



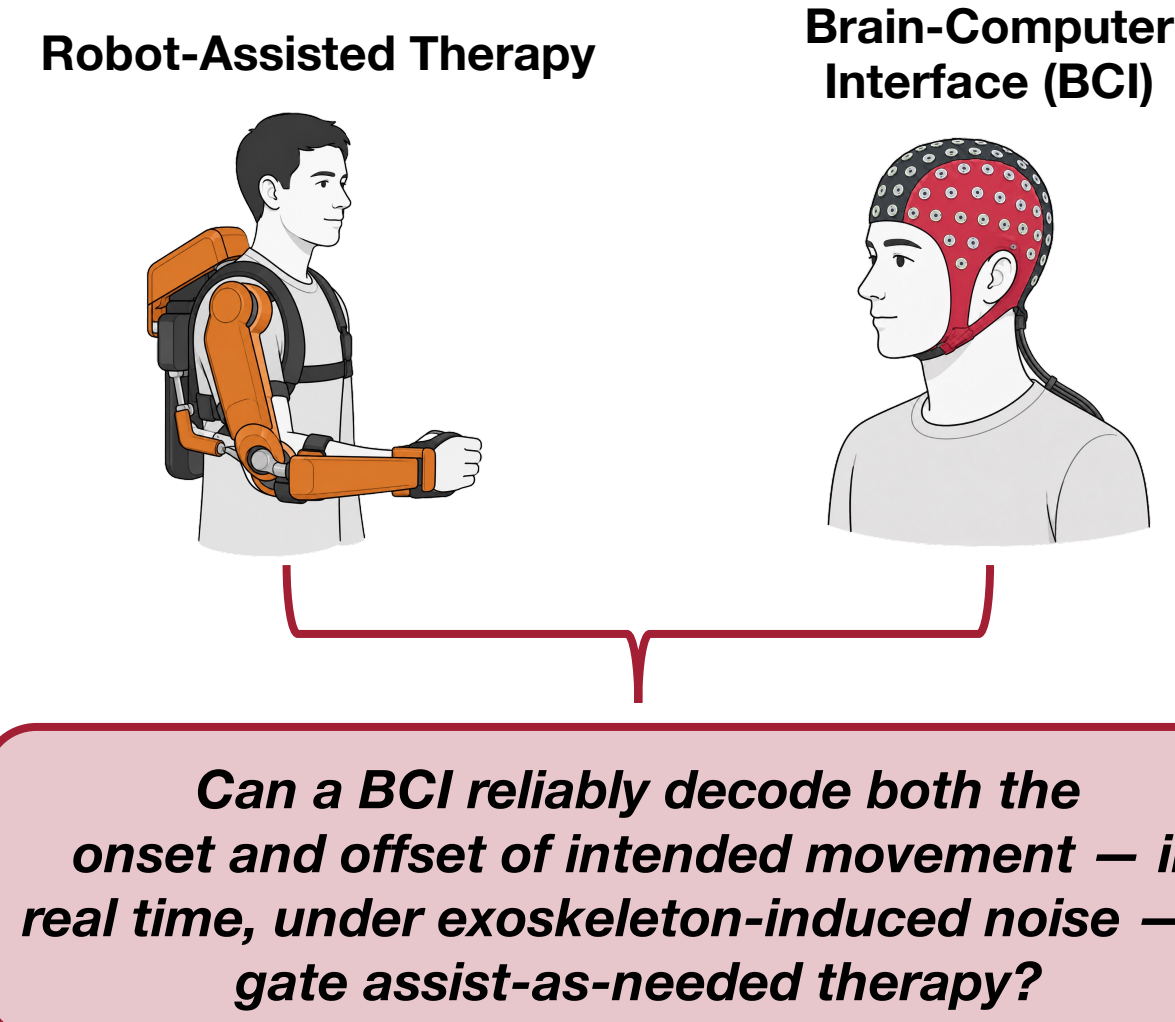
Kanishka Mitra<sup>1</sup>, Satyam Kumar<sup>2</sup>, Frigyes Samuel Rácz<sup>3</sup>, Deland Liu<sup>2</sup>, Ashish D. Deshpande<sup>4,5</sup>, José del R. Millán<sup>2,3,6</sup>

<sup>1</sup>Dept. of Electrical Engineering & Computer Science, Massachusetts Institute of Technology  
<sup>2</sup>Chandra Dept. of ECE · <sup>3</sup>Dept. of Neurology · <sup>4</sup>Walker Dept. of ME · <sup>5</sup>Dept. of BME, The University of Texas at Austin  
<sup>6</sup>Meta Reality Labs Research



