ROBUST ONLINE MULTIBAND DRIFT ESTIMATION IN ELECTROPHYSIOLOGY DATA

Charlie Windolf[1,2], *Angelique C. Paulk*[3], *Yoav Kfir*[4], *Eric Trautmann*[2], *Samuel Garcia*[5], *Domokos Meszéna*[3], *William Muñoz*[4], *Irene Caprara*[4], *Mohsen Jamali*[4], *Julien Boussard*[1,2], *Ziv M. Williams*[4], *Sydney S. Cash*[3], *Liam Paninski*[1,2], *Erdem Varol*[1,2,6]

[1] Department of Statistics, Columbia University
[2] Zuckerman Institute, Columbia University
[3] Department of Neurology, Center for Neurotechnology and Neurorecovery, Massachusetts General Hospital, Harvard Medical School
[4] Department of Neurosurgery, Massachusetts General Hospital, Harvard Medical School
[5] Centre national de la recherche scientifique, Centre de recherche en neuroscience de Lyon
[6] Department of Computer Science & Engineering, New York University

ABSTRACT

High-density electrophysiology probes have opened new possibilities for systems neuroscience in human and non-human animals, but probe motion (or drift) while recording poses a challenge for downstream analyses, particularly in human recordings. Here, we improve on the state of the art for tracking this drift with an algorithm termed **DREDge** (**D**ecentralized **R**egistration of **E**lectrophysiology **D**ata) with four major contributions. First, we extend previous decentralized methods to exploit *multiband* information, leveraging the local field potential (LFP), in addition to spikes detected from the action potentials (AP). Second, we show that the LFP-based approach enables registration at *sub-second* temporal resolution. Third, we introduce an efficient *online* motion tracking algorithm, allowing the method to scale up to longer and higher spatial resolution recordings, which could facilitate real-time applications. Finally, we improve the *robustness* of the approach by accounting for the nonstationarities that occur in real data and by automating parameter selection. Together, these advances enable fully automated scalable registration of challenging datasets from both humans and mice.

Index Terms— Decentralization, online optimization, electrophysiology, motion estimation, preprocessing, extracellular recording, neuropixels

1. INTRODUCTION

Dense electrophysiology via multi-channel microelectrode probes, such as the Neuropixels probe [1, 2], provide an unprecedented view of neural circuits in human and non-human animal brains at extremely high resolution both temporally (30 kHz) and spatially (20-400 µm). In contrast with older recording technologies using lower spatial resolution arrays, the latest probes that penetrate into the brain allow us to measure activity in large populations of neurons (several hundreds) and the local field potential (LFP) with high fidelity. Since their introduction and ongoing development, these probes have allowed testing a variety of novel scientific hypotheses, including those related to neural correlates of consciousness [3], motor planning [4] and visual choice tasks [5], cementing their role as a staple tool for systems neuroscience for the foreseeable future. Furthermore, Neuropixels probes have recently been successfully employed for high-quality intraoperative recordings in awake and anesthetized humans [6, 7], enabling us to directly answer fundamental questions about human brain physiology with possible clinical implications.

However, several biological and physical sources of noise and variability reduce the neural recording effectiveness of Neuropixels probes [8]. In the case of in vivo measurements, especially in human participants, the probe signal can be impacted by brain motion effects due to the heart rate and breathing of the patients as well as unexpected brain shifting during recording (such as when the participants start to talk in clinically indicated awake tasks) [6]. This motion results in drift in the voltage measurements across channels, potentially corrupting the ability to isolate single unit activity on a given channel, which may lead to undersampling of spikes or over-splitting of identified unit clusters [9, 10, 11]. Together, these errors reduce the ability to characterize the functional activity of neural populations that are measured.

While drift affects voltage measurements in any rigid electrode, it is both visible and fixable in probes with very high spatial resolution such as Neuropixels. There are two main approaches to solving the motion drift problem in dense electrophysiology probes. Experimental approaches involve designing hardware to steady the probe during measurement. For example, probe movement can be stabilized at the open craniotomy as done in non-human primate preparations using

O-rings or other materials pressing down on the brain [12]. However, in the human operating room, attempts to stabilize the probe using the same techniques could induce problematic capillary damage on the surface of the cortex and may require considerable careful non-human testing before implementing these approaches.

Two main computational methods have been developed to estimate motion drift in Neuropixels probes. Kilosort 2.5 [2] uses a template-based approach, computing a template signal, and then using cross-correlations to shift the drifting signal back to the template space. In contrast, Varol et al. [13] take a *decentralized* approach, measuring local signal shifts between all pairs of time-binned signals to learn a global displacement vector. These two techniques have been in wide deployment by several experimental groups [6, 14], however, they have three main shortcomings. Kilosort 2.5's template-based approach is plagued by model misspecification if the neural signal rapidly changes, eliminating the possibility that the recording would be described by a single template. The decentralized approach [13] circumvents this issue but is hindered by the computation of a $T \times T$ matrix that might be prohibitively large in chronic recordings. Furthermore, both of these methods have a variety of parameters such as time bin sizes and correlation cutoff parameters that need to be carefully tuned by practitioners to recover the tracked motion, reducing their robustness in high-throughput settings. Last and most importantly, these two techniques rely on the spike waveforms, probe locations, and discrete spike times of the high-frequency action potential (AP) band of the voltage signal, but ignore the smoother local field potential (LFP) band to estimate drift. Although AP band approaches require localizing spikes over time to estimate motion, the number of spikes in smaller time bins becomes too sparse to accurately compute motion-induced shifts. Hence, AP band approaches are limited to estimate drift on the order of $\sim$1 second temporal resolution. On the other hand, the LFP band possesses smoother and more continuous signal across the entire recording, which has the potential to capture drift with a temporal resolution that is only limited by the sampling rate of measurement ($\sim$2.5KHz).

To overcome these obstacles, we introduce a novel extension to the decentralized registra-

tion approach [13] with the following **main contributions**: 1) fast online GPU based-optimization of displacement estimates in large-scale recordings, 2) sub-second temporal resolution, 3) automatic statistical tuning of parameters, and 4) multiband (LFP and AP based) motion estimation. Together, these four new features provide us with a fast, scalable, and robust approach for motion-correcting electrophysiological data with minimal or no parameter tuning. We term our approach **DREDge** which stands for **D**ecentralized **R**egistration of **E**lectrophysiology **D**ata.

We validate **DREDge** on two human Neuropixels recordings [6] and a mouse recording [2] from two separate research groups. We measure its performance in terms of registration quality and run-time and compare it with state-of-the-art techniques [13] and [2]. See `https://github.com/evarol/DREDge` for open-source code.

2. METHODS

We first motivate the **DREDge** algorithm using the decentralized registration approach [13] that it builds on in section 2.1. Then we describe several additions that enable **DREDge** to estimate drift in multiband electrophysiology data in a robust, efficient, and online manner in sections 2.2,2.4,and2.5. In section 2.6, we describe the procedure for realigning electrophysiology data through interpolation, adjusting for the shifting signal after motion estimation. Finally, we summarize all of these steps in algorithmic pseudocode in section 2.7.

