

ERAASR: An algorithm for removing electrical stimulation artifacts from multielectrode array recordings

Daniel J. O'Shea^{1,2*}, Krishna V. Shenoy^{1,2,3,4}

1 Neurosciences Program

2 Departments of Electrical Engineering, Bioengineering, and Neurobiology

3 Bio-X Program, Stanford Neurosciences Institute

4 Howard Hughes Medical Institute

Stanford University, Stanford, CA, U.S.A.

* Correspondence: djoshea@stanford.edu

Abstract

Electrical stimulation is a widely used and effective tool in systems neuroscience, neural prosthetics, and clinical neurostimulation. However, electrical artifacts evoked by stimulation significantly complicate the detection of spiking activity on nearby recording electrodes. Here, we present ERAASR: an algorithm for Estimation and Removal of Artifacts on Arrays via Sequential principal components Regression. This approach leverages the similar structure of artifact transients, but not spiking activity, across simultaneously recorded channels on the array, across pulses within a train, and across trials. The effectiveness of the algorithm is demonstrated in macaque dorsal premotor cortex using acute linear multielectrode array recordings and single electrode stimulation. Large electrical artifacts appeared on all channels during stimulation. After application of ERAASR, the cleaned signals were quiescent on channels with no spontaneous spiking activity, whereas spontaneously active channels exhibited evoked spikes which closely resembled spontaneously occurring spiking waveforms. The ERAASR algorithm requires no special hardware and comprises sequential application of straightforward linear methods with intuitive parameters. Enabling simultaneous electrical stimulation and multielectrode array recording can help elucidate the causal links between neural activity and cognitive functions and enable the design and implementation of novel sensory protheses.

1 Introduction

Electrical stimulation is a method of modulating neural activity that is widely used within neuroscience and neuroengineering, as well as for treatment of chronic neurological pathologies. Within neuroscience, microstimulation is used to probe the functional organization of neural circuits. Whereas electrophysiological recordings provide correlative insights that neural activity in a brain region appears related to a specific behavior, electrical stimulation can demonstrate causal contributions of brain regions to specific cognitive functions (Salzman et al., 1990). In particular, intracortical microstimulation (ICMS) has been used to identify neural networks underlying perception, attention, cognition, and movement (see Cohen and Newsome, 2004; Clark et al., 2011; Histed et al., 2013, for a review). ICMS can also drive reliable percepts in humans and animal models, including somatosensory vibration (Romo et al., 2000) and visual phosphenes (Tehovnik and Slocum, 2007), which underscores its application for delivering artificial sensation in visual (Tehovnik et al., 2009) and

motor prosthetic systems (Berg et al., 2013; O'Doherty et al., 2011). In clinical contexts, stimulation, especially delivered to deeper brain structures, has demonstrated efficacy for the treatment of neurological and neuropsychiatric conditions (Wichmann and Delong, 2006; Holtzheimer and Mayberg, 2011) and recently as an avenue for providing sensory information as will be needed for a range of emerging brain-machine interface systems (Flesher et al., 2016).

However, as with all perturbations to the nervous system, accurate interpretation of the results relies on a solid understanding of the effects of the perturbation made (e.g., Jazayeri and Afraz, 2017; Otchy et al., 2015). With lesion and pharmacology studies, understanding the size of the affected region, completeness of the intended effect, the timecourse of the effect and recovery, and the influence of compensatory or homeostatic adaptation mechanisms are all critical to drawing correct conclusions from the behavioral impairments observed. Analogously, with ICMS, it is essential to understand the effect that electrical stimulation has on neural activity in order to draw correct inferences from causal experiments. In the context of neural prosthetics, researchers aim to improve the fidelity of artificial perception with sophisticated spatiotemporal patterning (e.g., Dadarlat et al., 2015) and to optimize stimulation parameters to maximize therapeutic benefit. Obtaining a clear picture of the effects of particular stimulation paradigms will likely prove essential to this goal, enabling closed-loop tuning of stimulation parameters relative to a desired effect on the local neural activity.

The effects on neural firing patterns evoked by ICMS have been difficult to characterize because the electrical artifact induced by stimulation interferes with the electrophysiological recording equipment used to record neuronal responses (Ranck, 1975; Merrill et al., 2005). A variety of approaches have been developed to work around these artifacts, which can be organized into several categories. Online approaches employ special hardware within the signal recording pathway to remove the stimulation artifact online (e.g., Wichmann and Devergnas, 2011; Brown et al., 2008; Müller et al., 2012). While effective, hardware approaches can be expensive or challenging to scale to high channel count recording arrays and typically require very stable artifacts over time for proper removal. Offline approaches employ standard electrophysiology collection hardware to record neural signals and then estimate and subtract the artifact post hoc. The most common approach for artifact estimation is to simply interpolate across the artifact (O'Keefe et al., 2001; Montgomery et al., 2005) and/or to average over repeated stimulation epochs (Hashimoto et al., 2002; Montgomery et al., 2005; Klink et al., 2017) and subtract, exploiting the larger stereotypy of the stimulus artifact relative to variable neural responses. Others have employed curve fitting approaches to estimate the artifact, exploiting the differences in shape of artifact and action potentials, thereby allowing for small variations in the observed artifact across trials (Wagenaar and Potter, 2002; Erez et al., 2010).

These previous approaches have two major shortcomings addressed in this work. First, these approaches only recover neural signals occurring after the stimulation pulse is completed, with a typical window of of several milliseconds during which signal is discarded. However, ICMS is typically delivered in trains of stimulation pulses occurring at frequencies approaching several hundred pulses per second (e.g., Churchland and Shenoy, 2007; Klink et al., 2017). This failure to recover neural activity concomitant with the stimulation pulses renders a large fraction of the peri-stimulation signal unusable for spike detection and waveform discrimination. Secondly, prior work has focused on artifacts detected at a single electrode, not exploiting the similar structure of contaminating artifacts across multiple electrodes using a multielectrode array. A notable exception to both of these shortcomings was developed by Mena and colleagues (2016), where they addressed both of these

limitations using a modern statistical framework to isolate electrical artifacts from 67
neural spiking signals in multielectrode array recordings of the retina *in vitro*. 68
However, this approach currently requires the availability of electrical images for each 69
neuron's spike waveform on multiple surrounding electrodes. Unfortunately, for the 70
multielectrode array configurations commonly employed in *in vivo* cortical 71
electrophysiology, the electrical image of most neurons is confined to a single electrode. 72
Additionally, since the set of recorded neurons is biased towards a highly active subset 73
of neurons (Barth and Poulet, 2012), we anticipate that some neurons that are driven 74
to spike by electrical stimulation may be quiet outside of the stimulation period, and 75
thus omitted in the dictionary of available electrical images. 76

In this paper, we present a fast method for removing electrical stimulation artifacts 77
from multielectrode array recordings. We term this approach ERAASR: Estimation 78
and Removal of Artifacts from multielectrode Arrays via Sequential principal 79
components Regression. ERAASR does not require any special hardware and can 80
recover the full timecourse of neural signals during and after stimulation, requiring 81
only that the electrical artifact does not saturate the amplifier on the recording 82
electrodes. Our algorithm exploits the similarity of artifacts across multiple electrodes, 83
multiple pulses within a pulse train, and multiple trials of repeated stimulation. The 84
ERAASR algorithm is structured as a sequence of cascaded, linear operations—PCA 85
and linear regression—across each axis of the data tensor. We validate our method on 86
neural signals collected during an arm-movement task with 24-channel linear 87
multielectrode arrays (“V probes”, Plexon Inc.) in macaque dorsal premotor cortex, 88
demonstrating that spiking waveforms extracted during the stimulation period closely 89
resemble spontaneous spiking waveforms on each electrode. 90

2 Methods 91

2.1 Subjects 92

Animal protocols were approved by the Stanford University Institutional Animal Care 93
and Use Committee. The subject was one adult male macaque monkey (*Macaca* 94
mulatta), monkey P. After initial training, we performed a sterile surgery during which 95
the macaque was implanted with a head restraint and a recording cylinder (NAN 96
Instruments) which was located over left, caudal, dorsal premotor cortex (PMd). The 97
cylinder was placed surface normal to the skull and secured with methyl methacrylate. 98
A thin layer of methyl was also deposited atop the intact, exposed skull within the 99
chamber. Before stimulation sessions began, a small craniotomy (3 mm diameter) was 100
made under ketamine/xylazine anesthesia, targeting an area in PMd which responded 101
during movements and palpation of the upper arm (17 mm anterior to interaural 102
stereotaxic zero). 103

2.2 Reaching task 104

For the purposes of this study, we engaged the monkey in a reaching task in order to 105
drive task related neural activity in premotor cortex where our electrodes were 106
located. We trained the monkey to use his right hand to grasp and translate a custom 107
3D printed handle (Shapeways, Inc.) attached to a haptic feedback device (Delta.3, 108
Force Dimension, Inc.). The other arm was comfortably restrained at the monkey's 109
side. The monkey was trained to perform a delayed reaching task by moving the 110
haptic device cursor towards green rectangular targets displayed on the screen. 111
Successful completion of each movement triggered a juice reward. 112

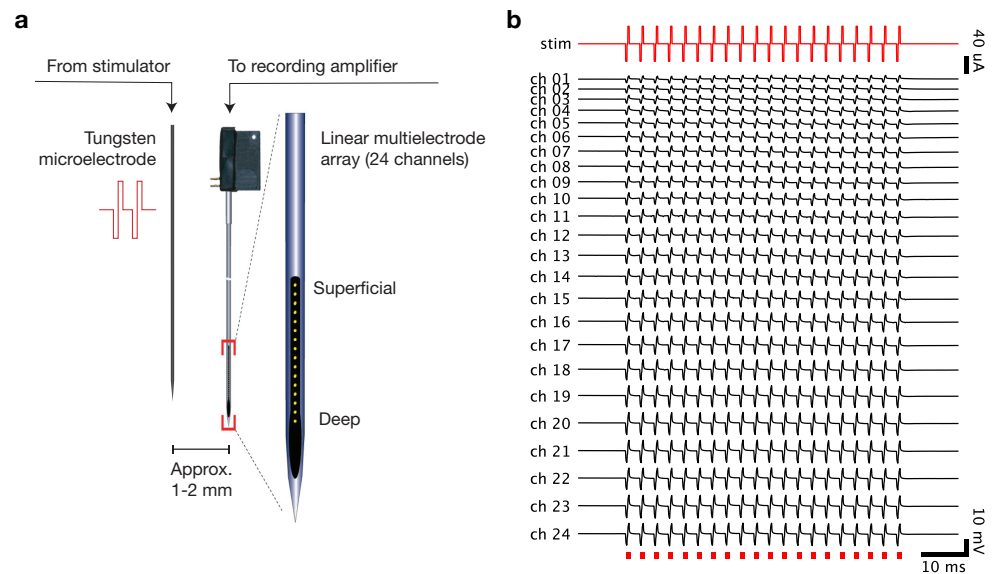


Figure 1. Our technique facilitates electrical recordings of local neural activity during and after nearby electrical stimulation, provided that the electrical artifact does not saturate the amplifier. (a) A schematic of the recording setup. A tungsten stimulating microelectrode and linear multielectrode array are inserted in parallel into dorsal premotor cortex. The linear array records the electrical activity of nearby neurons spanning the layers of cortex, while electrical current is introduced by the tungsten microelectrode. (b) An example of real electrical recordings on the linear multielectrode array (ch. 1-24) during delivery of a single 40 μ A pulse train (red, stim) through the stimulating electrode. The recording channels span 2.3 mm and are labeled from 1 (most superficial, typically above cortex) through 24 (deep, typically in white matter).

