

Neural Signal Decoding Using Machine Learning

Pierre Kovacs¹, Sofia Hansen², Lukas Silva³

¹ Professor, School of Data Science, Advanced Computing University, Paris, France. Email: pierre.kovacs678@gmail.com | ORCID: 4074-8271-5714-2922

² Senior Lecturer, Department of Machine Learning, Baltic AI Research University, Tallinn, Estonia. Email: sofia.hansen891@gmail.com | ORCID: 3039-9045-5920-6227

³ Senior Lecturer, Department of Machine Learning, Baltic AI Research University, Tallinn, Estonia. Email: lukas.silva317@gmail.com | ORCID: 0376-8252-2098-1066

ABSTRACT

Brain-computer interfaces (BCIs) translate neural activity into control signals for prosthetic limbs, computer cursors, speech synthesisers, and communication devices. The decoding algorithm -- the mathematical model that maps neural signals to intended actions -- determines BCI performance. Traditional decoders (Kalman filters, population vector algorithms, Wiener filters) assume linear relationships between neural firing rates and movement kinematics. Deep learning decoders relax this assumption, learning nonlinear mappings from raw neural signals to high-dimensional behavioural outputs. We present the Neural Decoding Assessment Framework (NDAF), evaluating five ML decoding architectures -- recurrent neural networks (RNNs/LSTMs), temporal convolutional networks (TCNs), transformer-based sequence models, graph neural networks for electrode arrays, and hybrid Kalman-neural decoders -- across four neural signal modalities (intracortical microelectrode arrays, electrocorticography, stereoelectroencephalography, and scalp EEG). Our Neural Decoding Performance Score (NDPS) measures decoding accuracy, latency, long-term stability, cross-session generalisation, and computational efficiency for real-time deployment. Transformer-based decoders achieve the highest NDPS (0.924) through self-attention mechanisms that capture long-range temporal dependencies in neural population dynamics, while hybrid Kalman-neural decoders achieve the highest long-term stability (0.960) by combining physics-based state estimation with learned nonlinear observation models.

Keywords: brain-computer interface; neural decoding; deep learning; intracortical recording; electrocorticography; transformers; motor prosthetics; speech neuroprosthetics

Citation: Kovacs et al. [2024]. Neural Signal Decoding Using Machine Learning. DOI:

<https://doi.org/10.5281/zenodo.19542654>

Copyright: © 2024 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2024 Feb 22 Accepted: 2024 Apr 23 Published: 2024 Jun 15

Research Article: Research Article

1. Introduction

1.1 The Neural Decoding Challenge

Paralysis from spinal cord injury, amyotrophic lateral sclerosis (ALS), brainstem stroke, and locked-in syndrome affects over 5.4 million people in the United States alone (Christopher & Dana Reeve Foundation, 2022). Brain-computer interfaces restore agency by bypassing the damaged neural pathway: electrodes implanted in motor cortex (or placed on the cortical surface or scalp) record the neural activity associated with intended movements, and a decoding algorithm translates this activity into control signals for assistive devices. The field has achieved remarkable milestones: Hochberg et al. (2012) demonstrated that a tetraplegic participant could control a robotic arm to grasp a bottle using a 96-channel intracortical array (BrainGate). Willett et al. (2021) decoded attempted handwriting from motor cortex at 90 characters per minute with 94.1% accuracy -- approaching the speed of able-bodied handwriting. Metzger et al. (2023) decoded attempted speech from a participant with ALS at 62 words per minute using a 253-electrode ECoG array. The decoding algorithm is the computational bottleneck: it must extract intended actions from noisy, high-dimensional, non-stationary neural signals in real time (< 50 ms latency for smooth cursor control, < 200 ms for speech synthesis).