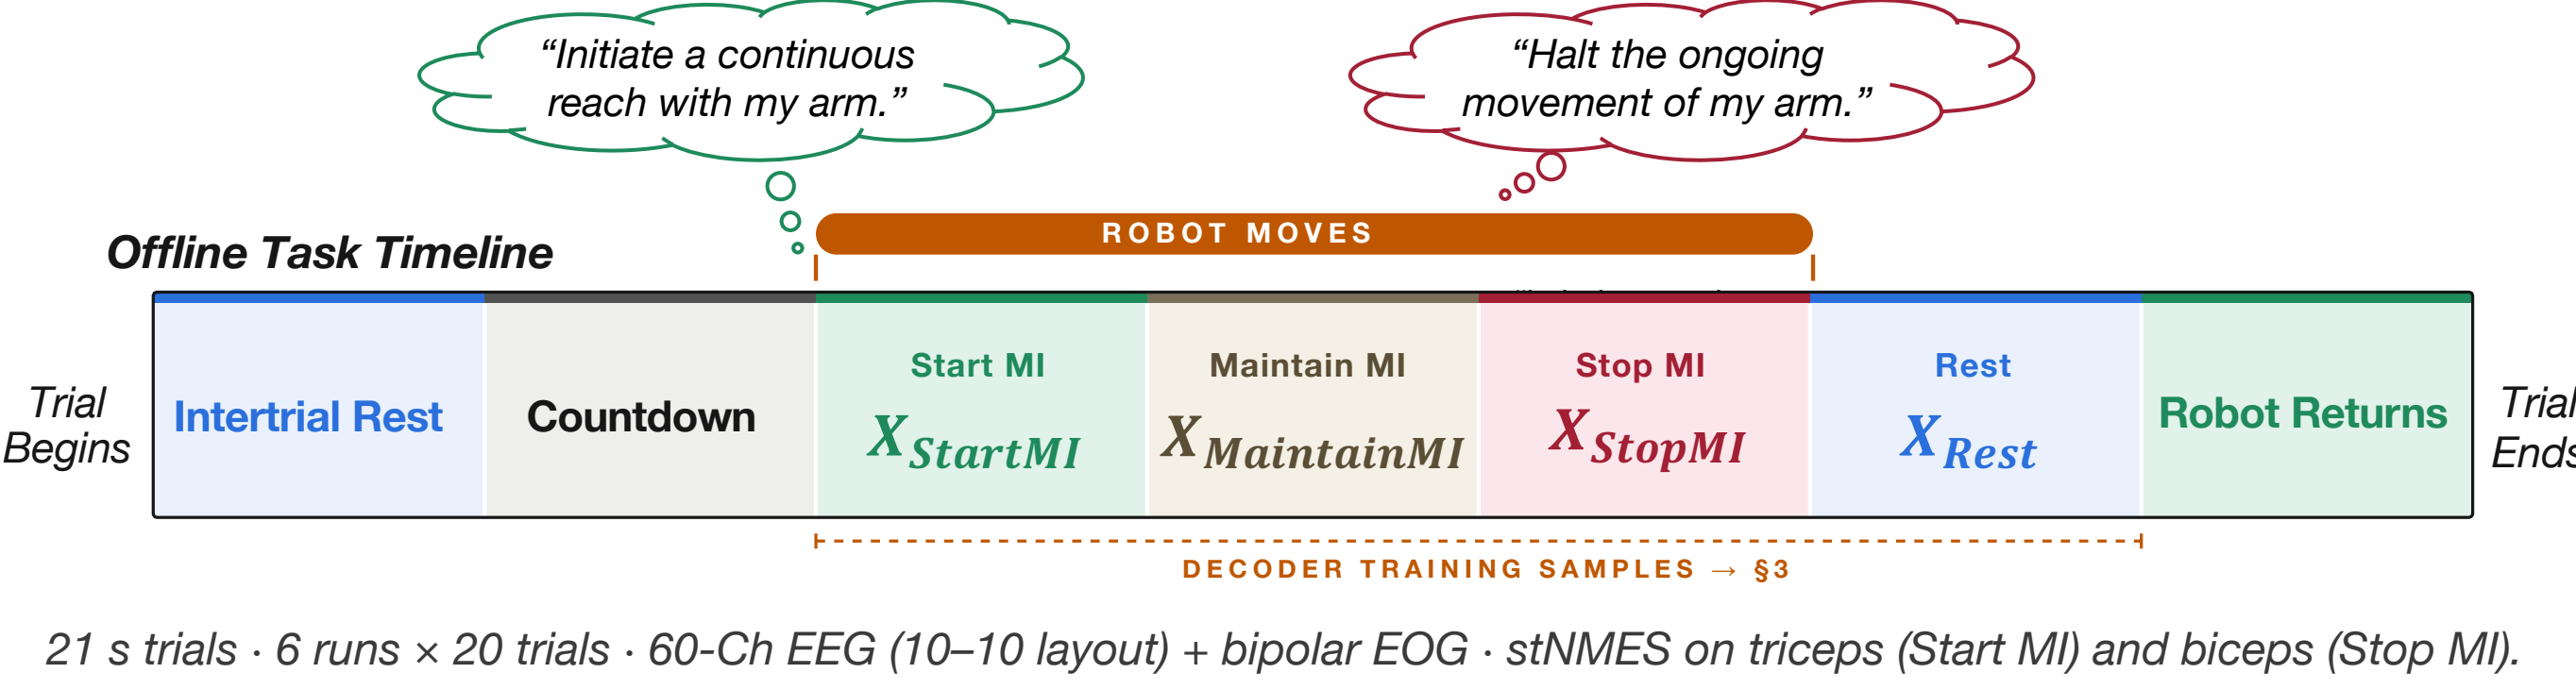
## 1 Motivation

- Robot-assisted therapy acts at the **limb level** – engaging impaired neural circuits only indirectly [1].
- Closing the loop between **intent** and **assistance** is the missing ingredient for neuroplasticity-targeted rehabilitation [2].
- Motor imagery (MI) gives a non-invasive route to intention-contingent control [3].
