# Adaptive decoding of temporally variable neural activity in single-trial time series

**Pablo Oyarzo (pablo@cfin.au.dk)**
Department of Clinical Medicine, Aarhus University
Aarhus, Denmark
Department of Education and Psychology, Freie Universität Berlin
Berlin, Germany

**Radoslaw M. Cichy (rmcichy@zedat.fu-berlin.de)**
Department of Education and Psychology, Freie Universität Berlin
Berlin, Germany

**Diego Vidaurre (dvidaurre@cfin.au.dk)**
Department of Clinical Medicine, Aarhus University
Aarhus, Denmark
Department of Psychiatry, Oxford University
Oxford, UK

## Abstract

**Cognitive processes, specially those involving higher-order functions, often unfold with temporal variability. This complicates the use of time-locked analysis techniques, including standard machine learning-based decoding methods. Although existing methods perform well in tasks with externally timed events, decoding covert processes -such as imagery or recall- remains difficult due to uncertainty in the timing of the underlying neural dynamics. In these cases, task-relevant neural signals may occur at variable latencies across trials, violating the temporal alignment assumptions of standard decoding models. We introduce the Adaptive Decoding Algorithm (ADA), a nonparametric framework for decoding under temporal uncertainty. ADA performs two coupled tasks: (i) it estimates, for each trial, the temporal window most likely to reflect task-relevant neural activity, and (ii) it uses this information to decode the trial label. Using controlled simulations, we show that ADA outperforms conventional methods that assume fixed temporal structure. These results demonstrate that explicitly modeling trial-specific timing can substantially improve decoding performance in scenarios where the timing of relevant neural activity is unknown.**

## Introduction

Inferring cognitive states from neural signals is a core goal in computational neuroscience. While multivariate decoding methods have advanced this goal (Haynes & Rees, 2006; Cichy, Pantazis, & Oliva, 2014; Grootswagers, Wardle, & Carlson, 2017), most still rely on a critical assumption: that task-relevant neural processes occur at fixed latencies across trials. This is particularly limiting in time-resolved modalities such as M/EEG, where trial averaging and time-locking to stimulus onset are standard practices to improve SNR (Pfurtscheller & Da Silva, 1999; Grootswagers, Robinson, & Carlson, 2019). However, these procedures risk conflating temporally variable components, potentially masking meaningful latency shifts and hindering interpretability (Stokes & Spaak, 2016; Vidaurre, Myers, Stokes, Nobre, & Woolrich, 2019).

To tackle this open problem, we introduce the Adaptive Decoding Algorithm (ADA), a nonparametric decoding framework that explicitly models trial-by-trial variability in the timing of task-relevant neural activity. ADA takes a two-step approach: first, identifying when informative neural patterns occur; then, using this information to classify behavioral or stimulus labels. Because of the first step, ADA does not assume temporal alignment across trials, enabling it to capture non-stereotyped and dynamically timed neural responses.

### Problem Formulation

Let $\mathbf{X}$ denote the recorded brain activity, $\mathbf{Y}$ the associated stimulus or behavioral labels, and $\mathcal{Z}$ the latent neural process that generates $\mathbf{X}$ and depends on $\mathbf{Y}$. We consider data from a subject across $N$ trials, where each trial $\mathbf{x}_n \in \mathbb{R}^{T\times p}$ is a multichannel time series, and each label $y_n \in \{-1,+1\}$. Although we focus on binary classification for simplicity, the method generalizes to multiclass and regression settings. Importantly, $\mathcal{Z}$ may not be active at the same time across trials. The observed signals $\mathbf{X}$ reflect not only $\mathcal{Z}$, but also unrelated activity and noise. These signals oscillate continuously across multiple frequencies (Buzsaki & Draguhn, 2004).

To capture the variability in $\mathcal{Z}$, we define an (unobserved) indicator variable $\mathbf{I} \in \{0,1\}^{T\times N}$, where $I_{tn} = 1$ if trial $n$ at time $t$ reflects that $\mathcal{Z}$ is active. For example, during visual processing, information typically reaches cortex $\sim$75 ms post-stimulus (Thorpe, Fize, & Marlot, 1996), so $I_{tn} = 0$ for $t < 75$ ms. Unlike standard decoding approaches, so we do not assume columns in $\mathbf{I}$ to be equal across trials.

Our goal is twofold: to decode $\mathbf{Y}$ from new trials and to estimate $\mathbf{I}$, thereby providing a trial-wise readout of when $\mathcal{Z}$ is active.

### Algorithm definition

ADA provides a practical approximation to the estimation of $I$, and uses this information to guide the decoding. The decoding itself builds on a $K$-nearest neighbors (KNN) classifier (Fix, 1985), incorporating temporal structure in a nonparametric framework.

