"Supersessioning": A hardware/software system for electrophysiology spanning multiple sessions in marmosets

Jens-Oliver Muthmann[1*], **Aaron J. Levi**[1], **Hayden C. Carney**[1], **Alexander C. Huk**[1*]

[1]The University of Texas at Austin

***For correspondence:**
huk@utexas.edu (ACH);
ollimuh@utexas.edu (JOM)

Abstract We introduce a straightforward, robust method for recording and analyzing spiking activity over timeframes longer than a single session, with primary application to the marmoset (*Callithrix jacchus*). Although in theory the marmoset's smooth brain allows for broad deployment of powerful tools in primate cortex, in practice marmosets do not typically engage in long experimental sessions akin to rhesus monkeys. This potentially limits their value for detailed, quantitative neurophysiological study. Here we describe chronically-implanted arrays with a 3D arrangement of electrodes yielding stable single and multi- unit responses, and an analytic method for creating "supersessions" combining that array data across multiple experiments. We could match units across different recording sessions over several weeks, demonstrating the feasibility of pooling data over sessions. This could be a key tool for extending the viability of marmosets for dissecting neural computations in primate cortex.

Introduction

The marmoset has drawn attention as a complementary nonhuman primate model system for visual neuroscience. While the dominant primate model system in neuroscience, the rhesus monkey (*Macaca mulatta*), has the advantage of (relatively) rich cognitive abilities, a large body and robust physiology, and an aggressive work ethic, their large and convoluted (gyrified) brains currently limit the number of techniques that can be applied for measurements of neural activity. Thus, despite their excellent trainability for complex tasks and willingness to engage in lengthy experimental sessions, the scale and variety of neurophysiological questions that can be addressed have been somewhat limited by practical constraints. Recently, the common marmoset (*Callithrix jacchus*) has emerged as a complementary primate model system because of their smooth (lissencephalic) cortex, opening up a much larger number cortical areas to the use of large-scale chronically implanted electrode arrays (in addition to other techniques). However, a major current concern for adopting the awake behaving marmoset for detailed quantitative studies is their tendency to perform far fewer trials per session compared to macaques. Such a behavioral limitation would result in correspondingly smaller amounts of neural data (and hence, statistical power) per experiment, undercutting the other advantages of the species, and likely limiting their applicability as a powerful neurophysiological complement to the sorts of quantitative neuroscience work done in macaques.

To redress this fundamental potential limitation, we have developed a straightforward, user-friendly tool for recording from large-scale arrays in marmosets while surmounting the relatively short behavioral sessions performed by this smaller (and gentler) species. First, we report success-

ful long-term electrophysiological recordings using a new type of multi-electrode array for which primate use has not yet been reported in publication to our knowledge, but which is commercially available. These "3D" arrays are available with customizable electrode spacing not just across a 2D grid, but also along the depth of individual shanks. The arrays yielded good quality single-unit (SUA) and multi-unit (MUA) activity, as demonstrated in two different marmoset cortical areas (area MT, and the posterior parietal cortex, PPC). Second, we introduce a transparent means for identifying activity recorded on these arrays, not just within individual sessions, but — importantly — *across* sessions. This integration of hardware and software solutions allowed for data from the same unit to be combined over multiple behavioral sessions, into what we termed "supersessions." This brings the statistical power of awake-behaving marmoset neurophysiology closer to that of macaques on a per-unit basis, while still allowing for larger scale recordings and/or powerful complementary tools that are more challenging to perform in macaques.

Here, we describe both the physiological and computational components of this tool and demonstrate its potential usefulness for transcending the behavioral limitations of marmosets into the realm of detailed, quantitative assessments of neural activity at large scales. Furthermore, the tool we introduce here is intentionally straightforward, meaning it can be readily implemented by others, as well as extended when ongoing updates to hardware and software emerge. We conclude by describing current limitations and how updates to this tool could further improve it.

To provide a bit more detail before delving into the results, we found that implanting commercially-available 3D "N-form arrays" (Modular Bionics) resulted in high quality, stable unit activity in marmosets. In our hands and experiences, this reflected a significant step forward in neural recording success, as prior attempts using more common types of 2D planar arrays (Utah, Black rock systems) yielded less reliable and lower-quality outcomes. Although our goal was simply to record neural activity and not to mechanistically understand why a particular array style works better or worse, our hypothesis is that the slow insertion style (***Nicolelis et al., 2003***) and sufficient spacing between shanks of the N-form array (compared to ***Rennaker et al.*** (***2005***); ***Karumbaiah et al.*** (***2013***)), produces less damage, inflammation, and/or gliosis, and still reduces chronic respiratory micromotion. (***Prodanov and Delbeke, 2016***).

Given the success of the neural recording hardware in yielding qualitatively impressive neural activity over long time periods, we designed a method to systematically compare and match (distributions of) spike waveforms across sessions. Our method identifies units from individual sessions independently, and then integrates spike clusters from new recordings into known, existing ones identified in prior sessions. Analyses of units can therefore be performed over multiple experimental sessions.

In order to achieve a representation of spike shapes that was robust to potentially varying noise levels and/or forms across experimental sessions, we extracted simple properties of spike shapes in a narrow window around their peak. This was achieved by matching a family of predefined templates on a GPU to yield a parametric representation of local excursions in the raw voltage traces, which included conventional unit spiking activity, spike events from weaker or more distant neural sources, and noise. Unit isolation was not a conventional detection problem, and was instead transformed into a multivariate classification problem to be solved by a clustering algorithm. The resulting clusters were then matched across recording sessions. Although we are not deeply attached to this particular spike sorting approach, we provide it as a robust, intuitive starting point, which we validated against a more sophisticated and complex spike-sorting package. Its simplicity also allows for online views of sorting results during experiments, which could be useful for experimental decisions even if more sophisticated sorting routines are employed post hoc.

Taken together, this work puts forth a synthesis of commercially-available hardware and intuitive software that allows experimenters to overcome one of the major limitations of the marmoset as a model species by introducing the concept of supersessions. More generally, this framework may support better integration of work done in marmosets and macaques, allowing these two awake-behaving primate preparations to have greater scientific overlap and thus to more solidly

allow for their relative strengths and weaknesses to be considered.

Results

Neural activity apparent for more than 9 months on chronically-implanted 3D arrays

We recorded single and multi-unit (hereafter, "unit") activity in the brains of 2 marmosets, one with a 3D N-form array in and around the middle temporal area (MT), the other with an identical array placed in posterior parietal cortex (PPC). For both arrays (Figure 1 A, B, respectively), we were able to record spiking activity starting a week after insertion. Activity lasted for a duration of at least 9 months, as depicted in Figure 1 (top rows). Figure 1 (second rows) show, in comparison, the relatively short durations of individual recording sessions (approximately a half hour to an hour). These durations likely reflect a lower bound on how long marmosets will work, as they were largely determined by the animal's preponent motivation to engage in various visual tasks with no fluid or food restriction.

Signal amplitudes (Figure 1, third rows) were fairly constant over long periods of time, perhaps with the first two weeks after implantation yielding smaller signals before stabilizing (i.e., first few recording sessions, visible at the very left of the plots). A gradual decline in signal amplitude was further apparent after about 7 months for marmoset J. Detected events (see Methods) had a wide amplitude range of relatively sparse (0.1 – 10 Hz) events, indicative of spiking activity (Figure 1, bottom rows). Taken together, these descriptions of the behavior of the animals and the signals from the electrode arrays lay the groundwork for attempting to stitch together data from multiple, subsequent recording sessions. The next critical step would be identifying unit activity that could conservatively be identified across such sessions.

Spike clusters overlap in consecutive sessions

Our goal was to identify spikes from the same units across recording sessions. This required measures that would be robust to noise, in the sense that spikes from other neurons would not perturb or distort characterization and identification of a given unit. To that aim, we focused our analysis on a very short temporal window, including only the depolarization phase of a spike, represented by a local minimum in the raw voltage traces.

For each local minimum (i.e., putative spike) in the raw voltage trace, we determined: (a) amplitude, measured as the dot product with a template (of unit power), expressed in standard deviations (σ), as calculated on the high-pass filtered voltage traces; (b) width, measured as the full width at half minimum; and (c) symmetry, measured as the ratio of its falling and rising phase durations (i.e., a 1 : 2 ratio means that recovering back to baseline took twice as long as reaching the voltage minimum).

These parameterized shape characterizations of the units were put into 3D-histograms (marginals shown in Figure 2 A) for each recording session, and clustered using a watershed algorithm (see Methods for details). This procedure yielded shape clusters (cyan markers in Figure 2 A) for every session in a common coordinate system to allow for cross-session comparisons of spike shapes. Shape clusters between consecutive sessions often looked very similar, and so we further tested whether they likely reflected spikes from the same or from different units.

Specifically, if the brain tissue was held in place by the 16 electrode shanks of the array such that relative movements between the electrodes and the sampled neurons rarely happened, we would always record from the same neurons and see identical spike shapes. Otherwise, if there were substantial shifts in relative position between brain and electrodes, both amplitude and spike shape would shift with movement, and we would be unable to track units across a large number of sessions.