2.3 Electrophysiology and stimulation configuration

At the start of each experimental session, both a stimulation electrode and a recording probe were secured to two independently controllable, motorized micromanipulators (NAN instruments). Both probes were lowered simultaneously through blunt, non-penetrating guide tubes into dorsal premotor cortex at 3 μ m/s to an approximate depth of 2 mm (Figure 1a). The stimulation electrode was a single tungsten microelectrode with approximately 1 M Ω impedance at 1 kHz (Part #UEWLGC-SECN1E, Frederick Haer Company). The recording probe was a linear electrode array consisting of 24 contact sites located at 100 μ m spacing along the length of the shank (V-probe PLX-VP-24-15ED-100-SE-100-25(640)-CT-500, Plexon Inc.). The recording probe penetration site was located at an approximate distance of 0.75 mm to 1.5 mm from the stimulation site. As best possible, the stimulating electrode was inserted at the same location and to the same depth, whereas we changed the location of the recording electrode for each recording session.

The stimulation was performed using a StimPulse electrical stimulation system used as a combined function generator and isolated current source (Frederick Haer Company). The electrical stimulation current flowed through the brain via the electrode to a “ground screw” located at the posterior pole of the implant whose tip was in contact with the dura below the skull. Microstimulation trains consisted of 20 biphasic pulses delivered at 333 Hz for 57 ms (150 μ s cathodic, 100 μ s pause, 150 μ s anodic). Stimulation amplitude was varied between 5 μ A to 40 μ A. Stimulation

delivered at the start of most experimental sessions, while the arm was in a passive, resting position, triggered brief movements of the contralateral upper arm and shoulder at thresholds typically $> 140 \mu\text{A}$. Stimulation onset was triggered via TTL pulse delivered by the task engine (Simulink Real-time Target, The Mathworks; NIDAQ digital to analog card, National Instruments). Stimulation was delivered on 20-40% of trials and randomly interleaved. Stimulation amplitude was fixed within each block of trials, where we collected at least 8-10 repetitions of each stimulation trial type for each block. Within each session, we typically began with blocks of lower amplitude stimulation before proceeding to higher amplitude stimulation.

The recording probe was connected to a 3-headed switching headstage, a component of the Blackrock StimSwitch system (Blackrock Microsystems). The probe connects to one bank of inputs, control lines from the StimSwitch control box connect to the second bank of inputs, and outgoing voltage signals leave via the third bank of outputs. During these experiments, this component was configured to act as a unity gain buffer (head stage) which simply relays the voltage signals to a shielded ribbon cable (Samtec Inc.) to the amplifier (Blackrock Microsystems). Each electrode is differentially amplified relative to a common reference line in the V-probe itself, which is also shorted to surrounding guide tube for better noise rejection.

Broadband voltages were recorded from all 24 electrode channels of the recording probe. We also recorded from the stimulation electrode during initial penetration to verify that we could see nearby neurons on the electrode. We then disconnected the stimulation electrode from the recording system and connected it to the StimPulse stimulator. Broadband signals were filtered at the amplifier (0.3 Hz 1 pole high-pass filter, 7.5 kHz 3 pole low-pass filter), digitized to 16 bit resolution over $\pm 8.196 \text{ mV}$ (resolution = $0.25 \mu\text{V}$), and sampled at 30 kHz.

2.4 Artifact amplitude model

The amplitude of the electrical artifact was measured on each recording channel as the peak to peak voltage concomitant with the first stimulus pulse, averaged across trials. These amplitude measurements were used to fit a model relating artifact amplitude to the distance of each site to the stimulation source. This model jointly optimizes a single scaling parameter for the distance-dependent amplitude relationship as well as parameters denoting the stimulation-relative coordinates of the recording probe in each session. We define x_s as the distance measured laterally from the probe to the closest point on the stimulating electrode, both of which are assumed to run parallel to the recording probe and normal to the cortical surface; we define y_s as the vertical distance along the penetration path from the most superficial electrode on the probe to the tip of the stimulating electrode. We begin by describing the location of each electrode e in each session s relative to the probe location $(x, y)_s$. Here

$$(x, y)_{e,s} = (x_s, y_s + e\Delta) \quad (1)$$

where $\Delta = 100 \mu\text{m}$ spacing between sites. This contact site is located at a distance

$$d_{e,s} = \sqrt{x_{e,s}^2 + y_{e,s}^2} \quad (2)$$

from the stimulation source defined as the origin. We then assume that the amplitude of the artifact on each electrode for a specific stimulation amplitude I_a , falls off as the reciprocal of this distance.

$$V_{e,s,a} = \frac{v_0 * I_a}{d_{e,s}} \quad (3)$$

where v_0 has units $\text{V}/(\mu\text{A mm})$. The parameters of this model are this single scaling parameter v_0 as well as a pair of location coordinates for the probe on each session.

We then fit this model using constrained non-linear curve fitting (lsqcurvefit in Matlab) to minimize the squared error between empirical and modeled artifact amplitude across electrodes, sessions, and amplitudes. Confidence intervals of the parameters were obtained using the Jacobian of the fitted model.

2.5 Precise temporal alignment and tensor construction

We begin with the raw, unfiltered voltage traces for each stimulation trial, collected simultaneously on the 24 channels, sampled at 30 kilosamples/sec. We first correct the small amount of jitter between triggering the stimulator to the beginning of the stimulation train, typically on the order of a few milliseconds. We begin by precisely aligning the stimulation trials to each other in time. This process is performed separately for each group of trials sharing a common stimulation amplitude. We extract the raw voltage data off of one electrode channel in a 60 ms window surrounding the triggering event, high-pass filter the signal with a fourth-order, 250 Hz corner frequency Butterworth filter, to remove slow drift. We then detect the onset of the artifact due to the first biphasic pulse on each trial using an appropriately low negative threshold, set to avoid spiking activity but reliably detect artifacts at the lowest stimulation amplitude. We then realigned the voltage data for each trial to this threshold crossing event. We then performed an outlier detection procedure to detect rare instances of stimulator failure, including trials where 19 or 21 pulses were delivered instead of 20. Outlier detection was performed by projecting the realigned, unfiltered voltage traces onto the first 10 principal components of that group, and detecting trials where the absolute value of the z-scored scores in any of the first 10 PCs exceeded a value of 5. Outlier trials (typically < 2% of trials) were excluded from subsequent analysis. The traces were then up-sampled 10x to 300 kHz using spline interpolation and then a maximum cross-correlation procedure was performed to determine the temporal offset between each trial and the first in the group. We manually verified the results of this procedure to ensure successful artifact removal and accurate alignment of the artifact pulse trains. The aligned traces are then resampled at the original 30 kHz sampling rate.

The ERAASR algorithm is summarized in alg. 1. We construct a data tensor for each group of trials sharing a common stimulation amplitude, as the artifact evoked across trials for a single stimulation amplitude were highly similar. We extract a 60 ms window of the trial-aligned broadband voltage data across the trials, yielding a data tensor \mathbf{X}_{raw} with size C (number of channels) \times T (number of timepoints per pulse) \times P (number of pulses) \times R (number of trials). In our datasets $C = 24$ channels, $T = 90$ timepoints per pulse (3 ms at 30 kHz sampling), $P = 20$ pulses, and R was typically on the order of 100-200 trials.

We lightly high pass filter the signals in time with a fourth-order, 10 Hz corner frequency Butterworth filter, which removes slow drifts from the traces and makes subsequent processing more robust. We first attempt to clean each channel by exploiting the similarities of the artifact across simultaneous channels (Figure 2a).

2.6 Removal of common structure over channels

We then unfold the data tensor into an $RTP \times C$ matrix \mathbf{M}^c (Figure 2b), and perform principal components analysis. This allows us to re-express each channel's response vector ($RTP \times 1$) as a linear combination of principal components, each $RTP \times 1$. As the artifact waveform is very large relative to the interesting spiking activity and shared across all channels, one would expect the top few principal components to capture preferentially shared variance due to the artifact. Empirically, we determined that the top $K_C = 4$ principal components (PCs) captured much of the artifact shape

Algorithm 1: ERAASR cleaning procedure

Input : Raw data tensor $\mathbf{X}_{\text{raw}}(c, t, p, r)$
with size C channels by T timepoints per pulse by P pulses by R trials

Output : Cleaned data tensor $\mathbf{X}_{\text{cleaned}}(c, t, p, r)$ with same size

Parameters: K_C, K_P, K_R - number of principal components to describe artifact structure over channels, pulses, trials
 $\lambda_C, \lambda_P, \lambda_R$ - the number of adjacent channels, pulses, trials excluded as regressors during artifact reconstruction
 $\beta_P, \beta_C \in \{\text{false}, \text{true}\}$ - Perform cleaning over pulses separately on each channel? Below we assume β_P is false and β_C is true for clarity.

$\mathbf{X} \leftarrow \mathbf{X}_{\text{raw}}$

Clean across channels

Unfold X into \mathbf{M}^C with size RPT by C .

$\mathbf{W}^C \leftarrow$ projection along leading K_C principal components of \mathbf{M}^C

for $c \in 1, \dots, C$ **do**

$\mathbf{A}_c \leftarrow \mathbf{M}^C \mathbf{W}_\phi^C$ with weights zeroed for channel c and channels within λ_C

$\mathbf{X}(c, :, :, :)$ \leftarrow projection of $\mathbf{X}(c, :, :, :)$ onto \mathbf{A}_c

end

Clean across pulses

Unfold X into \mathbf{M}^P with size TRC by P .

$\mathbf{W}^P \leftarrow$ projection along leading K_P principal components of \mathbf{M}^P

for $p \in 1, \dots, P$ **do**

$\mathbf{A}_p \leftarrow \mathbf{M}^P \mathbf{W}_\phi^P$ with weights zeroed for pulse p and pulses within λ_P

$\mathbf{X}(:, :, p, :)$ \leftarrow projection of $\mathbf{X}(:, :, p, :)$ onto \mathbf{A}_p

end

Clean across trials, separately for each channel

for $c \in 1, \dots, C$ **do**

 Unfold $X(c, :, :, :)$ into \mathbf{M}_c^R for with size TP by R .

$\mathbf{W}_c^R \leftarrow$ projection along leading K_R principal components of \mathbf{M}_c^R

for $r \in 1, \dots, R$ **do**

$\mathbf{A}_{c,r} \leftarrow \mathbf{M}_c^R \mathbf{W}_{c,r}^R$ with weights zeroed for trial r and trials within λ_R

$\mathbf{X}(c, :, :, r)$ \leftarrow projection of $\mathbf{X}(c, :, :, r)$ onto $\mathbf{A}_{c,r}$

end

end

return $\mathbf{X}_{\text{cleaned}} \leftarrow \mathbf{X}$

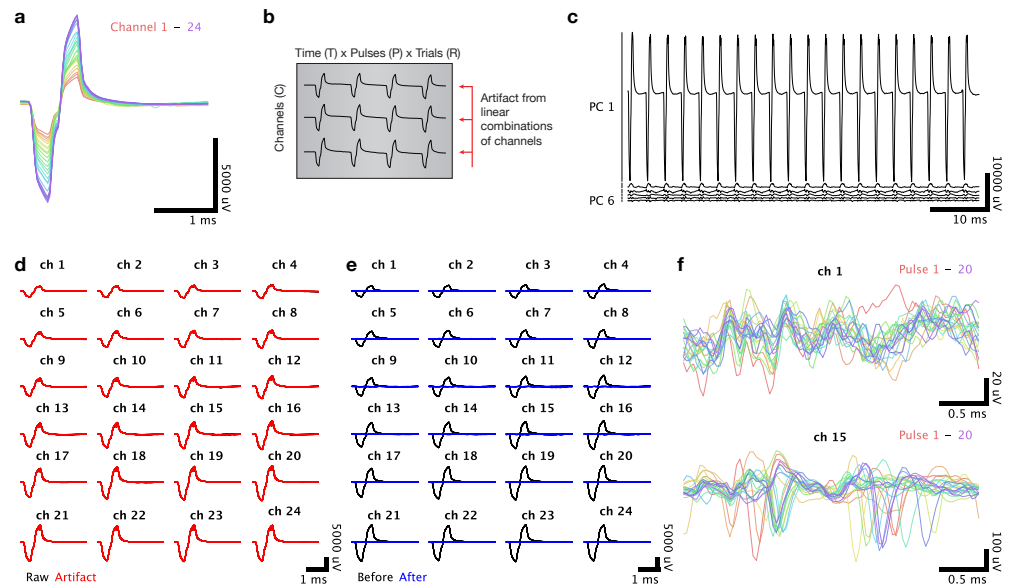


Figure 2. The first stage of artifact removal removes common structure recorded simultaneously across linear electrode array channels. (a) The electrical artifact evoked by a single pulse of a stimulus train is highly similar across the 24 channels of the array. (b) To identify the common artifact structure present across channels, we rearrange the artifact data into a matrix with dimensions time \times pulses \times trials ($T \times P \times R$) by channels (C), shown transposed. (c) The principal components of this matrix are linear combinations of the electrode channels. The projections of the recordings along these components clearly display artifact structure shared across channels. (d) Voltage recordings for the first pulse of the stimulus train, plotted separately for each channel with the responses on each trial superimposed. Raw voltage recordings are shown in black, with the artifact inferred via principal components regression shown in red. (e) Same presentation as in (d) but with the residual voltage recordings shown in blue, after the inferred artifact has been subtracted. (f) Cleaned voltage recordings for a single trial with the responses to each individual pulse superimposed. The top panel shows a superficial channel which had no detectable spontaneous neural activity, and the bottom panel shows a lower channel where a well-isolated single unit was spiking spontaneously.