2.1. Review: decentralized registration.

Given a $D \times T$ signal $\mathbf{R}$ with columns $\mathbf{r}_t$, our goal is to discover $\mathbf{p} \in \mathbb{R}^T$ such that $\mathbf{p}_t$ is the displacement of the tth time bin $\mathbf{r}_t$. Here, $\mathbf{R}$ may be the signal from the LFP band, or it may be a rasterized representation of spiking activity from the AP band 1. In the LFP case, in practice, rather than using LFP directly, we spatially filter it to obtain one form of what we call current source density (CSD) [15], which provides spatially sharpened local features for registration. In the AP case, we construct $\mathbf{R}$ by first estimating depth positions along the probe for all spikes,

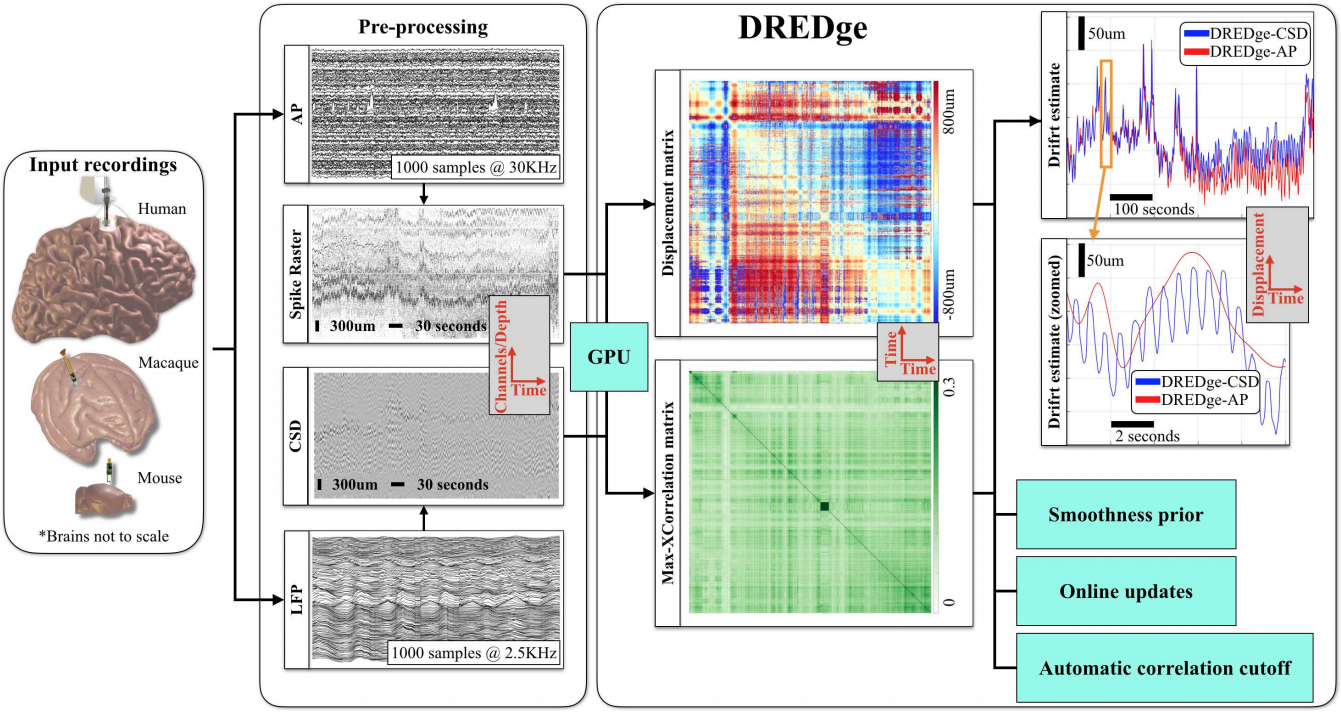

Fig. 1. Overview of the **DREDge** algorithm. AP or LFP data from electrophysiology recordings from human, non-human primate, or mouse data is first processed to yield depth (D) $\times$ time (T) features $\mathbf{R}$. Then each time-binned feature is cross-correlated with every other time-binned feature to generate a $T \times T$ displacement shift matrix and its corresponding maximum cross-correlation. This step is done efficiently on the GPU. The displacement matrix is filtered using an automatically derived correlation cutoff and the remaining terms are used to solve a *centralization* equation to estimate drift estimates for each time bin. This procedure is robustified using priors that ensure that nearby drift terms are close to each other. Furthermore, for large recordings, the entire routine is done in smaller overlapping time chunks in an online fashion to efficiently calculate drift without the need to store a large $T \times T$ displacement shift matrix in memory.

typically using a point-source localization model [16]. Then, we divide the depth domain into D bins and the time domain into T bins and set $\mathbf{R}_{dt}$ to the average amplitude of spikes falling in the dth depth and tth time bins, taking the average in empty bins to be 0 by convention. In the LFP case, we construct $\mathbf{R}$ by directly taking the voltage signal at time t at each channel location with depth d such that $\mathbf{R}_{dt}$ is the LFP value at depth d and time t (see Fig.1 for examples of $\mathbf{R}$ in the left side).

The decentralized approach [13] infers $\mathbf{p}$ using estimates of the displacements between all

pairs of time bins, represented in a $T \times T$ antisymmetric matrix $\mathbf{D}$ with

$$\mathbf{D}_{tt'} = \arg\max_{\Delta y} \operatorname{corr}(\mathbf{r}_t(y), \mathbf{r}_{t'}(y + \Delta y)), \tag{1}$$

the spatial offset which maximizes the correlation between the two time bins. For AP-band problems, T is measured in seconds and this matrix is relatively small; for longer recordings or LF bands, the online method below avoids large matrices. The "centralization" problem, then, is to find

$$\hat{\mathbf{p}} = \arg\min_{\mathbf{p}} \|\mathbf{D} - (\mathbf{p}\mathbf{1}^\top - \mathbf{1}\mathbf{p}^\top)\|_F^2. \tag{2}$$

The solution to this simple version of the problem is the row mean:

$$\hat{\mathbf{p}}_t = \frac{1}{T} \sum_{t'=1}^{T} D_{tt'}. \tag{3}$$

2.2. Extension: decentralized registration with correlation-based subsampling.

If due to nonstationarities in the neural activity or large amounts of drift, two time bins $\mathbf{r}_t$ and $\mathbf{r}_{t'}$ do not contain the same features, then their pairwise displacement $\mathbf{D}_{tt'}$ should be excluded from the objective in Eq. (2). For example, this could occur in two time bins that are temporally distant, and the subject has completely different neural subpopulations firing, preventing a good shift to be found that overlays these two patterns.

