Test-Time Adaptation of Spiking Neural Networks for Intracortical Neural Decoding using Membrane Potential Alignment

Guangzhi Tang[(1)]

[(1)] Maastricht University

Intracortical brain–computer interfaces degrade over days as electrode drift, firing-rate fluctuations, and glial encapsulation shift the statistics of recorded neural populations. Existing unsupervised adaptation methods such as NoMAD rely on variational autoencoders with gated RNNs, architectures too computationally demanding for the power budgets of implantable hardware. Spiking neural networks offer energy-efficient decoding but lack test-time adaptation methods suited to non-stationary neural recordings.

We introduce Membrane Potential Alignment (MPA), a test-time adaptation method that realigns a pretrained SNN decoder to shifted recordings by matching membrane potential distributions via KL divergence. Membrane potentials are modeled as multivariate Gaussians, and adaptation minimizes the divergence between source-session and target-session distributions. By restricting weight updates to low-rank (LoRA) matrices in the first hidden layer, MPA adapts fewer than 9% of parameters while keeping the rest of the decoder frozen. Two complementary compatibility metrics, representational similarity and Grassmann distance, flag sessions where adaptation is unlikely to succeed.

Evaluated on 23 sessions of a non-human primate reaching task spanning over one month, MPA improves mean R^2 from 0.09 (no adaptation) to 0.33, competitive with NoMAD while using a simpler fully connected SNN architecture operating at finer temporal resolution (4 ms vs. 20 ms). Effective adaptation is achieved with as little as 1.28 seconds of unlabeled target data. The compatibility metrics correctly identify sessions where both MPA and NoMAD fail, providing a practical safeguard against silent decoder degradation.

These results establish SNN-based test-time adaptation as a viable path toward recalibration-free brain–computer interfaces deployable on neuromorphic hardware. Future work targets on-device continual adaptation and multi-subject pretraining for more robust initial representations.

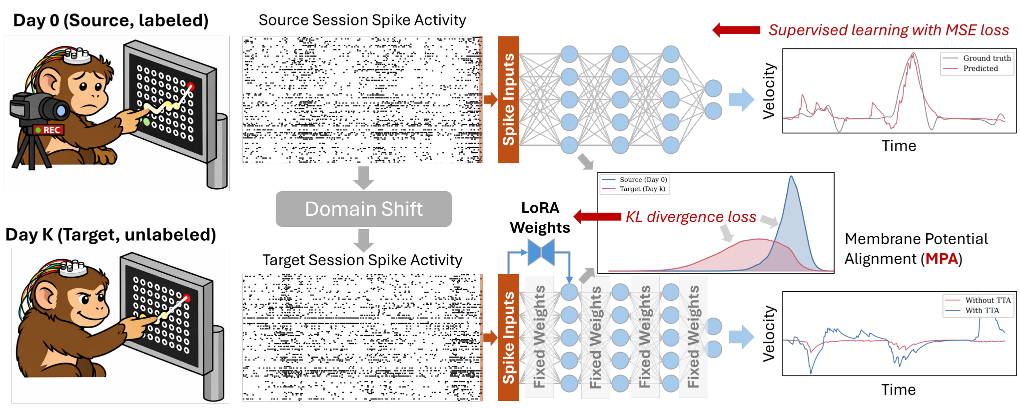

Figure 1: Overview of the proposed cross-day SNN test-time adaptation pipeline. An SNN pretrained on day 0 is applied to a later target session, day k, where cross-session drift causes a mismatch in neural inputs and labels are unavailable. We therefore keep the decoder fixed and adapt the SNN's LoRA weights by aligning source and target membrane potential distributions through membrane potential alignment (MPA), improving target-session decoding without supervised behavior labels.