Spike sorting: new trends and challenges of the era of high-density probes

Alessio P. Buccino^{*1}, Samuel Garcia^{†2}, and Pierre Yger^{‡3}

¹Department of Biosystems Science and Engineering, ETH Zurich, Switzerland

²Centre de Recherche en Neurosciences de Lyon, CNRS, France

³Sorbonne Université, INSERM, CNRS, Institut de la Vision, F-75012 Paris, France

Abstract

Recording from a large neuronal population of neurons is a crucial challenge to unravel how information is processed by the brain. In this review, we highlight the recent advances made in the field of “spike sorting”, which is arguably a very essential processing step to extract neuronal activity from extracellular recordings. We more specifically target the challenges faced by newly manufactured high-density multi-electrode array devices (HD-MEA), e.g. Neuropixels probes. Among them, we cover in depth the prominent problem of drifts (movements of the neurons with respect to the recording devices) and the current solutions to circumscribe it. In addition, we also review recent contributions making use of deep learning approaches for spike sorting, highlighting their advantages and disadvantages. Next, we highlight efforts and advances in unifying, validating, and benchmarking spike sorting tools. Finally, we discuss the spike sorting field in terms of its open and unsolved challenges, specifically regarding scalability and reproducibility. We conclude by providing our personal view on the future of spike sorting, calling for a community-based development and validation of spike sorting algorithms and fully automated, cloud-based spike sorting solutions for the neuroscience community.

keywords: spike sorting, extracellular electrophysiology, high-density neural probes

1 Introduction

Recording simultaneously from a large population of neurons is a crucial challenge to unravel how information is processed by the brain. This is especially striking if one assumes that fine neuronal correlations, that are ubiquitous in the brain, are a crucial key to such an understanding [4, 23]. Nowadays, only calcium imaging and extracellular electrophysiology can perform simultaneous large-scale recordings from hundreds of neurons and, at a first glimpse, it may seem that extracellular electrophysiology suffers from fewer intrinsic limitations than calcium imaging [86]. First, genetically-encoded calcium indicators (GECIs) are only an indirect sensor of neuronal spiking activity and have a slow kinetic leading to a poor temporal resolution compared to electrophysiological recordings. Second, calcium imaging requires indicators to be genetically expressed in the target cells, which is not the case with extracellular recordings. Finally, the imaged neurons need to be visible under a microscope (by removing the skull), which makes imaging more invasive than implanting an electrode. Therefore,

^{*}alessio.buccino@bsse.ethz.ch

[†]samuel.garcia@crns.fr

[‡]pierre.yger@inserm.fr

extracellular signals seem to be a much more flexible candidate to extract relevant information from a large ensemble of neurons.

The above-mentioned advantages are even more relevant when one considers the emergence of newly manufactured high-density multi-electrode array (HD-MEA) devices, both for *in vitro* [25, 7] and *in vivo* applications, e.g. Neuropixels probes [39, 77]. These devices introduced a tremendous increase in the number of extracellular channels that can be recorded simultaneously (from tens to thousands) and their spatial density, resulting in the capability to record from hundreds, or even thousands of neurons at once. Such HD-MEA probes can record numerous cells either across layers [77], with flexible designs across different structures [17, 32], or by optimizing electrode switching to increase the yield [44].

The procedure of “spike sorting” consists of various processing steps to extract single-neuron spiking activity from extracellular recordings. Mathematically, it belongs to the category of blind source separation problems, but despite the numerous possibilities given by these new HD extracellular recordings, the analysis of such data arguably creates new challenges [78]. Indeed, with the increased density of the recording channels, the probability of finding overlapping spikes (either in space or in time) is also increased as a direct consequence. These spike collisions are hard to resolve and complicate the spike sorting process [26]. In addition, the very dense spatial resolution of HD probes makes some known issues of spike sorting even more evident, especially regarding the prominent problem of drifts (movements of the tissue with respect to the recording devices). Finally, the amount of data generated by HD devices is now becoming so large that the need for fully automated and reproducible spike sorting solutions is pressing.

In this review, we will discuss these open issues in the field of spike sorting. On the algorithmic side, we will focus on describing the generic implementation principles that have been developed to handle drifts and reviewing recent contributions making use of deep learning approaches for spike sorting. Then we will discuss recent advances in benchmarking, validation, and comparison of available spike sorters. In particular, we will describe the SpikeInterface [12] and SpikeForest [48] projects, as common efforts to provide a unified framework for spike sorting and a quantitative comparison among a large number of well-established spike sorters on a large body of ground-truth datasets. Finally, we will discuss the open challenges of the spike sorting field and provide an overview on how, in our opinion, the future of spike sorting will look like.

2 The spike sorting pipeline

Spike sorting is a complex but necessary procedure to turn extracellular signals into spike trains. Over the years, several algorithms have been developed (see [45, 33, 15] for recent reviews), but despite mild subtleties and variations, they can all be described given the schematic shown in Figure 1. In a nutshell, the generic workflow of spike sorting can thus be described as follows. First, the raw extracellular signals (panel 1) are preprocessed to include only the relevant information for spike sorting. Usually, preprocessing involves filtering via a band-pass filter to exclude the low frequency Local Field Potential (LFP) contribution. In addition, a denoising step, such as whitening or common median reference, can also be used to reduce correlated noise among channels (panel 2). Once signals have been preprocessed, putative spike events are detected via threshold crossings (panel 3). Snippets of these putative spikes are then aligned in time (panel 4) and gathered so that they can all be projected into a lower dimensional space. This dimensionality reduction (or feature extraction) step is often achieved using Principal Component Analysis (PCA), but other methods can be considered (e.g. Independent Component Analysis or Wavelet Transform). In such a reduced space (panel 5), a clustering step splits the spike events into separate groups (see panel 6 and [84] for a review on all clustering algorithms used in the field). Every cluster centroid is a so-called “template” (panel 7), i.e., the average spatio-temporal motif visible on the extracellular channels when a given neuron is firing an action potential. Once the clustering step is completed, either the spike sorting algorithm stops (and the units are thus the labels of the clusters), or an extra template-matching step is added (see [45] for a review) to reconstruct the signals as a linear sum of the discovered templates and obtain a

Figure 1: **Overview of the generic spike sorting workflow.** The spike sorting pipeline is composed of multiple processing steps. Raw extracellular potentials (panel 1) are preprocessed via band-pass filtering (panel 2) before detecting putative spikes (panel 3) as thresholds crossings. Snippets around these spike times are then aligned (panel 4) and projected into a lower dimensional space (panel 5) in order to be clustered as groups (panel 6). Each group is a putative unit with a given template (panel 7), and spikes can be obtained either from the clustering or *via* an additional template-matching procedure (panel 8).

more accurate estimation of the spike trains (panel 8).

While template-matching is probably one of the most significant recent addition to the classical pipeline of spike sorting, since it enhances considerably the performance of the reconstruction [58, 26], this is not the only attempted modification. Indeed, several methods have been recently introduced to enhance the classical view of spike sorting. A full review of these methods is out of the scope of this article, but among those, one could cite works trying to improve the selection of the features (panel 5 of Figure 1) [16, 76], using Independent Component Analysis [11, 46] to identify the templates, tracking down the temporal evolution of the templates over time with [74, 80] with statistical assumptions, and consolidating the clustering step by consensus-based methods [22, 12].