Thus a simple heuristic approach is to include only those pairs of time bins whose maximal normalized cross-correlation exceeds a certain threshold, which indicates similar neural activity patterns up to a shift. To that end, we fix a correlation threshold θ, and let $\mathbf{C}_{tt'} = \operatorname{corr}(\mathbf{r}_t(y), \mathbf{r}_{t'}(y + D_{tt'}))$, to be the correlation corresponding to the displacement estimate $\mathbf{D}_{tt'}$. Let $\mathbf{S} \in \mathbb{R}^{T \times T}$ be the thresholded correlation matrix with $\mathbf{S}_{tt'} = \mathbb{1}_{\mathbf{C}_{tt'} > \theta}$. Now, we modify Eq. (2) to its subsampled form:

$$\hat{\mathbf{p}}_\theta = \arg\min_{\mathbf{p}} \left\| \mathbf{S} \circ [\mathbf{D} - (\mathbf{p}\mathbf{1}^\top - \mathbf{1}\mathbf{p}^\top)] \right\|_F^2, \tag{4}$$

where $\circ$ indicates the elementwise product.

To efficiently solve this problem, consider the set of pairs of times $\{(t_k, t'_k)\}_{k=1,\ldots,K}$, where $K = \sum_{t,t'} S_{tt'}$, such that $\mathbf{C}_{t_k t'_k} > \theta$. Now, let $\vec{\mathbf{D}}_\theta \in \mathbb{R}^K$ be the vector whose kth element is $\mathbf{D}_{t_k t'_k}$, and let the matrix $\mathbf{A} \in \mathbb{R}^{K \times T}$ have elements $\mathbf{A}_{kt} = [\mathbf{I} \otimes \mathbf{1}]_{t_k t} - [\mathbf{1} \otimes \mathbf{I}]_{t'_k t}$. Then, Eq. (4) can be rewritten in the least squares form

$$\min_{\mathbf{p}} \|\vec{\mathbf{D}}_\theta - \mathbf{A}\mathbf{p}\|_2^2, \tag{5}$$

which can be solved by a sparse least squares solver (e.g. LSMR [17]). Note that such least squares systems usually require $O(T^3)$ to solve, but since $\mathbf{A}$ only has $2K << T^2$ non-zero elements, the complexity is reduced to $O(K^2 T)$.

2.3. Adaptive choice of subsampling threshold:

Since different modalities and recordings will have different statistics, it is important to find a way to set the correlation threshold robustly. One simple solution is to set this threshold to a low percentile of the distributions of maximum cross-correlations between neighboring time bins. Since almost all neighboring time bins contain the same features, this method will discover a correlation threshold that is suitable for non-neighboring time bins in a way that adapts to the characteristics of the recording. In the online method below, a threshold can be chosen in this manner for each new batch in order to adapt to nonstationarities in the signal.

2.4. Extension: smoothing prior for robustness.

Since some time bins may be poorly correlated with the others, these bins may be separated from or only weakly linked to the rest through the subsampled cost function. This effect can lead to jumps in the resulting $\hat{\mathbf{p}}$. To resolve this issue, we place a standard Brownian prior on $\hat{\mathbf{p}}$: $\hat{\mathbf{p}}_{t+1} - \hat{\mathbf{p}}_t \sim N(0, 1)$. Since the difference operator is linear and sparse, it is simple to incorporate

this prior into the above least squares framework, effectively turning our objective function into:

$$\min_{\mathbf{p}} \|\vec{\mathbf{D}}_\theta - \mathbf{A}\mathbf{p}\|_2^2 + \|\Delta\mathbf{p}\|_2^2, \tag{6}$$

where Δ is the $T \times T$ finite difference matrix with elements $\Delta_{t,t} = 2$ and $\Delta_{t,t-1} = \Delta_{t-1,t} = -1$.

2.5. Extension: online motion tracking.

The methods in the previous two sections scale at least quadratically in T, since we compute $T \times T$ matrices $\mathbf{C}$ and $\mathbf{D}$ and solve a $T \times T$ system. When registering spiking data, where the time bins typically have lengths on the order of 1 second, this is no issue except in very long chronic recordings, or when registering sub-second resolution drift in human patients using the CSD or LFP signal.

To mitigate this, we estimate drift in chunks in an 'online' fashion. First, break the data $\mathbf{R}$ into C chunks of size at most $D \times T_0$, $\mathbf{R}^{(c)}$, $c = 1, \ldots, C$. We initialize the algorithm by using the batch version of **DREDge**(Algorithm 1) to find $\mathbf{p}^{(1)}$, the displacement in the first block. Then, given the previous chunk's displacement estimate $\mathbf{p}^{(c)}$, we can find the current chunk's displacement estimate $\mathbf{p}^{(c+1)}$ according to the problems in Eqs. (2) and (4). Proceeding through the recording chunk by chunk, we can recover the full displacement estimate by concatenating those in each chunk. Since the sizes of the chunks' sub-problems are bounded, this method will scale *linearly* in the total length of the recording. We present the online method in the case without correlation-based subsampling, since the notation for the subsampling will complicate things unnecessarily, but the extension is direct as in Sec. 2.2 and our results use subsampling.

Let $\mathbf{p} = [\hat{\mathbf{p}}^{(c)}; \mathbf{p}^{(c+1)}] \in \mathbb{R}^{2T_0}$ be the concatenation of the two displacement vectors, and define these chunks' $2T_0 \times 2T_0$ displacement matrix

$$\mathbf{D}^{(c+1)} = \begin{bmatrix} \mathbf{D}_{(c,c)} & \mathbf{D}_{(c,c+1)} \\ -\mathbf{D}_{(c,c+1)}^\top & \mathbf{D}_{(c+1,c+1)} \end{bmatrix} \tag{7}$$

where the $T_0 \times T_0$ blocks $\mathbf{D}_{(c,c)}$, $\mathbf{D}_{(c,c+1)}$, and $\mathbf{D}_{(c+1,c+1)}$ are pairwise displacement estimates in the previous block $\mathbf{R}^{(c)}$, between the two blocks $\mathbf{R}^{(c)}$ and $\mathbf{R}^{(c+1)}$, and within the current block $\mathbf{R}^{(c+1)}$, respectively. Consider Eq. (2), modified to hold $\hat{\mathbf{p}}^{(c)}$ fixed:

$$\hat{\mathbf{p}}^{(c+1)} = \underset{\mathbf{p}^{(c+1)}}{\arg\min} \left\| \mathbf{D}^{(c+1)} - (\mathbf{p}\mathbf{1}^\top - \mathbf{1}\mathbf{p}^\top) \right\|_F^2. \tag{8}$$

This online registration problem ensures that $\hat{\mathbf{p}}^{(c+1)}$ both aligns with the past estimate and centralizes the pairwise displacement estimates for the current time bins. Removing terms which do not include $\mathbf{p}^{(c+1)}$, Eq. (8) simplifies to

$$\min_{\mathbf{p}^{(c+1)}} \left\| \mathbf{D}_{(c+1,c+1)} - (\mathbf{p}^{(c+1)}\mathbf{1}^\top - \mathbf{1}\mathbf{p}^{(c+1)\top}) \right\|_F^2 + 2 \left\| (\hat{\mathbf{p}}^{(c)}\mathbf{1}^\top - \mathbf{D}_{(c,c+1)}) - \mathbf{1}\mathbf{p}^{(c+1)\top}) \right\|_F^2. \tag{9}$$

Here, the first term is the usual decentralized objective, ensuring the fidelity of the estimate in the current block. The cross term pushes $\mathbf{p}^{(c+1)}$ towards the column means of $\hat{\mathbf{p}}^{(c)}\mathbf{1}^\top - \mathbf{D}_{(c,c+1)}$, encouraging alignment with the past. As above, we can incorporate subsampling and solve this problem by rewriting it in OLS form and using a sparse least squares solver.