We were indeed able to systematically match units across recordings. This was done quantitatively, using the Jensen-Shannon divergence as a distance measure in the histogram shape space

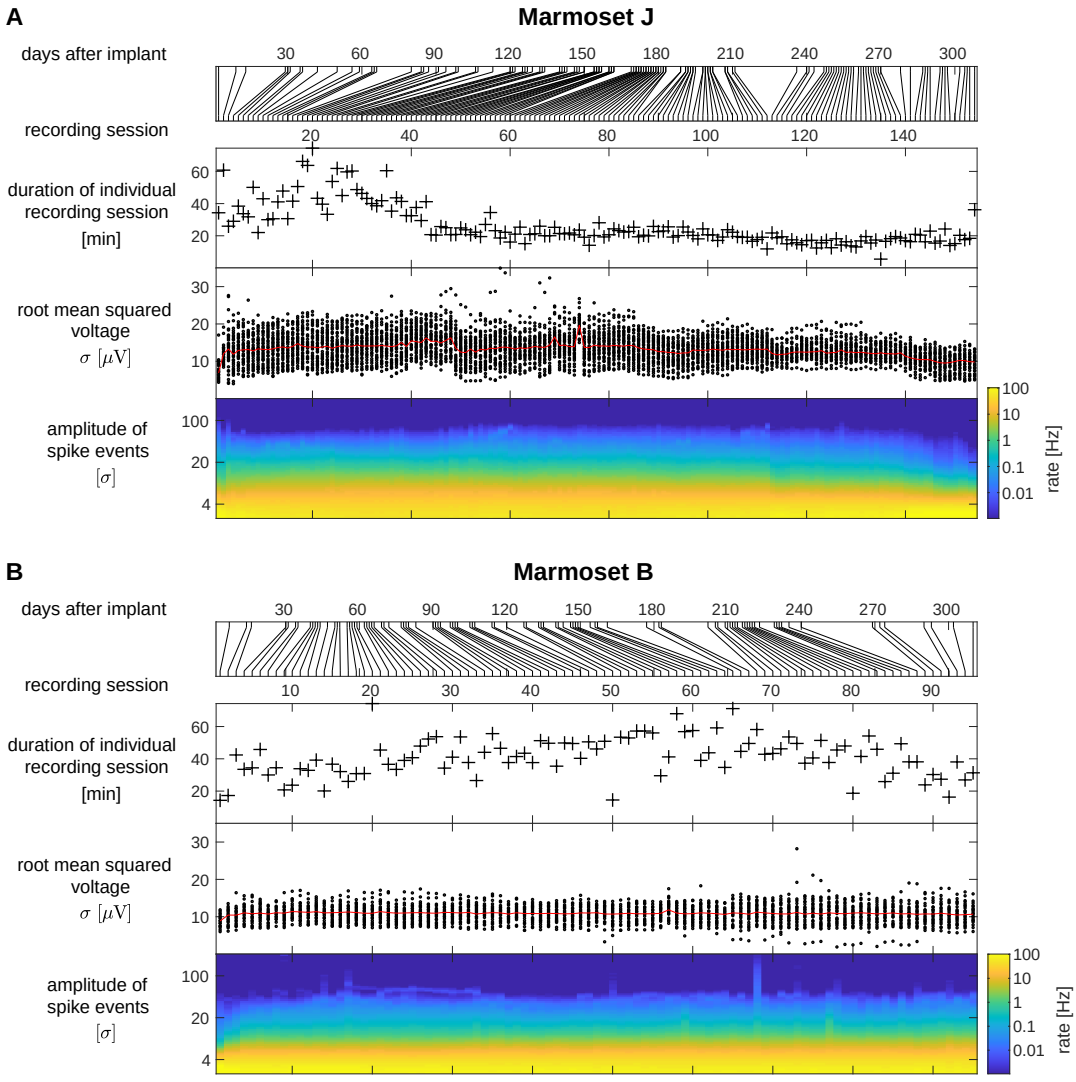

Figure 1. Long-term stability of arrays. (**A**) Marmoset J. Top panel: Illustration when individual recording sessions were performed. For clarity, the plots below and in subsequent Figures reflect individual recording sessions rather than time. Second row: Durations of electrophysiological recordings in individual sessions. Third row: Root-mean-squared voltage fluctuations of the common averaged, 300 Hz high-pass filtered data (scatter plots for active electrodes, average shown in red). Bottom row: Amplitude histograms of detected events, averaged across electrodes. (**B**) Same statistics for Marmoset B.

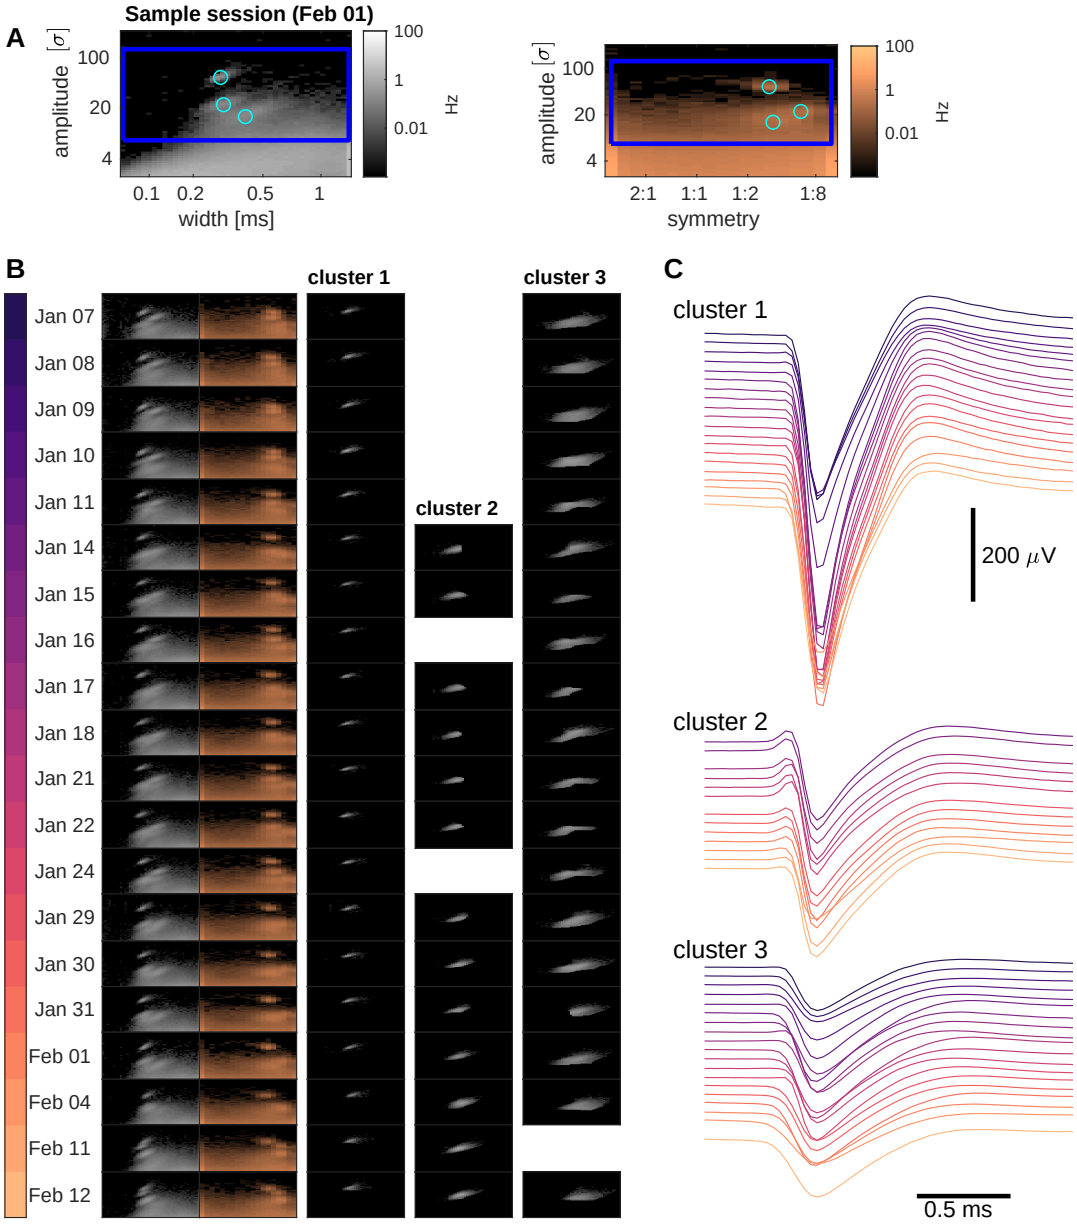

Figure 2. Example of merging clusters across sessions. (**A**) Histograms for amplitudes and widths (left panel) or symmetries right panes) of detected events on February 1. Regions outlined in blue are shown for a range of dates in (B), using the same color code and axes. Cyan circles mark the three clusters detected in this session. (**B**) Left: marginal histograms of local maxima for 20 consecutive recording sessions, labeled with dates. Right: temporal matches of the 3 clusters found on February 1. (**C**) Waterfall plots of average spike shapes, for dates as color-coded in (B). Data from marmoset B.

(allowing for small amplitude shifts under a penalty). Figure 2 B shows an example of tracking the 3 units observed on February 1 across multiple sessions. Cluster 1 provides an example of a clearly isolated unit with very large spikes, which lasted for about 5 weeks. For this cluster, averaged spike shapes were very similar across recording sessions, with smaller amplitudes for the initial and final recordings (Figure 2 C, cluster 1). Cluster 2 represents a cluster with decent amplitude spikes but relatively common spike shapes, resulting in highly variable sorting performance. While being reasonably well isolated from January 29 to February 1, it is contaminated to a variable degree with spikes from different units in other sessions and couldn't be separated from another cluster in two intermediate recording sessions. Cluster 3 had low spike amplitudes, but would be considered a decent multi-unit cluster from January 29 to February 1. For the other sessions there is a small local maximum in the shape histograms, but the cluster would be considerably contaminated with unclassified, smaller amplitude spikes. Given that larger amplitude clusters slowly (and independently) drift over time, we can assume that the same happens to units in this cluster, making it difficult to obtain exact matches across recordings. But, the relatively moderate firing rate of the cluster would suggest that few units with defined shapes were involved, distinguishing it from unclassified spikes.