(Figure 2c). We then want to reconstruct each channel's response omitting the contribution from the artifact. The simplest method to accomplish this would be to reconstruct the channel's response by using a linear combination of all PCs except the first K_C . However, in some cases, especially when lower currents were used and the magnitude of the artifact relative to the spiking activity was smaller, we observed that some of the top PCs would begin to incorporate a small amount of spiking activity from individual channels. By reconstructing those channels from the remaining PCs, the spiking activity itself would be partially distorted because a fractional portion of those spikes would be subtracted along with the artifact. We addressed this issue by again exploiting the locality of spiking activity, by assuming that spiking activity from any one channel would not be present on other channels except the immediately adjacent channels. However, the artifact is shared on a much larger spatial scale and can be separated from the spiking activity. Therefore, for each channel c , we used the following procedure: reconstruct the top K_c PCs using all channels except c and its

227
228
229
230
231
232
233
234
235
236
237
238
239
240

immediately adjacent neighbors $c - 1$ and $c + 1$, that is, using a modified version of the loading weights $w_{(k)}^C$ for the k th principal component in which we set the weights for $c, c - 1, c + 1$ to 0. More generally, we define parameter λ_C that dictates the number of adjacent channels excluded from the reconstruction. We refer to this vector of modified loading weights as $w_{(k,\phi)}^C$.

$$\left(w_{(k,\phi)}^C\right)_i = \begin{cases} 0 & \text{if } i \in \{c - \lambda_C, \dots, c, \dots, c + \lambda_C\} \\ \left(w_{(k)}^C\right)_i & \text{else} \end{cases} \quad (4)$$

Using these loading weights, we reconstruct the matrix of the top K_C artifact-capturing PCs, where each column is given by

$$\left(\mathbf{A}_c\right)_{(k)} = \mathbf{M}^C w_{(k,\phi)}^C \quad k = 1, \dots, K_C; \quad c = 1, \dots, C \quad (5)$$

We then regress the channel's response vector against these artifact components and subtract this reconstructed artifact from the channel's response:

$$\hat{\mathbf{M}}_c^C = \mathbf{M}_c^C - \left(\mathbf{A}_c^T \mathbf{A}_c\right)^{-1} \left(\mathbf{A}_c^T\right) \mathbf{M}_c^C \quad (6)$$

This cleaning across channels captured (Figure 2d) and removed (Figure 2e) much of the artifact from each pulse. However, when examining the responses on an individual channel across each of the 20 pulses superimposed, we observed consistent structure in the traces (Figure 2f).

2.7 Removal of common structure over pulses

This structure indicates that there remained artifact structure that is unique to individual channels but shared among the pulses in each trial. To address this, we rearrange $\hat{\mathbf{M}}^C$ into a new matrix \mathbf{M}^P with size $TRC \times P$, a set of responses to each individual pulse, concatenated over trials and channels (Figure 3a). Empirically, we determined that the top $K_P = 2$ principal components (PCs) captured much of the artifact shape. Consequently, we used the same PC regression procedure to estimate and subtract the artifact on each pulse, omitting the contribution of the pulse under consideration and optionally adjacent pulses within λ_P . For our dataset, $\lambda_P = 0$ was used. If $w_{(k)}^P$ is a vector of P loading weights describing the k th principal component, and for each pulse p , $w_{(k,p)}^P$ is given by:

$$\left(w_{(k,p)}^P\right)_i = \begin{cases} 0 & \text{if } i \in \{p - \lambda_P, \dots, p, \dots, p + \lambda_P\} \\ \left(w_{(k)}^P\right)_i & \text{else} \end{cases} \quad (7)$$

For each pulse, we reconstruct the matrix of the top K_P artifact-capturing PCs, whose columns are given by:

$$\left(\mathbf{A}_p\right)_{(k)} = \mathbf{M}^P w_{(k,p)}^P \quad k = 1, \dots, K_P; \quad p = 1, \dots, P \quad (8)$$

We then regress the channel's response vector against these artifact components and subtract this reconstructed artifact from the channel's response:

$$\hat{\mathbf{M}}_p^P = \mathbf{M}_p^P - \left(\mathbf{A}_p^T \mathbf{A}_p\right)^{-1} \left(\mathbf{A}_p^T\right) \mathbf{M}_p^P \quad (9)$$

This cleaning procedure over pulses captured much of the common artifact structure observed within each train (Figure 3b). It is also possible to perform this cleaning over pulses separately for each channel, following an approach like that described below for cleaning over trials.

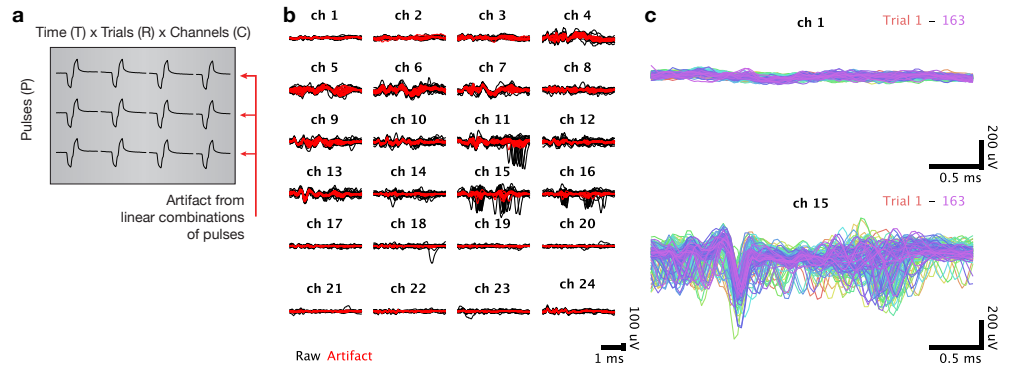


Figure 3. The second stage of artifact removal removes common artifact structure across the successive pulses of the stimulus train. (a) To isolate structure across pulses, we rearrange the artifact data into individual matrices for each channel, each with dimensions time (T) \times trials (R) \times channels (C) by pulses (P). (b) Voltage recordings for a single trial, plotted separately for each channel with the responses to each pulse superimposed. Raw voltage recordings are shown in black, with the artifact inferred via principal components regression shown in red. (c) Cleaned voltage recordings for a single pulse with the responses to all trials superimposed. The top panel shows a superficial channel which had no detectable spontaneous neural activity, and the bottom panel shows a lower channel where a well-isolated single unit was spiking spontaneously.

We next examined the set of cleaned traces for individual channels and pulses in the train, over the full set of trials. We observed that there remained common artifactual structure in these traces, which suggests that there is artifact structure unique to that channel and pulse number but shared among many trials (Figure 3c). In most cases, this common structure was similar among nearby trials but exhibited a slow drift in the shape of the artifacts over the experimental session.

2.8 Removal of common structure across trials

To remove these residual artifacts, we again employed a principal components regression approach, exploiting the shared structure of the artifact over multiple trials. We performed this operation separately for each channel. We rearranged the cleaned $\hat{\mathbf{M}}^P$ into a set of C data matrices \mathbf{M}_c^R for $c \in \{1, \dots, C\}$, each with size $TP \times R$ (Figure 4a). Empirically, we determined that the top $K_R = 4$ PCs captured much of the artifact shape over trials. If $w_{(c,k)}^R$ is a vector of R loading weights describing the k th principal component for channel c , we defined:

$$\left(w_{(c,k,r)}^R\right)_i = \begin{cases} 0 & \text{if } i \in \{r - \lambda_R, \dots, r, \dots, r + \lambda_R\} \\ \left(w_{(c,k)}^R\right)_i & \text{else} \end{cases} \quad (10)$$

where λ_R adjacent trials were excluded from the reconstruction. For our dataset, $\lambda_R = 0$ was used.

For each trial, we reconstruct the matrix of the top K_R artifact-capturing PCs, whose columns are given by:

$$\left(\mathbf{A}_{(c,r)}\right)_{(k)} = \left(\mathbf{M}_c^R\right) \left(w_{(c,k,r)}^R\right) \quad k = 1, \dots, K_R; \quad r = 1, \dots, R; \quad c = 1, \dots, C \quad (11)$$

We then regress the trial's response vector against these artifact components and

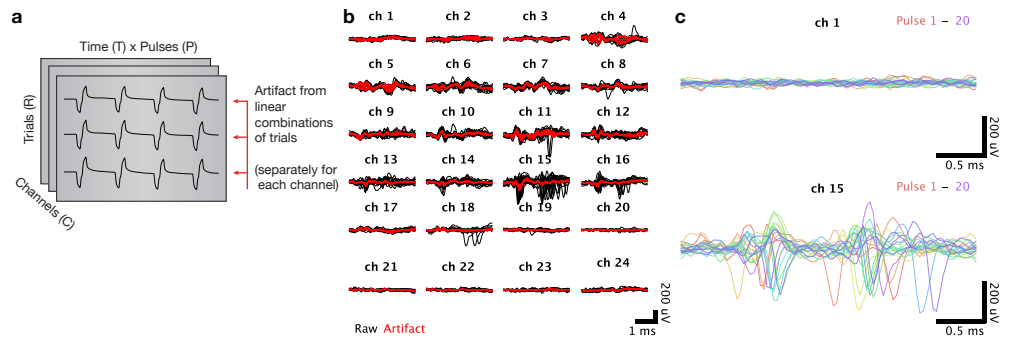


Figure 4. The third stage of artifact removal removes common artifact structure across trials in which identical electrical stimulus trains were delivered. (a) To isolate structure across trials, we rearrange the artifact data into individual matrices for each channel, each with dimensions trials (R) by time \times pulses ($T \times P$). (b) Voltage recordings for a single pulse, plotted separately for each channel with the responses to each trial superimposed. Raw voltage recordings are shown in black, with the artifact inferred via principal components regression shown in red. (c) Cleaned voltage recordings for a single trial with the responses to all pulses superimposed. The top panel shows a superficial channel which had no detectable spontaneous neural activity, and the bottom panel shows a lower channel where a well-isolated single unit was spiking spontaneously.

subtract this reconstructed artifact from the trial's response:

$$\hat{\mathbf{M}}_{(c,r)}^R = \mathbf{M}_{(c,r)}^R - \left(\mathbf{A}_{(c,r)}^T \mathbf{A}_{(c,r)} \right)^{-1} \left(\mathbf{A}_{(c,r)}^T \right) \mathbf{M}_{(c,r)}^R \quad (12)$$

We then rearranged the set of cleaned $\hat{\mathbf{M}}_c^R$ back into a $C \times T \times P \times R$ tensor $\mathbf{X}_{\text{cleaned}}$. This cleaning procedure captured much of the residual artifact (Figure 4c). The cleaned signals in $\mathbf{X}_{\text{cleaned}}$ did not display features obviously indicative of residual stimulation artifact, but did on some channels contain readily detectable spiking activity (Figure 4c, bottom panel). We then reinserted $\mathbf{X}_{\text{cleaned}}$ into the raw voltage traces for each stimulation trial, using appropriate offsets to preserve continuity at the first and last timepoints of the inserted segment.