2.6. Data interpolation after motion estimation

After estimating the motion trace, the SpikeInterface [11] library was used to interpolate the underlying raw data to correct for the motion. A simple method was used: first, a displaced coordinate is computed for each channel at each time according to the drift estimate, leading to a set of time-varying coordinates for a virtual probe. If the displacement at a given time was zero, then the original recording is used at that time. Otherwise, each virtual channel's value is computed by a weighted average of the three channels closest to its displaced location, where the weights are the inverse distances from the virtual channel.

2.7. Algorithmic details

We summarize the above pipeline to estimate motion drift in electrophysiology data as the **DREDge** algorithm whose pseudocode details can be found in Algorithm 1. Note that the spatiotemporal signal matrix $\mathbf{R} \in \mathbb{R}^{D \times T}$ can either be based on the AP band or the LFP band, resulting in two versions of the algorithm that we refer to below as **DREDge**-AP and **DREDge**-CSD. Algorithm 1 denotes the "batch" version of motion estimation. In large data cases such as in chronic recordings or in real-time applications, Algorithm 1 can serve as a *subroutine* for an online estimation of motion as described in section 2.5 where smaller time chunks of signal matrices $\mathbf{R}$ act as input in a streaming fashion.

Algorithm 1 DREDge (batch).

Input: signal $\mathbf{R} \in \mathbb{R}^{D \times T}$, neighbor correlation quantile q, true/false value `continuity-prior`
Output: motion estimate $\mathbf{p} \in \mathbb{R}^T$

Compute optimal displacements and correlations
Allocate $T \times T$ matrices $\mathbf{D}, \mathbf{C}$
for $1 \leq t, t' \leq T$ **do**
  $\mathbf{D}_{tt'} \leftarrow \arg\max_{\Delta y} \operatorname{corr}(\mathbf{r}_t(y), \mathbf{r}_{t'}(y + \Delta y))$
  $\mathbf{C}_{tt'} \leftarrow \operatorname{corr}(\mathbf{r}_t(y), \mathbf{r}_{t'}(y + D_{tt'}))$
end for

Compute the adaptive correlation threshold θ
$\theta \leftarrow \operatorname{quantile}_q\{\mathbf{C}_{t,t+1} : t = 1, \ldots, T-1\}$

Solve the centralization problem
Find pairs of times $\{(t_k, t'_k)\}_{k=1,\ldots,K}$ such that $C_{t_k,t'_k} \geq \theta$
Compute $\mathbf{d} \in \mathbb{R}^K$, where $\mathbf{d}_k = D_{t_k,t'_k}$
Compute the sparse matrix $\mathbf{A} \in \mathbb{R}^{K \times T}$, where $\mathbf{A}_{kt} = [\mathbf{I} \otimes \mathbf{1}]_{t_k t} - [\mathbf{1} \otimes \mathbf{I}]_{t'_k t}$
if `continuity-prior` **then**
  Grow $\mathbf{d}$ by appending a vector of $T-1$ 0s
  Grow $\mathbf{A}$ by appending $T-1$ new rows $\mathbf{A}_{K+1}, \ldots, \mathbf{A}_{K+T}$, where $\mathbf{A}_{K+t} = \delta_{t+1} - \delta_t$
end if
Find $\mathbf{p}$ by solving the sparse least squares problem $\min_{\mathbf{p}} \|\mathbf{d} - \mathbf{A}p\|_2^2$ using a solver like LSMR [17]