- Most MI-BCIs decode only **onset**, but reliable **offset** matters for assist-as-needed therapy [4].



## 2 Offline Data Collection

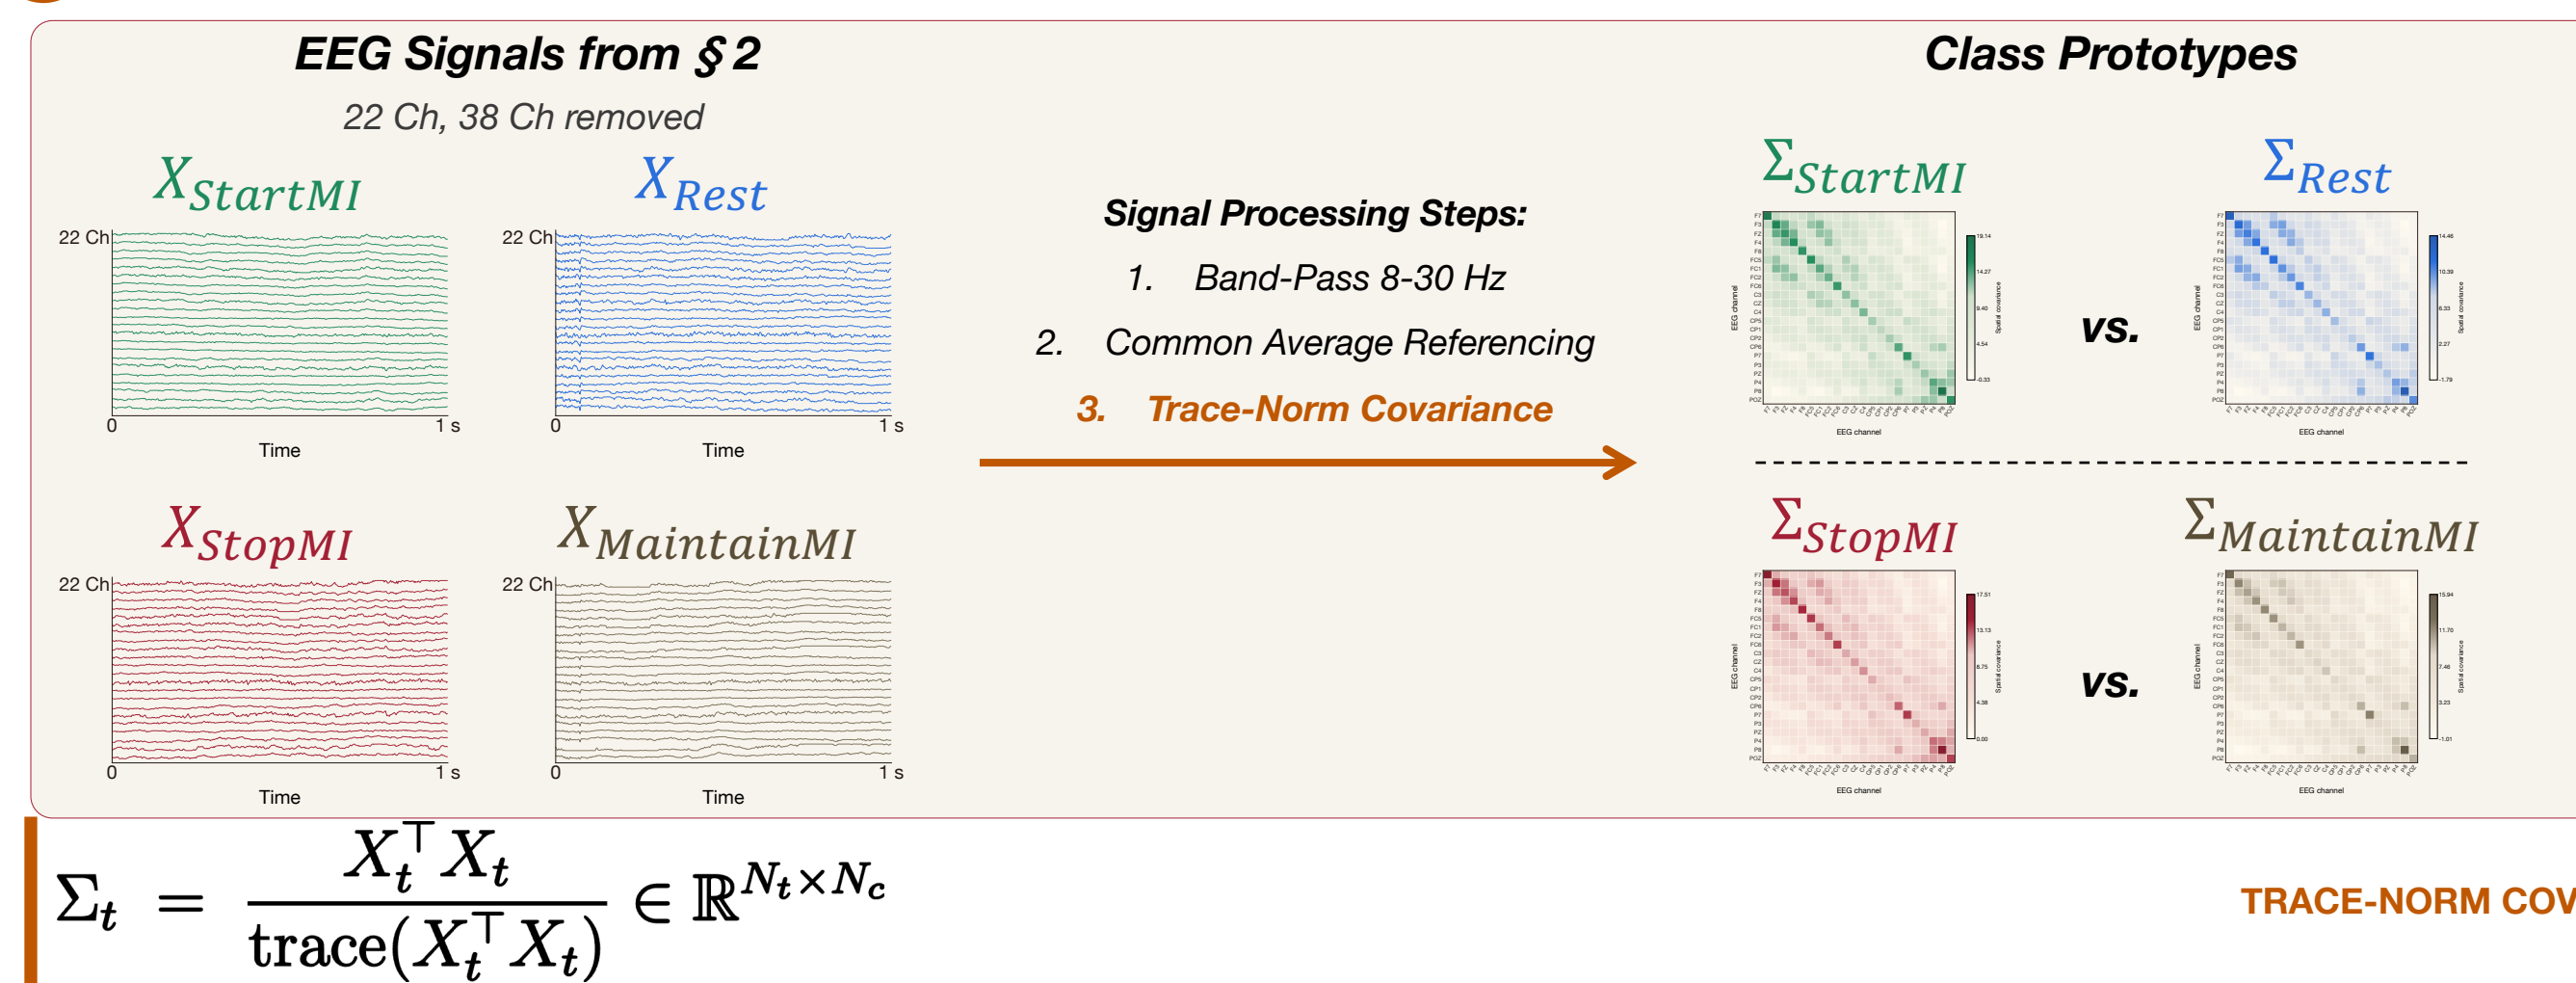
Eight healthy participants completed a single calibration session, performing MI **onset** and **offset** while donned in the Harmony SHR® exoskeleton [5].