In conclusion, our main result is that matching simple shape statistics of spike waveforms across several recording sessions using N-form arrays in marmosets is feasible, and for some units this consecutive recording is possible over notably long periods of time (> 1 month). This grants us the capacity to combine data from multiple experimental days, which we deem "supersessions".

Tuning properties on individual electrodes are stable across sessions

We further confirmed the stability of the measured "supersession" neuronal activity by evaluating the cross-session consistency of physiological tuning properties. This evaluation was done for the MT array implanted in marmoset J, where we were able to confirm that several sites on the array showed directionally-tuned activity in response to moving dots in the left visual field (as expected when recording from area MT in the right hemisphere).

The MT electrodes recorded strongly tuned multi-unit activity, so we focused on MUA supersessions for this analysis. We again used our parameterized representation of spike shapes to determine a region of interest (Figure 3 A, E, outlined in black) in spike shape space with strong directional tuning across recording sessions (Figure 3 A, E). This was feasible because tuning on a given electrode was consistent across a wide range of spike shapes (Figure 3 B, F). For the two MUA sites shown as examples, the direction tuning curves measured were stable over almost 3 weeks. This stability of physiological properties, built on top of the stability of spike shapes themselves, further strengthens the case for the validity and viability of supersessions.

We therefore created supersessions across these sessions that exhibited stable tuning and spike shapes, which allowed us to combine larger amounts of data for a single analysis. As an example here, we show that supersessions allow us to resolve the detailed time course of responses to individual motion directions at a high temporal resolution (Figure 3 C, G). Note that transient aspects of the motion-driven response were very short and consisted of only a few spikes per trial, such that averages across many trials were beneficial. To illustrate this effect, we show the same analysis for responses obtained in a single session (Figure 3 I-K). Averaging over the temporal responses, we then obtained tuning curves for individual sessions (Figure 3 D, H, L).

In this example, tuning was stable for considerably longer than one week. This demonstrates that not only were shape clusters with high amplitudes were stable across sessions, but also that functional properties of low-amplitude activity were conserved across many sessions. Furthermore, being able to combine 10 or more sessions provides an order-of-magnitude increase in trial count that, even assuming some degree of lower-quality unit isolation, should counterweight the relatively short individual behavioral sessions. We delve into this issue in more depth at the end of the results sections.

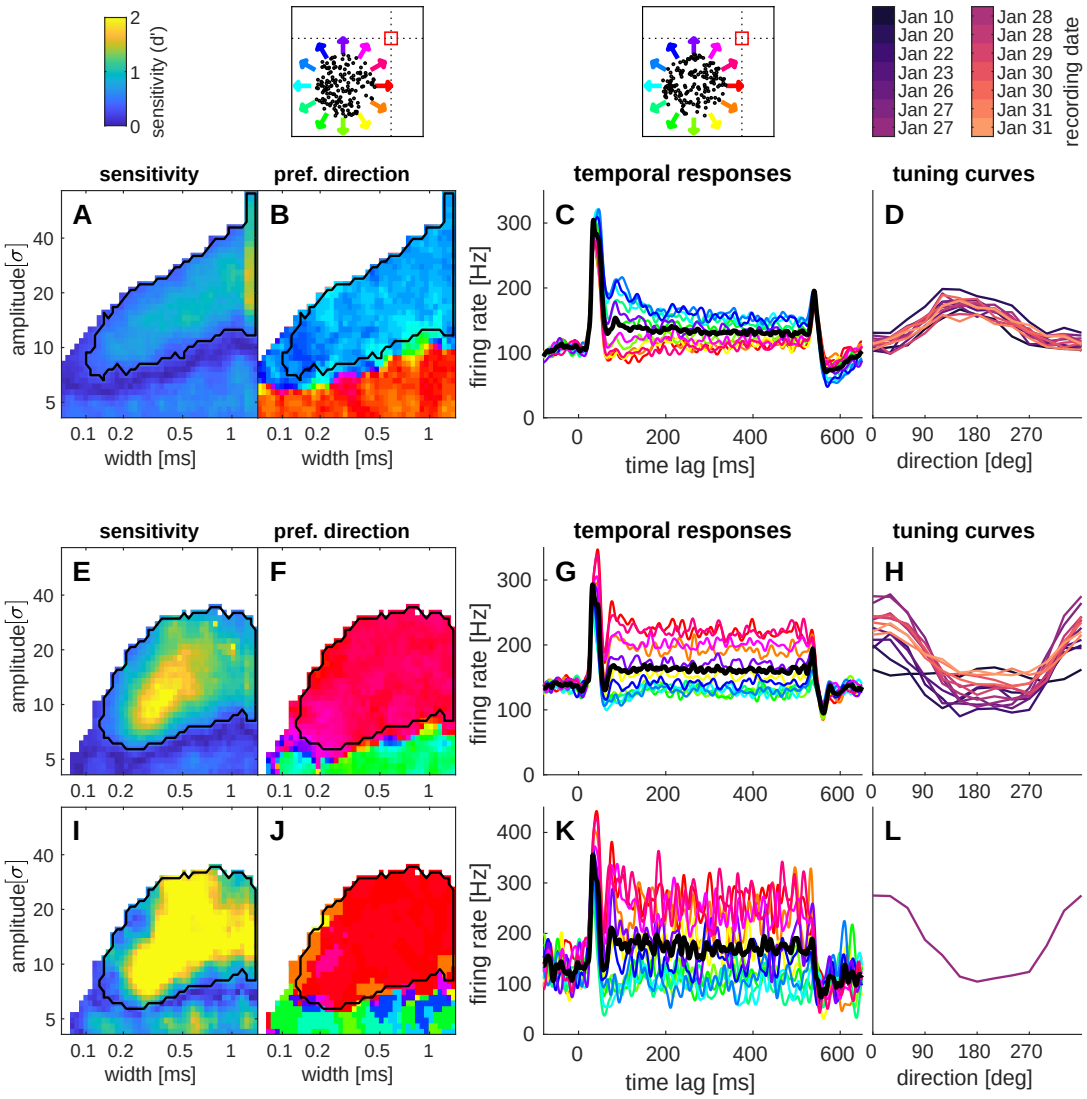

Figure 3. Examples of direction tuning on two electrodes. Top: Legends and stimuli for the examples below. Moving dots were presented at (-15,-15) degrees from the fixation point (red square).
(**A**) Sensitivity indices and (**B**) maximum response directions in dependence of spike shapes. (across sessions, corrected for a cross session baseline effect). The region outlined in black was used for further analysis. (**C**) Temporal firing rate responses, averaged across sessions and shown for individual tuning directions (colored lines, black line: avg. response, 4041 trials). (**D**) Tuning curves obtained for individual recording sessions (labeled above, some dates had a morning and afternoon session). (**E – H**) Same analysis for a second example electrode. (**I – L**) Tuning observed in a single session (January 27 afternoon session, 254 trials). Recordings in area MT (marmoset J).

Most units in a given recording were observed for several sessions

Having established stability of both spike waveforms and physiological tuning, we now turn to report a more comprehensive statistical description of recording stability and our ability to distinguish spike shape clusters (i.e., to isolate one unit from another). A summary of all tracked units across recording sessions is shown in Figure 4. Spike clusters were regions in 3D-shape-histograms, consisting of a set of voxels, which could be divided into boundary voxels (adjacent to a voxel outside the cluster) and center voxels. If the average spike count in boundary voxels was less than 3/4 of the average density in center voxels, clusters were considered as "better-isolated" and shown in darker colors in Figure 4.

We further distinguished clusters that lasted for shorter numbers of sessions (<5, orange) and longer numbers of sessions (blue, ≥ 5), as many of the short-lived units had low amplitudes and were less reliably detected. We found that a large proportion of units in a given recording survived for multiple recording sessions (histograms in Figure 4, blue vs. orange), especially when they were considered as better-isolated (Figure 4, darker colors).

A more detailed visualization of the survival of individual units is shown in the upper half of both panels in Figure 4. This plot can resolve whether the appearance or disappearance of units between two sessions happened locally (i.e., affecting only some individual units), or globally (i.e., affecting most, if not all, units across the array). To further see whether the temporal separation (i.e., number of days) between consecutive sessions was a major factor for the loss (/turnover) of units, we visualized the relation between the number of long lasting units lost and the temporal separation between the two sessions when the loss occurred (Figure 4, insets). Although larger temporal separations tended to correlate with a higher turnover of units, substantial unit turnover could also occur even with very short temporal separations between sessions.

This analysis also highlights a difference between the two animals: while there are several distinct time points of high turnover in marmoset J (Figure 4 A, likely indicative of discrete changes in electrode array position), no such events could be identified in marmoset B (Figure 4 B, likely indicative of only smaller and/or more gradual changes in array position within the brain). Although we are not sure why the array stability was different in the two animals, this does show that: (a) our analysis scheme is capable of revealing changes and differences in stability; and (b) regardless of whether an array was stable over longer or short terms with or without distinct temporal changes, it is possible to follow units across supersessions in both regimes.