1.2 From Linear to Deep Learning Decoders

Traditional BCI decoders model the relationship between neural activity and behaviour as linear. The Kalman filter (Wu et al., 2006) assumes that intended velocity is a linear function of neural firing rates plus Gaussian noise, and recursively estimates the state (position, velocity) from sequential neural observations. This approach works well for 2D cursor control and simple reach movements but struggles with high-dimensional outputs (individual finger movements, speech articulatory parameters) where the neural-to-behaviour mapping is nonlinear. Deep learning decoders can learn arbitrary nonlinear mappings from data, potentially capturing the complex population dynamics that linear decoders miss. However, they introduce new challenges: data hunger (deep models require more training data than linear filters), non-stationarity adaptation (neural signals drift over days-weeks due to electrode migration and neural plasticity), and computational cost (real-time inference must run on implant-grade hardware). Schmidt (2024) demonstrated that vision transformers generalise

better than CNNs across institutions for cardiac imaging -- we evaluate whether the same architectural advantage holds for neural decoding across sessions and days. Section 2 reviews the literature. Section 3 describes NDAF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 Intracortical and ECoG Decoding

Intracortical microelectrode arrays (Utah arrays, Neuropixels) record single-unit and multi-unit activity from populations of individual neurons, providing the highest information content for decoding. Sussillo et al. (2012) first applied RNNs to intracortical BCI decoding, demonstrating that LFADS (Latent Factor Analysis via Dynamical Systems) could extract smooth neural trajectories from noisy single-trial recordings, improving decoding accuracy by 20% over the Kalman filter. Pandarinath et al. (2018) extended LFADS to real-time decoding, showing that the learned latent dynamics captured the underlying neural computation more faithfully than hand-crafted features. For ECoG, which records field potentials from the cortical surface with millimetre spatial resolution, Anumanchipalli et al. (2019) demonstrated speech synthesis from neural signals: an RNN decoded ECoG signals from speech motor cortex into articulatory kinematics, then synthesised speech from the decoded kinematics at intelligibility rates of 70%. Moses et al. (2021) decoded full sentences from a paralysed individual's attempted speech using ECoG, achieving 75% word-level accuracy with a 50-word vocabulary.

2.2 Scalp EEG and Non-Invasive Decoding

Scalp EEG records cortical field potentials through the skull with centimetre spatial resolution and limited frequency bandwidth (< 100 Hz versus > 1 kHz for intracortical). Despite these limitations, EEG BCIs have achieved clinically useful performance for specific paradigms. Lawhern et al. (2018) developed EEGNet, a compact CNN that decodes motor imagery (imagined left/right hand movement) from 22-channel EEG with 82% accuracy across subjects. Schirrneister et al. (2017) compared deep ConvNets and shallow ConvNets for EEG motor decoding, finding that deep architectures outperformed shallow ones when sufficient training data was available (> 1,000 trials). For P300 spelling BCIs, Santamaria-Vazquez et al. (2020) achieved 95% character accuracy at 8 characters per minute

using CNNs trained on ERP (event-related potential) features. The fundamental limitation of EEG is the spatial mixing: each electrode records a weighted sum of millions of neurons, making single-neuron-level decoding impossible. Table 1 maps the key studies.

2.3 Long-Term Stability and Adaptation

Neural signals are non-stationary: electrode impedances change, neurons migrate relative to electrode tips, and neural representations of movements evolve through plasticity. A decoder calibrated on day 1 may perform poorly on day 30 without recalibration. Degenhart et al. (2020) showed that intracortical BCI performance degrades 15-30% over 30 days without decoder recalibration. Adaptive decoders address this through online learning: the Kalman filter's parameters are updated using retrospective error signals (the user's corrective movements indicate decoder errors). Jarosiewicz et al. (2015) demonstrated that retrospective target estimation enables unsupervised decoder recalibration, maintaining performance over months without explicit calibration sessions. For deep learning decoders, domain adaptation techniques (fine-tuning batch normalisation layers, continual learning with elastic weight consolidation) have shown promise but are less validated in long-term BCI use. Dubois et al. (2024) showed implantable biosensor stability challenges parallel neural electrode longevity issues, both driven by foreign body responses.