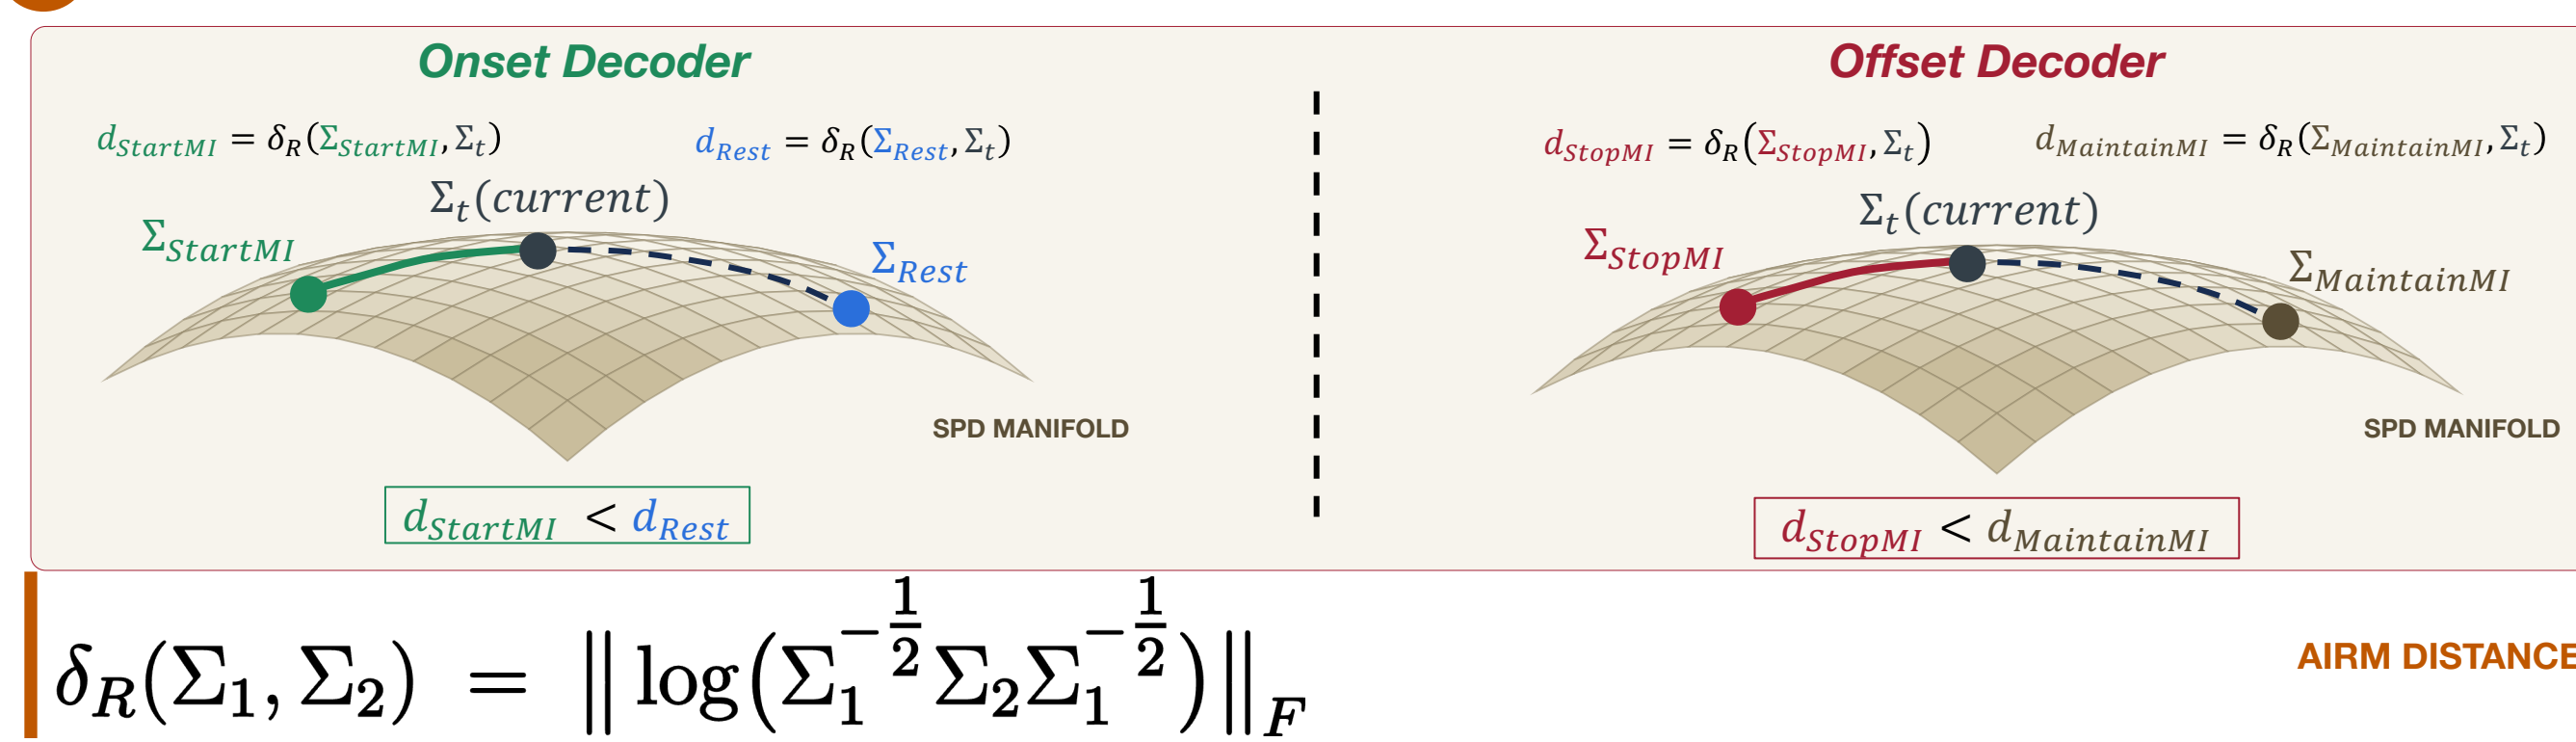
## 3 Riemannian Decoding Pipeline

Each 1s EEG window becomes a point on a curved space of covariance matrices. Hierarchical **minimum-distance** decoders ask, in real time: "Is the user trying to start?" then "Are they trying to stop?" [6]