Figure 5 shows descriptive histograms of the basic properties of all detected shape clusters (grayscale background). We distinguished clusters that survived short-term (upper row) and long-term (lower row). Several basic relations become apparent from visual inspection. First, the spread (avg. diameter) and firing rates of clusters tended to be larger for smaller amplitude waveforms, likely reflecting the effects of merging overlapping shapes from multiple units. Second, large amplitude waveforms were generally more skewed than those with low amplitudes, likely reflecting our descriptive approach's ability to identify the basic shape of individual unit waveforms. Third, waveforms from the array in MT tended to be narrower than those from the PPC array, perhaps revealing a biophysical difference that our approach is capable of picking up.

Viewing these basic descriptive plots, we also wondered whether long term matches of spike clusters might be a result of detecting different units that just happen to produce similar shapes. To test this, we estimated how likely a given cluster might be mistaken for a different cluster by counting the clusters with similar spike shapes from all recording sessions. We then ranked better-isolated clusters according to the number of similar shaped clusters. The resulting rank a cluster had in the sorted array is depicted in color in Figure 5. A low rank corresponds to isolated units and a low likelihood to detect the same cluster by chance (Figure 5, yellow/green circles), and a high rank means that the corresponding spike shapes were frequently observed (Figure 5, blue circles).

Sorting clusters in this way allows us to investigate whether clusters with commonly observed spike shapes would show a bias in long-term survival. We observed that many clusters with unique

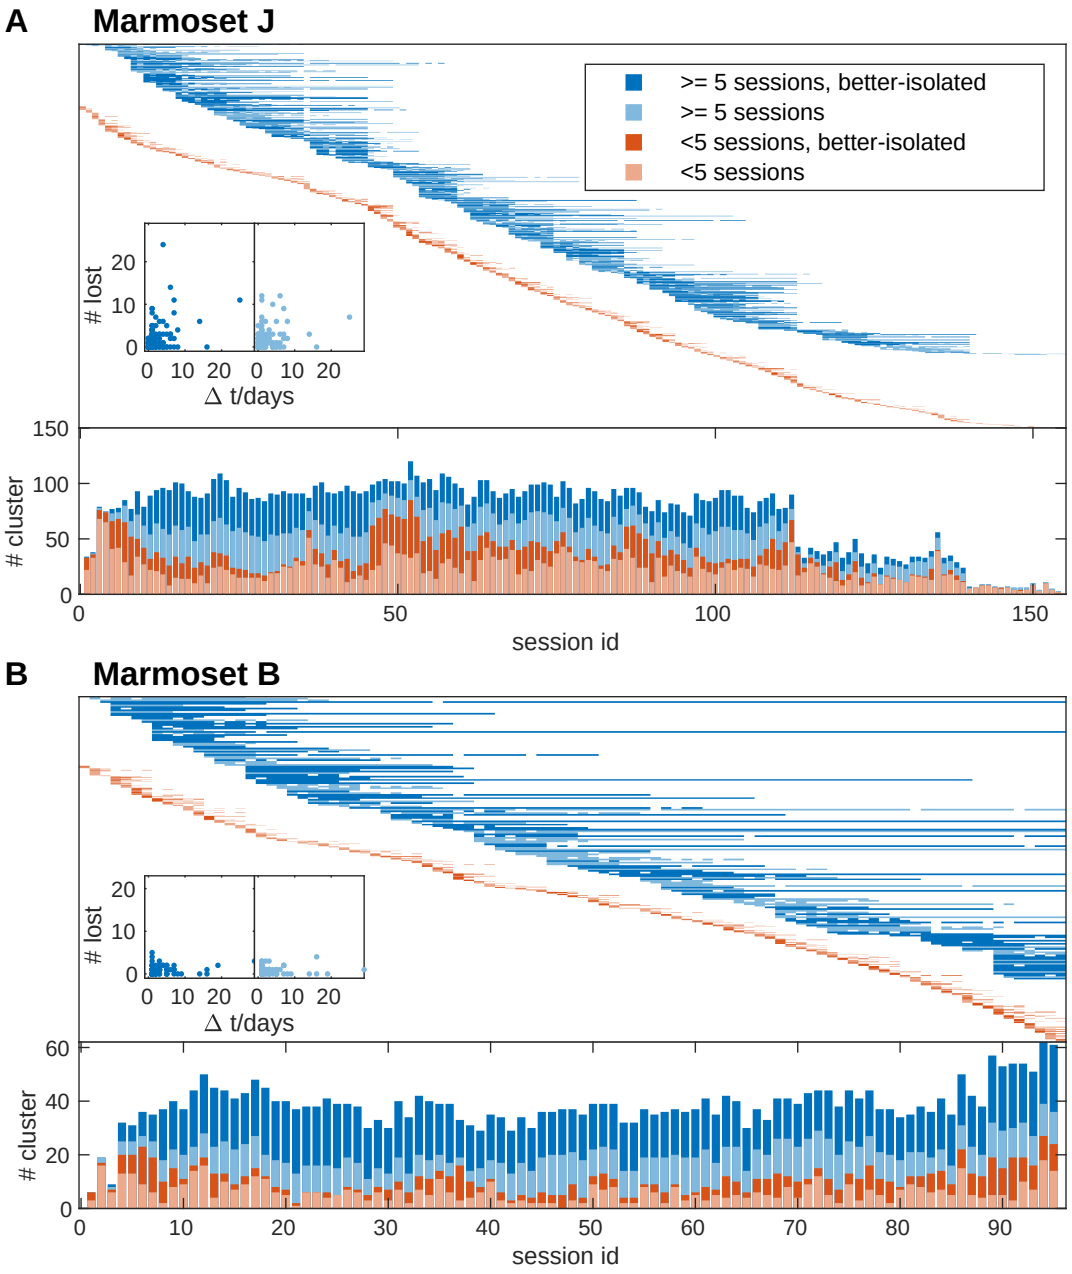

Figure 4. The majority of clusters survives for multiple sessions. (**A**) Clusters detected in recordings of area MT (marmoset J). Top: temporal pattern of long-term (at least 5 sessions, blue) and short lived (<5 sessions, orange) clusters. Better-isolated clusters are shown in darker shades. Inset: Number of disappearing units in dependence of the temporal gap between two recording sessions. Bottom: Number of clusters in each session. (**B**) Same plots for recordings in PPC (marmoset B).

shapes survived less than 5 sessions (Figure 5, yellow circles). However, we also noticed that many of these clusters had relatively low amplitudes and therefore might have been lost, not due to actual changes in the presence of the unit over a particular (brief) time frame, but due to insufficient signal-to-noise ratio relative to our spike-identification standards. We therefore re-focused our analysis of the relation between spike waveform uniqueness and lifetime using only clusters surviving for at least 5 sessions.

In order to assess whether clusters with more or less common waveform shapes might show

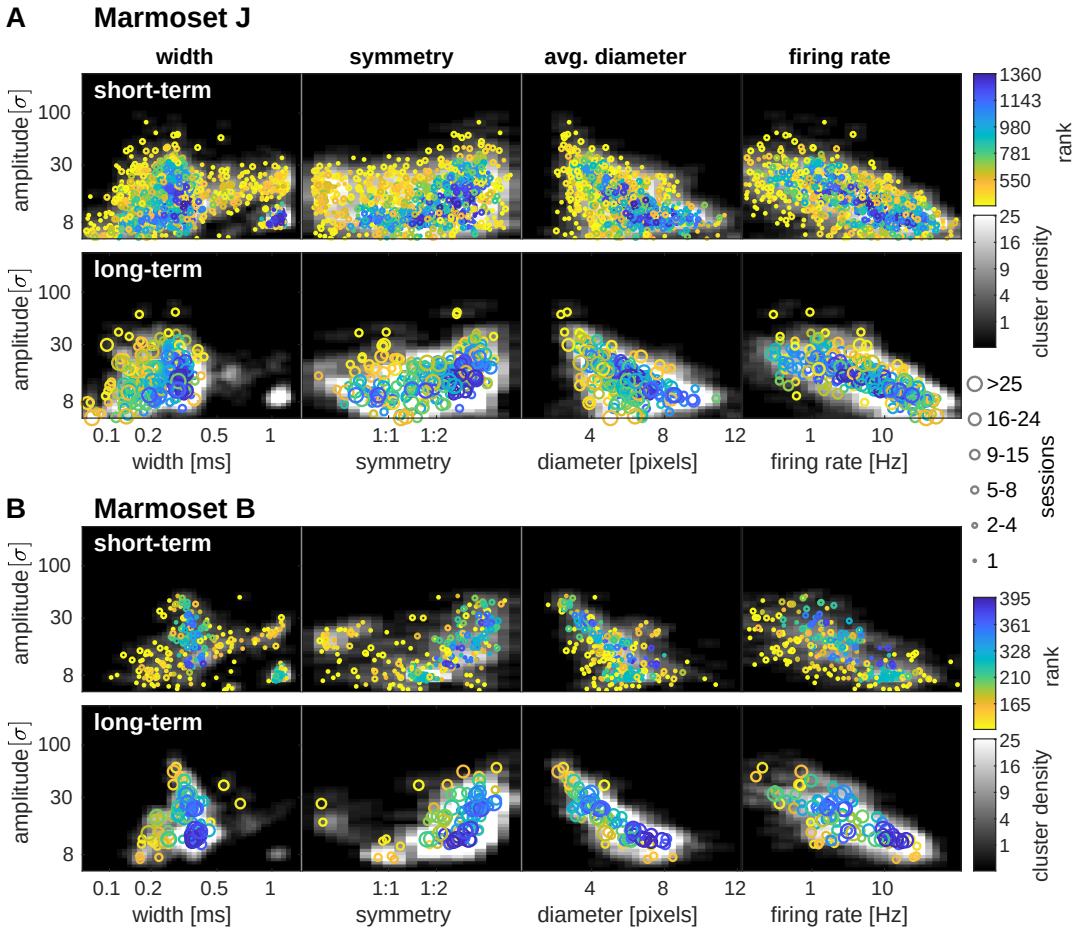

Figure 5. Detected shape clusters are similar (at a population level) when observed for multiple sessions. (**A**) Clusters detected in all recordings and electrodes of area MT (marmoset J). Grayscale represents the density of all detected clusters without merging them across sessions. Colored circles represent individual, better-isolated clusters, merged across sessions. These were ranked according to the corresponding overall density of clusters (i.e. grayscale background) and this ranking is shown in color. Specifically, properties of clusters depicted in yellow were rarely observed and those in blue were commonly found in the data. Clusters surviving less than (top row) and at least (bottom row) 5 sessions are plotted separately for clarity. (**B**) Same analysis for recordings in PPC (marmoset B).