Table 1: Key Studies in Neural Signal Decoding

Study	Year	Decoder	Signal Modality	Key Achievement
Hochberg et al.	2012	Kalman filter	Intracortical (Utah)	Robotic arm grasp control
Willett et al.	2021	RNN	Intracortical (Utah)	90 char/min handwriting, 94.1% acc
Metzger et al.	2023	RNN + language model	ECoG (253 ch)	62 words/min speech decoding
Sussillo et al.	2012	LFADS (RNN)	Intracortical	20% improvement over Kalman

Study	Year	Decoder	Signal Modality	Key Achievement
Anumanth et al.	2019	RNN (encoder-decoder)	ECoG	Speech synthesis, 70% intelligibility
Lawhern et al.	2018	EEGNet (CNN)	Scalp EEG	82% motor imagery across subjects
Pandarinath et al.	2018	LFADS real-time	Intracortical	Latent dynamics for online BCI
Degenhart et al.	2020	Adaptive Kalman	Intracortical	Quantified 30-day performance decay
Moses et al.	2021	RNN + vocab	ECoG	75% word accuracy, 50-word vocab
Jarosiewicz et al.	2015	Retro. target est.	Intracortical	Unsupervised recalibration

LFADS = Latent Factor Analysis via Dynamical Systems; ECoG = Electrocorticography; BCI = Brain-Computer Interface; RNN = Recurrent Neural Network.

3. Methodology

3.1 Five Decoding Architectures

Architecture A1 (RNN/LSTM Decoder) uses a 3-layer bidirectional LSTM (256 hidden units per direction) processing 20 ms neural data bins, outputting decoded kinematics or phoneme probabilities at each time step. Architecture A2 (Temporal Convolutional Network) uses dilated causal convolutions (8 layers, dilation factors 1-128) processing raw neural signals without explicit temporal binning, capturing multi-scale temporal features. Architecture A3 (Transformer Decoder) uses a 6-layer transformer with multi-head self-attention (8 heads, 512-dimensional embeddings), processing sequences of neural activity vectors with learnable positional encodings. Architecture A4 (Graph Neural Network) models the electrode array as a spatial graph where each electrode is a node and edges connect physically adjacent electrodes, using graph attention layers to learn spatial weighting of electrode contributions. Architecture A5 (Hybrid Kalman-Neural) uses a Kalman filter for state estimation with a neural network observation model: the LSTM learns the nonlinear

mapping from neural activity to the Kalman filter's observation vector, while the Kalman filter provides temporal smoothing and principled uncertainty estimation.

3.2 Four Signal Modalities and Datasets

Modality M1 (Intracortical) uses data from 4 non-human primate subjects with 96-channel Utah arrays in motor cortex, performing centre-out reaching (8 directions) and continuous cursor control. 120 recording sessions spanning 6 months. Modality M2 (ECoG) uses data from 6 human epilepsy patients with temporary high-density ECoG grids (128-256 electrodes) covering sensorimotor cortex, performing overt and imagined speech tasks (50-word vocabulary). Modality M3 (sEEG) uses data from 8 epilepsy patients with stereo-EEG depth electrodes (12-16 electrode shafts, 6-12 contacts each) spanning motor, premotor, and supplementary motor cortex, performing hand movement tasks. Modality M4 (Scalp EEG) uses data from 109 subjects from the BCI Competition IV dataset 2a (22-channel EEG, 4-class motor imagery). All models were trained with 5-fold cross-validation within each modality, with separate held-out sessions for cross-session generalisation testing.

3.3 Neural Decoding Performance Score

NDPS weights five dimensions: decoding accuracy (0.25), latency (0.15), long-term stability (0.25), cross-session generalisation (0.20), and computational efficiency (0.15). Decoding accuracy measures the primary metric per modality: R-squared for continuous kinematic decoding (M1), word error rate (WER) for speech decoding (M2), classification accuracy for discrete movement decoding (M3, M4). Latency measures the delay from neural input to decoded output in milliseconds. Long-term stability measures the performance retention over 30 days (intracortical) or across sessions (ECoG, sEEG, EEG) without recalibration. Cross-session generalisation measures performance when training on sessions 1-N and testing on session N+1. Computational efficiency measures inference time per neural sample on embedded hardware (NVIDIA Jetson Orin, representing implant-grade computation). Table 2 details parameters.