2.9 Post-stimulation transient cleaning

Following stimulation offset, a slower transient was observed in all channels. These transients are highly similar among trials. Consequently, we employed a similar principal components regression procedure to remove these slow transients. We performed the following procedure for each stimulation amplitude separately. We began by aligning the trials to stimulation offset and extracting voltage data in a 30 ms time window beginning after stimulation end, yielding a tensor \mathbf{X}_{post} with size R by T by C . We then rearrange this tensor into a set of C data matrices $\mathbf{M}_c^{\text{post}}$ for $c \in \{1, \dots, C\}$, each with size $T \times R$. Empirically, we determined that much of the transients temporal structure was present in each trial for a given channel and that the top $K_{\text{post}} = 2$ principal components (PCs) of these $\mathbf{M}_c^{\text{post}}$ matrices captured much of the transient shape over trials. Consequently, we repeated the same principal components regression approach described in eqs. (10) to (12) to clean the post-stimulation transients over trials. We then rearranged the set of cleaned $\hat{\mathbf{M}}_c^{\text{post}}$ back into a R by T by C tensor $\mathbf{X}_{\text{post-cleaned}}$. The signals in $\mathbf{X}_{\text{post-cleaned}}$ were then reinserted into the full voltage traces for each stimulation trial.

2.10 Spike thresholding and sorting

The cleaned voltage traces, along with the original voltage traces on non-stimulated trials were then high-pass filtered using a fourth-order 250 Hz corner frequency Butterworth filter, as is typically done online before spike detection. We then thresholded each signal at $-4.5 \times$ RMS voltage, using a greedy procedure in which the spike threshold crossings were detected in order of size (largest to smallest), and no further threshold crossings were considered within a lockout period extending 0.3 ms prior to and 1.0 ms after the current threshold crossing. We then hand-sorted the spiking waveforms using MKsort (Matthew Kaufman and Ripple, Inc., <https://github.com/ripple-neuro/mksort>).

2.11 Unit selection

During experiments, we recorded any units or multi-units that could be isolated on the multielectrode array, without regard for their responsiveness or modulation by the task. Before analysis, we applied a screening procedure which looked only at non-stimulated trials in order to remove neurons that were very unreliable or very weakly modulated by the task. Briefly, we filtered non-stimulated trials spike trains with a 30 ms Gaussian window, then aligned trials separately from 100 ms pre-target onset to 70 ms post go cue, and then from 300 ms pre-movement to 600 ms post, then averaged within groups of trials with the same reach target. We defined each unit's "signal" to be the range of firing rates over all times and conditions, and "noise" to be the maximum standard error of the mean firing rate. We included units where the ratio of signal to noise exceeded at least 2, a total of 138 units in monkey P. Note that this selection process looked only at non-stimulation trials.

3 Results

We delivered high-frequency (333 Hz) biphasic electrical stimulation in macaque dorsal premotor cortex through a tungsten microelectrode, while simultaneously recording neural spiking activity using a second 24-channel linear multielectrode array (Figure 1a). When a stimulus train was delivered through the electrode, a highly stereotyped stimulation artifact appeared on all channels of the recording array (Figure 1b). These broadband recordings were corrupted by a very large electrical stimulation artifact. Amplitudes could exceed $\pm 7500 \mu\text{V}$ which is nearly two orders of magnitude larger than spiking waveform amplitudes ($\sim 100 \mu\text{V}$). The artifact traces resembled filtered versions of the original current stimulus. with similar transients evoked for each of the 20 biphasic pulses. These artifacts increased in amplitude with larger stimulation current and smoothly varied over the 24 channels of the recording probes (Figure 5a,b), consistent with a distance-dependent attenuation of the electrical stimulus.

We used these amplitude profiles over channels for each session to fit a model relating artifact amplitude and the reciprocal of the distance of each electrode to the stimulation source. The model fitting process jointly optimizes a scaling parameter of the amplitude-distance relationship as well as parameters specifying the location of the recording probe on each session (See methods). We fit this model to data collected in 9 recording and stimulation sessions in macaque prefrontal cortex. Figure 5c shows the empirical and predicted artifact amplitude (normalized by stimulation current) under the model against distance from the stimulation site (squared correlation $R^2 = 0.91$), demonstrating that this model well-describes the distant-dependent falloff of the electrical artifact. The fitted location of each probe relative to the stimulation

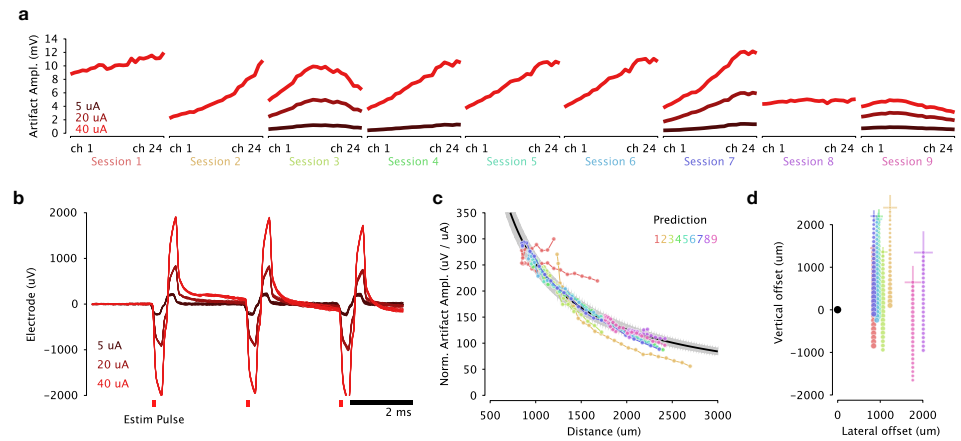


Figure 5. Electrical artifacts recorded on the electrode array resemble scaled, filtered versions of the original stimulus pulse train. The artifact amplitude varies smoothly with stimulation current and electrode distance, enabling *post hoc* inference of the locations of the recording probes in each session. (a) Electrical artifact amplitude (peak to peak) across multielectrode array channels. Each plot corresponds to a recording session in which the multielectrode array was inserted in a new location. (b) Electrical artifacts recorded on several trials on superficial electrode with increasing electrical stimulation current. (c) A simple model relating the electrical artifact amplitude was fit to each channels on every session. The model assumes that the artifact scales linearly with stimulation current and falls off with the reciprocal of Euclidean distance from the stimulating electrode tip. These distances are themselves inferred from the data as parameters describing the position of the electrode array in each session. The fitted model predicted amplitudes (black curve flanked by gray 95% CI) align well with the empirical amplitudes observed on each session (groups of connected dots, each color corresponds to a session). (d) The fitted locations of the recording probe on each session relative to the stimulating electrode tip. Marker size is scaled according to artifact amplitude. Cross-hairs indicate 95 % CI for probe location parameters.

source are depicted in Figure 5d, which closely aligns with the noted approximate probe insertion locations at 0.7 mm to 2 mm from the recording probe. This simple relationship suggests that under certain circumstances, electrical artifacts could be used to reconstruct recording locations *in vivo*.

Having characterized the artifacts, we applied ERAASR to remove the artifact from these traces. The ERAASR algorithm leverages several features of these recording datasets. First, it assumes that the true neural signal is corrupted by an additive stimulation artifact, which depends critically on the assumption that no information is lost due to amplifier saturation during the stimulation period. For the 5 μ A to 40 μ A range of stimulation amplitudes used, the stimulation artifact produced did not saturate the amplifier's range ($\pm 8196 \mu$ V) on any of the recording channels in any experimental session. Therefore, if the artifact shape can be properly estimated, it can be removed via simple subtraction. Second, because we recorded simultaneously on multiple, closely spaced channels, we have multiple simultaneous observations of the electrical stimulation artifact at different points in space, whereas the individual channels do not share the same spiking waveforms, especially if we exclude immediately adjacent channels. Third, the stimulation train contains 20 repeats of identical biphasic pulses, which produced highly similar artifact transients. Finally, we deliver the stimulation train over many trials, yielding similar artifacts on each

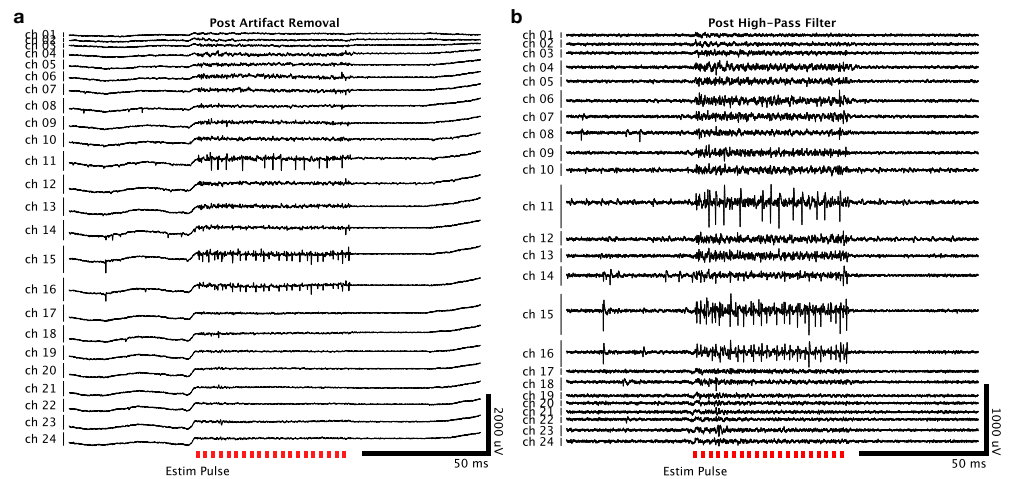


Figure 6. The cleaned single trial voltage recordings exhibit very little residual stimulation artifact. (a) Electrical recordings on the linear multielectrode array for the same trial shown in 1(b). Spiking activity during the stimulus is readily apparent on multiple channels, including several where spontaneous neural activity was observed. (b) The same recording example as in (a), following high-pass filtering to isolate spiking activity.

repetition. Because spiking activity evoked during the stimulation period is unlikely to occur at precisely the same time relative to each pulse and on each trial, we can use these repeated artifact measurements to build an estimate of the artifact shape for subsequent removal.

WeFigure 6a shows a set of traces for a single trial across the array on one of the recording sessions, before high-pass filtering is performed; Figure 6b shows the same data after high-pass filtering before spike detection. The raw, pre-cleaned voltage traces for this same trial is shown in Figure 1b, demonstrating that the vast majority of the artifact initially present has been removed. These high-pass filtered traces were then thresholded to extract spikes and then manually sorted into different units.

To evaluate the artifact removal process, we use prior knowledge about the multielectrode array channels. The most superficial (low numbered) and deepest (high numbered) channels were located above cortex and in white matter respectively, which we inferred from the signals observed on these channels, the density of neural signals on intermediate channels presumed to lie within cortex, and the known thickness of premotor cortex relative to the depth span of the probe. During non-stimulated trials, no detectable spiking waveforms were observed on these superficial and deep channels, therefore the presence of threshold crossings in the cleaned stimulation voltage traces on these channels would be surprising, and suggest that the cleaning procedure had failed to fully remove the stimulation artifact. However, as expected, the presence of evoked spiking activity during the stimulation window is limited to the intermediate channels, where spontaneous spiking activity was also observed (Figure 6). For closer inspection, we superimposed the raw voltage traces with cleaned, artifact-removed, high-pass filtered voltage traces for individual trials. Figure 7a-b shows such a comparison for the most superficial channel, where neither spontaneous nor evoked spiking activity was detectable. Figure 7c-d shows the same comparison for a channel with spontaneous spiking as well as robust evoked spiking activity.