a difference in their lifespans, we analyzed cluster survival, excluding different amounts of the most common cluster shapes. Due to the limited amount of data, we visualized the inverse of the expected lifetime at a given age, which is assuming a constant probability to lose a cluster in each session. Figure 6 shows that this assumption is reasonable, as the fraction of clusters lost per session does not change dramatically after 5 sessions. Importantly, except for clusters with the 10% most uncommon shapes, the rate at which spike clusters were lost over time did not depend on how common the spike shapes of that cluster were. This is good news, as it does not appear that the longevity of units over sessions is strongly confounded by the appearance and disappearance of units which happen to have similar spike shapes.

This analysis also revealed an interesting difference between the two animals: For the array in PPC, cluster survival was about twice as long as for the array in area MT. Although there were more clusters observed for the MT array, we also observed greater variations in signal amplitude and we gradually lost signal in the later recordings of that array (Figure 1 A). We therefore infer that the observed effect could have been due to a higher degree of general instability of the MT array over

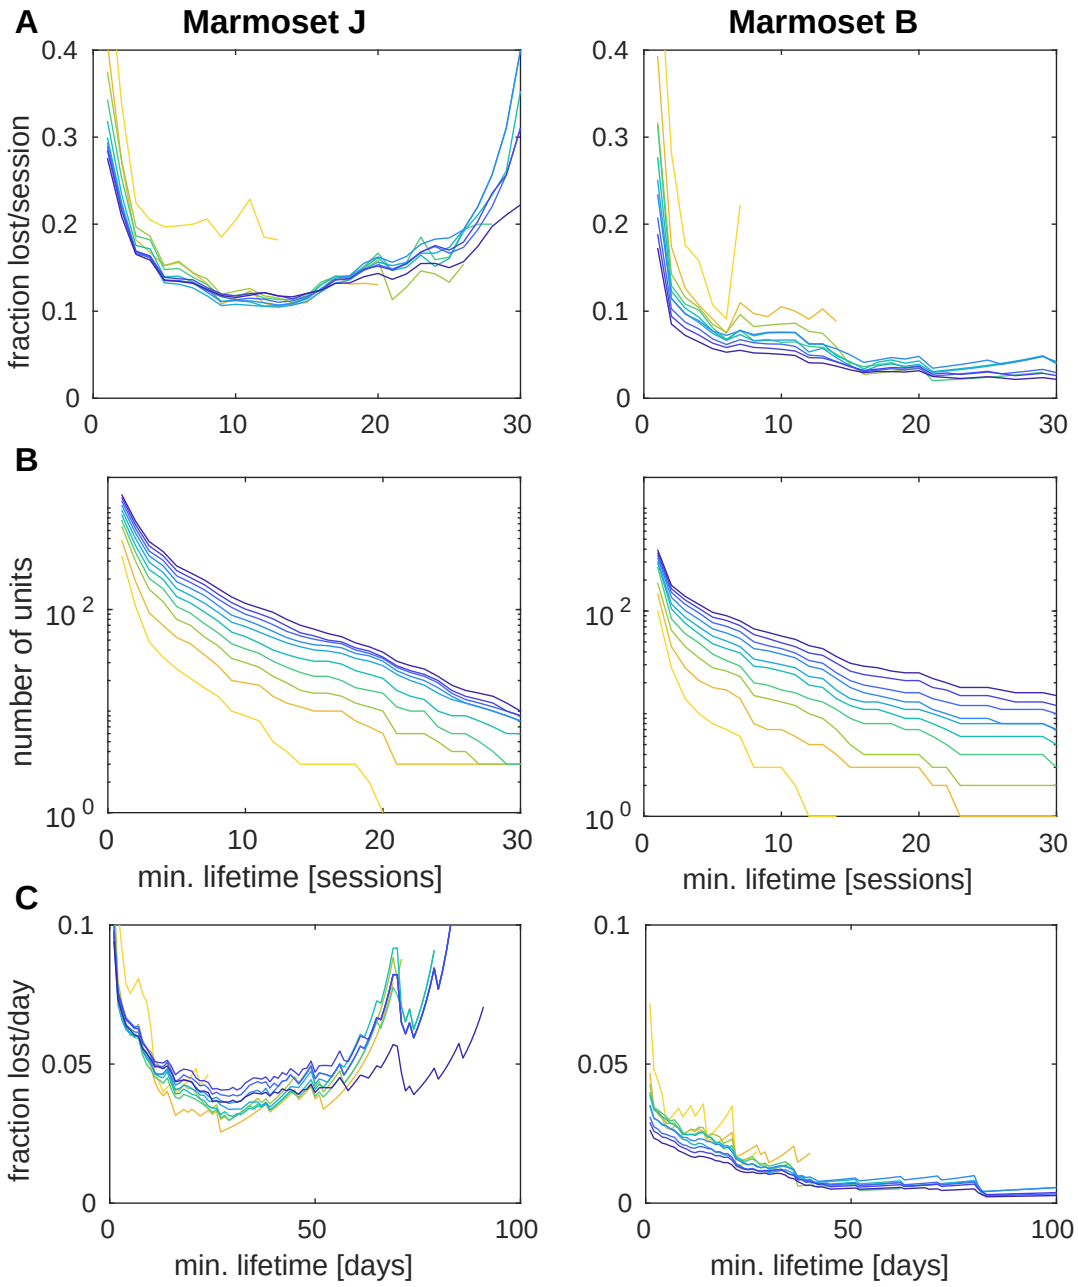

Figure 6. Cluster survival is not an effect of common spike shapes. (**A**) Estimated fraction of the clusters that would be lost after surviving the number of sessions indicated on the x-axis. Coloured lines correspond to the fraction of clusters included in the analysis (steps of 10%, as in Figure 5), where the most yellow curve corresponds to only including the 10% most uncommon shapes. (**B**) Number of units observed for a minimum lifetime. (**C**) Same as in (A) when measured in days rather than sessions. Recordings in area MT (marmoset J, left column) and PPC (marmoset B, right column).

time.

Supersessions provide the power to estimate spatial and temporal aspects of responses across sessions

Finally, we tested whether clearly isolated units could be matched across multiple sessions to assess their spatial and temporal properties. We therefore performed generic receptive field map-

ping assays at regular intervals over multiple experimental sessions. As proof of concept, here, we describe an example in which both spatial receptive fields and temporal dynamics of responses were estimated using supersession data.

Figure 7 shows two example units. The first unit had well isolated, high amplitude spike shapes (Figure 7 C,E) and a pronounced refractory period (Figure 7 F) for at least 6 recording sessions (firing rate $(1.7 \pm 0.2)\,\mathrm{Hz}$; avg. spike count per trial (400 ms) 0.7 ± 0.4 overall and 1.5 ± 0.5 for stimuli in the receptive field). It consistently responded transiently to stimuli in the left visual field, 50-80 ms after stimulus onset. The second example ((Figure 7 G-L) shows a unit with an amplitude gradually increasing and decreasing across sessions. Corresponding to an increase in SNR and lower contamination by false detections averaged spike shapes became sharper for sessions with large spikes (Figure 7 K). This unit had a much faster response around 40 ms, consisting of about 1 spike per trial (and eventually a slightly elevated sustained activity during stimulus presentation). In both of these cases, the response properties of the unit would have been difficult to determine using only a single session's worth of data, due to the low absolute number of spikes recorded. For example, the total number of spikes recorded in the first 400 ms in the receptive field of the unit in a single session was just 20-80 spikes, the total number of spikes across all trials about twice that amount. But by evaluating data across sessions, the supersession data shows that these units had clearly-localized receptive fields.

Discussion

Modern neurophysiological studies in primates require increasingly large amounts of data, either because the parameter space of relevant stimuli or behaviors grows richer (and hence, data are distributed across a larger number of conditions), or because the goal of the experiment itself is to measure more detailed aspects of population activity (and hence, more data are required to estimate higher order statistics). Here, we established the potential of chronically-implanted 3D electrode arrays, coupled with a simple unit identification scheme, to allow for the creation of supersession datasets that transcend the standard limitations of marmoset behavior within individual experimental sessions. We found that high quality activity was evident on this type of array for many months, that a mixture of stable SUA and MUA data could be collected spanning multiple individual sessions, and that these supersessions yielded stable physiological characterizations that were more detailed than those from single sessions.