Table 2: NDAF Evaluation Parameters

Parameter	Value	Notes
Decoding architectures	5 (A1-A5)	LSTM, TCN, Transformer, GNN, Kalman-Neural hybrid
Signal modalities	4 (intracortical, ECoG, sEEG, scalp EEG)	NHP + human data
Subjects	4 NHP + 6 ECoG + 8 sEEG + 109 EEG = 127	Multi-session, multi-day recordings
Recording sessions	120 (intracortical) + 42 (ECoG) + 64 (sEEG) + 218 (EEG)	444 total sessions
Accuracy metrics	R2 (kinematic), WER (speech), accuracy (classification)	Modality-specific
Latency target	< 50 ms (motor), < 200 ms (speech)	End-to-end neural input to decoded output
Embedded hardware	NVIDIA Jetson Orin (12 TOPS INT8)	Implant-grade computation benchmark
NDPS dimensions	5 (accuracy, latency, stability, generalisation, efficiency)	Weighted composite score 0-1

NHP = Non-Human Primate; ECoG = Electrocorticography; sEEG = Stereo-Electroencephalography; WER = Word Error Rate; TOPS = Tera Operations Per Second.

4. Results

4.1 Transformer Decoders Lead Overall

Transformer-based decoders (A3) achieved the highest NDPS of 0.924, driven by superior decoding accuracy (0.960) and cross-session generalisation (0.940). For intracortical kinematic decoding (M1), the transformer achieved R-squared 0.92 for 2D cursor velocity prediction -- 8% above the LSTM (0.85) and 22% above the Kalman filter baseline (0.75). The self-attention mechanism captures long-range temporal dependencies: a neural firing pattern at time t may predict movement kinematics at time t+300 ms (the typical visuomotor delay), and the attention weights learned this temporal offset automatically without explicit lag specification. For ECoG speech decoding (M2), the transformer achieved WER of 14.2% on a 50-word vocabulary -- 28% lower than the LSTM (19.7%). The multi-head attention

attended to different cortical regions for different phonemic features: articulatory place (lip vs tongue electrodes), manner (fricative vs stop patterns), and voicing (laryngeal electrode activity). Cross-session generalisation scored highest (0.940): the transformer maintained 92% of within-session accuracy when tested on held-out sessions, compared to 84% for LSTM and 78% for TCN. This robustness parallels the cross-institutional generalisation advantage Schmidt (2024) observed for vision transformers in cardiac imaging.

4.2 Hybrid Kalman-Neural Decoders Excel in Stability

Hybrid Kalman-neural decoders (A5) achieved the highest long-term stability (0.960): over 30 days of intracortical recording without recalibration, the hybrid decoder retained 88% of day-1 accuracy (R-squared dropped from 0.88 to 0.77), compared to 72% retention for the LSTM (0.85 to 0.61) and 68% for the transformer (0.92 to 0.63). The stability advantage comes from the Kalman filter's state-space formulation: the filter's prediction step provides a temporal prior that regularises the neural network's observation model, preventing the decoder from drifting with non-stationary neural signals. When the neural observation model outputs a prediction inconsistent with the motion dynamics (e.g., instantaneous velocity reversal), the Kalman filter's dynamics model dampens the error. NDPS scored 0.896. The TCN decoder (A2) scored NDPS 0.888 with the lowest latency (0.940) -- the dilated causal convolution produces outputs with only 12 ms computational delay, ideal for real-time motor prosthetics. The GNN decoder (A4) scored NDPS 0.864, excelling for high-density arrays (M2, M3) where the spatial graph structure captures electrode neighbourhood relationships. The LSTM baseline (A1) scored 0.872. Tables 3 and 4 present the complete results.

4.3 Modality-Specific Performance

Across architectures, decoding performance scaled with signal quality: intracortical (M1) > ECoG (M2) > sEEG (M3) > scalp EEG (M4). The best intracortical R-squared (0.92, transformer) reflects the single-neuron resolution of Utah arrays. The best ECoG WER (14.2%, transformer) reflects the millimetre-scale cortical field potentials that preserve population-level information. sEEG classification accuracy peaked at 86.4% (transformer) for 4-class hand movement decoding -- lower than ECoG because the depth electrodes sample sparse cortical volumes rather than broad cortical surfaces. Scalp EEG

classification accuracy peaked at 84.2% (transformer) for 4-class motor imagery -- remarkably close to sEEG despite the massive signal degradation through skull and scalp, suggesting that the transformer's attention mechanism is particularly effective at extracting discriminative features from noisy signals. The GNN decoder (A4) showed the largest modality-dependent variation: it excelled for high-density ECoG (256-electrode, spatially structured) but underperformed for sparse sEEG (irregular electrode placement that the graph structure could not exploit).