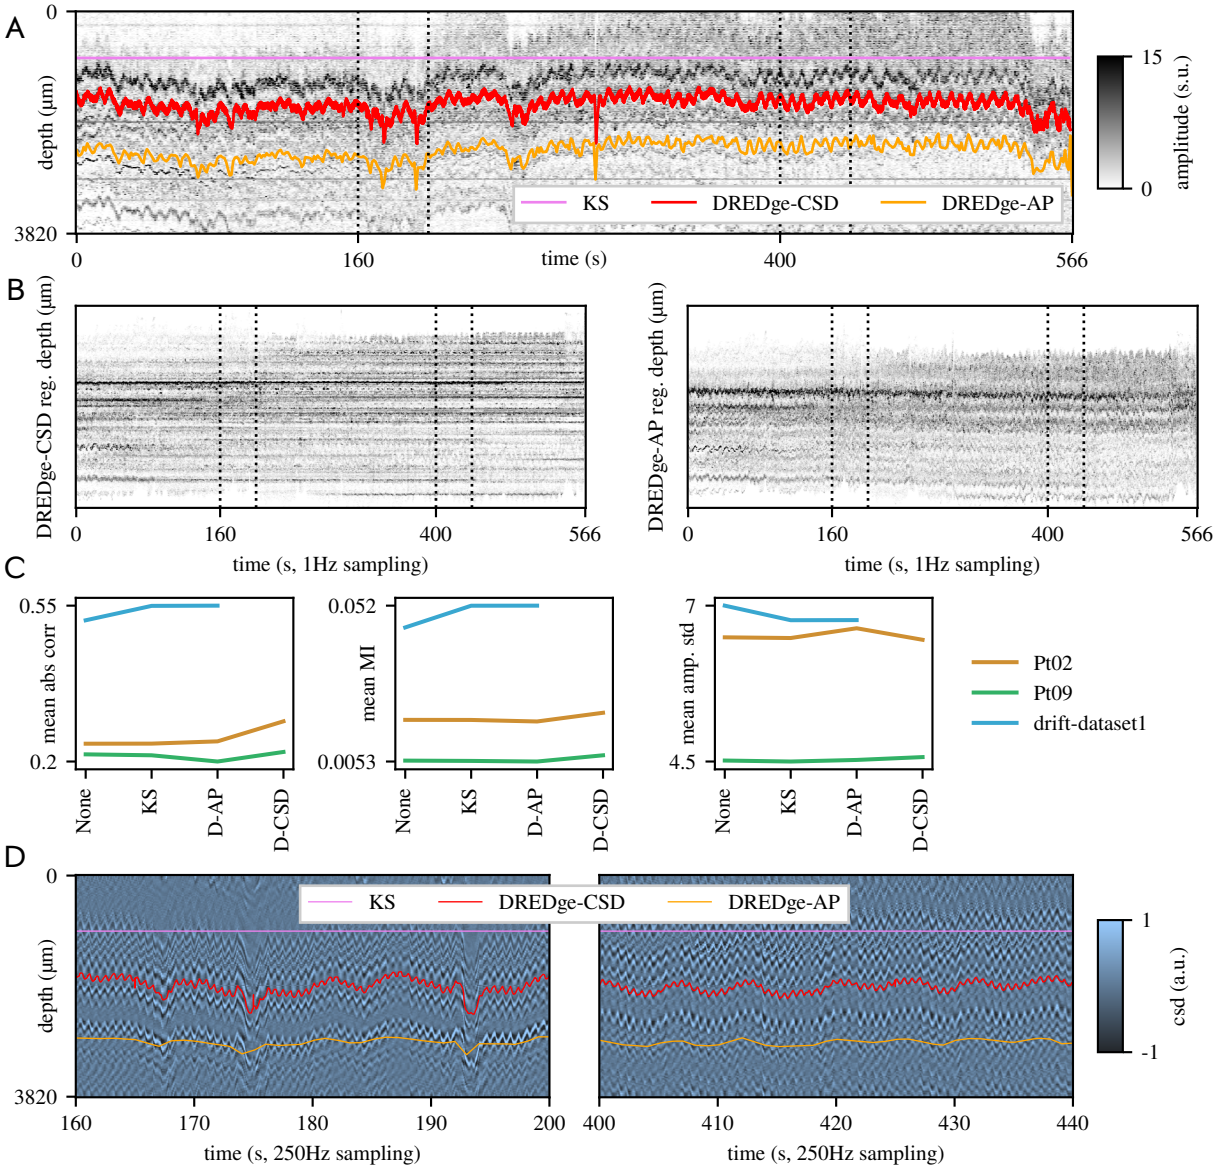

Fig. 2. (a) Rasterized spike amplitudes from the Pt. 02 dataset [6], with estimated motion traces from Kilosort 2.5 and **DREDge**; regions of interest starting at 160 and 400s displayed in (d) show that the large jumps do correspond to the observed spatially filtered LFP which we refer to as CSD (d). (b) Spike amplitude maps after shifting the spike depths according to **DREDge** displacement estimates (CSD, top, and AP, below). (c) Metrics on three datasets, two human (Pts. 02 and 09, [6]) and one mouse (drift-dataset1, [2]). The first two metrics are the mean correlation/mutual information across pairs of rasterized spike activity time bins; the last is the mean standard dev. of the amplitude in-depth bins. (d) Motion traces for 40s regions of interest in the CSD. Color scale shared in (a) and (b).

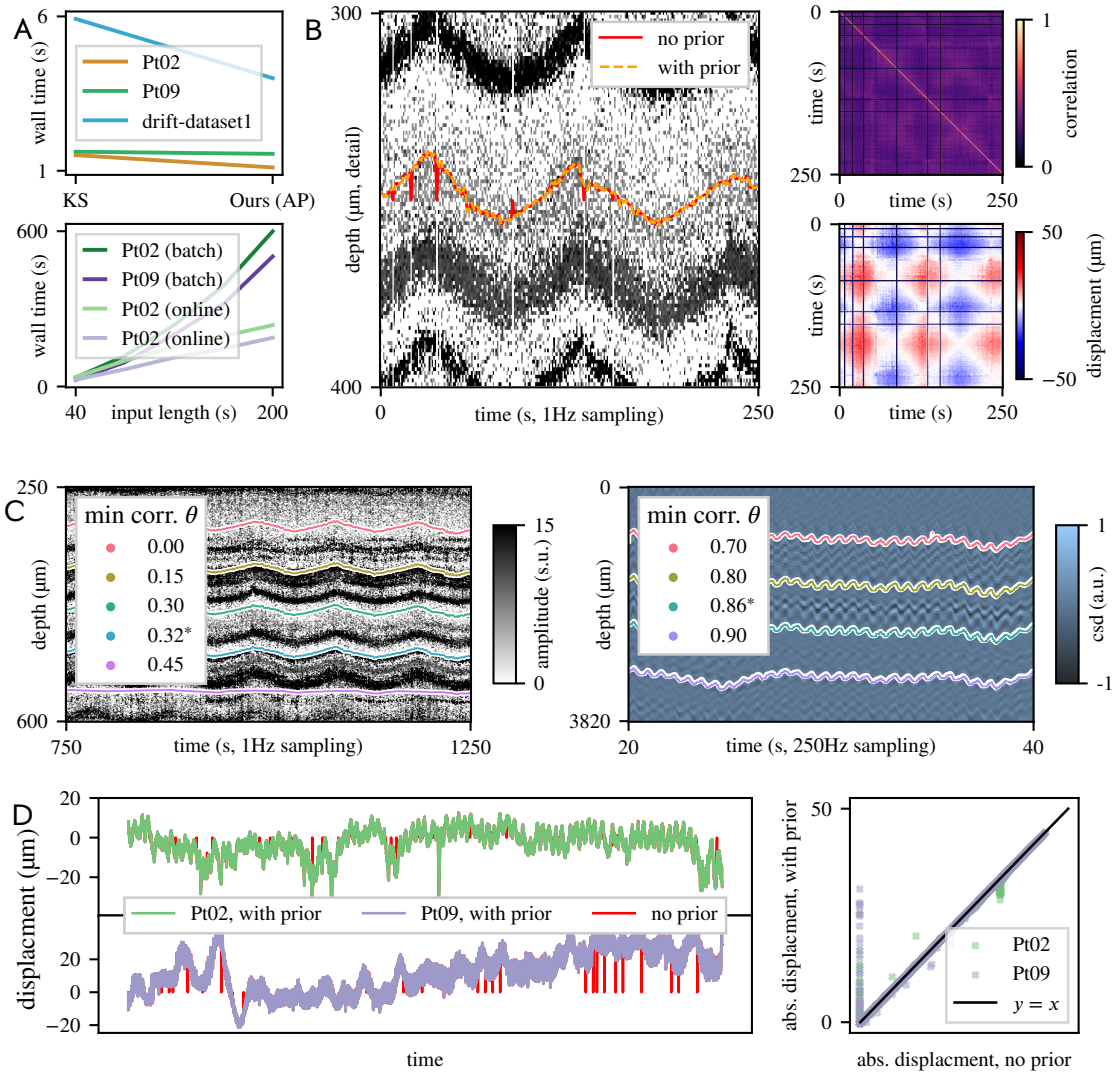

Fig. 3. (a) Timing comparisons: top, **DREDge**-AP method vs. KS; below, online vs. offline in CSD. **DREDge**-AP is on par with KS thanks to a fast GPU implementation. (b) Comparing results with no prior and with prior on a test version of drift-dataset1 [2] with 5% of time bins erased at random: left, estimated motion traces over rasterized spike amplitudes; right, correlation and displacement matrices; horizontal and vertical stripes in these reflect the erased time bins. (c) Motion estimates using adaptive correlation threshold choices (marked *) vs. a grid of other choices: top, in real mouse AP data (θ^* =5 percentile of neighbor correlations) and bottom, in human CSD (0.1 percentile). (d) Left, full CSD displacement estimates with prior (purple, green) and without (red); right, adding the prior pushes estimates for isolated frames away from 0 without causing shrinkage.