Recording performance

With the goal of making the marmoset more strongly viable for detailed quantitative studies, we aimed to develop an analysis pipeline that would be robust to different levels of recording quality, measuring single-unit activity where possible, but at the same time considering multi-unit activity. When applying this analysis to data recorded from implanted electrode arrays over the course of more than 9 months and averaging across all recording sessions, we obtained 28 better-isolated units/array/session. For individual arrays, these averages were 32 and 23 for marmoset J and B, respectively, 20 and 18 of which would be seen across a span of five or more sessions. In addition, we found another 40 and 16 multi-unit clusters per array per session for marmosets J and B, respectively; 18 and 10 sessions lasting for five sessions or more).

In comparison, previous reports of recording stability using planar (2D) 'Utah' arrays in macaques (***Dickey et al., 2009***; ***Vaidya et al., 2014***; ***Fraser and Schwartz, 2011***) focused on single unit activity, which strengthened their claims to be able to track individual units, but at the cost of discarding multi-unit activity. Values reported in those prior studies were at most 137 units/array/session, but with large variations across arrays and with decreasing number over time, the average values were closer to 30 units/array/session. In addition, most recordings were done in the first two months after implantation, possibly implying a quicker falloff in signal quality than we encountered with different arrays, and making the comparison to our unit identification and quality less direct.

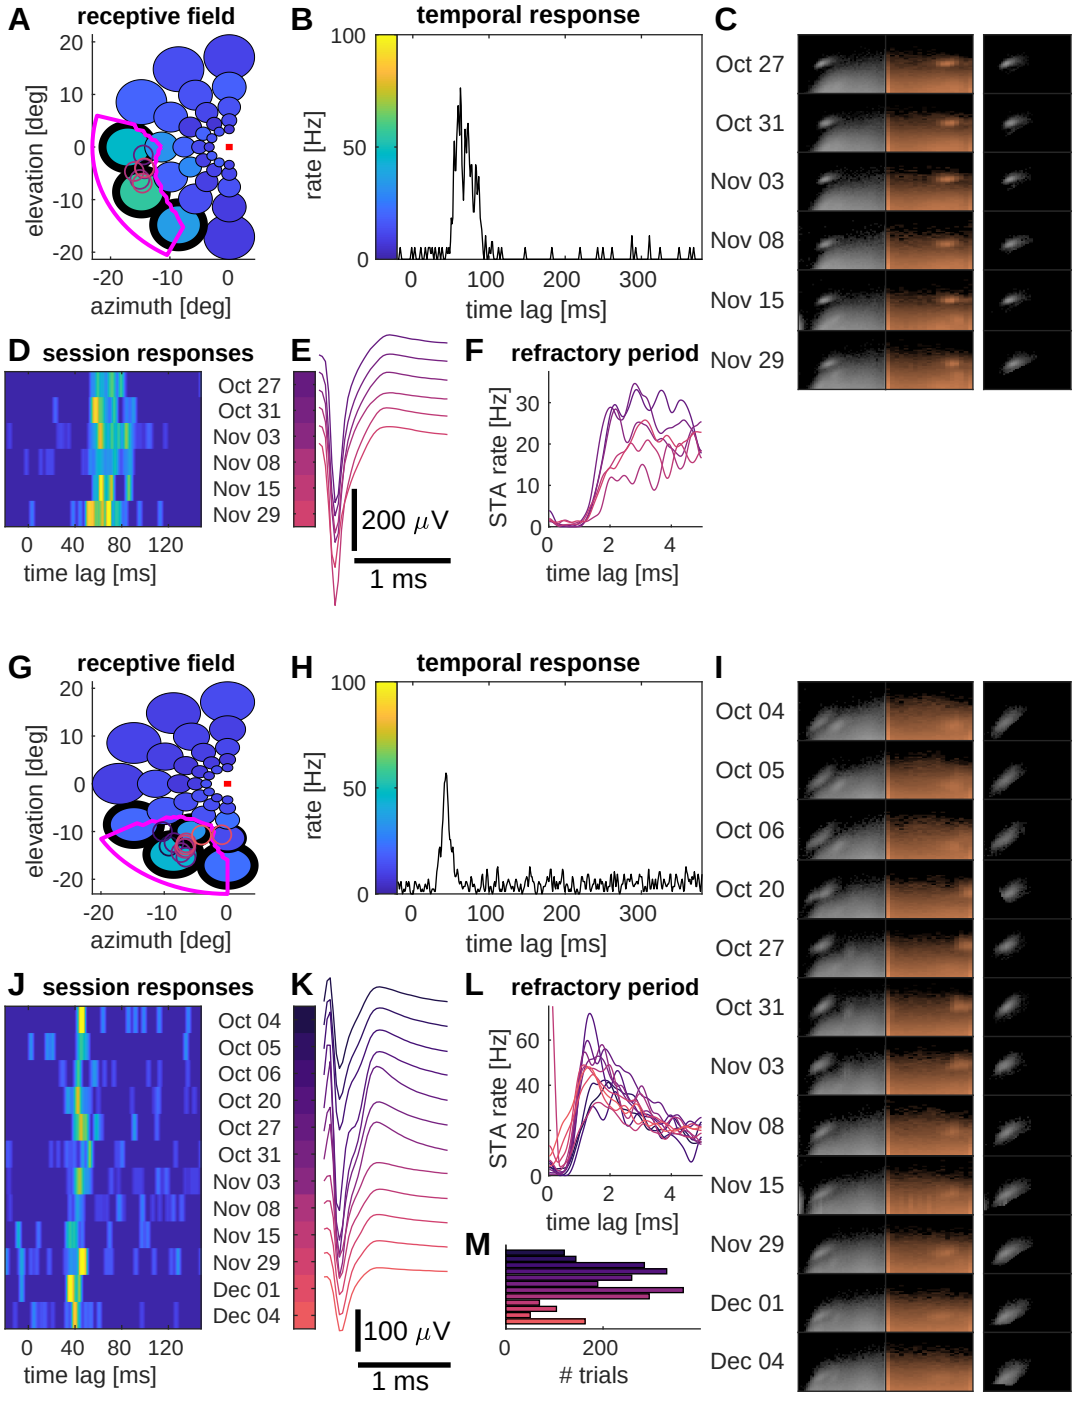

Figure 7. Examples of receptive fields of two units near area MT. (**A**) Maximum firing rates in response to presentation of a disk of moving dots (diameter scaled by 1/2 for clarity; colors indicates firing rate) at a given location in the visual field (fixation spot indicated by a red square). The receptive field (region where the interpolated firing rate exceeded a threshold; see Methods) is outlined in magenta. Colored circles represent estimates of receptive field locations for individual recording sessions. (**B**) Average firing rate for the three conditions (around the RF) outlined in black in (A). (**C**) Marginal shape histograms (as in Figure 2). (**D**) Close-up for firing rates shown in (B) for each recording session. (**E**) Averaged spike shapes. (**F**) Spike triggered averaged firing rates show a refractory period after spikes. (**G-L**) Same as (A-F) for a different unit. (**M**) Total number of trials per session. Colors indicate recording dates (sessions) and firing rates, respectively, and are matched across panels. Recordings near area MT (marmoset J).

Although a complete comparison between these types of array is beyond the scope of this proof-of-concept tool introduction, we believe it is likely that the variations in performance observed with 'Utah' arrays in macaques were larger than for the 3D arrays we used. In fact, in marmosets, arrays with similar sizes as the ones used in this study (but with fewer electrode contacts) have been reliably implanted and often measured spiking activity for months (***Debnath et al., 2018***).

We conclude this comparison by noting that we recorded from a similar number of units as reported for the larger 96 channel 'Utah' arrays (***Dickey et al., 2009***; ***Vaidya et al., 2014***; ***Fraser and Schwartz, 2011***), but from a smaller region of the brain, largely thanks to the denser 3D geometry of the arrays. This is another advantage on the hardware side of this tool, as it allows for larger scale recordings within small brain areas in the marmoset– arrays built for larger primate brains will often sparsely sample within a single area, spanning their footprint over many adjacent areas.

Long-term stability of units

The 3D array recordings had excellent long-term stability, which is a novel and important result for studies using marmosets. The feasibility of long term recordings is itself not totally unprecedented, as there are multiple approaches that align with our observations in a number of species. Here we review some examples, not just to bolster the case that long term stable recordings can be made in a number of species, but to point to the broader potential adoption of the supersession analysis approach we have introduced.

For example, ***Jackson and Fetz*** (***2007***) used microwires and studied stability of single units in continuous recordings using a window discriminator, and found single units surviving for up to 17 days in a one year experiment, where microwires were moved periodically to different neurons to improve signal quality. More systematic experiments addressing long-term stability of individual units were done with 'Utah' arrays by matching spike waveforms and inter-spike interval histograms across recording sessions (***Dickey et al., 2009***; ***Vaidya et al., 2014***), eventually in combination with correlations and firing rates (***Fraser and Schwartz, 2011***) to increase statistical power. While comprising relatively small numbers of units and recording sessions, these studies demonstrated a few single units being recorded for months, suggesting that there was likely no relative movement between the electrodes and the neural tissue. ***Linderman et al.*** (***2006***) used continuous recordings were used to study short-term changes of spike amplitudes and reported moderate amplitude fluctuations in two example units.

The N-form arrays we used had the same spacing between shanks as the 'Utah' type of array — albeit with a higher density of recording sites along a shank, and far fewer total shanks. Even though the N-form arrays comprised only 16 shanks, we found a similar long-term stability for well-isolated single units, suggesting that this number of shanks is sufficient to mitigate substantial array drift. The smaller "bed of nails" also permits a slow insertion method, which we hypothesize is important for avoiding damage associated with ballistic insertion methods, especially important in the smaller and more delicate marmoset brain.