3 Recent advances in spike sorting algorithms

3.1 Drift correction

One of the major disadvantages of extracellular recordings over calcium imaging is that the tissue in the vicinity of the electrodes remains *invisible*. Thus the number, the types, and more importantly the positions of the neurons in the surrounding of the neural probe is unknown. The latest is currently the most challenging problem in the field of spike sorting, since one wants to guarantee the identities of the cells over times. While recording, as shown in Figure 2A, it is possible for cells to drift over time and change their positions with respect to the recording probe. These drifts can introduce distortions of the waveforms recorded extracellularly. In order to ensure that units are recognized correctly over the course of a recording, one needs to track down these waveform changes over time. One way to visualize drifts is to represent the distribution of the spiking activity (e.g. the firing rates) as function of the electrode position over time. In Figure 2B, for example, one can see the activity of the neurons as function of time, recorded by a Neuropixels probe *in vivo* [77]. More precisely, the panel displays a *positional* raster plot, where the peak positions are estimated just after spike detection, before any spike sorting *per se*. Usually, the position of the peaks is estimated as its center of mass, but more

Figure 2: **Neuronal drifts.** **A)** Illustration of the problems raised by neuronal drifts. Two neurons (red and orange) have two distinct waveforms, but after some non rigid drifts (movements), the templates are changed. **B)** Illustration of the drift *in vivo*. A Neuropixels probe in the cortex of the mice is moved up and down with an imposed triangular movement of amplitude 50 μm (top trace) [77]. The row shows the depth of all the peaks detected via signal threshold (*positional* raster plot, see text) as function of the time. As one can see, when the probe starts moving (red dotted line), so does the depth of the peaks, since neurons are drifting across channels. While the drift here is imposed, spontaneous drifts can happen locally and non homogeneously over the whole probe (see inside light blue boxes before and after movements).

refined methods can be applied [9, 13, 69]. During the recording, the probe is moved along the vertical axis with a triangular movement of 50 μm amplitude. The oscillations visible in the raster plot here reflects the fact that neurons are moving up and down with respect to the recording channels.

The causes and the temporal extent of these drifts can be numerous. One the main source is mechanical: *in vivo*, cells are likely to slowly drift from initial positions because of the pressure release in the tissue after the insertion of a probe. These drifts can have arbitrary time constants, and thus either be abrupt and discontinuous or slow and continuous. Another source of waveform distortion, which is not related to relative displacement, can be functional: if a pharmacological drug is added to the media, or if neurons are forced to fire in a non-physiological manner (for example with optogenetics). Such external perturbations are likely to distort the templates, in a possibly non-reversible manner, and thus are another real challenge for spike sorting.

There is yet no consensus on how to deal with drifts in extracellular recordings. However, before reviewing the different methods used by modern algorithms, it is worthwhile noticing that solutions can be grouped into two main categories: drifts can either be handled during the pre- or post-processing stages. Both approaches have advantages and disadvantages. On one hand, handling drifts at the pre-processing stage could ease the task of the spike sorting algorithms, at the cost of (possibly) distorting the input data (for example by introducing unwanted artifacts in the signals). On the other hand, handling drifts at the post-processing stage does not modify the input data, but it might result in erroneous results depending on the extent and nature of the drifts.

The basic idea behind correcting for drifts at the pre-processing stage is to mimic drift-correction procedures that have been applied in the field of calcium imaging [28], and thus perform a non-rigid spatial registration of the extracellular signals directly (*pre-registration*). While such approach, to date, has only been applied to Neuropixels devices to estimate the drifts along the depth dimension of the electrodes [77], in theory it could be generalized to drifts in several dimensions [9]. The algorithm starts from the *positional* raster plot, where each detected event is plotted against its estimated position; next, a binning procedure allows to create an image of the drifts as function of time, i.e., a picture of the activity as function of the depth (see Figure 3A). Each column of this image is termed activity

histogram. Assuming the activity of neurons in the recording is stationary enough, the pre-registration method tries to find the (non) rigid transformations over time that would maintain the histogram of the peaks, for every chunk of duration d , as constant as possible.

The fine details of the registration might vary from algorithm to algorithm. For example, in [77], an average template is obtained as the mean of activity histograms computed by chunk (see Figure 3B) and used as a reference for motion correction. However, this *average template* approach assumes that the spiking activity is relatively constant over the time course of the recording. To relax this assumption, recent work have proposed a decentralized method [83] that computes all the pairwise displacements between chunks to directly estimate a motion signal to register the traces (see Figure 3C). These pre-registration methods can be either rigid (motion is constant over depth) or non rigid (motion depends on the depth). The extension to the non-rigid case can easily be performed by considering only local spatial neighbourhood along the depth axis. However, note that irregardless of the strategy, the registration procedure, acting at the pre-processing stage, is distorting the input data and, in doing so, might introduce unwanted artifacts that might affect, to some extent, the overall spike sorting performance.

A different approach to handle drifts is to try to perform spike localization and motion correction in order to ease the job of downstream clustering. An accurate estimate of spike position with respect to the probe layout can be used as an additional feature for clustering and reduce spurious noisy fluctuations in the feature space. While the simplest estimation of spike positions can be done via the centers of mass of the detected spike waveforms, more refined methods using variational inference [37] or point-source models [9] have been proposed. As it can be seen in Figure 3D, using estimated positions as a spike feature for the clustering step can enhance the the recognition and tracking of

Figure 3: Strategies for drift correction **A)** Peaks position as function of the depth, estimated via the center of mass procedure, and for the recording shown in Figure 2. A binning procedure can be used to create an image of the peaks distribution, over space, as function of time. **B)** Registration of the extracellular recordings via a common activity template, computed as the average activity profile over all chunks [77]. **C)** Same as **B)**, but with a template computed in a decentralized manner [83]. **D)** Illustration of the motion clouds methods, trying to refine the estimation of the peaks' position before clustering. Spikes (black dot) during three different consecutive intervals of time δt are shown, superimposed with the probe layout. Red and blue circles show spikes from two different cells, drifting along the depth axis over time. Tracking these movements (bottom lines) can be used to resolve drifts. **E)** Illustration of the post-processing methods. Several independent sorting are launched, with overlapping time windows. Overlapping results can be used to track drifts over units.

different clusters over time [9]. In addition, the estimated spike positions can be used as an alternative method for image-based registration approaches (Figure 2B-C), using a point-cloud registration to estimate motion signals and correct for drifts.

Finally, the last option is to handle drifts at the post-processing stage (Figure 2E). To do so, the approach followed by [18, 55] is to divide the recording into small and overlapping temporal intervals and to perform the sorting in each of these chunks individually. Using the information in the overlapping regions, one could then try to establish the juncture and track the cell identities over time despite small changes in waveforms. For example, if neuron A is firing during $[t_0, t_1]$ and neuron B is firing during $[t_2, t_3]$ with $t_1 \geq t_2$, then if spikes of A and B are the same in $[t_2, t_1]$, the two cells are likely to be the same. The main disadvantage of such an approach is that one needs to launch several instances of the spike sorting procedure, at the cost of slowing down the whole processing pipeline. In addition, the duration of the overlapping window is *ad hoc*, and should reflect a prior expectation of the temporal evolution of the drifts.

3.2 Deep learning approaches

The last few years have witnessed a boom in the efforts of using deep learning approaches to tackle the spike sorting problem [60, 20, 62, 71, 43, 92, 91, 81, 89, 47]. Deep learning methods, in fact, have proven so powerful and accurate in many complicated applications, ranging from image classification to natural language processing. It is therefore natural to turn to deep learning for such a complex task as spike sorting.

Deep Learning (DL) refers to machine learning techniques using very *deep* Artificial Neural Networks (ANN), with many more layers, parameters and artificial neurons. In addition to the higher complexity with respect to standard feed-forward ANN models, DL can also employ a wide variety of architectures to tackle different problems. *Convolutional Neural Networks* (CNNs), for example, make use of convolutional layers that combine the input features using convolutions in various dimensions (1D, 2D, and even 3D). They are commonly used in image classification tasks, where 2D convolutional layers allow the model to capture 2-dimensional features of the data [40]. Another class of DL models that have been applied to spike sorting are *Long Short-Term Memory* (LSTM) networks, which are a type of recurrent neural networks well suited to process time series data, as their recurrent structure enables them to keep a *memory* of the previous times [35]. A third DL concept that is worth introducing is *autoencoders* [5]. Autoencoders (AE) are a family of unsupervised models that aim at encoding the input data into a low-dimensional latent space. This is usually achieved by an *encoder* network, that projects the input data onto the latent space, and a *decoder* network that goes from the latent space to the output.