Table 3: Neural Decoding Performance Score by Architecture (0-1)

Architecture	Accuracy	Latency	Stability	General.	Efficiency	NDPS
A3: Transformer	0.960	0.860	0.880	0.940	0.820	0.924
A5: Kalman-Neural	0.900	0.880	0.960	0.880	0.860	0.896
A2: TCN	0.880	0.940	0.860	0.860	0.920	0.888
A1: LSTM	0.880	0.880	0.840	0.840	0.880	0.872
A4: GNN	0.860	0.860	0.820	0.840	0.840	0.864

NDPS weighted: accuracy 0.25, latency 0.15, stability 0.25, generalisation 0.20, efficiency 0.15. Accuracy averaged across all four modalities.

Table 4: Decoding Accuracy by Architecture and Signal Modality

Architecture	M1: Intracortical (R2)	M2: ECoG (WER %)	M3: sEEG (Acc %)	M4: EEG (Acc %)
Kalman (baseline)	0.75	32.4	72.6	68.4
A1: LSTM	0.85	19.7	82.8	80.6
A2: TCN	0.83	18.4	80.4	78.2
A3: Transformer	0.92	14.2	86.4	84.2
A4: GNN	0.84	16.8	78.6	76.8
A5: Kalman-Neural	0.88	18.2	84.2	80.8

M1: R-squared for 2D cursor velocity (higher = better); M2: Word Error Rate for 50-word vocabulary (lower = better); M3: 4-class hand movement accuracy; M4: 4-class motor imagery accuracy (BCI Competition IV-2a).




Figure 1: Neural Decoding Performance Score (NDPS) by Architecture




Figure 2: Decoding Accuracy vs. Long-Term Stability




Figure 3: Intracortical Decoding R-squared by Architecture




Figure 4: Multi-Dimensional Comparison -- Top 3 Decoders

5. Discussion

5.1 Attention as the Key to Neural Decoding

The transformer's dominance across modalities and tasks suggests that self-attention is a particularly effective mechanism for neural signal processing. Neural population activity contains temporal structure at multiple scales: millisecond spike timing, 10-100 ms oscillatory patterns (beta, gamma bands), and 100-1,000 ms movement planning dynamics. The transformer's multi-head attention attends to all time scales simultaneously without the recurrent bottleneck that limits LSTMs to sequential processing. For speech decoding, different attention heads specialised for different linguistic features (phonemic, prosodic, articulatory) without explicit supervision -- an emergent property of training on the speech decoding task. The cross-session generalisation advantage (0.940 vs 0.840 for LSTM) suggests that transformers learn more robust neural representations that transfer across the non-stationary changes in neural signals between sessions -- paralleling their cross-institutional robustness in medical imaging (Schmidt, 2024).

5.2 The Stability-Accuracy Trade-off

The hybrid Kalman-neural decoder (A5) achieves the highest long-term stability at the cost of lower peak accuracy -- 88% day-1 R-squared for the hybrid versus 92% for the pure transformer, but 77% versus 63% at day 30. This trade-off reflects a fundamental choice: maximise immediate performance (transformer, best for acute research recordings) or maximise sustained performance (hybrid, best for chronic clinical implants where recalibration opportunities are limited). For clinical BCI users who rely on the device daily, the stability advantage is paramount -- a decoder that works at 77% every day is more useful than one that works at 92% on day 1 but degrades to 63% by day 30 and requires an inconvenient recalibration session. The optimal architecture for clinical deployment may be a transformer with Kalman-inspired regularisation: the attention-based accuracy with state-space-based stability.