In assessing the usefulness of supersession unit data, we used relatively relaxed criteria for unit selection. Given this liberal approach, we did not focus on comparing session-scale average spike waveforms (as these are sensitive to varying amounts of other-spike contamination and , but rather distributions of a parametric representation of spikes, where contamination could be considered as a mostly flat, additive component. Likewise, we dropped the comparison of inter-spike interval histograms, firing rates and correlations. While these can provide useful information about unit identity, they rely on a high SNR and good isolation of units in every single session and might even depend on the animal's engagement in experiments. To avoid discarding large amounts of good data without further inspection, we argue that these measures might best be used for post-hoc tests. Spike shapes themselves proved to be reasonably informative about cluster identity, and for short experimental sessions and low firing rates, multiple sessions may be required to obtain useful second order estimates.

Recent studies in rodents have been very successful in long-term tracking of neuronal activ-

ity. However, this performance was in large part made possible by increasing the density of electrode contacts, and therefore the number of observables available for spike sorting. Specifically, ***Okun et al.*** (***2016***) successfully sorted concatenated data for a small number of sessions and immobile NeuroNexus silicon probes with 4-8 tetrodes (slow insertion). Tetrode recordings in mouse (***Dhawale et al., 2017***) have been used for continuous tracking over weeks. Continuous tracking seems required here due to larger fluctuations in electrical coupling of neurons to electrodes. Recent work with high density arrays (***Chung et al., 2019***) in rats showed smaller fluctuations and allowed sorting segments of data and linking these together. Other recent high-density recording techniques using ultraflexible mesh electronics (***Fu et al., 2016***, ***2017***) and silicon high-density arrays (***Jun et al., 2017b***) have not yet been systematically studied for unit longevity. In primates, heptodes have been used in acute recordings, in marmoset cerebellum (***Sedaghat-Nejad et al., 2019***) and in macaques ***Kaneko et al.*** (***2007***), and single unit tracking was done in the latter case.

In terms of stability of units, the following general picture emerges: wires and tetrodes drift within days, but stability is better when they are left in place without an attached micromanipulator ***Okun et al.*** (***2016***) or when they are continuously tracked (***Dhawale et al., 2017***), approaches which can yield stability for days to weeks. Multiple shanks likely reduce electrode drift and units can be tracked for weeks to months ('Utah' arrays potentially for months if no degrading signal quality, ***Vaidya et al.*** (***2014***); ***Fraser and Schwartz*** (***2011***)), while ultraflexible, polymer based electrodes might remain stable even longer. Our results fit well into this picture.

Implications for experimental planning and spike sorting methods

Long-term stability offers the potential to generate detailed characterizations of neuronal behavior, but it also requires more careful experimental planning. In the two sections below, we highlight conceptual differences for experimental planning and spike sorting compared to the classical single-session approach.

Experimental Planning

While the general long-term stability and the observation of single- and multi-unit activity did support more data-rich analyses than would have been possible from a single session, the fashion in which units ended up being sampled across recordings crucially affects the planning of possible experiments. If, at one extreme, we had recorded from a different set of neurons in every recording session, we would have ended up with a large sample of recorded neurons, but not more data per unit. Such a scenario would be allowing us to estimate distributions of neuronal behavior in a given area. At the other extreme, if we were to always record from the same set of neurons, we would end up with a small sample, but would be able to measure their responses in many different conditions and further quantify the higher-order statistical interactions between them.

In reality, we found ourselves in a fruitful middle regime: Units were recorded for variable durations, in which a small fraction of units both appeared and was lost between recording sessions. This process was not entirely random, as we saw that most units disappeared during the initial sessions after their appearance. This means that the chance for a unit to survive for another session increased with the number of sessions that this neuron had already been observed. Hence, if we were to ask which of the units we would most likely observe in a future session, the best bet would be those units that were already observed for the most sessions in the past.

The variable lifetimes of units also provide an additional tool for raising the standard for isolation. Restricting an analysis to only long-lasting units would likely reduce the chance of including less clearly isolated units. Such units may not be found in some of the recordings due to variations in signal amplitude.

The exact timescales at which units were lost between sessions varied slightly across our two test arrays/animals. However, there may be two different mechanisms involved: while we found a relatively low, constant turnover of units on both arrays, in marmoset J we additionally saw a few events where a large fraction of units was lost between subsequent recordings (Figure 4). These

events could not be explained by a long temporal gap between the recordings, suggesting a relatively fast mechanism for that, with a timescale of hours to days (as opposed to weeks and months).

We believe that these findings can impact the planning of experiments using chronic arrays. In the classical single session approach, experimenters devote part of the experimental time for general characterization of receptive fields and tuning of neurons, in order to target a neuron and adapt the stimulus properties to efficiently sample responses, avoiding stimuli without an expected effect on the neuron's firing behavior. In the case of chronic array recordings, we record from many neurons with potentially different receptive fields and tuning properties, suggesting the use of more general stimuli, e.g. sampling a larger visual area and different tuning directions. Especially when studying interactions between a small number of units, one should keep in mind that some of these units may disappear during the course of an experiment and it would be advisable to start with a larger group of candidate units. In this regard, chronic arrays would be ideally suited for continuous tasks and naturalistic stimuli (e.g. ***Huk et al.*** (***2018***); ***Knöll et al.*** (***2018***)), which efficiently sample a large parameter space, allowing for simultaneous characterization of units with different tuning properties.

If, however, an experimental design requires finding persistent units in order to adapt focused studies to suit their tuning, we recommend choosing units that have already been observed for at least 3 sessions, as these units have a high chance to survive the next sessions. In our experiments, such units had a conditional (additional) lifespan of 6 and 14 sessions (for marmoset J and B, respectively, cf. Figure 6 A). Likewise, studies of changes in firing behaviour of single units across sessions (e.g. while an animal is learning a task, or after drug treatment) are in principle feasible. However, such experiments can usually not be repeated in the same animal, and few units will be clearly isolatable, resulting in a rather inefficient use of the acquired data. In this case, the suggested approach is to perform several consecutive studies on an animal, which is possible given the longevity of the arrays used here.

Importantly, we have shown that it is feasible to combine data across multiple sessions to infer tuning properties of neurons from multiple sessions. The same should be possible for interneuronal correlations. Our results also highlight that, in many cases, it would be incorrect to assume that units with similar spike shapes recorded on the same electrode in subsequent sessions would correspond to different neurons.

We conclude that chronically implanted electrode arrays allow for both sampling of a large set of neurons and detailed analysis of a few long-term units, but different timescales need to be considered when planning experiments. If the objective is to sample the population of neurons across a brain area, experimental sessions could be separated by a month to take advantage of appearance and disappearance of neurons on the array. If instead the objective is a detailed analysis of a smaller set of neurons and their interactions, daily recordings for 2-4 weeks are ideal.

Features of the spike sorting method

We adopted a modular strategy for spike sorting, where individual sessions were processed independently and could be iteratively merged to form 'supersessions'. In this way, experimenters can perform preanalyses as data are generated and determine receptive fields and tuning properties of neurons to guide stimulus selection as well as monitor recording quality. This modular approach further facilitates excluding particularly noisy segments in individual sessions, which might impair or bias the clustering algorithm.

The primary reason for eschewing existing spike sorting methods was a general concern about robustness when stationarity assumptions were not met across recording sessions. This is a known challenge to even cutting-edge algorithms (***Jun et al., 2017a***). We instead chose a simple parametric representation that was designed to be robust to noise and artifacts, which can differ from session to session. Our focus was on characterizing the peak of the depolarization phase using unimodal templates where the SNR would be highest. While spike shapes can be strongly bimodal, depending on the relative position of the electrode and neuron, the shapes for spikes with highest

amplitudes near the soma have been shown to be largely unimodal in theoretical studies (***Lindén et al., 2011***; ***Quian Quiroga, 2009***; ***Camuñas-Mesa and Quiroga, 2013***). As we recorded spikes on single electrodes and could expect a large number of neurons in the vicinity of an electrode (***Pedreira et al., 2012***), high amplitude spikes would be easiest to separate from other units. This situation would certainly be different for high-density probes. The process of estimating parameters of the spike shapes was essentially an optimization. We would shift a template temporally at sub-sampling resolution and change its width and symmetry to best match a local minimum in the raw voltage traces. In practice, this step was implemented by running the raw data through a large filter bank on a GPU.

Our spike sorting approach did not solve the problem of overlapping spikes. However, it greatly reduced the problem as the time interval needed for detection was reduced to the width of the spike and thus, due to zero padding, much smaller than the the width of the templates in the filter bank. In addition, for cases where overlapping spikes exist, we should see them in the shape histograms as somewhat isolated shapes that are a bit wider and of higher amplitude than an adjacent cluster. In our data, we did not find evidence for significant numbers of overlapping spikes near isolated clusters. Shape clusters were either nicely separated in the sense that overlapping spikes had at least half an order of magnitude lower amplitudes, or we would be unable to separate clusters in the first place, due to low amplitudes and a large number of sources, with a combined firing rate beyond 100 Hz (as in Figure 3). In the latter case, peak firing rates in single trials in response to a stimulus can be an order of magnitude higher and we necessarily detect overlapping spikes. Hence, firing rate estimates for low amplitude spikes should be read as a lower bound, providing useful (slightly distorted) information about tuning in sustained responses, while truncating transient responses.