Secondly, it is essential to ensure that threshold crossings detected during the stimulation period are likely due to real neural spiking activity, as opposed to residual

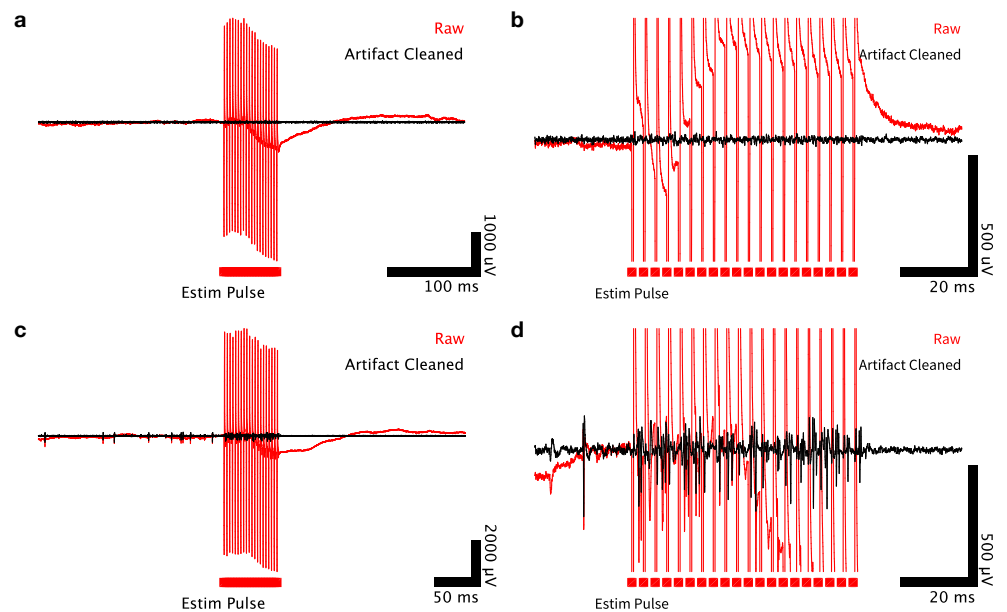


Figure 7. Comparison of raw and artifact cleaned signals. (a) Raw (red) and cleaned (black) voltage traces for a single trial on the most superficially located recording channel where no spontaneous activity was recorded, and no spiking activity is detected during stimulation. (b) Zoomed view of (a). (c) Same as (a) but for a channel located within cortex where spontaneous activity was present. (d) Zoomed view of (c) demonstrating evoked spikes visible during the stimulation period.

transients due to stimulation artifact that persist through the cleaning procedure. We 411
reasoned that threshold crossings created by residual artifact on a specific channel 412
would not be expected to resemble the spontaneous spiking waveforms collected 413
during non-stimulated trials on a per-trial basis. For each session, we generated a 414
visual comparison of the spiking waveforms detected on non-stimulated trials with 415
those detected within the stimulation window (putatively “evoked” spikes). These 416
evoked spiking waveforms detected during the stimulation period were highly similar 417
to the spiking waveforms detected during non-stimulation trials. Figure 8 418
demonstrates this similarity for a representative pair of experimental sessions. We 419
note that these waveforms need not be identical, as evoked spiking activity from other 420
neurons located further from the electrode could superimpose to corrupt the 421
waveforms from nearby neurons, especially when this faraway spiking activity is highly 422
synchronized by stimulation. 423

Lastly, we can quantify the amount of residual artifact directly by again utilizing 424
channels that displayed no spontaneous spiking activity. Figure 9a compares the RMS 425
voltage of these individual channels sampled from the post-cleaning stimulation 426
window against a time window taken before stimulation. This metric captures the 427
amount of baseline noise (largely thermal noise) on each channel, thereby providing an 428
expectation for the variance of a given channel if the entirety of the corrupting artifact 429
were successfully removed. For low amplitudes, the RMS voltage is often slightly lower 430
than the pre-stimulation RMS, which is possible due to the greedy nature of the 431
cleaning procedure. For larger amplitudes, the stimulation RMS remains acceptably 432
close to pre-stimulation RMS. Figure 9b summarizes the similarity of spiking 433
waveforms on each channel observed both spontaneously and during the stimulation 434

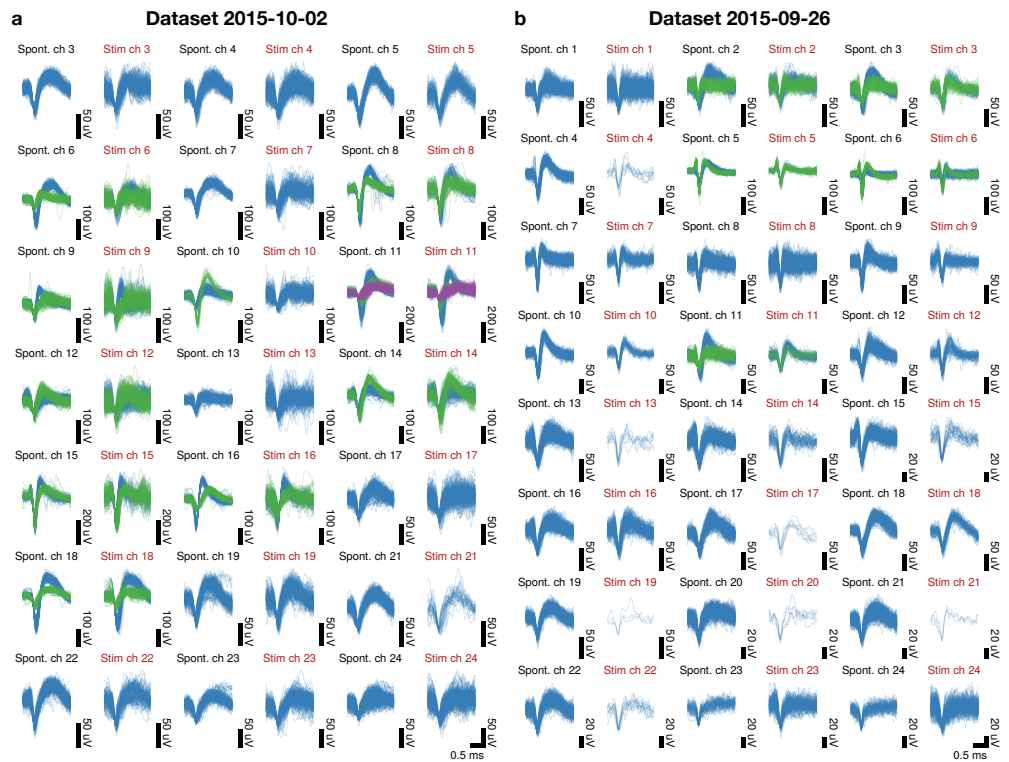


Figure 8. Recovered spiking waveforms detected during the stimulation period closely resemble spontaneously occurring waveforms recorded on the same channel. (a-b) Superimposed waveforms from each channel with spiking activity of two representative recording sessions. Each pair of columns compares spontaneously occurring (spont.) vs. peri-stimulus period (stim.) spiking waveforms. Waveform colors denote hand-sorted unit identity.

window.

435

4 Discussion

436

We discuss the main features of our method for estimating and removing electrical stimulation artifacts in comparison with existing approaches, as well as limitations of our approach and possible improvements. We also highlight a set of interesting neuroscientific arenas where the ability to observe electrically perturbed neural activity might be particularly illuminating, underscoring the utility of artifact removal methods.

437

438

439

440

441

442

4.1 Comparison to other methods

443

ERAASR exploits similarities in the electrical artifact observed across multiple channels, pulses within a stimulus train, and trials with repeated delivery of the same stimulus. Of these the shared structure of simultaneously recorded artifact across multiple channels was most useful. Most previously described attempts exploit only the common structure across trials, typically using an averaging or moving average approach (Hashimoto et al., 2002; Montgomery et al., 2005; Klink et al., 2017) to

444

445

446

447

448

449

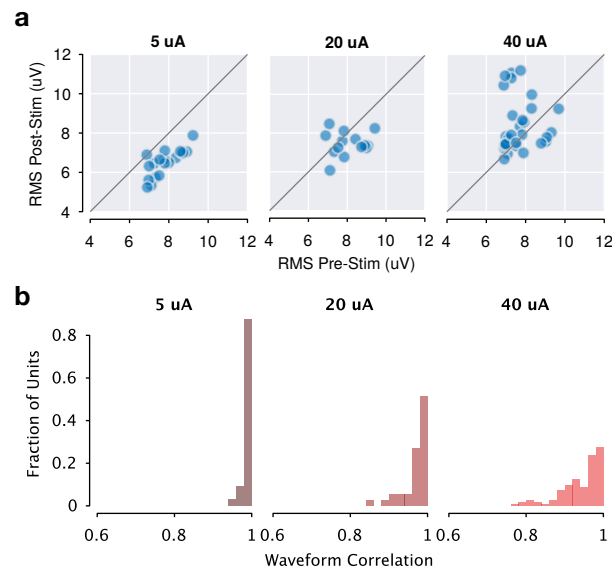


Figure 9. Summary of artifact removal quality across recording sessions. (a) RMS voltage on electrode channels which displayed no spontaneous spiking activity, across all recording sessions. Post-cleaning peri-stimulation RMS voltage is similar to spontaneous pre-stim RMS on each channel. (b) Summary of correlation coefficients between spontaneous spiking waveforms and recovered waveforms during stimulation, across all recovered units in all recording sessions.

estimate artifacts on an individual channel. However, these approaches are necessarily quite sensitive to variability in artifact shape or amplitude across trials. While this variability may be small in relative terms, the much larger amplitude of artifacts relative to neural spiking waveforms can create meaningfully large errors in the residual signal, requiring that several milliseconds of signal in the vicinity of the stimulation pulse be discarded. Another common approach is to exploit differences in the shape of stimulation artifacts relative to spiking waveforms, by using curve-fitting algorithms to capture artifacts in the voltage signal but exclude spiking activity (Wagenaar and Potter, 2002; Erez et al., 2010). Similarly, these fits reliably capture artifact shape at a fixed delay from the stimulation pulse, but struggle to capture earlier portions of the transient (Erez et al., 2010), requiring peri-stimulus signal to be discarded. Motivated by the popularity of high-frequency pulse trains in ICMS experiments, our method addresses problem of observing neural activity through the entire stimulation pulse without discarding signal.

Our algorithm combines a series of simple, intuitive, linear operations along the channels, pulses, and trials axes of the data tensor. The core operation is principal components regression, in which PCA is used to identify a template describing common structure across a given axis of the tensor, followed by a regression step in which the artifact on a given channel (or pulse, or trial) is reconstructed by excluding that channel (pulse, trial) and its neighbors. This approach mitigates the risk of removing spiking signals during the cleaning process, exploiting the local nature of spiking responses across channels and the inherent biological variability in spiking responses on each pulse and each trial. Furthermore, each of these stages is intuitive and the results easily understood. A small set of key parameters, including the number of principal components and the set of neighbors excluded from the

reconstruction at each stage, can be adjusted to trade-off between more aggressive cleaning and more veridical preservation of spiking waveform shape.

Our method operationally separates the process of artifact estimation from spike detection and extraction, which is performed as usual using typical high-pass filtering and thresholding on all voltage data after the cleaning process is complete. In contrast, a promising alternative method for multi-channel artifact removal is presented by Mena and colleagues (2016) for multielectrode array recordings of the retina. This approach employs a statistical model describing both the artifact and the spike generation process, jointly estimating artifact and spikes present in the voltage signals. This method exploits common structure of the artifact across a local group of channels around the stimulation source, but also requires that the electrical image of each neuron (the shape of spiking waveforms over several nearby electrodes) be known as an input to the cleaning algorithm. While this approach is highly effective and appropriate in the context of retinal recordings, where individual neurons are recorded on many densely spaced electrodes, the electrical images of neurons in primate cortex using typical multielectrode arrays are often limited to one or two electrodes.