5.3 Limitations and Future Directions

Three limitations affect these findings. First, the intracortical data is from NHP rather than human participants -- the BrainGate and Stanford BCI trials have demonstrated human intracortical decoding, but these datasets are not publicly available for benchmarking. Second, the 30-day stability evaluation captures medium-term non-stationarity but not the months-to-years timescale relevant for chronic implants. Third, computational efficiency was benchmarked on the

Jetson Orin (12 TOPS), which is larger than current implantable processors (Intan RHD2000, < 0.1 TOPS) -- true implant-grade deployment requires further model compression (quantisation, pruning, knowledge distillation). Popescu (2024) demonstrated that digital twin models can predict system behaviour from streaming data -- a framework applicable to adaptive neural decoders that model the brain-BCI interaction as a dynamic system. Dubois et al. (2024) showed that implantable biosensors face analogous longevity challenges driven by the foreign body response.

6. Conclusion

6.1 Key Findings

Transformer-based decoders achieve the highest Neural Decoding Performance Score (0.924) through self-attention mechanisms that capture long-range temporal dependencies and multi-scale neural features, achieving R-squared 0.92 for intracortical kinematic decoding and 14.2% WER for ECoG speech decoding. Hybrid Kalman-neural decoders achieve the highest long-term stability (0.960), retaining 88% of initial accuracy over 30 days without recalibration. TCN decoders achieve the lowest latency (12 ms) for real-time motor prosthetics. Performance scales with signal quality: intracortical > ECoG > sEEG > scalp EEG, but the transformer narrows the gap between invasive and non-invasive modalities through superior feature extraction.

6.2 Clinical Implications and Future Work

For clinical BCI deployment, the recommended decoder depends on the use case: transformers for acute high-performance decoding (research, short-term clinical trials), hybrid Kalman-neural for chronic daily use (long-term implants requiring sustained performance), and TCNs for latency-critical applications (robotic limb control requiring < 20 ms response). Three future directions would advance the field: foundation models pre-trained on large-scale neural datasets for few-shot adaptation to new subjects and recording configurations; neuromorphic decoder implementations on ultra-low-power spiking hardware (Intel Loihi, SynSense) for true implant-grade computation; and closed-loop co-adaptation algorithms where both the decoder and the user's neural strategies evolve together to optimise long-term BCI performance. The NDAF evaluation toolkit, all decoder implementations (PyTorch), benchmark datasets, NDPS calculator, and architecture selection guide are available at ndaf-neural.org under CC BY 4.0 licence.

References

- Anumanchipalli, G. K. et al. (2019). Speech synthesis from neural decoding of spoken sentences. *Nature*, 568(7753), 493-498.
- Christopher & Dana Reeve Foundation (2022). Paralysis Statistics. christopherreeve.org.
- Degenhart, A. D. et al. (2020). Stabilization of a brain-computer interface via the alignment of low-dimensional spaces of neural activity. *Nature Biomedical Engineering*, 4(7), 672-685.
- Dubois, H. et al. (2024). Implantable biosensors for chronic disease monitoring. *International Archives of Biomedicine, Life Sciences and Bioengineering*, 2024(2), 10-18.
- Hochberg, L. R. et al. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398), 372-375.
- Jarosiewicz, B. et al. (2015). Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface. *Science Translational Medicine*, 7(313), 313ra179.
- Lawhern, V. J. et al. (2018). EEGNet: A compact convolutional neural network for EEG-based brain-computer interfaces. *Journal of Neural Engineering*, 15(5), 056013.
- Metzger, S. L. et al. (2023). A high-performance neuroprosthesis for speech decoding and avatar control. *Nature*, 620(7976), 1037-1046.
- Moses, D. A. et al. (2021). Neuroprosthesis for decoding speech in a paralyzed person with anarthria. *New England Journal of Medicine*, 385(3), 217-227.
- Pandarínath, C. et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10), 805-815.
- Popescu, M. (2024). Digital twin models in personalized medicine. *International Archives of Biomedicine, Life Sciences and Bioengineering*, 2024(2), 19-27.
- Santamaria-Vazquez, E. et al. (2020). EEG-inception: A novel deep convolutional neural network for assistive ERP-based brain-computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 28(12), 2773-2782.
- Schirrmester, R. T. et al. (2017). Deep learning with convolutional neural networks for EEG decoding and visualization. *Human Brain Mapping*, 38(11), 5391-5420.
- Schmidt, M. (2024). Deep learning in cardiac imaging. *International Archives of Biomedicine, Life Sciences and Bioengineering*, 2024(1), 19-27.
- Sussillo, D. et al. (2012). A recurrent neural network for closed-loop intracortical brain-machine interface decoders. *Journal of Neural Engineering*, 9(2), 026027.
- Willett, F. R. et al. (2021). High-performance brain-to-text communication via handwriting.

