION · MU yield at F1 ≥ 0.95, 70- / 320-ch Per-channel affine correction instead of identity + low rank: −39% / −41%

Decoupled training: keeping ICA's identifiability

An ICA loss through a network loses identifiability, so the two losses never share a gradient.



$$\mathcal{L}_I(\mathbf{w}_j) = -\,\mathbb{E}\left[\mathbf{sign}(u_j(t))\,|u_j(t)|^e\right]$$

trains w only, compensator off: classical linear ICA, so the guarantee holds.

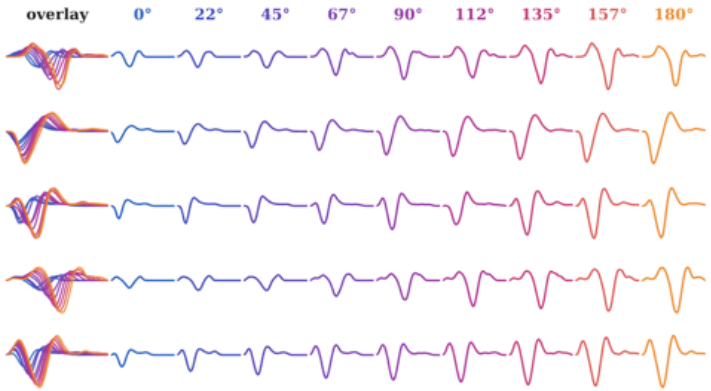
$$\mathcal{L}_N(\boldsymbol{\phi}_j) = \left(\sigma^2_{\text{spike}} + \max(\sigma^2_{\text{null}}, \bar{\sigma}^2_{\text{null}})\right)^{1/2}$$

two-class GMM fit to the compensated source. Trains φ only, w detached: tightens the spike and noise clusters of the chosen unit, never picks it.

ABLATION · MU yield at F1 ≥ 0.95, 70- / 320-ch Coupled gradients (both losses train both parts) instead: −20% / −19%

Neural peel-off: templates that move with the unit

Average where the unit looks stationary, then map back through the closed-form inverse at each spike.



One unit's waveform across the sweep: no single template fits all of them

1 Average spikes in the compensated signal A(t)x̃(t), where the unit looks the same every time

$$\mathrm{STA}_j^{\mathrm{comp}}(\tau) = \frac{1}{n_j}\sum_{i=1}^{n_j} \tilde{\mathbf{x}}^{\mathrm{comp}}(t_{j,i}+\tau)$$

2 Map it back through the inverse at each spike's own time

$$\hat{\mathbf{x}}_j^{(i)}(t_{j,i}+\tau) = A_j(t_{j,i}+\tau)^{-1}\,\mathrm{STA}_j^{\mathrm{comp}}(\tau)$$

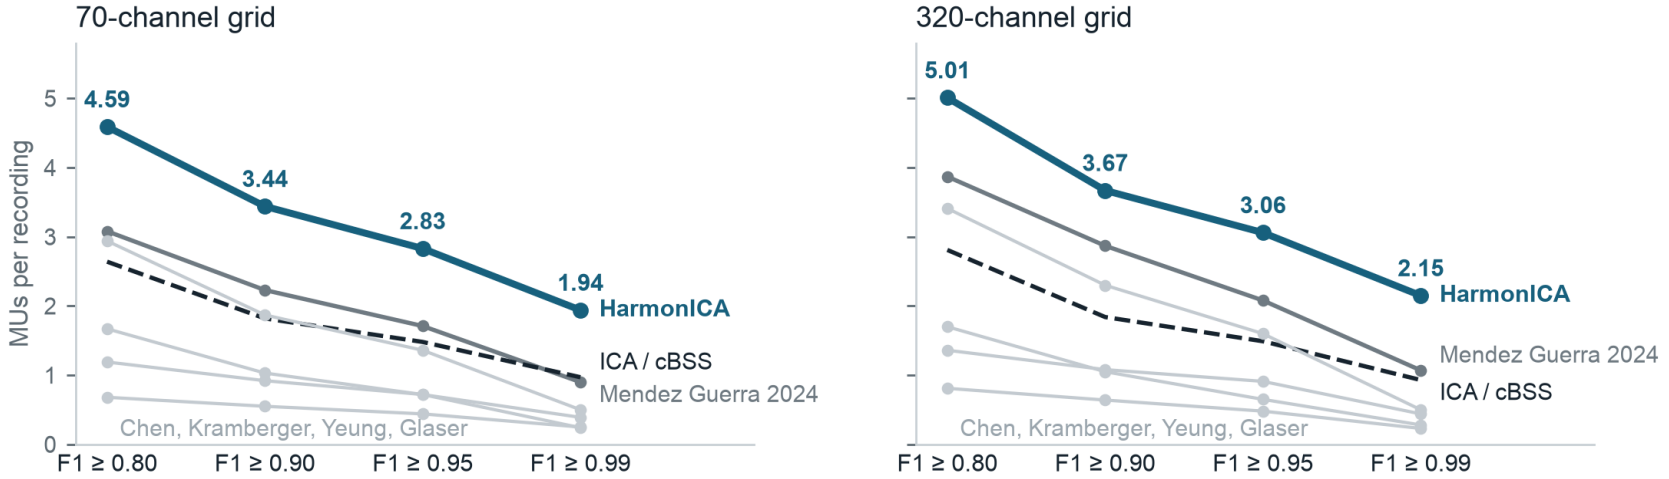
3 Subtract, then look for the next unit

$$\tilde{\mathbf{x}}(t_{j,i}+\tau) \leftarrow \tilde{\mathbf{x}}(t_{j,i}+\tau) - \hat{\mathbf{x}}_j^{(i)}(t_{j,i}+\tau)$$

ABLATION · MU yield at F1 ≥ 0.95, 70- / 320-ch No neural peel-off: −16% / −5%

Results: ahead of every adaptive baseline

Mean MUs recovered per recording at each F1 floor, MUniverse dynamic benchmark: 168 recordings per grid × 5 seeds.



At F1 ≥ 0.95: about 2× the ICA seed (2.83 vs 1.48; 3.06 vs 1.49), ahead of the best adaptive baseline (1.71; 2.08).
At F1 ≥ 0.99: about 2× the best adaptive baseline, Mendez Guerra 2024 (1.94 vs 0.90; 2.15 vs 1.07).
With the joint angle: 515 → 571 MUs in total on the 320-channel grid, against 249 for ICA.

IMPERIAL

Thank you

github.com/AlexKClarke/harmonica_neurips