4.2 Limitations

First, our approach relies critically on the assumption that stimulation artifact linearly superimposes with neural spiking signals. Therefore, its utility is limited to stimulation configurations and amplitudes which occupy the linear, non-saturated regime of the amplifier. Practically, this sets a minimum distance between the electrode array and the stimulation source for a given stimulation amplitude. This distance could be designed such that the return path of the electrical current steers the electric field so as to minimize the recorded artifact amplitude (Rattay and Resatz, 2004). This problem can also be effectively managed in hardware, employing special circuitry to estimate and partially cancel artifacts online (Wichmann and Devergnas, 2011; Brown et al., 2008; Müller et al., 2012). A hybrid approach employing multi-channel artifact removal circuitry to prevent saturation using a predictable transformation of the recorded signal, supplemented with *post hoc* artifact cleaning procedure like the one described here could be effective across a much larger range of stimulation amplitudes.

Second, our algorithm sequentially removes shared structure across channels, pulses, and trials. We accomplish this by reshaping the voltage data tensor into a matrix (or set of matrices) so that familiar methods like PCA can be employed. This problem of finding shared structure naturally lends itself to tensor decomposition (Kolda and Bader, 2009), which could identify shared structure jointly over each of these axes and potentially improve the artifact estimation. Tensors naturally arise in neuroscientific data collection contexts, and decomposition methods are becoming increasingly useful for identifying population structure (Seely et al., 2016; Elsayed and Cunningham, 2017; Williams et al., 2017).

Lastly, our algorithm operates by greedily removing any shared structure as artifact, which can inadvertently remove or distort spiking waveform signals as well. In our dataset, we did not observe this while removing shared structure across channels, as we explicitly excluded adjacent channels from the reconstruction process. However, removing shared structure across pulses and trials assumes that spiking responses evoked by a stimulation pulse are variable in time. Moreover, with increasing stimulation amplitude, the temporal precision of evoked spikes may increase (Ranck, 1975; Butovas and Schwarz, 2003), which would lead to waveform distortion due to overzealous reconstruction and removal of spiking signals as artifact. We experimented with using matched-filters for robust spike detection while relaxing the artifact estimation technique (data not shown), though we did not explore whether this would be applicable at higher currents and more synchronous evoked spiking. In these

circumstances, we expect that a joint estimation of artifact and spiking as proposed by Mena et al. (2016) could be adapted to more effectively recover highly regular evoked spiking by exploiting other differences in the structure of spikes from artifact.

4.3 Motivation for direct observation of electrically perturbed neural activity

ICMS is a well-established and popular technique for perturbing neural activity and probing the causal contributions of a certain brain region to specific cognitive functions (Cohen and Newsome, 2004; Histed et al., 2013; Clark et al., 2011; Tehovnik, 1996). Here we argue that recent results suggest that the effect of electrical stimulation on neural activity may be significantly more complex than previously realized, which implies that direct observations of the net effect on neural activity in the perturbed region as well as upstream and downstream areas could reveal the precise mode by which microstimulation modulates behavior.

The initial conception of ICMS is that stimulation activates most neurons within a sphere surrounding the electrode tip (Stoney et al., 1968; Tehovnik, 1996; Tehovnik et al., 2006). This idea derives originally from the findings of Stoney et al. (1968), who cleverly side-stepped the issue of stimulation artifacts by using collisions of antidromic and orthodromic spikes in a dual-stimulation paradigm as an indirect measure of local neuronal activation. Using this technique, they estimated that 10 μ A and 100 μ A ICMS currents would activate most pyramidal cells in a local ball extending 100 μ m and 450 μ m in radius, respectively. A wealth of additional research has also attempted to carefully characterize the sensitivity of various elements of the CNS to stimulation, as a function of current amplitude, pulse duration and shape, electrode configuration, etc. (e.g. Ranck, 1975; Asanuma et al., 1976; Tehovnik, 1996; Rattay, 1999; Tehovnik and Slocum, 2007; Marcus et al., 1979; Nowak and Bullier, 1996; Nowak and Bullier, 1998b; Nowak and Bullier, 1998a; Kimmel and Moore, 2007). The primary findings of this body of work are that stimulation primarily evokes spikes at axons, in particular the excitable axon initial segment, and that the likelihood of evoking a spike in a neuron depends on distance, pulse duration, and current according to a relatively simple power law. The spatial profile of ICMS has been bolstered by multiple behavioral studies as well (e.g. Murasugi et al., 1993; Tehovnik et al., 2004; Tehovnik et al., 2005; Bartlett et al., 2005; Tehovnik and Sommer, 1997). These studies derive an estimate of the effective activation volume and dependency on stimulation parameters by combining behavioral readouts with known features of cortical spatial organization. For example, Murasugi et al. (1993) found that low current ICMS delivered to monkey area MT could bias perceived dot motion direction, but that ICMS currents above 20 μ A less effectively biased perception. The authors conclude that below this threshold current, electrical activation is confined to a single “directional” column in MT, about 0.2 mm in diameter (Albright et al., 1984).

However, recent evidence challenges this traditional view of localized, dense activation. Using two-photon calcium imaging, Histed et al. (2009) demonstrated that low-intensity ICMS activated a sparse collection of neurons distributed in a relatively large volume extending several hundred μ m. In two experiments in cats, the authors also found neurons as far as 4 mm away activated by 10 μ A stimulation, perhaps resulting from lateral axons extending multiple millimeters within layer 2/3 of visual cortex in cats but not rodents. The widely distributed collection of neurons is likely activated by antidromic stimulation of their axons. By advancing the electrode 30 μ m, the authors found that an entirely non-overlapping sparse set of cells was activated, suggesting that ICMS at low currents (10-25 μ A) activates axonal processes within \approx 15 μ m of the electrode tip. At higher currents, more cells within the distributed

volume were activated, though all experiments were conducted with currents $\leq 30 \mu\text{A}$. 576

Additionally, many mapping studies have focused on direct activation in response 577
to a single isolated current pulse. Nonetheless, it has long been known that ICMS 578
induces transsynaptic spiking in connected neurons (Stoney et al., 1968; Butovas and 579
Schwarz, 2003). As larger currents activate more neurons, the effect of transsynaptic 580
activation may be enhanced due to the summation of many from this larger directly 581
activated population. Cortically evoked spiking may also depend on subthreshold 582
activity in neurons, which is continuously modulated by ongoing task or 583
stimulus-evoked responses (Kara et al., 2002). When multiple pulses are delivered, as 584
is typical in experiments applying ICMS in behaving animals, short-term synaptic 585
depression commonly seen in cortex would also modulate the effectiveness of 586
transsynaptically evoked spiking (Thomson and Deuchars, 1994; Stratford et al., 1996; 587
Deisz and Prince, 1989; Varela et al., 1997). When longer pulse trains are used, 588
particularly at high currents, electrical stimulation might readily engage a large 589
collection of neural circuits while driving temporally complex neural responses (Strick, 590
2002; Graziano et al., 2002). 591

It is also well-established that larger stimulation currents can recruit inhibitory 592
interneurons and result in inhibition of cortical responses lasting hundreds of 593
milliseconds after stimulation (Berman et al., 1991; Watanabe et al., 1966; Li et al., 594
1960; Chung and Ferster, 1998; Creutzfeldt et al., 1966; Kara et al., 2002; Masse and 595
Cook, 2010; Seidemann et al., 2002; Butovas et al., 2006; Butovas and Schwarz, 2003). 596
Additionally, in an experiment combining cortical stimulation with fMRI and 597
extracellular recording, Logothetis et al. (2010) demonstrated that high-frequency 598
ICMS of the thalamic lateral geniculate nucleus (LGN) initially excites (but later 599
suppresses) monosynaptically connected primary visual cortex (V1), but suppresses 600
responses in downstream extrastriate cortices. This suppression was abolished by 601
bicuculine infusion. The authors concluded that, high-frequency ICMS pulse trains 602
disrupt normal information flow along cortico-cortical pathways, by recruiting strong 603
inhibition that prevents the spread of activation beyond the regions whose direct 604
afferents were stimulated. 605

Taken together, the net effects of ICMS on the firing patterns of cortical neuronal 606
populations may be quite complicated. The directly activated population is likely 607
sparse, widely distributed, and stochastic due to a strong dependence on axonal 608
proximity to the stimulating electrode. Moreover, the transsynaptically activated 609
population of cells undergoes a complex, time-varying pattern of modulation that 610
reflects circuit micro-architecture, local circuit dynamics, short-term synaptic 611
plasticity, recruitment of inhibition, and an interaction with ongoing task-related 612
activity. With this evident complexity, we feel that arriving at an accurate 613
interpretation of the behavioral consequences of electrical microstimulation first 614
requires careful measurement of its effects on local neuronal population activity 615
(Jazayeri and Afraz, 2017; Otchy et al., 2015). Furthermore, an accurate picture of 616
electrically modulated neural activity could elucidate the circuit mechanisms 617
underlying intriguing aspects of the interaction between microstimulation and 618
behavior, such as how electrically evoked eye movements interact with visual guided 619
eye movements (Sparks and Mays, 1983), how precise timing of electrical pulses 620
modulates saccadic effects (Kimmel and Moore, 2007), how perceptual decisions are 621
biased by stimulation of LIP (Hanks et al., 2006), how motor preparation recovers 622
subsequent to electrical disruption (Churchland and Shenoy, 2007; Shenoy et al., 623
2011), and how motor cortical stimulation can evoke complex movements converging 624
on a specific endpoint (Graziano et al., 2011) and effectively nullify the contribution of 625
goal-directed behavior in the motor cortical output (Griffin et al., 2011). A detailed 626
comparison with the effect of optogenetic stimulation could clarify the differences 627

(Gerits and Vanduffel, 2013; Diester et al., 2011; Ohayon et al., 2013) and similarities (Dai et al., 2013) observed between the two modalities' effects on behavior. Lastly, a detailed model of microstimulation's interaction with cortical dynamics could facilitate delivery of more reliable or realistic artificial sensory perception in visual, auditory, and motor prostheses (Otto et al., 2005; Tehovnik et al., 2009; O'Doherty et al., 2011; Bensmaia and Miller, 2014; Dadarlat et al., 2015).

5 Conclusion

We developed ERAASR, an algorithm for estimating and removing artifacts caused by electrical stimulation on multielectrode array experiment. ERAASR assumes that artifact is shared over many channels and that evoked transients are highly repeatable across pulses and trials, whereas spiking activity is highly local and temporally jittered. Shared structure in the voltage signals is identified and removed sequentially across channels, across pulses in a stimulus train, and across trials, using straightforward linear methods. We believe our technique will be useful to neuroscientists in drawing precise causal links between perturbations and the effect of stimulation on neural activity and behavior and aid in the design and implementation of enhanced neural prosthetic systems capable of restoring lost sensation and facilitating precise control.

6 Acknowledgments

We thank Tirin Moore, Rob Franklin, and Paul Venable for helpful discussions. This article was supported by the following grants: Christopher and Dana Reeve Paralysis Foundation, Burroughs Wellcome Fund Career Awards in the Biomedical Sciences,, Defense Advanced Research Projects Agency Reorganization and Plasticity to Accelerate Injury Recovery N66001-10-C-2010, US National Institutes of Health Institute of Neurological Disorders and Stroke Transformative Research Award R01NS076460, US National Institutes of Health Director's Pioneer Award 8DP1HD075623-04, US National Institutes of Health Director's Transformative Research Award (TR01) from the NIMH 5R01MH09964703, Defense Advanced Research Projects Agency NeuroFAST award from BTO W911NF-14-2-0013, Howard Hughes Medical Institute, NSF GRFP, and NSF IGERT.