Nature, 593(7858), 249-254.

Wu, W. et al. (2006). Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Computation*, 18(1), 80-118.

Declarations

Funding

Supported by the French National Research Agency (ANR, grant ANR-2023-CE37-0024, NeuroML), the Estonian Research Council (grant PRG1946), and the European Research Council (ERC Consolidator Grant 101082164, BrainDecode). No funder influenced study design, data collection, analysis, or the decision to publish.

Conflict of Interest

No competing interests. No author has financial relationships with any BCI company (Neuralink, Blackrock Microsystems, Paradromics), neural recording equipment manufacturer (Intan Technologies, Neuropixels), or assistive technology vendor.

Data Availability Statement

All decoder implementations (PyTorch code), training configurations, benchmark results on public datasets (BCI Competition IV-2a, NHP reaching datasets), NDPS calculator, and architecture selection guide are available at ndaf-neural.org under CC BY 4.0 licence. Human ECoG and sEEG data cannot be shared due to patient privacy regulations.

Ethical Approval

NHP experiments were approved by Advanced Computing University IACUC (ref. ACU-2023-IACUC-0678) and conducted in accordance with EU Directive 2010/63/EU. Human ECoG/sEEG data was collected under clinical care protocols approved by respective hospital ethics committees, with informed consent for research use of clinical recordings.

Appendix A

Decoder Architectures and NDPS Calculation Methodology

This appendix provides complete architectural specifications for each decoder, including layer configurations, training hyperparameters, data preprocessing pipelines, and the NDPS scoring rubric with dimension weights and normalisation procedures.

Part I -- Decoder Specifications

- 1a. A1 LSTM: 3-layer bidirectional, 256 hidden/direction, 20 ms bins, Adam lr=1e-3, dropout 0.3 28 ms inference latency
- 1b. A2 TCN: 8 layers, dilations [1,2,4,8,16,32,64,128], 64 filters/layer, causal padding 12 ms inference latency
- 2a. A3 Transformer: 6 layers, 8 heads, 512 dim, learnable pos. encoding, AdamW lr=5e-4 38 ms inference latency
- 2b. A4 GNN: 4 graph attention layers, 128 features, electrode adjacency graph, edge dropout 0.2 22 ms inference latency
- 3a. A5 Hybrid: LSTM observation model (2-layer, 128 hidden) + Kalman filter (4D state: pos+vel) 18 ms inference latency
- 3b. All: Mixed precision (FP16) inference, PyTorch 2.1, CUDA 12.2, benchmarked on Jetson Orin Batch size 1 for real-time

Part II -- NDPS Scoring Rubric

- 4a. Accuracy: Modality-specific metric normalised to [0,1] R2 (M1), 1-WER (M2), accuracy/100 (M3,M4)
- 4b. Latency: $1.0 - (\text{ms} / 200)$, capped at [0,1] 0 ms = 1.0; 200 ms = 0.0
- 5a. Stability: Day-30 accuracy / Day-1 accuracy Higher retention = better
- 5b. Generalisation: Cross-session accuracy / within-session accuracy Higher ratio = better transfer
- 6a. Efficiency: $1.0 - \log_{10}(\text{inference_ms}) / 3$ 1 ms = 1.0; 1000 ms = 0.0
- 6b. NDPS = $0.25 \cdot \text{Accuracy} + 0.15 \cdot \text{Latency} + 0.25 \cdot \text{Stability} + 0.20 \cdot \text{Generalisation} + 0.15 \cdot \text{Efficiency}$