In detail, each trial is divided into $W$ overlapping windows of length $L$, forming a matrix $\mathbf{D} \in \mathbb{R}^{WN\times Lp}$, where each row is a window. A label vector $\mathbf{r} \in \mathbb{R}^{WN}$ is constructed by assigning to each window the label $y_n$ of its parent trial.

To represent when $\mathcal{Z}$ is active, we define a binary matrix $\mathbf{H} \in \{0,1\}^{N\times W}$, where each row marks the window(s) used for decoding in that trial. $\mathbf{H}$ is thus a temporally-coarser approximation to $\mathbf{I}$.

ADA applies a weighted KNN classifier at the window level in a leave-one-trial-out fashion. Each window's prediction $\hat{r}_j$ is a weighted average of its $K$ nearest neighbors, based on cosine similarity. Accuracy is computed as $a_j = \hat{r}_j \cdot r_j$, and the top-scoring window(s) per trial define the estimated matrix $\hat{\mathbf{H}}$.

We then fit a ridge regression model to predict the window-level accuracies from data features:

$$\hat{\beta} = \arg\min_{\beta} \sum_{j=1}^{WN} (a_j - \mathbf{d}_j\beta)^2 + \alpha\|\beta\|_2^2, \tag{1}$$

where $\alpha$ is a regularization parameter.

In testing, each unseen trial is segmented into $W$ windows. Each is scored using the estimated $\hat{\beta}$. The top $\kappa$ windows are selected, and a final prediction is made using the same weighted KNN rule, restricted to the training windows indexed by $\hat{\mathbf{H}}$. If $\kappa > 1$, the predictions can be integrated as:

$$\hat{y} = \text{sign} \sum_{l=1}^{\kappa} f(\mathbf{d}_l, \mathbf{D}_{\hat{\mathbf{H}}}, \mathbf{r}_{\hat{\mathbf{H}}}, K). \tag{2}$$

where $f(\mathbf{d}_l, \mathbf{D}_{\hat{\mathbf{H}}}, \mathbf{r}_{\hat{\mathbf{H}}}, K)$ denotes the weighted K-nearest neighbors prediction for window $l$. This design enables ADA to