In the following, we review the recent contributions that used machine/deep learning strategies to tackle the spike sorting process. We identified four main strategies in which DL solutions have been employed for spike sorting (Figure 4): *i*) end-to-end solutions, *ii*) DL approaches for different spike sorting steps, *iii*) SNN-based solutions, and *iv*) *DL for automatic spike sorting curation*.

3.2.1 End-to-end solutions

The first strategy to employ DL models for spike sorting is to build a so-called *end-to-end* solution. End-to-end solutions ideally take the raw (or filtered) traces as input to the DL model and output directly the spike trains of the different spike sorted neurons (Figure 4A). A few different solutions have been proposed to this end. In the work from *Racz et al.* [60], the authors implemented separate networks for spike detection and spike sorting for a 128-channel high-density device [21] using a combination of CNN and LSTM. *Li et. al* [47] developed a similar approach focused on recordings from single channels. However, instead of splitting the detection and sorting phase, the used a single 1D CNN architecture to directly detect and sort spikes in the input recording. *Racz et al.* [62] proposed a different architectures that uses an CNN+LSTM-based autoencoder architecture to find a latent space and dense output networks to perform the spike sorting task.

While end-to-end solutions arguably represent a *tempting* strategy to tackle the spike sorting problem, there are a few apparent limitations that need to be highlighted. Supervised learning methods, and in particular DL models with many parameters, require a very large amount of *labeled* or *ground-truth* data to be trained. In the three solutions described above, the authors mainly used manually curated or synthetic recordings to train their models, and each model was trained separately on each recording. This allowed them to know *a priori* the number of output neurons to be used, for example, but this information is normally not known in advance when spike sorting, which is an *unsupervised* problem in nature. In addition, it is still unclear how end-to-end strategies could generalize over different recordings from different brain regions and states, as well as from different probe geometries. For example, in the cerebellum, Purkinje cells produce both simple and complex spikes, which have very different spike shapes. Complex spikes represent a well-known problem for spike sorting, so that specific cerebellum-specific spike sorters have been developed [72]. To tackle the same problem, Markanday and colleagues [49] implemented an end-to-end solution to detect complex spikes data manually labeled by experts.

3.2.2 Deep learning for spike sorting steps

A second strategy to use of DL models for spike sorting is to utilize specific networks for different spike sorting steps, rather than as an end-to-end solution (Figure 4B). As introduced in Section 2, spike sorting methods can be split into several and relatively isolated steps, some of which may be tackled and improved with the aid of a DL approach.

Recording denoising: Extracellular recordings contain several sources of independent noise components, that may arise from the readout electronic circuit and also from the biological sample itself (e.g. the activity of far away neurons). Before further processing, it could be therefore useful to apply some

Figure 4: **Deep learning strategies for spike sorting.** **A)** End-to-end solutions use DL network to predict spike times and labels from the signal traces. **B)** Alternatively, different DL approaches can be used to for separate steps, such as denoising, spike detection, and feature extraction. **C)** Unsupervised spiking neural networks (SNN) use spiking neuronal models and plasticity rules to learn to recognize different units in the recordings. **D)** DL can also be used to automatically curate and clean the raw output of an automated spike sorter.

denoising techniques, e.g. common average/median reference or whitening of the signals. Lecoq and colleagues [42] approached the problem of denoising by using an autoencoder approach to *interpolate* the input signal (*DeepInterpolation*). The idea is to train a deep network to predict the signal at time t using samples before and after t . This allowed the authors to remove independent noise sources in the input signals (the same approach was also applied to calcium imaging and fMRI data). When applied as a pre-processing step to extracellular recordings from a Neuropixels probe, the spike sorting output (using Kilosort2 [56]), yielded a larger number of high quality units.

Spike detection and waveform denoising: A second processing step that is essential to spike sorting is spike detection. Normally, this step is performed channel-wise by simply detecting intervals of the signals that cross a certain threshold that is set based on some dispersion measures (e.g median absolute deviation [61]). One of the major problems of spike detection is the presence of colliding spiking events, that result in distorted waveforms. These collision events can strongly degrade downstream processing steps, especially clustering, and therefore need to be removed. In order to ameliorate the performance of spike detection, Lee and colleagues, in developing the YASS spike sorter, used a DL strategy for spike detection and waveform cleaning/denoising. First, the authors implemented a temporal and spatial CNN architecture to detect spikes in the signals, similarly to [60]. The detection network was trained using synthetic snippets on channel neighborhoods that either included a spike event (manually curated templates with additive noise) or did not (noise snippets). After a spike event was detected, a second denoising network was used to clean collision-corrupted spikes. In this case, the model was trained using simulated spike events corrupted by a collision to output the clean template. The combination of the spike detection and waveform denoising network showed to greatly improve the separation of different units for the clustering space. Saif et al. [70] implemented a 1D-CNN model for single-channel data to remove noisy spike events before clustering. In this case, the model was trained on a large dataset of manually labeled data, in which experts flagged single spikes as being of neuronal origin or noise events.

Feature extraction: After spikes are detected, the aligned waveforms need to be projected onto a low-dimensional space before the clustering step; this *feature extraction* step is normally performed by means of principal component analysis (PCA) or wavelet decomposition. As briefly mentioned before, autoencoders have emerged as optimal candidates to find low-dimensional latent spaces and have therefore been proposed to reduce the dimensionality of waveforms. Wu and colleagues [91], for example, used deep compressive autoencoders to extract features from detected spikes and showed that they outperform several other dimensionality reduction techniques in compressing waveforms. Similarly, Eom et al. [20] proposed an ensemble of autoencoders with different complexities to project waveforms onto a low-dimensional space and showed that their method provides a better separability of different units in the latent space in comparison to other dimensionality reduction strategies.

In addition to the use of autoencoders to extract features, Wouters et al. [89] introduced a DL strategy to tackle the problem of spike collisions in spike sorting [26]. When two spikes are perfectly overlapping in time, the observed collision event will be a linear sum of the waveforms from the two neurons; if the feature extraction method is linear (e.g. PCA), the collision spikes from two neurons will form a third cluster with a centroid in correspondence to the sum of the two neurons' clusters. However, when spikes from two neurons have small jitter in time, the linearity in the feature space is lost. The authors introduced a deep neural network for dimensionality reduction and trained it with synthetic data to enforce linearity in the feature space with respect to spike collisions even when a jitter between the colliding spikes is present. The promising results showed that the proposed feature map is capable of achieving a performance close to optimal template matching without a template-matching phase.

3.2.3 Spiking Neural Networks (SNN)-based solutions

Deep neural networks utilize a very simplistic model of artificial neurons, i.e., the *perceptrons* [64], which integrate their weighted input to produce an output. The output of an artificial neuron can be thought as a mean firing rate, as there is not concept of *spikes* in these kind of neural networks. Conversely, in Spiking Neural Networks (SNN) the artificial neurons communicate via spikes and synaptic inputs. Leaky Integrate-and-Fire (LIF) models of single neuron temporally integrate the synaptic inputs from received from other neurons and they produce a spike whenever their membrane potential crosses a threshold. By employing models of synaptic plasticity, SNNs have shown to be able to perform several unsupervised learning tasks [30, 19].

In the context of spike sorting, Werner et al. [87] introduced a 2-layer SNN implementation capable to spike sort extracellular data from a single channel in real-time. The proposed method first decomposed the input signal using bank filters whose output modulated the firing of the corresponding input neurons. The input neurons were fully connected to 5 output neurons (the number of output neurons needs to be greater than the number of expected units), which connected to each other with inhibitory synapses to implement a winner-take-all approach (only one output neuron can fire at a time). Using spike-timing dependent plasticity (STDP), the SNN was able to learn to recognize spikes from different neurons in the recordings. Importantly, the spike trains generated by the output layer correspond directly to the sorted spike trains (Figure 4C). In a follow-up study, Bernert and colleagues [8] used a different SNN-based solution to spike sort data both from single channel recordings and from tetrode data. The authors introduced an additional hidden layer uses plasticity to learn different waveform fragments in the input signal. The output neurons learn to recognize combinations of different fragments as different units and to output the spike trains. Additionally, an attention neuron is employed to gate the learning only when an action potential is occurring.