References

- Albright, T D, R Desimone, and C G Gross (1984). “Columnar organization of directionally selective cells in visual area MT of the macaque.” en. In: *Journal of neurophysiology* 51.1, pp. 16–31. ISSN: 0022-3077.
- Asanuma, H, A Arnold, and P Zarzecki (1976). “Further study on the excitation of pyramidal tract cells by intracortical microstimulation.” en. In: *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale* 26.5, pp. 443–461. ISSN: 0014-4819.
- Barth, Alison L and James F a Poulet (2012). “Experimental evidence for sparse firing in the neocortex.” In: *Trends in neurosciences* 35.6, pp. 345–355. ISSN: 0166-2236. DOI: 10.1016/j.tins.2012.03.008.
- Bartlett, John R, Edgar A DeYoe, Robert W Doty, Barry B Lee, Jeffrey D Lewine, Nubio Negrão, and William H Overman Jr (2005). “Psychophysics of electrical stimulation of striate cortex in macaques.” en. In: *Journal of neurophysiology* 94.5, pp. 3430–3442. ISSN: 0022-3077. DOI: 10.1152/jn.00406.2005.
- Bensmaia, Sliman J and Lee E Miller (2014). “Restoring sensorimotor function through intracortical interfaces: progress and looming challenges.” In: *Nature reviews. Neuroscience* 15.5, pp. 313–325. ISSN: 1471-003X, 1471-0048. DOI: 10.1038/nrn3724.
- Berg, J A, J F Dammann 3rd, F V Tenore, G A Tabot, J L Boback, L R Manfredi, M L Peterson, K D Katyal, M S Johannes, A Makhlin, R Wilcox, R K Franklin, R J Vogelstein, N G Hatsopoulos, and S J Bensmaia (2013). “Behavioral demonstration of a somatosensory neuroprosthesis.” en. In: *IEEE transactions on neural systems and rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society* 21.3, pp. 500–507. ISSN: 1534-4320, 1558-0210. DOI: 10.1109/TNSRE.2013.2244616.
- Berman, N J, R J Douglas, K A Martin, and D Whitteridge (1991). “Mechanisms of inhibition in cat visual cortex.” en. In: *The Journal of physiology* 440, pp. 697–722. ISSN: 0022-3751. DOI: 10.1113/jphysiol.1991.sp018731.
- Brown, E A, J D Ross, R A Blum, Y Nam, B C Wheeler, and S P DeWeerth (2008). “Stimulus-Artifact Elimination in a Multi-Electrode System.” In: *IEEE transactions on biomedical circuits and systems* 2.1, pp. 10–21. ISSN: 1932-4545. DOI: 10.1109/TBCAS.2008.918285.
- Butovas, Sergejus, Sheriar G Hormuzdi, Hannah Monyer, and Cornelius Schwarz (2006). “Effects of Electrically Coupled Inhibitory Networks on Local Neuronal Responses to Intracortical Microstimulation.” en. In: *Journal of neurophysiology* 96.3, pp. 1227–1236. ISSN: 0022-3077, 1522-1598. DOI: 10.1152/jn.01170.2005.
- Butovas, Sergejus and Cornelius Schwarz (2003). “Spatiotemporal effects of microstimulation in rat neocortex: a parametric study using multielectrode recordings.” In: *Journal of neurophysiology* 90.5, pp. 3024–3039. ISSN: 0022-3077. DOI: 10.1152/jn.00245.2003.
- Chung, S and D Ferster (1998). “Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression.” en. In: *Neuron* 20.6, pp. 1177–1189. ISSN: 0896-6273.
- Churchland, Mark M and Krishna V Shenoy (2007). “Delay of movement caused by disruption of cortical preparatory activity.” In: *Journal of neurophysiology* 97.1, pp. 348–359. ISSN: 0022-3077. DOI: 10.1152/jn.00808.2006.
- Clark, Kelsey L, Katherine M Armstrong, and Tirin Moore (2011). “Probing neural circuitry and function with electrical microstimulation.” en. In: *Proceedings of the Royal Society of London B: Biological Sciences* 278.1709, pp. 1121–1130. ISSN: 0962-8452, 1471-2954. DOI: 10.1098/rspb.2010.2211.

- Cohen, Marlene R and William T Newsome (2004). “What electrical microstimulation has revealed about the neural basis of cognition.” In: *Current opinion in neurobiology* 14.2, pp. 169–177. ISSN: 0959-4388. DOI: 10.1016/j.conb.2004.03.016.
- Creutzfeldt, Otto D, Satoru Watanabe, and Hans D Lux (1966). “Relations between EEG phenomena and potentials of single cortical cells. I. Evoked responses after thalamic and epicortical stimulation.” In: *Electroencephalography and clinical neurophysiology* 20.1, pp. 1–18. ISSN: 0013-4694. DOI: 10.1016/0013-4694(66)90136-2.
- Dadarlat, Maria C, Joseph E O’Doherty, and Philip N Sabes (2015). “A learning-based approach to artificial sensory feedback leads to optimal integration.” In: *Nature neuroscience* 18.1, pp. 138–144. ISSN: 1097-6256. DOI: 10.1038/nn.3883.
- Dai, Ji, Daniel I Brooks, and David L Sheinberg (2013). “Optogenetic and Electrical Microstimulation Systematically Bias Visuospatial Choice in Primates.” In: *Current biology: CB* 24.1, pp. 63–69. ISSN: 0960-9822, 1879-0445. DOI: 10.1016/j.cub.2013.11.011.
- Deisz, R A and D A Prince (1989). “Frequency-dependent depression of inhibition in guinea-pig neocortex in vitro by GABAB receptor feed-back on GABA release.” en. In: *The Journal of physiology* 412, pp. 513–541. ISSN: 0022-3751.
- Diester, Ilka, Matthew T Kaufman, Murtaza Mogri, Ramin Pashaie, Werapong Goo, Ofer Yizhar, Charu Ramakrishnan, Karl Deisseroth, and Krishna V Shenoy (2011). “An optogenetic toolbox designed for primates.” In: *Nature neuroscience* 14.3, pp. 387–397. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn.2749.
- Elsayed, Gamaleldin F and John P Cunningham (2017). “Structure in neural population recordings: an expected byproduct of simpler phenomena?” en. In: *Nature neuroscience* 20.9, pp. 1310–1318. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn.4617.
- Erez, Yaara, Hadass Tischler, Anan Moran, and Izhar Bar-Gad (2010). “Generalized framework for stimulus artifact removal.” en. In: *Journal of neuroscience methods* 191.1, pp. 45–59. ISSN: 0165-0270, 1872-678X. DOI: 10.1016/j.jneumeth.2010.06.005.
- Flesher, Sharlene N, Jennifer L Collinger, Stephen T Foldes, Jeffrey M Weiss, John E Downey, Elizabeth C Tyler-Kabara, Sliman J Bensmaia, Andrew B Schwartz, Michael L Boninger, and Robert A Gaunt (2016). “Intracortical microstimulation of human somatosensory cortex.” en. In: *Science translational medicine* 8.361, 361ra141. ISSN: 1946-6234, 1946-6242. DOI: 10.1126/scitranslmed.aaf8083.
- Gerits, Annelies and Wim Vanduffel (2013). “Optogenetics in primates: a shining future?” In: *Trends in genetics: TIG* 29.7, pp. 403–411. ISSN: 0168-9525. DOI: 10.1016/j.tig.2013.03.004.
- Graziano, Michael S A, Tyson N S Aflalo, Dylan F Cooke, Darcy M Griffin, Heather M Hudson, Abderraouf Belhaj-saïf, Paul D Cheney, Simon Sponberg, Andrew J Spence, Chris H Mullens, Robert J Full, Michael S A Graziano, Tyson N S Aflalo, and Dylan F Cooke (2011). “Arm Movements Evoked by Electrical Stimulation in the Motor Cortex of Monkeys Arm Movements Evoked by Electrical Stimulation in the Motor Cortex of Monkeys.” In: August 2005, pp. 4209–4223. DOI: 10.1152/jn.01303.2004.
- Graziano, Msa, Csr Taylor, and Tirin Moore (2002). “Complex Movements Evoked by Microstimulation of Precentral Cortex.” In: *Neuron* 34, pp. 841–851. ISSN: 0896-6273.

- Griffin, D M, H M Hudson, A Belhaj-Saif, and P D Cheney (2011). “Hijacking Cortical Motor Output with Repetitive Microstimulation.” In: *Journal of Neuroscience* 31.37, pp. 13088–13096. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.6322-10.2011.
- Hanks, Timothy D, Jochen Ditterich, and Michael N Shadlen (2006). “Microstimulation of macaque area LIP affects decision-making in a motion discrimination task.” en. In: *Nature neuroscience* 9.5, pp. 682–689. ISSN: 1097-6256. DOI: 10.1038/nn1683.
- Hashimoto, Takao, Christopher M Elder, and Jerrold L Vitek (2002). “A template subtraction method for stimulus artifact removal in high-frequency deep brain stimulation.” en. In: *Journal of neuroscience methods* 113.2, pp. 181–186. ISSN: 0165-0270. DOI: 10.1016/S0165-0270(01)00491-5.
- Histed, Mark H, Vincent Bonin, and R Clay Reid (2009). “Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation.” In: *Neuron* 63.4, pp. 508–522. ISSN: 0896-6273, 1097-4199. DOI: 10.1016/j.neuron.2009.07.016.
- Histed, Mark H, Amy M Ni, and John H R Maunsell (2013). “Insights into cortical mechanisms of behavior from microstimulation experiments.” In: *Progress in neurobiology* 103, pp. 115–130. ISSN: 0301-0082. DOI: 10.1016/j.pneurobio.2012.01.006.
- Holtzheimer, Paul E and Helen S Mayberg (2011). “Deep Brain Stimulation for Psychiatric Disorders.” In: *Annual review of neuroscience* 34.1, pp. 289–307. ISSN: 0147-006X. DOI: 10.1146/annurev-neuro-061010-113638.
- Jazayeri, Mehrdad and Arash Afraz (2017). “Navigating the Neural Space in Search of the Neural Code.” en. In: *Neuron* 93.5, pp. 1003–1014. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2017.02.019.
- Kara, Prakash, John S Pezaris, Sergey Yurgenson, and R Clay Reid (2002). “The spatial receptive field of thalamic inputs to single cortical simple cells revealed by the interaction of visual and electrical stimulation.” en. In: *Proceedings of the National Academy of Sciences of the United States of America* 99.25, pp. 16261–16266. ISSN: 0027-8424. DOI: 10.1073/pnas.242625499.
- Kimmel, Daniel L and Tirin Moore (2007). “Temporal patterning of saccadic eye movement signals.” en. In: *The Journal of neuroscience: the official journal of the Society for Neuroscience* 27.29, pp. 7619–7630. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.0386-07.2007.
- Klink, P Christiaan, Bruno Dagnino, Marie-Alice Gariel-Mathis, and Pieter R Roelfsema (2017). “Distinct Feedforward and Feedback Effects of Microstimulation in Visual Cortex Reveal Neural Mechanisms of Texture Segregation.” en. In: *Neuron* 95.1, 209–220.e3. ISSN: 0896-6273, 1097-4199. DOI: 10.1016/j.neuron.2017.05.033.
- Kolda, T and B Bader (2009). “Tensor Decompositions and Applications.” In: *SIAM Review* 51.3, pp. 455–500. ISSN: 0036-1445. DOI: 10.1137/07070111X.
- Li, C L, A Ortiz-Galvin, S N Chou, and S Y Howard (1960). “Cortical intracellular potentials in response to stimulation to lateral geniculate body.” en. In: *Journal of neurophysiology* 23, pp. 592–601. ISSN: 0022-3077.
- Logothetis, Nikos K, Mark Augath, Yusuke Murayama, Alexander Rauch, Fahad Sultan, Joziën Goense, Axel Oeltermann, and Hellmut Merkle (2010). “The effects of electrical microstimulation on cortical signal propagation.” en. In: *Nature neuroscience* 13.10, pp. 1283–1291. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn.2631.
- Marcus, S, P Zarzecki, and H Asanuma (1979). “An estimate of effective spread of stimulating current.” In: *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale*. ISSN: 0014-4819.