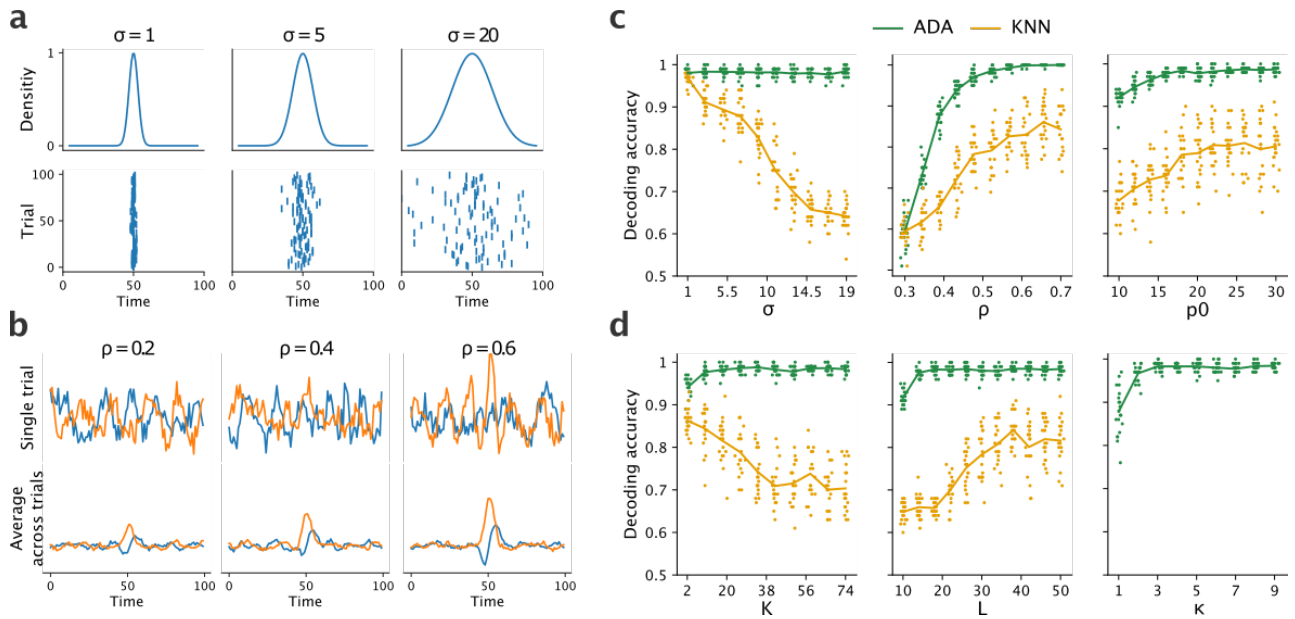


Figure 1: **Simulation experiments. a)** Trial-wise variability is introduced by sampling effect latencies from a distribution $\eta$, shaped by the dispersion parameter $\sigma$. Top: three examples of $\eta$; bottom: corresponding samples of s. **b)** Single-trial and averaged signals for a representative channel at varying $\rho$ levels. **c)** Accuracy of ADA (green) and KNN (yellow) as a function of $\sigma$ (left), $\rho$ (center), and $p_0$ (right). **d)** Accuracy as a function of $K$ (left), $L$ (center), and $\kappa$ (right; ADA only). Defaults: $K = 20$, $L = 30$, $\kappa = 4$.

flexibly adapt to trial-specific temporal variability in the neural process underlying behavior.

In summary, ADA's main hyperparameters are: $K$ (neighbors), $L$ (window length), and $\kappa$ (windows per trial for prediction).

## Simulation Experiments

We benchmarked ADA against KNN, a baseline that does not model between-trial temporal variability, using synthetic data produced with Genephys, a generative model of electrophysiological signals (Vidaurre, 2024) that can simulate multiple stimulus-induced effects. Each dataset included 40 channels, 100 time points per trial, and 200 training/testing trials.

### Sensitivity to Signal Properties

We probed the influence of three signal-level parameters: dispersion $\sigma$ (between-trial temporal variability; Fig. a), the signal-to-noise ratio $\rho$ (effect strength; Fig. b), and the number of task-relevant channels $p_0$ (effect sparsity); see (Vidaurre, 2024) for details on these.

As observed, ADA remained robust across increasing temporal variability, while KNN performance degraded markedly (Fig. c, left). Both methods improved with higher $\rho$, but ADA consistently outperformed KNN across all levels (Fig. c, center). As sparsity increased (i.e., lower $p_0$), accuracy declined for both methods, yet ADA retained a clear performance margin (Fig. c, right).

### Sensitivity to Algorithm Parameters

We then assessed the influence of algorithmic hyperparameters: number of neighbors ($K$), window length ($L$), and the number of windows ($\kappa$, specific to ADA). Simulations were run with fixed $\sigma = 10$, $\rho = 0.5$, and $p_0 = 20$.

ADA performed robustly across a wide range of $K$, except at very low values. KNN favored smaller $K$, likely due to its internal averaging across windows (Fig. d, left). Accuracy improved with longer windows $L$, with ADA consistently outperforming KNN ($L$; Fig. d, center). For ADA, increasing $\kappa$ beyond one improved robustness, reducing sensitivity to misestimations in $\mathbf{H}$ (Fig. d, right).

## Conclusion

Unraveling temporal uncertainty is important to study higher-order cognition. We introduced ADA, a framework that explicitly models trial-specific variability in the timing of task-relevant signals. Controlled simulations demonstrated that ADA outperforms conventional methods when temporal alignment cannot be assumed. This underscores the importance of integrating temporal flexibility into decoding models, particularly for covert or internally driven processes.

## References


Buzsaki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *science*, *304*(5679), 1926–1929.

Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature neuroscience*, *17*(3), 455–462.

Fix, E. (1985). *Discriminatory analysis: nonparametric discrimination, consistency properties* (Vol. 1). USAF school of Aviation Medicine.

Grootswagers, T., Robinson, A. K., & Carlson, T. A. (2019). The representational dynamics of visual objects in rapid serial visual processing streams. *NeuroImage*, *188*, 668–679.

Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of cognitive neuroscience*, *29*(4), 677–697.

Haynes, J.-D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature reviews neuroscience*, *7*(7), 523–534.

Pfurtscheller, G., & Da Silva, F. L. (1999). Event-related eeg/meg synchronization and desynchronization: basic principles. *Clinical neurophysiology*, *110*(11), 1842–1857.

Stokes, M., & Spaak, E. (2016). The importance of single-trial analyses in cognitive neuroscience. *Trends in Cognitive Sciences*, *20*(7), 483–486.

Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *nature*, *381*(6582), 520–522.

Vidaurre, D. (2024). A generative model of electrophysiological brain responses to stimulation. *Elife*, *12*, RP87729.

Vidaurre, D., Myers, N. E., Stokes, M., Nobre, A. C., & Woolrich, M. W. (2019). Temporally unconstrained decoding reveals consistent but time-varying stages of stimulus processing. *Cerebral Cortex*, *29*(2), 863–874.