SNN-based solutions represent an emerging strategy for spike sorting, especially given their intrinsic capability of functioning in real-time and of being implemented in neuromorphic hardware solutions [87]. This characteristics makes them optimal candidates for brain-machine interface applications, where the raw signals need to be processed and decoded in real-time. However, in the context of precise spike sorting for systems neuroscience, some limitations exist. Unsupervised learning rules for SNN (such as Spike Timing Dependent Plasticity) are very sensitive to initial conditions, and quite often unstable without metaplasticity (see [93] for review). Thus, very precise fine tuning of the parameters is often necessary. Moreover, questions such as how the extracellular traces should be turned into spikes, or what is the optimal size of the SNN required to capture all the spiking dynamics, still remain open.

3.2.4 Deep learning for spike sorting curation

The last strategy to utilize DL in the spike sorting process is to build models that automatically curate a spike sorting output. Differently from all the strategies described above, in this case the DL model will take a raw output from a spike sorter as input and output a curated version of the spike sorted units (Figure 4D). Manual curation of spike sorting output is still a very common and required step to improve the results of automatic spike sorters before further analysis. Spike sorters, in fact, do make mistakes [12, 48] and an expert electrophysiologist can correct such mistakes by inspecting features of the spike sorted units. Several software tools (Phy [66], Neuroscope [31], CellExplorer [57], JRClust [38]) are used by the community to remove bad or noisy units, merge oversplit units, and split unit clusters (when overmerged or when contaminated by noise events). The former two curation actions (removing bad units and merging) could be well suited for a supervised DL approach, since they do not modify the underlying spike trains. However, this solution is still relatively unexplored. One option could be to gather all the annotated curation steps and train a network with using these community-labeled data. However, since it is well known that manual curation is prone to bias [88, 66], a second or complementary option could be to use ground-truth datasets to train a network to perform the best curation steps given the ground-truth spiking activity.

4 Benchmarking spike sorting results

Spike sorting tools require extensive and quantitative validation to assess their performance. To provide such validation, one needs to utilize so-called ground-truth (GT) data, i.e., recordings in which the spiking activity of the underlying neurons is known. There exist three different strategies to obtain ground-truth recordings, which differ in their biophysical detail and amount of units available as ground truth (Figure 5A).

Paired recordings: these GT data are obtained by recording the same neuron simultaneously with an extracellular probe and a second measuring modality, such as a glass micro-pipette inserted in the target neuron (intracellular) [85], attached to the cell membrane (juxtacellular) [54], or via whole-cell patch-clamp [94, 2, 50] (Figure 5.1). The second modality enables to record spikes of the target neuron unequivocally and reliably, so that they can be used as ground-truth. While paired recordings are the most realistic source of ground-truth data, the acquisition of such paired dataset is arguably challenging, especially *in vivo* [54, 50, 2]. A second limitation of this approach is its very low-yield: in most cases only a single ground-truth neuron is available for each dataset and this prevents a full characterization of the spike sorting performance. In addition, recording

Figure 5: Spike sorting validation. **A)** Different strategies for GT data: 1. paired extracellular and intra-/juxtacellular recordings; 2. hybrid recordings with injected artificial neurons; 3. fully synthetic recordings with biophysical models. **B)** The **SpikeInterface** project showed how largely spike sorters disagree. 1. Snippet of a Neuropixels trace with its *activity map* on the top; 2. Number of units found by six different sorters (HS - Herding Spikes, KS - Kilosort2, IC - IronClust, TDC - Tridesclous, SC - Spyking Circus, HDS - HDSort); 3. Units in agreement with k different sorters. Panels are adapted from [12] – CC BY. **C)** The **SpikeForest** project systematically benchmarks spike sorters against a large body of ground-truth data. Example of the result tables, with the accuracy of spike sorters on the left (blue) and the number of units found by each sorter on the right (green). Figure from [48] – CC BY.

via a micro-pipette could be inherently biased in terms of the cell-type of the target neuron, as interneurons are less numerous and smaller than pyramidal cells, and therefore less likely to target.

Hybrid recordings: a second type of GT data is constructed by adding artificial spiking activity to experimental recordings [66, 90] (Figure 5.2). The templates of the artificial neurons are usually extracted from previously sorted and curated data and are re-injected into the recording by adding modulations to resemble physiological variability, for example due to bursting [2]. While hybrid recordings provide a higher yield with respect to paired recordings, they still do not allow a full characterization of the spike sorting output. In practice, only a limited number of artificial units can be injected in the original recording in order to avoid an *overcrowded* spiking activity that could compromise the intrinsic performance of spike sorters.

Synthetic recordings: a third strategy to obtain GT data is to use simulations of extracellular signals and corresponding spiking activity [14, 29, 53, 38, 10] (Figure 5.3). In this case, the GT dataset allows for an exhaustive and full benchmark of spike sorting results, since all neurons making up the spiking activity are known. However, the debate on whether synthetic recordings are realistic enough to capture subtleties of experimental data is still open. Nevertheless, simulations can provide a powerful and highly controllable tool to develop and benchmark spike sorting tools. As an example, controlled simulations were used to specifically assess the performance of spike sorters against spike collisions, which are a well known problematic for accurate spike sorting [26].

With these different ground-truth recordings, one can start approaching the task of systematically benchmarking the performance of multiple spike sorting tools. However, this step is challenging for several reasons: first of all, in extracellular electrophysiology, there is a large variety of proprietary formats and acquisition systems; second, the community has produced several spike sorters, each coming with its own requirements in terms of input format, parameters, configuration files, and output format. **SpikeInterface** is an open-source software tool developed to unify existing technologies into a single framework [12]. Providing a standardized access to the multiple acquisition formats and spike sorters, **SpikeInterface** allows one to run and assess the performance of many sorters on the same recordings. Strikingly, when six modern spike sorters were run on the same Neuropixels dataset [75] (Figure 5B.1 shows a snippet of the traces) there was a large disagreement among them (Figure 5B.2 and 3). Out of a total of $\sim 1'400$ units found, only 33 were in agreement between all sorters and 263 in agreement between at least two sorters [12]. Clearly, this highlights the need of more systematic benchmarking efforts in order to *trust* the results of automatic spike sorters.

In this direction, the **SpikeForest** project [48] attempted to gather a large body of ground-truth data, systematically run spike sorters on them, and report the results in an interactive and intuitive manner. The resulting website (<https://spikeforest.flatironinstitute.org/>) reports the performance tables (Figure 5C) of spike sorting results from over 1.3 TB of ground-truth data coming from diverse sources and strategies (paired, hybrid, and synthetic recordings). Such effort provides unprecedented guidance to end-users, willing to explore and understand the performance of different tools on different types of recordings (e.g. tetrodes vs high-density probes) and it arguably represented a push for the spike sorting community, as now developers have a massive and curated ground-truth repository to test their tools, and new features, against. Moreover, the **SpikeForest** team had to overcome intrinsic software engineering challenges that we cover in more details in the next section.

5 Open challenges and the future of spike sorting

In this review, we covered the recent advances in the field of spike sorting, especially all the recent pre and post-processing techniques applied by spike sorting algorithms to track the neurons over time when they are drifting with respect to the recording sites. We also gave an overview of all the attempts to use

deep learning approaches in order to boost the performances of these algorithms, highlighting different strategies with their advantages and limitations. Finally, we covered latest progress in generating ground-truth data for systematic benchmarking and validation of spike sorting tools. Despite the tangible advances in spike sorting development and validation efforts over the last few years, the neurotechnology *revolution* that we are experiencing with the advent and broad distribution of high-density probes is setting novel challenges for spike sorting pipelines.