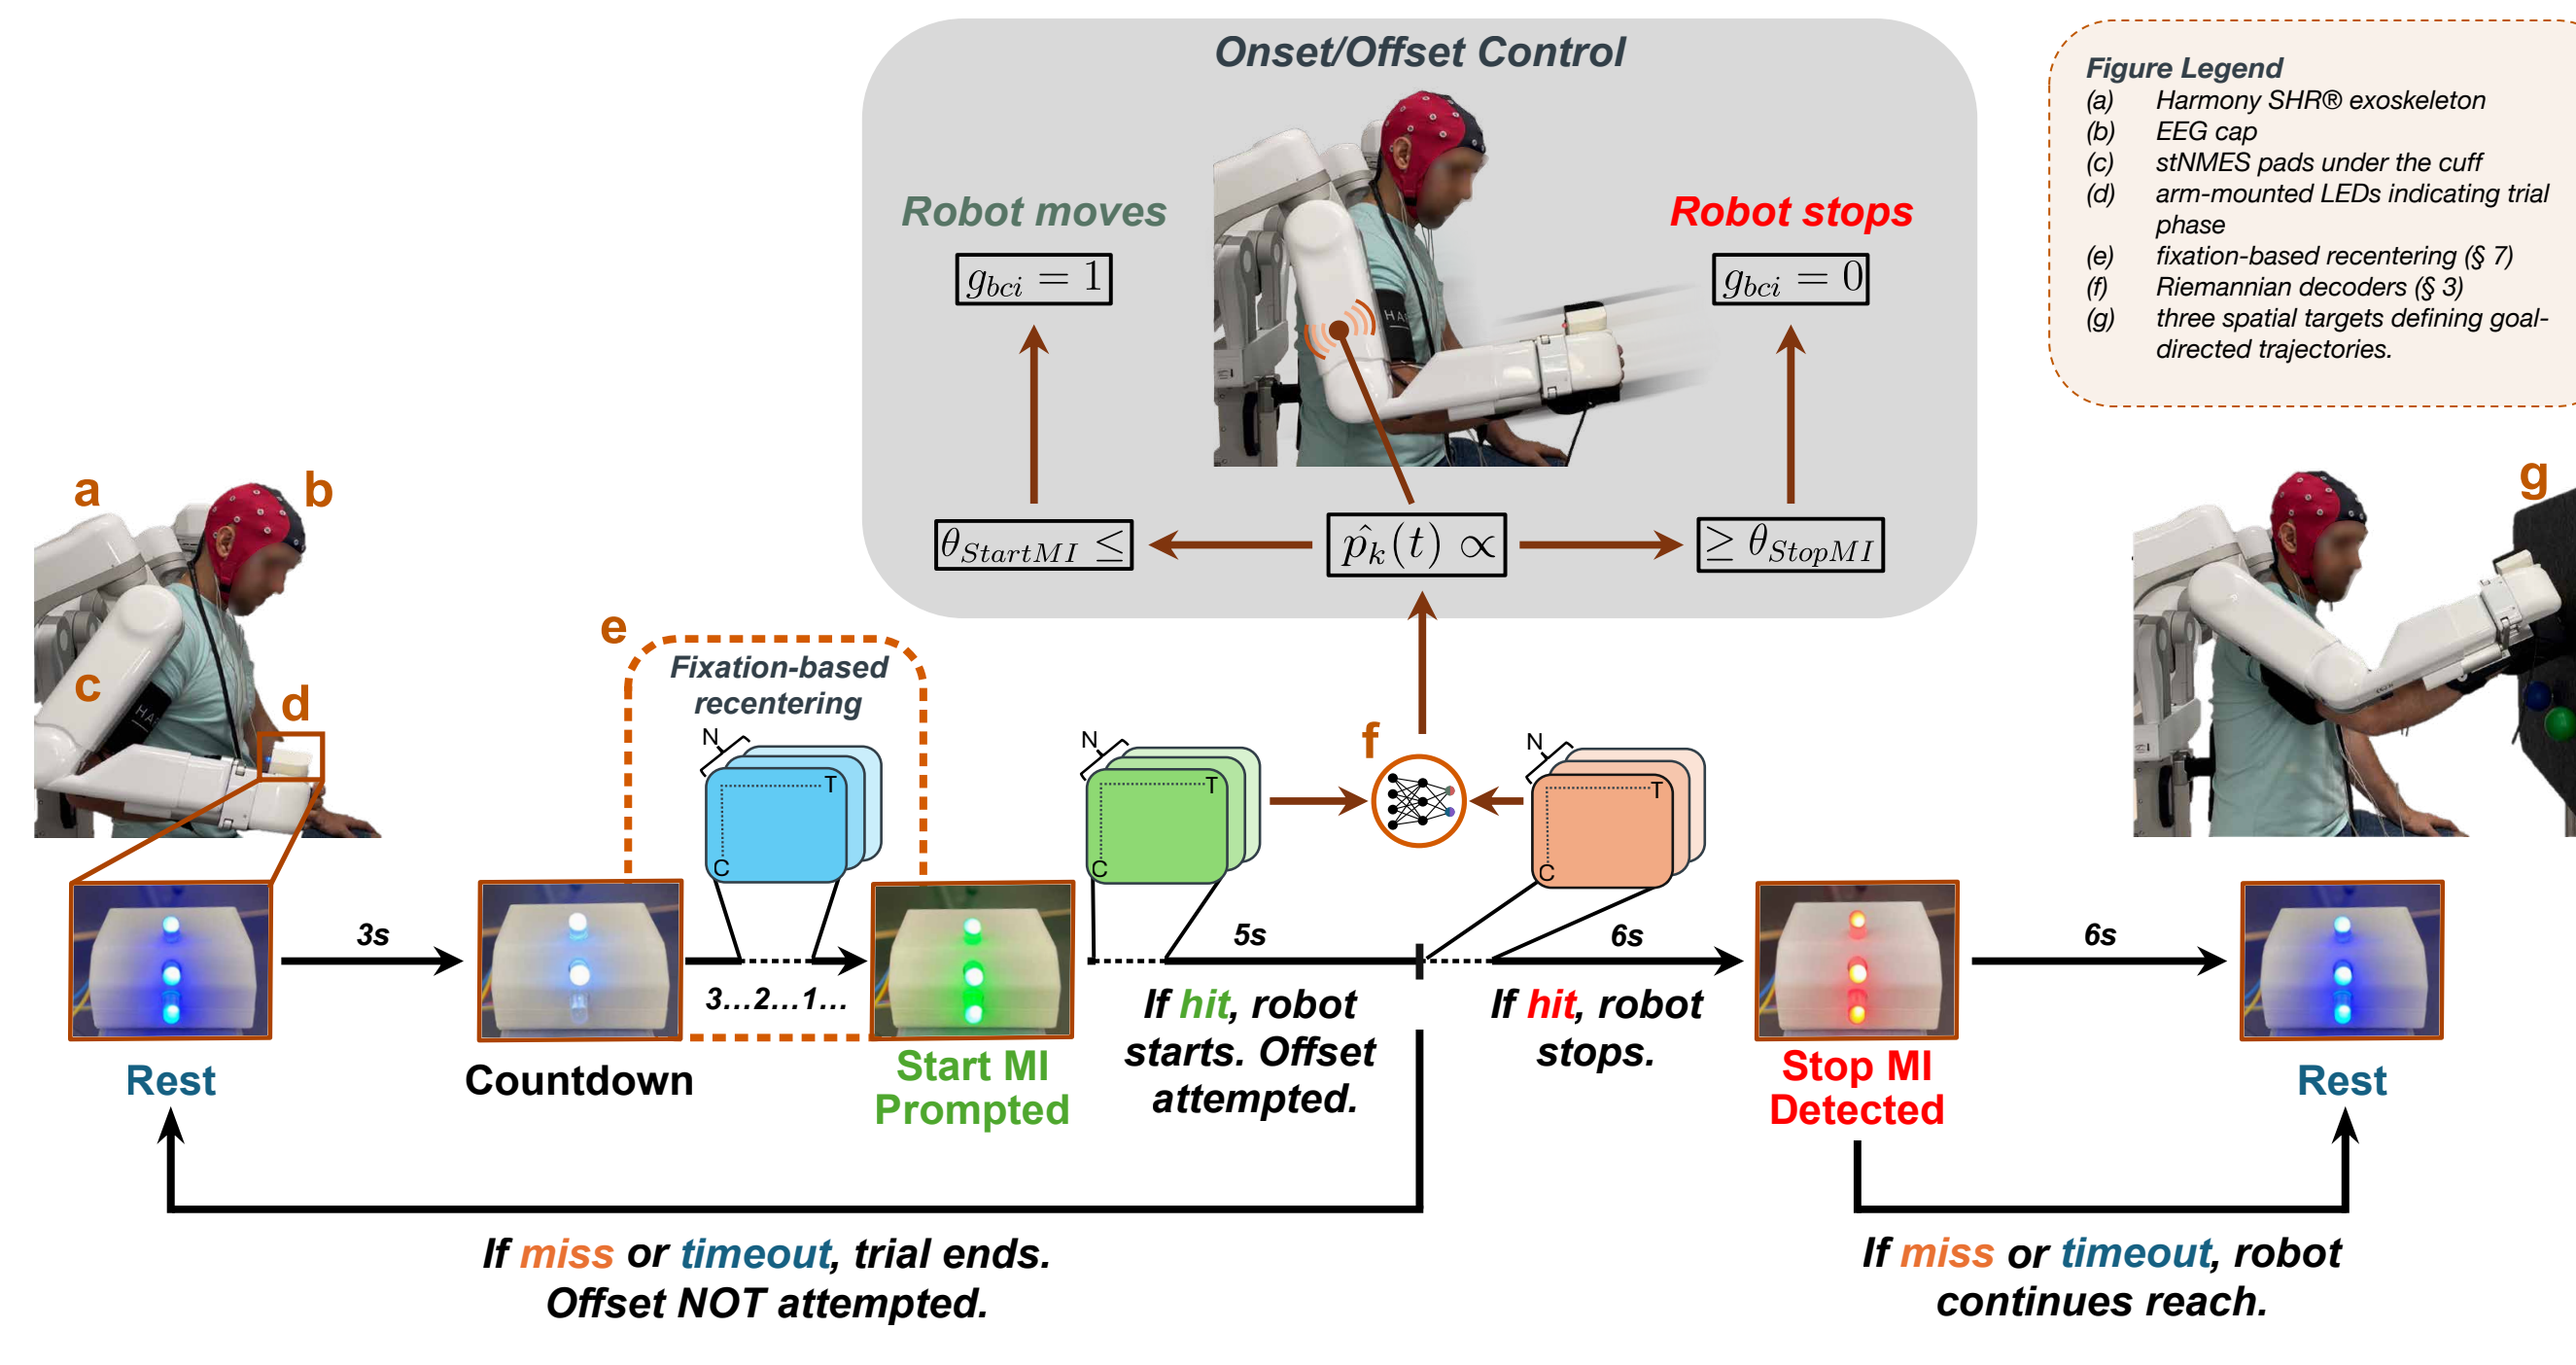
### 1 From EEG signals to spatial covariances