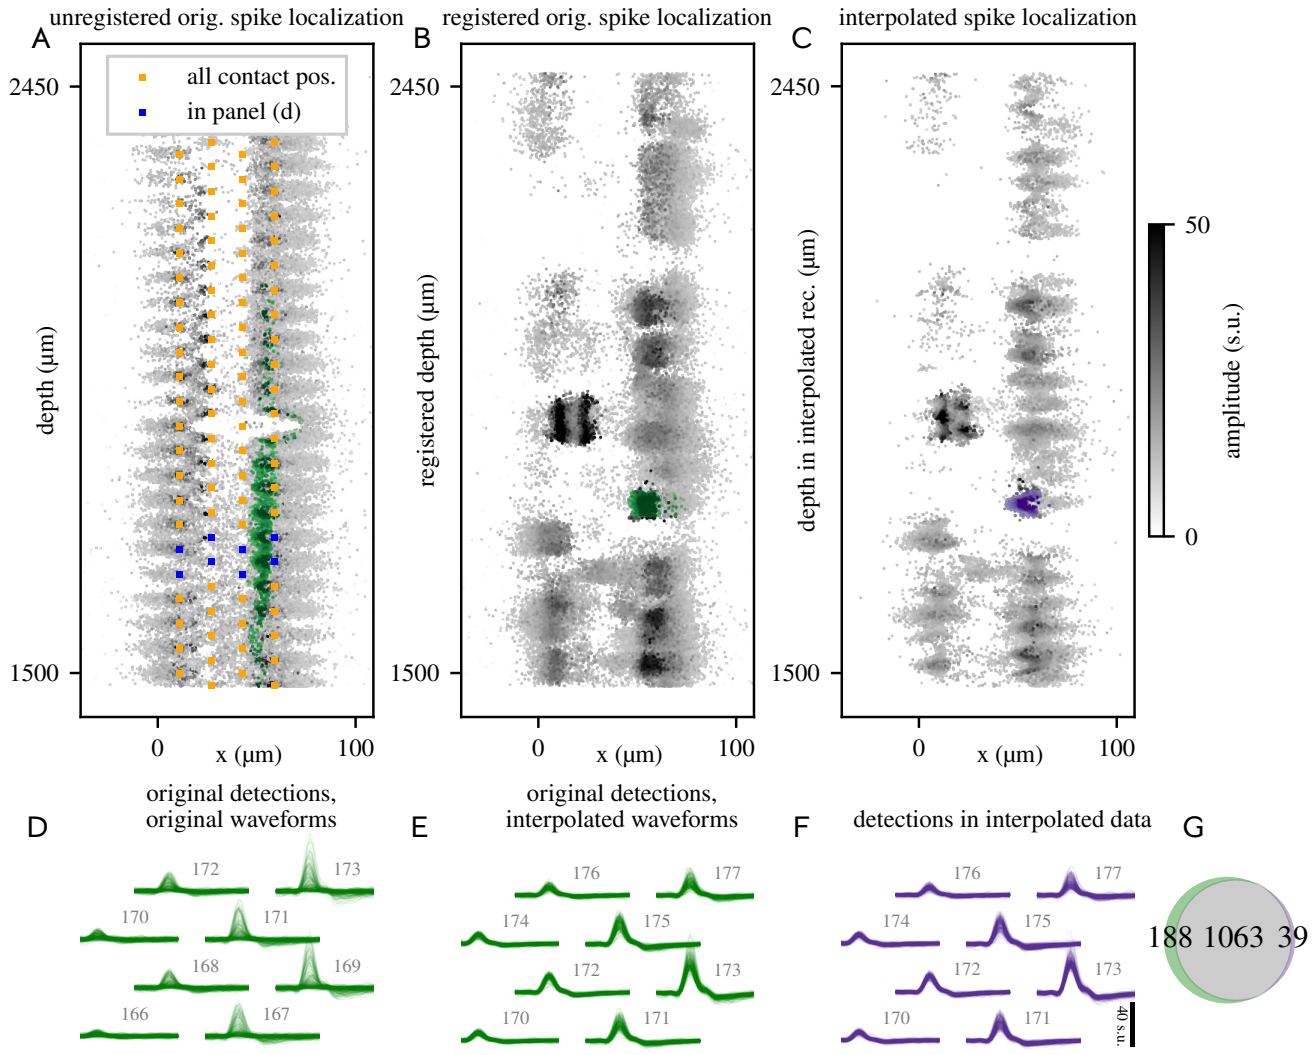

Fig. 4. Checking the drift estimate using drift-corrected spike localizations [16] and interpolated waveforms. (a) Localizations of spikes detected in the Pt. 02 dataset. (b) Spike localizations shifted according to a drift estimate (**DREDge**-CSD in Fig. 2.a); green dots in (a) and (b) represent the same events, corresponding to spikes in a well-isolated cluster in (b) isolated by manually thresholding amplitude and selecting a rectangular region in x, z. (c) Spike localizations computed from events detected in a version of the recording which was interpolated to correct for drift using the SpikeInterface framework [11]. Purple dots lie in a region matching the region used to select the green dots in (b). (d) Waveforms corresponding to spikes highlighted in the green box in (a-b) read from the original recording, shown on high-amplitude channels (contacts shown in blue in (a)). (e) Waveforms corresponding to the same spike times, read from the motion-corrected interpolated binary. The cluster which emerged after correcting the spike depths corresponds to a unit with a well-stereotyped waveform shape in the interpolated binary, providing a validation of the drift estimate. (f) Waveforms corresponding to spikes highlighted in purple in (c); these waveforms match (e) in appearance. (g) A Venn diagram comparing spike times found in the green and purple clusters in (b) and (c); most (1063) spikes are shared between the two spike trains, and few spikes in each cluster do not match (188 and 39 spikes).

3. RESULTS

Experiments were carried out in three datasets: a mouse dataset with induced motion (dataset1 from [2]), and two human datasets (Pts. 02 and 09 from [6]). The mouse dataset is characterized by a slow (period $\sim$ 100s) and shallow (tens of microns) induced triangle wave drift, while the human datasets are characterized by large (hundreds of microns) and fast (sub-second) drift driven by heartbeat and breathing patterns. Comparisons were carried out against the registration algorithm in Kilosort 2.5, which performs well on the same mouse dataset as previously shown in [2], but fails on the human datasets (see Fig. 2.(a, c)). Online methods used 10^4 samples per chunk.