The first new challenge is *scalability*. Since the first extracellular recording from a tungsten microelectrode [36] neural devices for both *in vivo* and *in vitro* application have experienced a steady increase in the number of electrodes that can be recorded simultaneously (Figure 6A). For *in vitro* applications, where probe size is less of a constraint than *in vivo*, micro-electrode array technologies have allowed to reach over 1000 simultaneously recorded electrodes [25, 7] packed at very high spatial densities. Recently, some prototypes have been proposed that can even record from over $\sim 20'000$ electrodes simultaneously [95]. Neural devices with several hundreds of electrodes have also been developed for *in vivo* applications [39, 3, 77, 17, 21], with the Neuropixels technology being arguably a game changer in the field of systems neuroscience. Spike sorting algorithms need to be able to cope with this steadily increasing number of channels. A 1-hour Neuropixels 1.0 probe recording, with its 384 channels sampled at 30kHz and with 2 bytes per sample, will produce over ~ 80 GB of unprocessed raw data. With similar calculation, a 1-hour recording of a Neuropixels 2.0 device [77] will produce ~ 160 GB of data, and 1-hour recording of the 20'000-channel device in [95] (sampled around 10 kHz) will give a 1 TB file to analyze. Spike sorters, therefore, need to implement algorithmic solutions with scalability, speed, and performance in mind, for example by leveraging multiprocessing, distributed, and GPU-accelerated solutions. In addition, spike sorting developers need to be aware of the trends in neurotechnology and try to predict what neural devices will in the next years to come¹. Finally, it is also essential that the computational time of spike sorting is faster than 1x (i.e., it takes less for spike sorting than the duration of the recording), otherwise, as a field, we will never finish our spike sorting *pre*-processing and move on with downstream analysis.

A second challenge for spike sorting is *reproducibility*. Although the more and more common usage of automatic tools for an initial spike sorting, a manual curation step to *clean* the spike sorting output is still very common, if not necessary. A number of tools have been developed to allow users to remove, merge, and split units [31, 65, 38, 57], but this manual curation is still very subjective and can provide different outputs [88, 66]. Clearly this hinders full reproducibility of the results and all downstream analysis that depends on spike sorting. In Section 3.2.4 we discussed the possibility to use deep learning methods trained on ground-truth data to address automatically curate spike sorting outputs. Alternatively, or complementary, a wide range of unsupervised quality metrics have been proposed to assess the quality of spike sorted units [34, 6]. However, the debate is still open on which quality metrics are most informative and should be used and how they will generalize over different brain areas, species, and brain states (or even across sorters). A third but computationally expensive option to avoid manual curation could be running several spike sorters on the same recording and aggregating the results in an *ensemble* fashion [22, 12]. In addition to reducing reproducibility, manual curation of sorting results from HD devices is also very time consuming because of the many units that these probes can record from.

Reproducibility can also be affected by the intrinsic life cycle of spike sorting software tools, which can undergo many modifications and releases as bugs are fixed, methods are improved, and new features are added. While version controlled systems (VCS) can keep track of all these changes, from a user perspective it can be confusing and even complicated to know which exact version of a spike sorter was used on specific datasets and to report correctly. A possible solution to the *version* problem is to encapsulate spike sorters of different versions in containers, i.e., pieces of software that package all codes and dependencies for an application to run independent of the environment. Containers can be run through platforms such as **Docker** [52] and **Singularity** [41] and ensure reproducibility on the software side. The **SpikeForest** project first introduced containerized and version-tagged containers

¹Strikingly, the **Klusta** spike sorter [66], developed in 2016 for “large and dense” probes (up to 64 channels), has become outdated in less than 4 years

of 11 different spike sorters which we recently collected and distributed through **SpikeInterface**². In addition to improve reproducibility of the results, containers can also ease the process of installing and setting up different spike sorters, which can be annoying as the number and complexity of dependencies increases. The combination of containerized solutions and automatic curation would enable a fully reproducible spike sorting pipeline, which does not require any human intervention (and possibly human bias) in the loop (Figure 6B).

Despite not highlighted in this review, it is worthwhile mentioning some other open challenges for spike sorting, that remain to be fully solved. Among them, a major problem is how to automatically handle waveforms modulation due to bursting, which causes an amplitude and shape distortion of the extracellular waveforms during the firing of consecutive spikes [2]. Several algorithmic attempts have been proposed to tackle the issue [73, 59], but to date there is no definitive solution implemented in the most commonly used spike sorting tools and correction of spike sorting errors due to bursting modulation is still mainly addressed via post-hoc manual curation. Spike collisions (spike events overlapping in space and/or time), which have been proven the spike sorting performance [26], represent a second well-known issue for spike sorting methods to address. Although recent spike sorting solutions based on template-matching appear to ameliorate the problem [26], automatic spike sorters are still more error-prone when spikes are colliding (especially from neurons with similar templates). Spike sorting becomes particularly challenging during network bursts, which are very common in *in vitro* cell cultures [82], due to the concomitant bursting modulation and high spike collision rates. Finally, while we have mainly focused on offline applications, online spike sorting for HD probes has the potential to tremendously boost the possibilities for neuroprosthetics or brain-machines interfaces. While several attempts have been made [68, 24], all the aforementioned bottlenecks (drifts, bursts, collisions, scalability) are even more striking in an online context.

In addition to the increased number of electrodes and size of single files, the large availability and distribution of HD probes is producing, worldwide, an unprecedented amount of data. Current solutions utilized by single laboratories or institutions usually rely on local machines or institutional servers, which might fall short in terms of both storage capacity and computational load, given the sharp increase of number of electrodes and overall number of recordings. Additionally, there is still a large variation among groups in the way that spike sorting is carried out, both in terms of the chosen software and custom post-processing and manual curation steps. A possible solution forward is to move both data storage and spike sorting fully remotely, on cloud computing services [63, 1]. A remote and centralized solution for spike sorting could solve several issues that we have introduced so far. As new probes with even more electrodes are developed and more and more recordings are generated by the neuroscience community, cloud-based solutions (e.g. the DANDI archive <https://gui.dandiarchive.org/#/> and the EBRAINS platform <https://ebrains.eu/>) could provide on-demand storage as required and would prevent individual institutions to constantly upgrade their servers and storage capabilities. In addition, when uploading data onto a remote server, standardized file formats could be enforced, e.g. NWB [79, 67], NIX [51], and BIDS [27], and as a by-product several standardized tools for data processing and visualization can be developed. Plus, data sharing and reuse would be drastically facilitated. Leveraging cloud computing providers, we could move the heavy computational burden of spike sorting remotely and use processing resources as needed. In this regard, the containerization of spike sorting tools is already an important step toward the deployment of spike sorting on the cloud. However, another benefit of moving processing on the cloud is to allow a centralized and collaborative spike sorting development and deployment, instead of the current *individual* state where many different groups are aiming to develop and improve their own tools. Given the modular nature of the spike sorting processing, where several well identified and circumscribed steps are needed to go from extracellular traces to spike trains, it is very likely that different sorters are providing optimal solutions for different steps. Following this idea, the **SpikeInterface** project has started to break down different sorters and re-implement them in a modular fashion. By doing so, it calls upon the spike sorting developer community to collaborate in this effort and collaborate to

²on GitHub <https://github.com/SpikeInterface/spikeinterface-dockerfiles> - and DockerHub <https://hub.docker.com/orgs/spikeinterface>

Figure 6: Challenges and future of spike sorting **A)** Scalability. The constant increase of electrode counts both *in vivo* (blue) and *in vitro* (orange) represents a challenge for spike sorting tools. **B)** The reproducibility challenge can be addressed by using containerized software (via Docker or Singularity) and automatic curation. **C)** The future of spike sorting, in our view, is to make storage and processing remotely. After acquisition, recordings are converted to standardized formats, such as NWB, NIX, or BIDS, and uploaded to cloud-based archives, e.g. DANDI or EBRAINS. Analysis pipelines, including spike sorting, can be then deployed on cloud computing services.

discuss, define, implement, and test different *sorting components* available from multiple sorters³. In such a view, a modular and collaborative spike sorting development phase, where different steps of the pipeline are implemented in a standardized, efficient, accessible, and continuously tested framework by the entire spike sorting developer community, will represent a tangible progress in making spike sorting tools more reliable, efficient, and trustworthy.