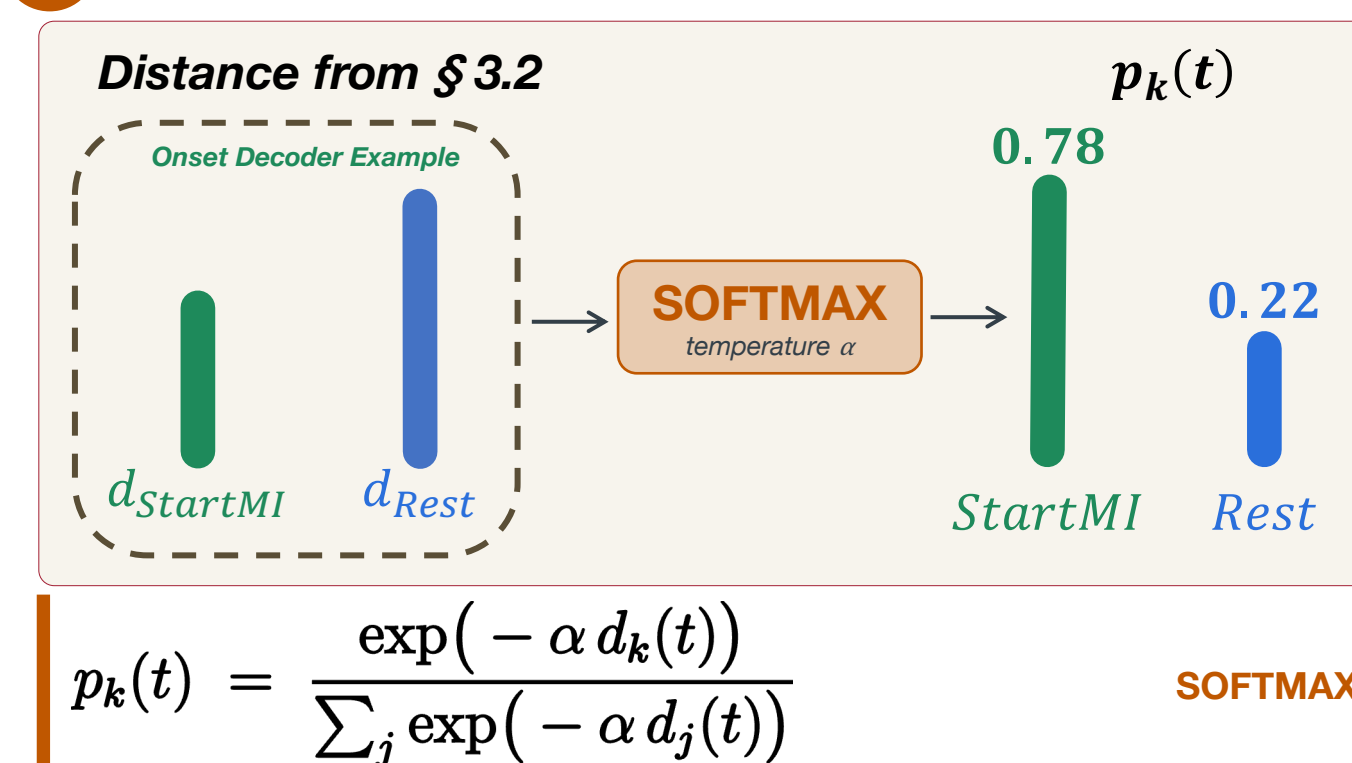
### 2 Distance on the SPD manifold



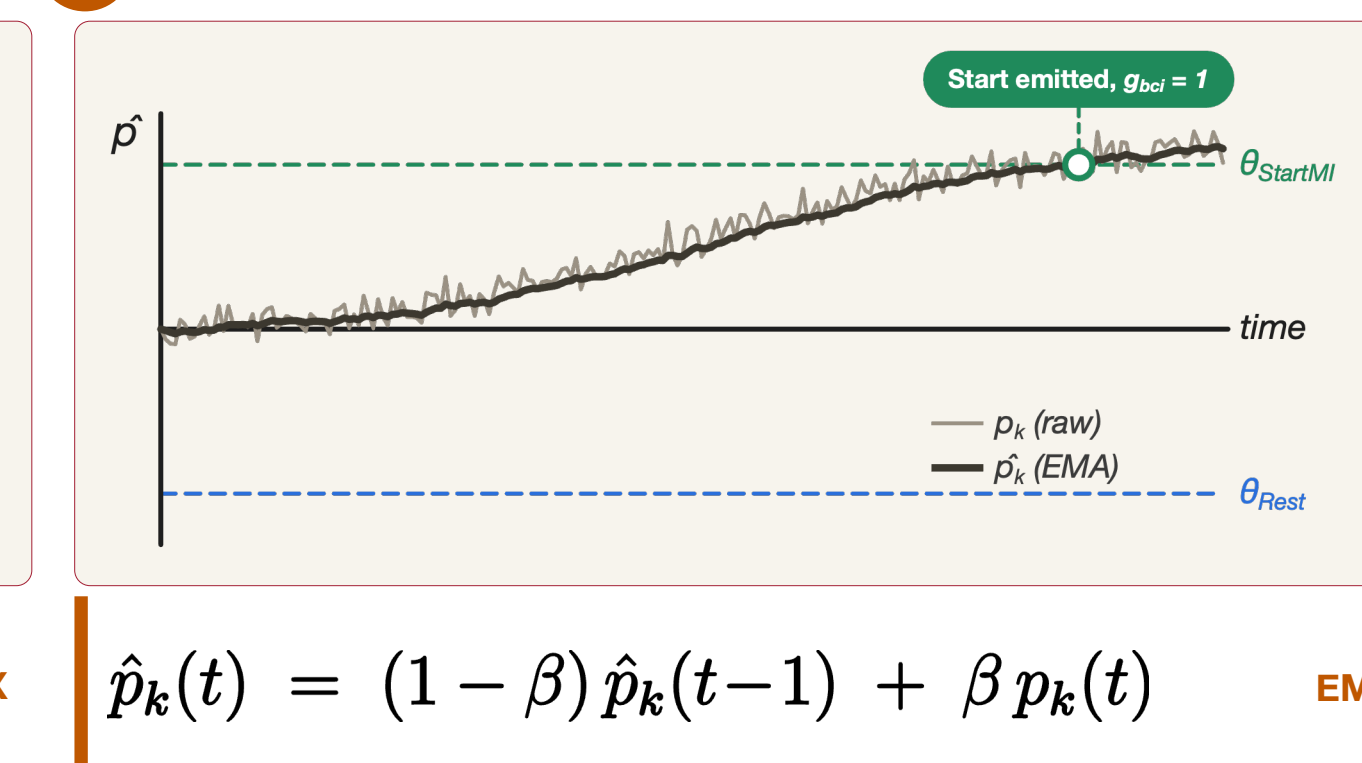
## 4 Experimental Setup & Online Control