Due to the characteristic noise present in unsorted spiking data, the AP methods (KS and **DREDge**-AP) shown in Fig. 2 cannot be pushed to a sub-second resolution and thus cannot capture the fine drift discovered by the **DREDge**-CSD method (Fig. 2.c), reflected in a well-stabilized spike amplitude raster (Fig. 2.a, top) when compared to the AP method (Fig. 2.a, bottom). The CSD method's improvements over the AP method are reflected in its performance on simple metrics (Fig. 2.c). While the time complexity of the AP-based method is on par with Kilosort and both are just a few seconds (Fig. 3.a, top), when running on CSD sampled at 250Hz the runtime increases dramatically, an effect which is mitigated by the online method (Fig. 3.a, bottom). By using the prior and automatic correlation thresholding, the CSD method can robustly register full recordings (Fig. 3.d), without incurring shrinkage from the prior (right). Intuition for the effect of the prior is examined in a simulation study (Fig. 3.b), in which randomly selected time bins in the mouse AP data are zeroed out, resulting in "glitches" in the correlation and displacement matrices (right). With no prior, these glitches carry through to the drift estimate, but the prior is able to remove them. Throughout these figures, adaptive correlation thresholds were used. The 5th percentile of correlations of neighboring frames is used in AP, where the time scale is faster and thus more nonstationarities appear, while the 0.1 percentile is used in the smoother CSD. These thresholds are among the best performing when looking at a grid of choices (Fig. 3.c).

To further validate our drift estimate, we shifted localization features [16] of detected spikes according to the drift. Clusters that were not visible at first (Fig. 4.a) appear after correction (Fig. 4.b). In particular, spikes from a particular cluster shown in green in panel (b) have localizations with a broad spread before motion correction. To see if these particular spikes truly correspond to a single unit, we interpolated the binary file to correct for the drift using the open-source SpikeInterface framework [11]. As expected, waveforms corresponding to spikes in the cluster were spread across many channels in the uncorrected recording (Fig. 4.d), but appear as a stable unit after correction (Fig. 4.e), validating that the drift correction was accurate. For another perspective, we check to see if the cluster can be identified in the localizations of spikes detected in the interpolated binary. Indeed, after localizing events from that binary, shown in (Fig. 4.c), we identified the same cluster manually (purple dots). Waveforms loaded at times corresponding to these spikes (Fig. 4.f) match those shown in (Fig. 4.e). Further, after using SpikeInterface [11] to identify matching spike times across these clusters, the Venn diagram in (Fig. 4.g) shows that these clusters have high spike time agreement.

4. DISCUSSION

Here, we present an online decentralized algorithm, **DREDge**, to track drift in electrophysiological recordings from the brain using Neuropixels probes. We demonstrate the applicability of tracking the movement of the neural signal relative to the probe with this approach in three different data sets, two recordings in the human cortex [6] and one in mouse [2]. **DREDge** is shown to work using as input both action potentials, representing the activities of single neurons, as well as the local field potential, representing the summed activity of hundreds to thousands of neurons. We develop methods to adapt this algorithm to the structure of individual datasets, improving the robustness and ease of use of this tool.

Finally, we suggest that **DREDge** could be used to improve the clustering of spike localization features for applications in spike sorting. Furthermore, as a fast online estimation of the motion of the probe relative to the neural signal, **DREDge** has potential applications in recording sessions

with real-time neural feedback and brain-computer interfaces (BCI) [18, 19].

The latter point may be a key feature in gathering neural data using Neuropixels in the human cortex. Unlike in mouse preparations [1, 2], recordings in humans, so far, have presented a challenge in that the brain, by virtue of the scale and size, moves considerably due to the heartbeat, breathing, the participants talking, etc. This has required researchers in this domain to rely on manual tracking or only use data where the spike rate is high enough to track the neural signal reliably using Kilosort tools [6, 7]. Such other approaches might not fully correct for the observed motion, with implications for lower yields of unit quality and/or number. **DREDge** allows researchers to move forward in using such precious data to understand human brain activity.

A major question is whether after applying the tracked motion to the signal, there are changes that are not present in the original signal. Our data suggest that applying **DREDge** correction to motion-degraded spike and local field potential data has the potential to accurately reconstruct these electrophysiological signals, while not altering their spatiotemporal profiles in datasets without significant motion artefacts. However, the power of our approach is that we use both the LFP and the spiking activity in parallel to estimate the motion and, alternately, can use each, particularly the spike waveforms, to confirm that we are not adding in artefacts. Further validation could therefore make use of 2D cross-correlation to determine if new spiking waveforms were found in the corrected data sets [6]. An additional question is whether ongoing large-scale neural changes such as evoked potentials due to a task, clinically indicated neural signals such as epileptiform activity, anesthesia-induced burst suppression [6], or other types of stimulation could alter the motion registration and, therefore, the tracking. Further analyses of additional data sets from humans, non-human primates, rodents, and other species where we know specific events are occurring will be necessary to parse this effect in a future study.

Finally, a true test of this approach is how it can handle different types of probes, probe designs, and layouts, both Neuropixels and other silicon probes, and how this can alter spike sorting [1, 2, 20]. To this end, we were able to apply our method to the staggered Neuropixels 1.0 version (the two human data sets) and the Neuropixels 2.0 probe (the mouse data set).

Future work would involve applying the method to other probe layouts, particularly to ask whether this approach could be applied to low spatial resolution probes and the spatial and temporal limitations in tracking motion relative to the neural signal.

To facilitate this sort of portability, **DREDge** is being developed into a part of the SpikeInterface [11] framework. In future applications, and as **DREDge** is applied to further data sets using different probe layouts and reaching deeper brain structures other than the cortex, we propose that the underlying algorithm may be applicable to a number of different conditions and paradigms within and beyond electrophysiology.

5. ACKNOWLEDGMENTS

We would like to thank Yangling Chou, Aaron Tripp, Brian Coughlin, and Fausto Minidio for their help in the original data collection. We would like to especially thank the patients for their willingness to participate in this research. EV is supported by 1K99MH128772-01A1. This research was supported by the ECOR and K24-NS088568 (to SSC) and the Tiny Blue Dot Foundation (to SSC and ACP) and NIH grant U01NS121616 (to ZMW). CW, EV and LP is funded by Simons Foundation 543023, NSF Neuronex Award DBI-1707398 and the Gatsby Charitable Foundation. The views and conclusions contained in this document are those of the authors and do not represent the official policies, either expressed or implied, of the funding sources. The funders had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript.

6. COMPETING INTERESTS STATEMENT

The MGH Translational Research Center has clinical research support agreements with Neuralink, Paradromics, and Synchron, for which SSC provide consultative input. None of these entities listed are involved with this research or the Neuropixels device. The remaining authors declare no competing interests.