- Masse, Nicolas Y and Erik P Cook (2010). “Behavioral time course of microstimulation in cortical area MT.” en. In: *Journal of neurophysiology* 103.1, pp. 334–345. ISSN: 0022-3077, 1522-1598. DOI: 10.1152/jn.91022.2008.
- Mena, Gonzalo E, Lauren E Grosberg, Pawel Hottowy, Alan Litke, John Cunningham, E J Chichilnisky, and Liam Paninski (2016). “Electrical Stimulus Artifact Cancellation and Neural Spike Detection on Large Multi-Electrode Arrays.” en.
- Merrill, Daniel R, Marom Bikson, and John G R Jefferys (2005). “Electrical stimulation of excitable tissue: design of efficacious and safe protocols.” en. In: *Journal of neuroscience methods* 141.2, pp. 171–198. ISSN: 0165-0270. DOI: 10.1016/j.jneumeth.2004.10.020.
- Montgomery Jr, Erwin B, John T Gale, and He Huang (2005). “Methods for isolating extracellular action potentials and removing stimulus artifacts from microelectrode recordings of neurons requiring minimal operator intervention.” en. In: *Journal of neuroscience methods* 144.1, pp. 107–125. ISSN: 0165-0270. DOI: 10.1016/j.jneumeth.2004.10.017.
- Müller, Jan, Douglas J Bakkum, and Andreas Hierlemann (2012). “Sub-millisecond closed-loop feedback stimulation between arbitrary sets of individual neurons.” en. In: *Frontiers in neural circuits* 6, p. 121. ISSN: 1662-5110. DOI: 10.3389/fncir.2012.00121.
- Murasugi, C M, C D Salzman, and W T Newsome (1993). “Microstimulation in visual area MT: effects of varying pulse amplitude and frequency.” en. In: *The Journal of neuroscience: the official journal of the Society for Neuroscience* 13.4, pp. 1719–1729. ISSN: 0270-6474.
- Nowak, L G and J Bullier (1996). “Spread of stimulating current in the cortical grey matter of rat visual cortex studied on a new in vitro slice preparation.” en. In: *Journal of neuroscience methods* 67.2, pp. 237–248. ISSN: 0165-0270.
- Nowak, L G and J Bullier (1998a). “Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter. I. Evidence from chronaxie measurements.” en. In: *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale* 118.4, pp. 477–488. ISSN: 0014-4819.
- Nowak, L G and J Bullier (1998b). “Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter. II. Evidence from selective inactivation of cell bodies and axon initial segments.” en. In: *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale* 118.4, pp. 489–500. ISSN: 0014-4819.
- O’Doherty, Joseph E, Mikhail A Lebedev, Peter J Ifft, Katie Z Zhuang, Solaiman Shokur, Hannes Bleuler, and Miguel A L Nicolelis (2011). “Active tactile exploration using a brain-machine-brain interface.” en. In: *Nature* 479.7372, pp. 228–231. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature10489.
- O’Doherty, Joseph E, Mikhail a Lebedev, Peter J Ifft, Katie Z Zhuang, Solaiman Shokur, Hannes Bleuler, and Miguel a L Nicolelis (2011). “Active tactile exploration using a brain-machine-brain interface.” In: *Nature*, pp. 1–5. ISSN: 0028-0836. DOI: 10.1038/nature10489.
- Ohayon, Shay, Piercesare Grimaldi, Nicole Schweers, and Doris Y Tsao (2013). “Saccade modulation by optical and electrical stimulation in the macaque frontal eye field.” en. In: *The Journal of neuroscience: the official journal of the Society for Neuroscience* 33.42, pp. 16684–16697. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.2675-13.2013.
- O’Keefe, D T, G M Lyons, A E Donnelly, and C A Byrne (2001). “Stimulus artifact removal using a software-based two-stage peak detection algorithm.” en. In: *Journal of neuroscience methods* 109.2, pp. 137–145. ISSN: 0165-0270. DOI: 10.1016/S0165-0270(01)00407-1.

- Otchy, Timothy M, Steffen B E Wolff, Juliana Y Rhee, Cengiz Pehlevan, Risa Kawai, Alexandre Kempf, Sharon M H Gobes, and Bence P Ölveczky (2015). “Acute off-target effects of neural circuit manipulations.” en. In: *Nature* 528.7582, pp. 358–363. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature16442.
- Otto, Kevin J, Patrick J Rousche, and Daryl R Kipke (2005). “Microstimulation in auditory cortex provides a substrate for detailed behaviors.” In: *Hearing research* 210.1-2, pp. 112–117. ISSN: 0378-5955. DOI: 10.1016/j.heares.2005.08.004.
- Ranck Jr, J B (1975). “Which elements are excited in electrical stimulation of mammalian central nervous system: a review.” en. In: *Brain research* 98.3, pp. 417–440. ISSN: 0006-8993.
- Rattay, F (1999). “The basic mechanism for the electrical stimulation of the nervous system.” en. In: *Neuroscience* 89.2, pp. 335–346. ISSN: 0306-4522.
- Rattay, Frank and Susanne Resatz (2004). “Effective electrode configuration for selective stimulation with inner eye prostheses.” en. In: *IEEE transactions on bio-medical engineering* 51.9, pp. 1659–1664. ISSN: 0018-9294. DOI: 10.1109/TBME.2004.828044.
- Romo, R, A Hernández, A Zainos, C D Brody, and L Lemus (2000). “Sensing without touching: psychophysical performance based on cortical microstimulation.” In: *Neuron* 26.1, pp. 273–278. ISSN: 0896-6273.
- Salzman, C D, Kenneth H Britten, and William T Newsome (1990). “Cortical microstimulation influences perceptual judgements of motion direction.” In: *Nature* 346, pp. 174–177. ISSN: 0028-0836.
- Seely, Jeffrey S, Matthew T Kaufman, Stephen I Ryu, Krishna V Shenoy, John P Cunningham, and Mark M Churchland (2016). “Tensor Analysis Reveals Distinct Population Structure that Parallels the Different Computational Roles of Areas M1 and V1.” en. In: *PLoS computational biology* 12.11, e1005164. ISSN: 1553-734X, 1553-7358. DOI: 10.1371/journal.pcbi.1005164.
- Seidemann, Eyal, Amos Arieli, Amir Grinvald, and Hamutal Slovin (2002). “Dynamics of Depolarization and Hyperpolarization in the Frontal Cortex and Saccade Goal.” en. In: *Science* 295.5556, pp. 862–865. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1066641.
- Shenoy, Krishna V, Matthew T Kaufman, Maneesh Sahani, and Mark M Churchland (2011). “A dynamical systems view of motor preparation Implications for neural prosthetic system design.” In: *Progress in brain research* 192, pp. 33–58. ISSN: 0079-6123, 1875-7855. DOI: 10.1016/B978-0-444-53355-5.00003-8.
- Sparks, D L and L E Mays (1983). “Spatial localization of saccade targets. I. Compensation for stimulation-induced perturbations in eye position.” en. In: *Journal of neurophysiology* 49.1, pp. 45–63. ISSN: 0022-3077.
- Stoney Jr, S D, W D Thompson, and H Asanuma (1968). “Excitation of pyramidal tract cells by intracortical microstimulation: effective extent of stimulating current.” en. In: *Journal of neurophysiology* 31.5, pp. 659–669. ISSN: 0022-3077.
- Stratford, K J, K Tarczy-Hornoch, K A Martin, N J Bannister, and J J Jack (1996). “Excitatory synaptic inputs to spiny stellate cells in cat visual cortex.” en. In: *Nature* 382.6588, pp. 258–261. ISSN: 0028-0836. DOI: 10.1038/382258a0.
- Strick, Peter L (2002). “Stimulating research on motor cortex.” In: *Nature neuroscience* 5.8, pp. 4–5. ISSN: 1097-6256.
- Tehovnik, E J (1996). “Electrical stimulation of neural tissue to evoke behavioral responses.” In: *Journal of Neuroscience Methodsuroscience methods* 65.1, pp. 1–17. ISSN: 0165-0270.
- Tehovnik, E J and M A Sommer (1997). “Electrically evoked saccades from the dorsomedial frontal cortex and frontal eye fields: a parametric evaluation reveals

- differences between areas.” en. In: *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale* 117.3, pp. 369–378. ISSN: 0014-4819.
- Tehovnik, E J, a S Tolias, F Sultan, W M Slocum, and N K Logothetis (2006). “Direct and indirect activation of cortical neurons by electrical microstimulation.” In: *Journal of neurophysiology* 96.2, pp. 512–521. ISSN: 0022-3077. DOI: 10.1152/jn.00126.2006.
- Tehovnik, Edward J and Warren M Slocum (2007). “Phosphene induction by microstimulation of macaque V1.” In: *Brain research reviews* 53.2, pp. 337–343. ISSN: 0165-0173. DOI: 10.1016/j.brainresrev.2006.11.001.
- Tehovnik, Edward J, Warren M Slocum, and Peter H Schiller (2004). “Microstimulation of V1 delays the execution of visually guided saccades.” en. In: *The European journal of neuroscience* 20.1, pp. 264–272. ISSN: 0953-816X. DOI: 10.1111/j.1460-9568.2004.03480.x.
- Tehovnik, Edward J, Warren M Slocum, and Peter H Schiller (2005). “Delaying visually guided saccades by microstimulation of macaque V1: spatial properties of delay fields.” en. In: *The European journal of neuroscience* 22.10, pp. 2635–2643. ISSN: 0953-816X. DOI: 10.1111/j.1460-9568.2005.04454.x.
- Tehovnik, Edward J, Warren M Slocum, Stelios M Smirnakis, and Andreas S Tolias (2009). “Microstimulation of visual cortex to restore vision.” en. In: *Progress in brain research* 175, pp. 347–375. ISSN: 0079-6123, 1875-7855. DOI: 10.1016/S0079-6123(09)17524-6.
- Thomson, A M and J Deuchars (1994). “Temporal and spatial properties of local circuits in neocortex.” en. In: *Trends in neurosciences* 17.3, pp. 119–126. ISSN: 0166-2236.
- Varela, J A, K Sen, J Gibson, J Fost, L F Abbott, and S B Nelson (1997). “A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex.” en. In: *The Journal of neuroscience: the official journal of the Society for Neuroscience* 17.20, pp. 7926–7940. ISSN: 0270-6474.
- Wagenaar, Daniel A and Steve M Potter (2002). “Real-time multi-channel stimulus artifact suppression by local curve fitting.” en. In: *Journal of neuroscience methods* 120.2, pp. 113–120. ISSN: 0165-0270. DOI: 10.1016/S0165-0270(02)00149-8.
- Watanabe, S, M Konishi, and O D Creutzfeldt (1966). “Postsynaptic potentials in the cat’s visual cortex following electrical stimulation of afferent pathways.” en. In: *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale* 1.3, pp. 272–283. ISSN: 0014-4819.
- Wichmann, Thomas and Mahlon R DeLong (2006). “Deep brain stimulation for neurologic and neuropsychiatric disorders.” en. In: *Neuron* 52.1, pp. 197–204. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2006.09.022.
- Wichmann, Thomas and Annaelle Devergnas (2011). “A novel device to suppress electrical stimulus artifacts in electrophysiological experiments.” en. In: *Journal of neuroscience methods* 201.1, pp. 1–8. ISSN: 0165-0270, 1872-678X. DOI: 10.1016/j.jneumeth.2011.06.026.
- Williams, A, H Kim, F Wang, S Vyas, K V Shenoy, M Schnitzer, T Kolda, and S Ganguli (2017). *Dimension reduction of multi-trial neural data by tensor decomposition. Frontiers in Neuroscience (T-28)*.