### 3 Distance to Posterior Estimation



### 4 Temporal Smoothing & Decision Rule



## 5 Results

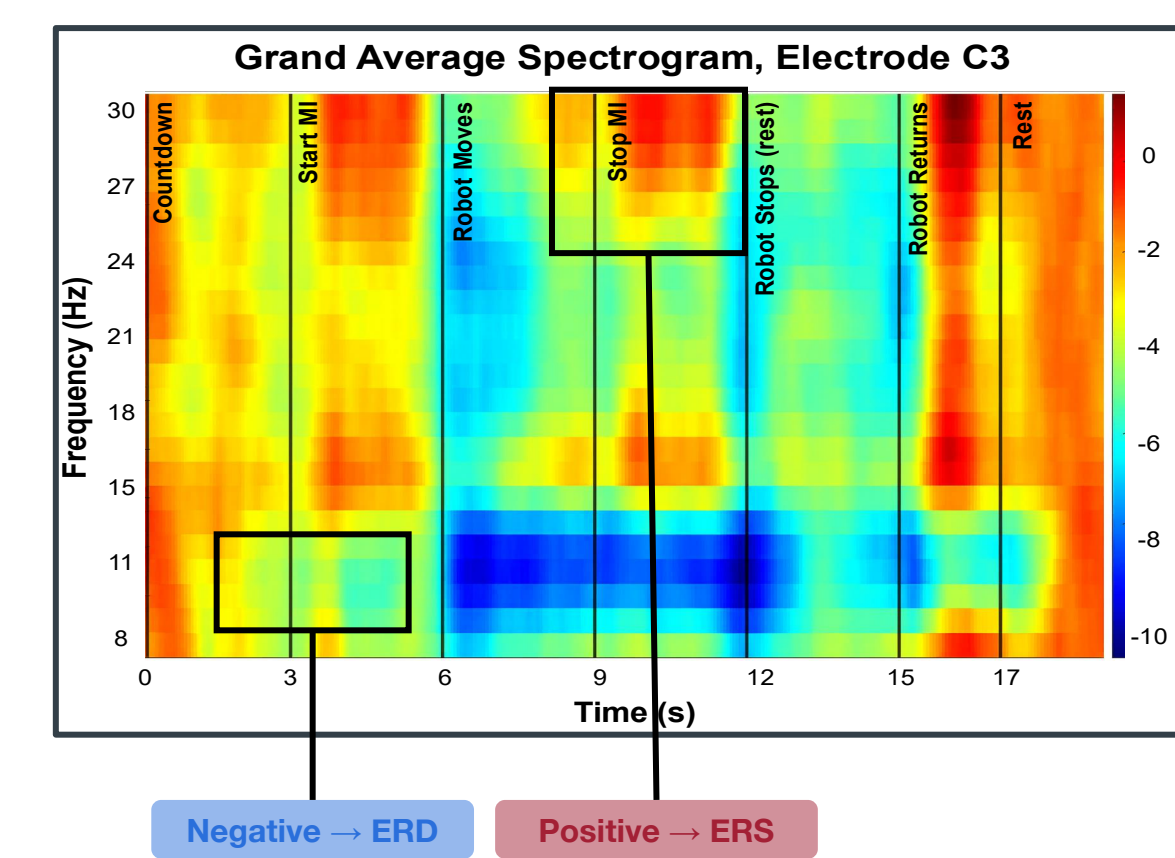
### Offline Results

To see **when** and **where** the brain modulates, we track each frequency's power:


$$ERD/ERS(f, t) = \log_{10} A(f, t) / B(f)$$

$A(f, t)$  = live spectral power ·  $B(f)$  = baseline power before the countdown.

Robust  **$\mu$ -band ERD** locked to Start MI and a clear  **$\beta$ -band rebound** after Stop MI, canonical MI signatures persist despite movement [4], evoked sensory feedback and exoskeleton-induced noise.



### Online Results



A single non-invasive decoder delivers both reliable start and stop commands to a rehabilitation exoskeleton – closing the loop in real time, under real instrumental noise and passive arm motion.

## 6 Recentering Bias

Recentering corrects day-to-day EEG drift, but our online protocol only ever sees the **target class** (Start or Stop MI). Building the reference from those samples alone **pulls them toward the offline mean**.

1. A margin that scores each sample  
For every online covariance, take its two Riemannian distances (§ 3.2) and form a single margin:

$$m = d_{neg} - d_{pos}$$

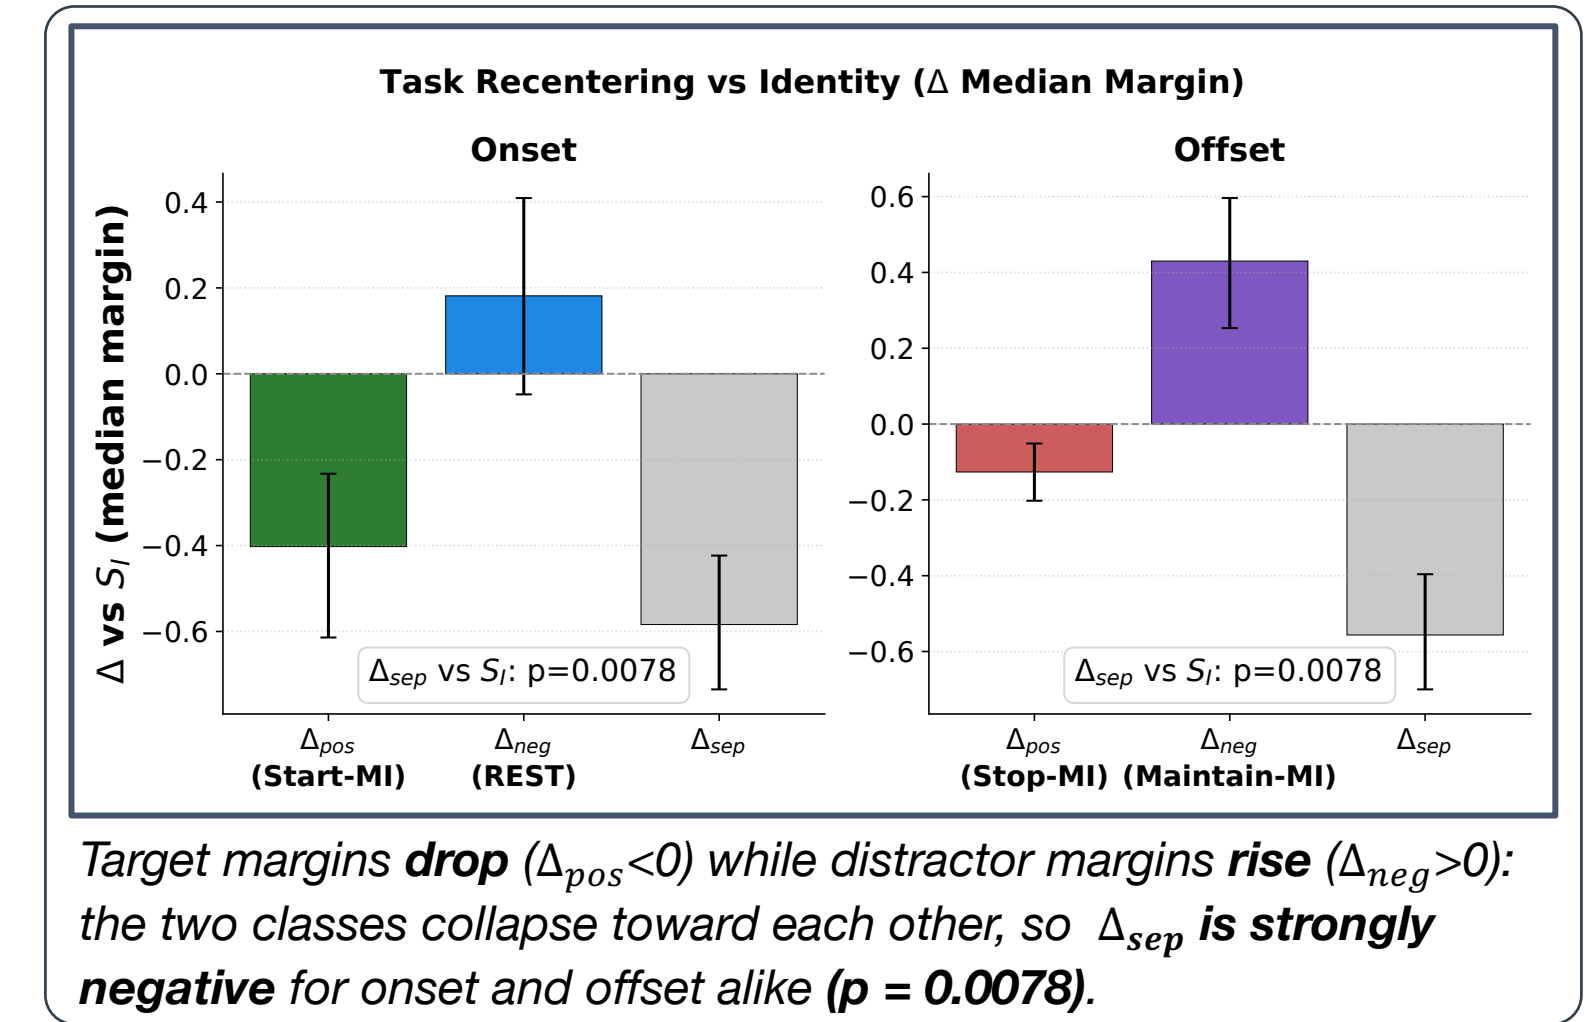
$m > 0$  closer to the positive class (Start / Stop)  
 $m < 0$  closer to the negative class (Rest / Maintain)