7. REFERENCES

[1] James J Jun, Nicholas A Steinmetz, Joshua H Siegle, Daniel J Denman, Marius Bauza, Brian Barbarits, Albert K Lee, Costas A Anastassiou, Alexandru Andrei, Çağatay Aydın, et al., "Fully integrated silicon probes for high-density recording of neural activity," *Nature*, vol. 551, no. 7679, pp. 232–236, 2017.

[2] Nicholas A. Steinmetz, Cagatay Aydin, Anna Lebedeva, Michael Okun, Marius Pachitariu, Marius Bauza, Maxime Beau, Jai Bhagat, Claudia Böhm, Martijn Broux, Susu Chen, Jennifer Colonell, Richard J. Gardner, Bill Karsh, Dimitar Kostadinov, Carolina Mora-Lopez, Junchol Park, Jan Putzeys, Britton Sauerbrei, Rik J. J. van Daal, Abraham Z. Vollan, Marleen Welkenhuysen, Zhiwen Ye, Joshua Dudman, Barundeb Dutta, Adam W. Hantman, Kenneth D. Harris, Albert K. Lee, Edvard I. Moser, John O'Keefe, Alfonso Renart, Karel Svoboda, Michael Häusser, Sebastian Haesler, Matteo Carandini, and Timothy D. Harris, "Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings," *bioRxiv*, 2020.

[3] Janis Karan Hesse and Doris Y Tsao, "A new no-report paradigm reveals that face cells encode both consciously perceived and suppressed stimuli," *Elife*, vol. 9, pp. e58360, 2020.

[4] Hidehiko K Inagaki, Susu Chen, Margreet C Ridder, Pankaj Sah, Nuo Li, Zidan Yang, Hana Hasanbegovic, Zhenyu Gao, Charles R Gerfen, and Karel Svoboda, "A midbrain-thalamus-cortex circuit reorganizes cortical dynamics to initiate movement," *Cell*, vol. 185, no. 6, pp. 1065–1081, 2022.

[5] Nicholas A Steinmetz, Peter Zatka-Haas, Matteo Carandini, and Kenneth D Harris, "Distributed coding of choice, action and engagement across the mouse brain," *Nature*, vol. 576, no. 7786, pp. 266–273, 2019.

[6] Angelique C Paulk, Yoav Kfir, Arjun R Khanna, Martina L Mustroph, Eric M Trautmann, Dan J Soper, Sergey D Stavisky, Marleen Welkenhuysen, Barundeb Dutta, Krishna V

Shenoy, et al., "Large-scale neural recordings with single neuron resolution using Neuropixels probes in human cortex," *Nature Neuroscience*, vol. 25, no. 2, pp. 252–263, 2022.

[7] Jason E Chung, Kristin K Sellers, Matthew K Leonard, Laura Gwilliams, Duo Xu, Maximilian E Dougherty, Viktor Kharazia, Sean L Metzger, Marleen Welkenhuysen, Barundeb Dutta, et al., "High-density single-unit human cortical recordings using the Neuropixels probe," *Neuron*, 2022.

[8] Nicholas A Steinmetz, Christof Koch, Kenneth D Harris, and Matteo Carandini, "Challenges and opportunities for large-scale electrophysiology with neuropixels probes," *Current opinion in neurobiology*, vol. 50, pp. 92–100, 2018.

[9] Marius Pachitariu, Nicholas Steinmetz, Shabnam Kadir, Matteo Carandini, et al., "Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels," *BioRxiv*, p. 061481, 2016.

[10] JJ Jun, JF Magland, C Mitelut, and AH Barnett, "IronClust: Scalable and drift-resistant spike sorting for long-duration, high-channel count recordings; 2020," .

[11] Alessio P Buccino, Cole L Hurwitz, Samuel Garcia, Jeremy Magland, Joshua H Siegle, Roger Hurwitz, and Matthias H Hennig, "SpikeInterface, a unified framework for spike sorting," *Elife*, vol. 9, pp. e61834, 2020.

[12] Eric M Trautmann, Sergey D Stavisky, Subhaneil Lahiri, Katherine C Ames, Matthew T Kaufman, Daniel J O'Shea, Saurabh Vyas, Xulu Sun, Stephen I Ryu, Surya Ganguli, et al., "Accurate estimation of neural population dynamics without spike sorting," *Neuron*, vol. 103, no. 2, pp. 292–308, 2019.

[13] Erdem Varol, Julien Boussard, Nishchal Dethe, Olivier Winter, Anne Urai, The International Brain Laboratory, Anne Churchland, Nick Steinmetz, and Liam Paninski, "Decentralized motion inference and registration of neuropixel data," in *ICASSP 2021-2021 IEEE*

International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2021, pp. 1085–1089.

[14] Gaelle Chapuis, Mayo Faulkner, Kenneth D Harris, Julia M Huntenburg, Cole Hurwitz, Hyun Dong Lee, Liam Paninski, Cyrille Rossant, Noam Roth, Nicholas A Steinmetz, Charlie Windolf, et al., "Spike sorting pipeline for the International Brain Laboratory," *channels*, vol. 10, pp. 6, 2022.

[15] Daniel K. Wójcik, "Current source density (CSD) analysis," in *Encyclopedia of Computational Neuroscience*, Dieter Jaeger and Ranu Jung, Eds., pp. 1–10. Springer New York, New York, NY, 2013.

[16] Julien Boussard, Erdem Varol, Hyun Dong Lee, Nishchal Dethe, and Liam Paninski, "Three-dimensional spike localization and improved motion correction for neuropixels recordings," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, Eds. 2021, vol. 34, pp. 22095–22105, Curran Associates, Inc.

[17] David Chin-Lung Fong and Michael Saunders, "LSMR: An iterative algorithm for sparse least-squares problems," *SIAM Journal on Scientific Computing*, vol. 33, no. 5, pp. 2950–2971, 2011.

[18] Leigh R Hochberg, Daniel Bacher, Beata Jarosiewicz, Nicolas Y Masse, John D Simeral, Joern Vogel, Sami Haddadin, Jie Liu, Sydney S Cash, Patrick Van Der Smagt, et al., "Reach and grasp by people with tetraplegia using a neurally controlled robotic arm," *Nature*, vol. 485, no. 7398, pp. 372–375, 2012.

[19] Francis R Willett, Donald T Avansino, Leigh R Hochberg, Jaimie M Henderson, and Krishna V Shenoy, "High-performance brain-to-text communication via handwriting," *Nature*, vol. 593, no. 7858, pp. 249–254, 2021.

[20] Réka Barbara Bod, János Rokai, Domokos Meszéna, Richárd Fiáth, István Ulbert, and Gergely Márton, "From end to end: Gaining, sorting, and employing high-density neural single unit recordings," *Frontiers in Neuroinformatics*, vol. 16, 2022.