Acknowledgments

This work was supported by the ETH Zurich Postdoctoral Fellowship 19-2 FEL-17 (APB). In addition, we would like to thank Julien Boussard for fruitful conversations on the spike localization and drift correction topic. We are thankful to Nick Steinmetz for acquiring and sharing the Neuropixels dataset with imposed drift movements that we used to generate Figures 2-3. Finally, we wish to thank the spike sorting community for sharing and discussing exciting updates over the past year: Jeremy Magland, Liam Paninski, Matthias Hennig, Jeff Soules, Fernando Chaure, Hernan Rey, Julien Boussard, Erdem Varol, Cole Hurwitz, Ben Dichter, and Olivier Winter.

³https://spikeinterface.readthedocs.io/en/latest/module_sorting_components.html

References

- [1] T. Abe, I. Kinsella, S. Saxena, E. K. Buchanan, J. Couto, J. Briggs, S. L. Kitt, R. Glassman, J. Zhou, L. Paninski, et al. Neuroscience cloud analysis as a service. *bioRxiv*, pages 2020–06, 2021.
- [2] B. D. Allen, C. Moore-Kochlacs, J. G. Bernstein, J. P. Kinney, J. Scholvin, L. F. Seoane, C. Chronopoulos, C. Lamantia, S. B. Kodandaramaiah, M. Tegmark, et al. Automated in vivo patch-clamp evaluation of extracellular multielectrode array spike recording capability. *Journal of neurophysiology*, 120(5):2182–2200, 2018.
- [3] G. N. Angotzi, F. Boi, A. Lecomte, E. Miele, M. Malerba, S. Zucca, A. Casile, and L. Berdoncini. Sinaps: An implantable active pixel sensor cmos-probe for simultaneous large-scale neural recordings. *Biosensors and Bioelectronics*, 126:355–364, 2019.
- [4] B. B. Averbeck, P. E. Latham, and A. Pouget. Neural correlations, population coding and computation. *Nature reviews neuroscience*, 7(5):358–366, 2006.
- [5] P. Baldi. Autoencoders, unsupervised learning, and deep architectures. In *Proceedings of ICML workshop on unsupervised and transfer learning*, pages 37–49. JMLR Workshop and Conference Proceedings, 2012.
- [6] A. H. Barnett, J. F. Magland, and L. F. Greengard. Validation of neural spike sorting algorithms without ground-truth information. *Journal of neuroscience methods*, 264:65–77, 2016.
- [7] L. Berdoncini, K. Imfeld, A. Maccione, M. Tedesco, S. Neukom, M. Koudelka-Hep, and S. Martinoia. Active pixel sensor array for high spatio-temporal resolution electrophysiological recordings from single cell to large scale neuronal networks. *Lab on a Chip*, 9(18):2644–2651, 2009.
- [8] M. Bernert and B. Yvert. An attention-based spiking neural network for unsupervised spike-sorting. *International journal of neural systems*, 29(08):1850059, 2019.
- [9] J. Boussard, E. Varol, H. D. Lee, N. Dethe, and L. Paninski. Three-dimensional spike localization and improved motion correction for neuropixels recordings. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [10] A. P. Buccino and G. T. Einevoll. Mearec: a fast and customizable testbench simulator for ground-truth extracellular spiking activity. *Neuroinformatics*, 19(1):185–204, 2021.
- [11] A. P. Buccino, E. Hagen, G. T. Einevoll, P. D. Häfliger, and G. Cauwenberghs. Independent component analysis for fully automated multi-electrode array spike sorting. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 2627–2630. IEEE, 2018.
- [12] A. P. Buccino, C. L. Hurwitz, S. Garcia, J. Magland, J. H. Siegle, R. Hurwitz, and M. H. Hennig. Spikeinterface, a unified framework for spike sorting. *Elife*, 9:e61834, 2020.
- [13] A. P. Buccino, M. Kordovan, T. V. Ness, B. Merkt, P. D. Häfliger, M. Fyhn, G. Cauwenberghs, S. Rotter, and G. T. Einevoll. Combining biophysical modeling and deep learning for multielectrode array neuron localization and classification. *Journal of neurophysiology*, 120(3):1212–1232, 2018.
- [14] L. A. Camuñas-Mesa and R. Q. Quiroga. A detailed and fast model of extracellular recordings. *Neural computation*, 25(5):1191–1212, 2013.
- [15] D. Carlson and L. Carin. Continuing progress of spike sorting in the era of big data. *Current opinion in neurobiology*, 55:90–96, 2019.

- [16] F. J. Chaure, H. G. Rey, and R. Quian Quiroga. A novel and fully automatic spike-sorting implementation with variable number of features. *Journal of neurophysiology*, 120(4):1859–1871, 2018.
- [17] J. E. Chung, H. R. Joo, J. L. Fan, D. F. Liu, A. H. Barnett, S. Chen, C. Geaghan-Breiner, M. P. Karlsson, M. Karlsson, K. Y. Lee, et al. High-density, long-lasting, and multi-region electrophysiological recordings using polymer electrode arrays. *Neuron*, 101(1):21–31, 2019.
- [18] J. E. Chung, J. F. Magland, A. H. Barnett, et al. A fully automated approach to spike sorting. *Neuron*, 95(6):1381–1394, 2017.
- [19] P. U. Diehl and M. Cook. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Frontiers in computational neuroscience*, 9:99, 2015.
- [20] J. Eom, I. Y. Park, S. Kim, H. Jang, S. Park, Y. Huh, and D. Hwang. Deep-learned spike representations and sorting via an ensemble of auto-encoders. *Neural Networks*, 134:131–142, 2021.
- [21] R. Fiáth, B. C. Raducanu, S. Musa, A. Andrei, C. M. Lopez, C. van Hoof, P. Ruther, A. Aarts, D. Horváth, and I. Ulbert. A silicon-based neural probe with densely-packed low-impedance titanium nitride microelectrodes for ultrahigh-resolution in vivo recordings. *Biosensors and Bioelectronics*, 106:86–92, 2018.
- [22] J. Fournier, C. M. Mueller, M. Shein-Idelson, M. Hemberger, and G. Laurent. Consensus-based sorting of neuronal spike waveforms. *PloS one*, 11(8):e0160494, 2016.
- [23] F. Franke, M. Fiscella, M. Sevelev, B. Roska, A. Hierlemann, and R. A. da Silveira. Structures of neural correlation and how they favor coding. *Neuron*, 89(2):409–422, 2016.
- [24] F. Franke, M. Natora, C. Boucsein, M. H. Munk, and K. Obermayer. An online spike detection and spike classification algorithm capable of instantaneous resolution of overlapping spikes. *Journal of computational neuroscience*, 29(1-2):127–148, 2010.
- [25] U. Frey, U. Egert, F. Heer, S. Hafizovic, and A. Hierlemann. Microelectronic system for high-resolution mapping of extracellular electric fields applied to brain slices. *Biosensors and Bioelectronics*, 24(7):2191–2198, 2009.
- [26] S. Garcia, A. P. Buccino, and P. Yger. How do spike collisions affect spike sorting performance? *bioRxiv*, 2021.
- [27] K. J. Gorgolewski, T. Auer, V. D. Calhoun, R. C. Craddock, S. Das, E. P. Duff, G. Flandin, S. S. Ghosh, T. Glatard, Y. O. Halchenko, et al. The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific data*, 3(1):1–9, 2016.
- [28] D. S. Greenberg and J. N. Kerr. Automated correction of fast motion artifacts for two-photon imaging of awake animals. *Journal of neuroscience methods*, 176(1):1–15, 2009.
- [29] E. Hagen, T. V. Ness, A. Khosrowshahi, C. Sørensen, M. Fyhn, T. Hafting, F. Franke, and G. T. Einevoll. Visapy: a python tool for biophysics-based generation of virtual spiking activity for evaluation of spike-sorting algorithms. *Journal of neuroscience methods*, 245:182–204, 2015.
- [30] H. Hazan, D. Saunders, D. T. Sanghavi, H. Siegelmann, and R. Kozma. Unsupervised learning with self-organizing spiking neural networks. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–6. IEEE, 2018.
- [31] L. Hazan, M. Zugaro, and G. Buzsáki. Klusters, neuroscope, ndmanager: a free software suite for neurophysiological data processing and visualization. *Journal of neuroscience methods*, 155(2):207–216, 2006.