2. Compare task vs identity reference  
Measure how each class's median margin shifts when we switch from a neutral identity reference to task-based recentering:

$$\Delta_{pos} = \bar{m}_{pos}(task) - \bar{m}_{pos}(identity)$$

$$\Delta_{neg} = \bar{m}_{neg}(task) - \bar{m}_{neg}(identity)$$

$$\Delta_{sep} = \Delta_{pos} - \Delta_{neg}$$



Task-based recentering makes **Start look more like Rest** (and Stop like Maintain) – squeezing the classes into the ambiguous middle, where a single fixed threshold can no longer tell them apart.

## 7 Fixation-Based Recentering

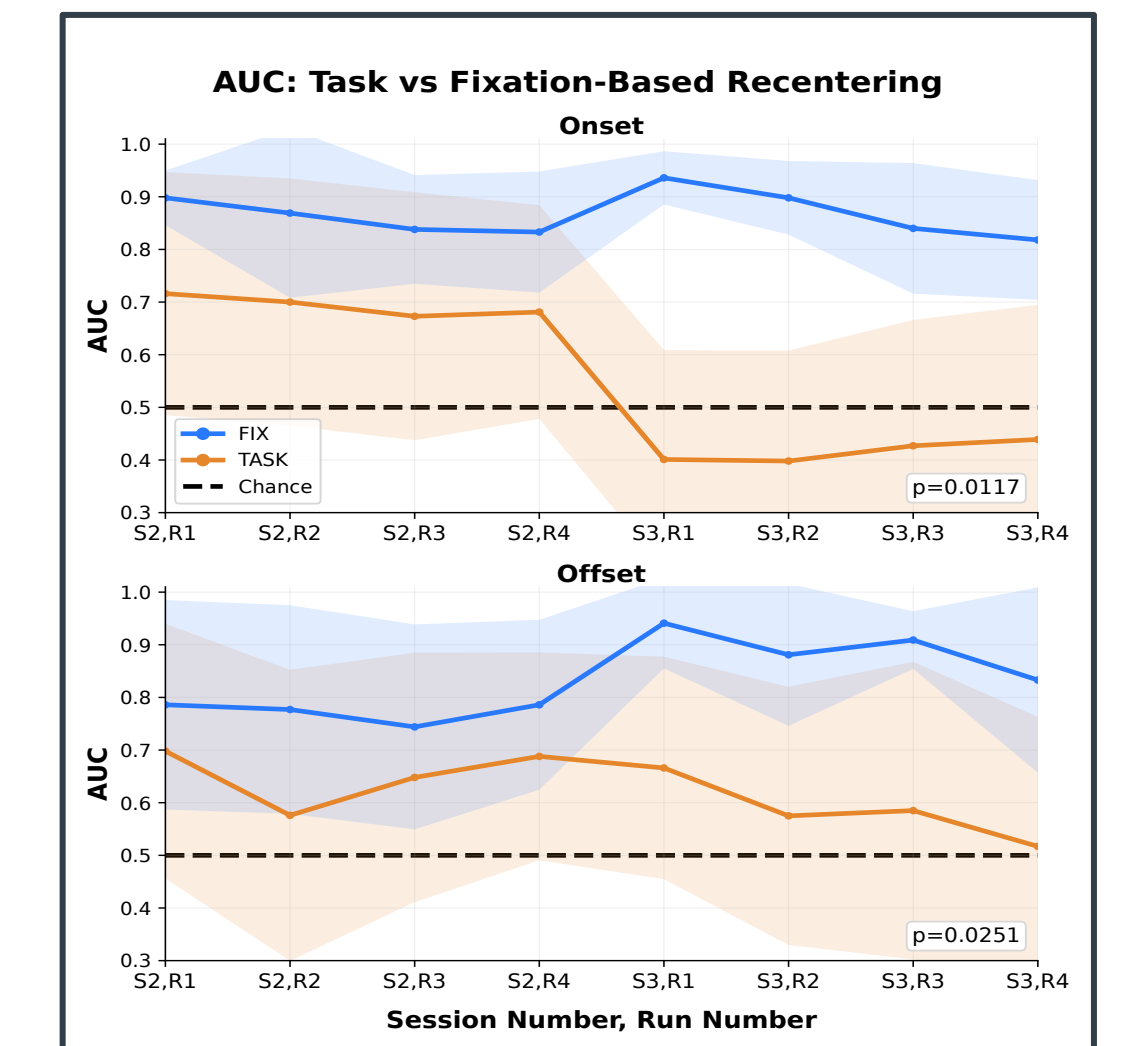
The fix: build the drift reference from the **pre-cue fixation period** - EEG that belongs to neither class. This tracks day-to-day drift without ever touching the command classes, so it can't pull them together [6].

KEY IDEA: CLASS-AGNOSTIC REFERENCE  
Conventional recentering whitens using **target-class** samples. We instead whiten using only **fixation** samples, neutral by construction, preserving the Start→Stop geometry.

- Covariances + **log-Euclidean mean** of the fixation window.
- Trim outliers** on the SPD manifold: drop the farthest 20% (blinks, motion), recompute.
- Shrink toward identity** for a neutral reference  
 $\log S_{fix} = (1 - \alpha) \log S_{raw}$
- Eigenvalue shrinkage** for conditioning, so no channel dominates  
 $S_{fix} = V \text{diag}((1 - \lambda) \ell + \lambda \ell^T) V^T$
- Smooth across runs** in the log domain to slowly track the drift.

ONSET · THRESHOLD-FREE AUC  
0.55 → 0.87  
+56% · p = 0.0117

OFFSET · THRESHOLD-FREE AUC  
+34% gain  
p = 0.0251



A class-agnostic reference restores separability and suppresses the unilateral bias - a more symmetric, more stable decision boundary that holds within runs and across days, with less recalibration.

## 8 Takeaways & Future Work

- Intention-gated start & stop**  
First online demonstration of both initiating assistance and terminating it **mid-trajectory** on a clinical upper-limb exoskeleton from non-invasive motor imagery, in real time.
- Reliable across sessions**  
A geometry-aware Riemannian decoder with subject-specific thresholds and proprioceptive feedback held up under robot motion across multiple days.
- Diagnosed & fixed a hidden bias**  
Revealed and quantified the task-recentering bias; a class-agnostic fixation reference restored separability and cut unilateral bias within & across days.

An important feasibility step toward MI-BCI rehabilitation, bridging motor-imagery transition decoding and real-time robot control, before clinical translation.

## 9 References

- [1] L. Tonin & J. d. R. Millán, *Annu. Rev. Control Robot. Auton. Syst.*, 2021. [2] A. Biasucci et al., *Nat. Commun.*, 2018. [3] G. Pfurtscheller & C. Neuper, *Proc. IEEE*, 2001. [4] B. Orset et al., *bioRxiv*, 2021. [5] K. Mitra et al., *IEEE/RSJ IROS*, 2023. [6] S. Kumar, F. Yger & F. Lotte, *Int. Winter Conf. BCI*, 2019.