- [32] F. He, R. Lycke, M. Ganji, C. Xie, and L. Luan. Ultraflexible neural electrodes for long-lasting intracortical recording. *IScience*, page 101387, 2020.
- [33] M. H. Hennig, C. Hurwitz, and M. Sorbaro. Scaling spike detection and sorting for next-generation electrophysiology. *In Vitro Neuronal Networks*, pages 171–184, 2019.
- [34] D. N. Hill, S. B. Mehta, and D. Kleinfeld. Quality metrics to accompany spike sorting of extracellular signals. *Journal of Neuroscience*, 31(24):8699–8705, 2011.
- [35] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [36] D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology*, 160(1):106–154, 1962.
- [37] C. L. Hurwitz, K. Xu, A. Srivastava, A. P. Buccino, and M. H. Hennig. Scalable spike source localization in extracellular recordings using amortized variational inference, 2019.
- [38] J. J. Jun, C. Mitelut, C. Lai, S. Gratiy, C. Anastassiou, and T. D. Harris. Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction. *bioRxiv*, page 101030, 2017.
- [39] J. J. Jun, N. A. Steinmetz, J. H. Siegle, D. J. Denman, M. Bauza, B. Barbarits, A. K. Lee, C. A. Anastassiou, A. Andrei, Ç. Aydın, et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679):232–236, 2017.
- [40] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- [41] G. M. Kurtzer, V. Sochat, and M. W. Bauer. Singularity: Scientific containers for mobility of compute. *PloS one*, 12(5):e0177459, 2017.
- [42] J. Lecoq, M. Oliver, J. H. Siegle, N. Orlova, P. Ledochowitsch, and C. Koch. Removing independent noise in systems neuroscience data using deepinterpolation. *Nature Methods*, pages 1–8, 2021.
- [43] J. Lee, C. Mitelut, H. Shokri, I. Kinsella, N. Dethe, S. Wu, K. Li, E. B. Reyes, D. Turcu, E. Batty, et al. Yass: Yet another spike sorter applied to large-scale multi-electrode array recordings in primate retina. *bioRxiv*, 2020.
- [44] K. H. Lee, Y.-L. Ni, J. Colonell, B. Karsh, J. Putzeys, M. Pachitariu, T. D. Harris, and M. Meister. Electrode pooling can boost the yield of extracellular recordings with switchable silicon probes. *Nature communications*, 12(1):1–14, 2021.
- [45] B. Lefebvre, P. Yger, and O. Marre. Recent progress in multi-electrode spike sorting methods. *Journal of Physiology-Paris*, 110(4):327–335, 2016.
- [46] C. Leibig, T. Wachtler, and G. Zeck. Unsupervised neural spike sorting for high-density microelectrode arrays with convolutive independent component analysis. *Journal of neuroscience methods*, 271:1–13, 2016.
- [47] Z. Li, Y. Wang, N. Zhang, and X. Li. An accurate and robust method for spike sorting based on convolutional neural networks. *Brain Sciences*, 10(11):835, 2020.
- [48] J. Magland, J. J. Jun, E. Lovero, A. J. Morley, C. L. Hurwitz, A. P. Buccino, S. Garcia, and A. H. Barnett. Spikeforest, reproducible web-facing ground-truth validation of automated neural spike sorters. *Elife*, 9:e55167, 2020.

- [49] A. Markanday, J. Bellet, M. E. Bellet, J. Inoue, Z. M. Hafed, and P. Thier. Using deep neural networks to detect complex spikes of cerebellar purkinje cells. *Journal of neurophysiology*, 123(6):2217–2234, 2020.
- [50] A. Marques-Smith, J. P. Neto, G. Lopes, J. Nogueira, L. Calcaterra, J. Frazão, D. Kim, M. G. Phillips, G. Dimitriadis, and A. R. Kampff. Recording from the same neuron with high-density cmos probes and patch-clamp: a ground-truth dataset and an experiment in collaboration. *bioRxiv*, page 370080, 2018.
- [51] M. Martone, R. Gerkin, R. Moucek, S. Das, W. Goscinski, J. Hellgren-Kotaleski, D. Kennedy, T. Leergaard, J. Boline, and M. Abrams. Nix–neuroscience information exchange format. *F1000Research*, 9, 2020.
- [52] D. Merkel et al. Docker: lightweight linux containers for consistent development and deployment. *Linux journal*, 2014(239):2, 2014.
- [53] S. L. Mondragón-González and E. Burguière. Bio-inspired benchmark generator for extracellular multi-unit recordings. *Scientific reports*, 7(1):1–13, 2017.
- [54] J. P. Neto, G. Lopes, J. Frazão, J. Nogueira, P. Lacerda, P. Baião, A. Aarts, A. Andrei, S. Musa, E. Fortunato, et al. Validating silicon polytrodes with paired juxtacellular recordings: method and dataset. *Journal of neurophysiology*, 116(2):892–903, 2016.
- [55] J. Niediek, J. Boström, C. E. Elger, and F. Mormann. Reliable analysis of single-unit recordings from the human brain under noisy conditions: tracking neurons over hours. *PloS one*, 11(12):e0166598, 2016.
- [56] M. Pachitariu, N. A. Steinmetz, S. N. Kadir, M. Carandini, and K. D. Harris. Fast and accurate spike sorting of high-channel count probes with kilosort. *Advances in neural information processing systems*, 29:4448–4456, 2016.
- [57] P. C. Petersen, J. H. Siegle, N. A. Steinmetz, S. Mahallati, and G. Buzsáki. Cellexplorer: A framework for visualizing and characterizing single neurons. *Neuron*, 109(22):3594–3608, 2021.
- [58] J. W. Pillow, J. Shlens, E. Chichilnisky, and E. P. Simoncelli. A model-based spike sorting algorithm for removing correlation artifacts in multi-neuron recordings. *PloS one*, 8(5):e62123, 2013.
- [59] C. Pouzat, M. Delescluse, P. Viot, and J. Diebolt. Improved spike-sorting by modeling firing statistics and burst-dependent spike amplitude attenuation: a markov chain monte carlo approach. *Journal of neurophysiology*, 91(6):2910–2928, 2004.
- [60] M. Rácz, C. Liber, E. Németh, R. Fiáth, J. Rokai, I. Harmati, I. Ulbert, and G. Márton. Spike detection and sorting with deep learning. *Journal of neural engineering*, 17(1):016038, 2020.
- [61] H. G. Rey, C. Pedreira, and R. Q. Quiroga. Past, present and future of spike sorting techniques. *Brain research bulletin*, 119:106–117, 2015.
- [62] J. Rokai, M. Rácz, R. Fiáth, I. Ulbert, and G. Márton. Elvisort: encoding latent variables for instant sorting, an artificial intelligence-based end-to-end solution. *Journal of Neural Engineering*, 18(4):046033, 2021.
- [63] A. Rokem, B. Dichter, C. Holdgraf, and S. Ghosh. Pan-neuro: interactive computing at scale with brain datasets. 2021.
- [64] F. Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.

- [65] C. Rossant and K. D. Harris. Hardware-accelerated interactive data visualization for neuroscience in Python. *Front Neuroinform*, 7:36, 2013.
- [66] C. Rossant, S. N. Kadir, D. F. Goodman, J. Schulman, M. L. Hunter, A. B. Saleem, A. Grosmark, M. Belluscio, G. H. Denfield, A. S. Ecker, et al. Spike sorting for large, dense electrode arrays. *Nature neuroscience*, 19(4):634–641, 2016.
- [67] O. Rübel, A. J. Tritt, R. Ly, B. K. Dichter, S. S. Ghosh, L. Niu, I. Soltesz, K. Svoboda, L. M. Frank, and K. Bouchard. The neurodata without borders ecosystem for neurophysiological data science. *bioRxiv*, 2021.
- [68] U. Rutishauser, E. M. Schuman, and A. N. Mamelak. Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo. *Journal of neuroscience methods*, 154(1-2):204–224, 2006.
- [69] I. D. Ruz and S. R. Schultz. Localising and classifying neurons from high density mea recordings. *Journal of neuroscience methods*, 233:115–128, 2014.
- [70] M. Saif-ur Rehman, O. Ali, S. Dyck, R. Lienkämper, M. Metzler, Y. Parpaley, J. Wellmer, C. Liu, B. Lee, S. Kellis, et al. Spikeddeep-classifier: A deep-learning based fully automatic offline spike sorting algorithm. *Journal of Neural Engineering*, 18(1):016009, 2021.
- [71] M. Saif-ur Rehman, R. Lienkämper, Y. Parpaley, J. Wellmer, C. Liu, B. Lee, S. Kellis, R. Andersen, I. Iossifidis, T. Glasmachers, et al. Spikedeeper: A deep-learning based method for detection of neural spiking activity. *Journal of neural engineering*, 16(5):056003, 2019.
- [72] E. Sedaghat-Nejad, M. A. Fakharian, J. Pi, P. Hage, Y. Kojima, R. Soetedjo, S. Ohmae, J. F. Medina, and R. Shadmehr. P-sort: an open-source software for cerebellar neurophysiology. *Journal of Neurophysiology*, 126(4):1055–1075, 2021. PMID: 34432996.
- [73] P. S. Shabestari, A. P. Buccino, S. S. Kumar, A. Pedrocchi, and A. Hierlemann. A modulated template-matching approach to improve spike sorting of bursting neurons. In *2021 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 1–4. IEEE, 2021.
- [74] K. Q. Shan, E. V. Lubenov, and A. G. Siapas. Model-based spike sorting with a mixture of drifting t-distributions. *Journal of neuroscience methods*, 288:82–98, 2017.
- [75] J. H. Siegle, X. Jia, S. Durand, S. Gale, C. Bennett, N. Graddis, G. Heller, T. K. Ramirez, H. Choi, J. A. Luviano, et al. A survey of spiking activity reveals a functional hierarchy of mouse corticothalamic visual areas. *Biorxiv*, page 805010, 2019.
- [76] B. C. Souza, V. Lopes-dos Santos, J. Bacelo, and A. B. Tort. Spike sorting with gaussian mixture models. *Scientific reports*, 9(1):1–14, 2019.
- [77] N. A. Steinmetz, C. Aydin, A. Lebedeva, M. Okun, M. Pachitariu, M. Bauza, M. Beau, J. Bhagat, C. Böhm, M. Broux, et al. Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, 372(6539), 2021.
- [78] N. A. Steinmetz, C. Koch, K. D. Harris, and M. Carandini. Challenges and opportunities for large-scale electrophysiology with neuropixels probes. *Current opinion in neurobiology*, 50:92–100, 2018.
- [79] J. L. Teeters, K. Godfrey, R. Young, C. Dang, C. Friedsam, B. Wark, H. Asari, S. Peron, N. Li, A. Peyrache, et al. Neurodata without borders: creating a common data format for neurophysiology. *Neuron*, 88(4):629–634, 2015.

- [80] R. Toosi, M. A. Akhaee, and M.-R. A. Dehaqani. An automatic spike sorting algorithm based on adaptive spike detection and a mixture of skew-t distributions. *Scientific Reports*, 11(1):1–18, 2021.
- [81] D. Valencia and A. Alimohammad. Neural spike sorting using binarized neural networks. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 29:206–214, 2020.
- [82] J. Van Pelt, I. Vajda, P. S. Wolters, M. A. Corner, and G. J. Ramakers. Dynamics and plasticity in developing neuronal networks in vitro. *Progress in brain research*, 147:171–188, 2005.
- [83] E. Varol, J. Boussard, N. Dethe, O. Winter, A. Urai, T. I. B. Laboratory, A. Churchland, N. Steinmetz, and L. Paninski. Decentralized motion inference and registration of neuropixel data. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1085–1089. IEEE, 2021.
- [84] R. Veerabhadrapa, M. Ul Hassan, J. Zhang, and A. Bhatti. Compatibility evaluation of clustering algorithms for contemporary extracellular neural spike sorting. *Frontiers in systems neuroscience*, 14:34, 2020.
- [85] M. Wehr, J. S. Pezaris, and M. Sahani. Simultaneous paired intracellular and tetrode recordings for evaluating the performance of spike sorting algorithms. *Neurocomputing*, 26:1061–1068, 1999.
- [86] Z. Wei, B.-J. Lin, T.-W. Chen, K. Daie, K. Svoboda, and S. Druckmann. A comparison of neuronal population dynamics measured with calcium imaging and electrophysiology. *PLoS computational biology*, 16(9):e1008198, 2020.
- [87] T. Werner, E. Vianello, O. Bichler, D. Garbin, D. Cattaert, B. Yvert, B. De Salvo, and L. Perniola. Spiking neural networks based on oxram synapses for real-time unsupervised spike sorting. *Frontiers in neuroscience*, 10:474, 2016.
- [88] F. Wood, M. J. Black, C. Vargas-Irwin, M. Fellows, and J. P. Donoghue. On the variability of manual spike sorting. *IEEE Transactions on Biomedical Engineering*, 51(6):912–918, 2004.
- [89] J. Wouters, F. Kloosterman, and A. Bertrand. A data-driven spike sorting feature map for resolving spike overlap in the feature space. *Journal of Neural Engineering*, 18(4):0460a7, 2021.
- [90] J. Wouters, F. Kloosterman, and A. Bertrand. Shybrid: A graphical tool for generating hybrid ground-truth spiking data for evaluating spike sorting performance. *Neuroinformatics*, 19(1):141–158, 2021.
- [91] T. Wu, A. Rátkai, K. Schlett, L. Grand, and Z. Yang. Learning to sort: Few-shot spike sorting with adversarial representation learning. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 713–716. IEEE, 2019.
- [92] T. Wu, W. Zhao, E. Keefer, and Z. Yang. Deep compressive autoencoder for action potential compression in large-scale neural recording. *Journal of neural engineering*, 15(6):066019, 2018.
- [93] P. Yger and M. Gilson. Models of metaplasticity: a review of concepts. *Frontiers in computational neuroscience*, 9:138, 2015.
- [94] P. Yger, G. L. Spampinato, E. Esposito, B. Lefebvre, S. Deny, C. Gardella, M. Stimberg, F. Jetter, G. Zeck, S. Picaud, et al. A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo. *Elife*, 7:e34518, 2018.
- [95] X. Yuan, M. Schröter, M. E. J. Obien, M. Fiscella, W. Gong, T. Kikuchi, A. Odawara, S. Noji, I. Suzuki, J. Takahashi, et al. Versatile live-cell activity analysis platform for characterization of neuronal dynamics at single-cell and network level. *Nature communications*, 11(1):1–14, 2020.