In this work, we used the parametric representation of local mimina as a spike sorting method. But we could certainly perform spike sorting with an existing method and obtain these parametric representations for spikes in order to subsequently match spike clusters across recording sessions. Likewise, as current sorting techniques are validated with respect to stability over long time frames, it would be straightforward to replace our sorting approach. However, our sorting approach could still be used for fast, online assessments of recording quality, neuronal yield and tuning properties as it does not require manual curation.

Application to data

In many cases, we observed that shape clusters appeared and disappeared gradually over time, such that the observed spike amplitudes were highest around the middle of their lifetime. We could thus have a situation where some shape clusters of a given unit were clearly isolated single unit activity, and others were contaminated (e.g. Figure 7 I). Although this effect means that some of the unit data from 'supersessions' is less well-isolated than conventional singe-session data, the framework can also be used to estimate the impact of contamination for a given analysis, and hence to determine in a principled manner how high an isolation standard is required.

To give an example how such analysis could look like, assume that we have a number of sessions (W) where a unit was well-isolated, and some sessions (C), where the same unit was contaminated with low amplitude spikes from other neurons and some of its spikes were lost due to low amplitudes. We would then pool data from each group (W and C) of sessions to obtain a larger sample size and estimate firing rates and interspike interval histograms.

Assuming that low amplitude spikes from other neurons are uncorrelated (alternatively, the interspike interval distribution of low amplitude spikes could be estimated with sufficient data) and uniformly distributed, we would fit the ISI histograms of group C as a linear combination of the ISI histogram of group W and a uniform distribution. The component explained by the uniform distribution could then be translated into an estimate of the spike count for the low amplitude spikes from other neurons (i.e., dividing the rate of the uniform component by spike count of group C and multiply with the total recording duration of group C). To obtain an estimate of the

number of spikes missed in group C due to low spike amplitudes, one can multiply the difference in firing rates between group W and C with the total recording duration of group C and add the spike count for the low amplitude spikes determined above. After doing a given analysis separately for groups W and C, one could then compare the results and see how they are affected for a known contamination and signal loss.

Furthermore, if one looked into the datasets of group W, one would likely find spikes that are statistically similar to the contaminating spikes in group C, simply by identifying identically shaped spikes at much lower amplitudes. Therefore, it is possible to create surrogate datasets with known contamination (and, by removing spikes, signal loss) and treat them as a model to predict effects on a given analysis. The above analysis would then provide independent data to test this model.

Apart from spike clusters, our sorting approach also gives access to low amplitude spikes that do show tuned responses to visual stimulation, but likely arise from a multitude of units with a continuum of corresponding spike shapes (e.g. Figure 3). For the purpose of decoding neural activity, such low amplitude spikes can be of great value. In fact, results from other groups indicate that lowering the detection threshold increased the performance of a decoder despite losing information about the neuronal identity (***Trautmann et al., 2019***; ***Kloosterman et al., 2013***; ***Todorova et al., 2014***). Our work suggests that we can define a detection threshold (or region of interest) post-hoc, based on responsiveness to stimuli known to drive neural activity. We refer to this activity as multi-unit hash (MUH), creating a third category alongside with MUA, which should form clusters that are separable from MUH, and SUA which would additionally show a clear refractory period. We need to stress here that MUH is distinct from the 'unsorted spikes' often left behind by most sorting algorithms.

In summary, we were able to create 'supersessions' for individual units on a timescale of several days to a few weeks. This allows for more statistical power than a single session's worth of data can provide, and hence could put the awake marmoset preparation more on par with that of macaques. This is important because the marmoset is also a "pivot species" to richer and more powerful techniques that are more difficult to apply to the macaque. Such supersessions do require reconsidering the design of experiments to handle the comings-and-goings of identified units. Such experiments will likely have a long term structure in which where basic characterization of neural response properties is performed approximately once a week, with the remainder of experimental data collection being dedicated to more sophisticated experiments.

Methods and Materials

Electrophysiology preparation

Two marmosets were with implanted N-Form arrays (Modular Bionics, Berkeley, CA, USA) in area MT (marmoset J) or PPC (marmoset B). Prior to placing the chronically implanted array, we drilled a grid of 9 burr-holes over and surrounding the desired brain area based on stereotaxic coordinates from ***Paxinos et al.*** (***2012***). We performed extracellular recordings using single tungsten electrodes in each burr-hole to fine tune the placement of the array based on the physiological response. The MT array was placed based on high response to direction of motion, while the LIP array was placed based on high eye-movement related activity. A small craniotomy and duratomy were made surrounding the desired area for array placement.

The N-form array was mounted on a stereotax arm and manually lowered till tips of the shanks had entered the brain. The brain dimpled slightly, then the tissue relaxed around the implant. The array was then slowly lowered until the baseplate was just above the brain's surface. The array was stabilized and sealed with KwikCast before being closed entirely with dental cement and acrylic. The array connectors were enclosed in a custom 3D-printed box embedded in the acrylic implant.

Animal procedures described in this study were approved by the UT Austin Institutional Care and Use Committee (IACUC, Protocol AUP-2017-00170). All of the animals were handled in strict

accordance with this protocol.

The N-form arrays (Modular Bionics, Berkeley, CA, USA) consisted of a 4x4 grid of electrode shanks, spaced by $400\,\mu m$. Each shank was $1.5\,mm$ long and had 4 electrode contacts, one at its tip, and three more at $250\,\mu m$, $375\,\mu m$ and $500\,\mu m$ distance from the tip. Extracellular signals were recorded at all 64 electrode contacts with sampling rate of $30\,kHz$, using the OpenEphys recording system (***Siegle et al., 2017***). For marmoset J, seven of the electrode contacts were found damaged after the surgery and ignored for further analyses.

Visual tasks and stimuli

All stimuli were presented using custom MATLAB (Mathworks) code with the Psychophysics Toolbox (***Brainard, 1997***) and a Datapixx I/O box (Vpixx) for precise temporal registration of stimulus, behavioral, and electrophysiological events (***Eastman and Huk, 2012***).

Marmosets were trained to fixate a central dot in the presence of peripheral visual stimuli. The animals fixated the dot within a window of 1.5 degree radius for the whole trial to obtain liquid reward in the form of marshmallow juice. If the marmoset broke fixation, the trial was aborted. Fixation was acquired and held for $200\,ms$ before a stimulus appeared.

To measure MT receptive fields, we presented a circular cloud of randomly moving dots for $350\,ms$ at one of 35 different screen locations during controlled fixation. The diameter of the stimulus aperture scaled with the eccentricity of its center.

To measure direction tuning, we presented coherent motion in 12 possible directions at a fixed location based on previously measured receptive fields. Each trial contained motion in one direction for a duration of $500\,ms$.

For PPC recordings, marmosets were trained to perform a memory guided saccade task. The animals fixated the central dot while a target dot was briefly flashed at a random location in the periphery. After a delay of 400-1000 ms, the central dot was extinguished and the marmosets received liquid reward for saccades to the remembered location of the target. Memory guided saccades are well known to generate PPC activity in primates (***Andersen et al., 1990***). The task itself was not part of the investigations in this work. We outline it here as context for the behavioral engagement of the animal in the experiments and to emphasize its potential to drive neuronal activity in PPC.

On average, recording durations of individual sessions were $(26 \pm 13)\,min$ for marmoset J and $(41 \pm 12)\,min$ for marmoset B.

Pre-processing

We filtered a $60\,Hz$ component out of the raw data for each electrode using a custom made algorithm. We also performed common average referencing by subtracting (projections onto) the median of high-pass filtered signals over all electrodes from each channel. We further up-sampled data to $60\,kHz$ before feeding into Kilosort (***Pachitariu et al., 2016***). For this, values between samples were obtained by linear interpolation and values at samples were smoothed with a [1/6 2/3 1/6] smoothing kernel to obtain a uniform variance across data points for the case of Gaussian white noise.

Spike sorting

We aimed at jointly sorting spike data from tens of recording sessions (marmoset J: N=154, marmoset B: N=95) under the following constraints:

1. Marmosets were head-fixed, but able to move their bodies within the chair, creating temporally variable amounts of noise in the data.
2. Electrodes were separated by at least $\geq 125\,\mu m$ and spikes were not generally expected to be seen on multiple electrodes.
3. We observed only few separable units (0-3) per electrode.
4. There was no apparent electrode drift within recording sessions.
5. Spike clusters needed to be matched across recordings.

If spike shapes are known, then template matching would be the best way to detect spikes. However, if spikes are to be sorted, information in the raw data needs to be used to separate spike clusters, and especially to separate them from fluctuations in the background noise level and low-amplitude events of neuronal origin. A good sorting algorithm therefore needs to make estimates that are maximally invariant when subjected to noise. Potential issues are:

1. Baseline estimate: errors could change the match of bimodal templates. This may especially become a problem when the noise level is temporally varied.
2. Sampling frequency and temporal resolution for peak detection: Misaligned spikes differ in shape. This can be resolved by upsampling the data, but results in longer templates.
3. Temporally overlapping spikes: Need to be detected and fitted.

To address these three issues, we generated a bank of unimodal templates (essentially triangles with a tip rounded off by a cosine function) which varied in phase (to effectively yield 180 kHz sampling frequency), width and symmetry (see examples in Figure 8 B), covering a wide range of possible shapes. Each template was normalized to have an energy (sum of squared entries) of one. Using this bank of templates in a template matching strategy reduces baseline errors, temporal misalignment and the chance of fitting overlapping spikes, but does sacrifice some detection power (when compared to using templates generated from the data, about 10% of the signal power).

We determined local maxima (in time and width, but global in symmetry to avoid double detections) for the match (dot product) between our templates and the preprocessed voltage traces. In this setting, we were fitting the peak of the depolarization phase of a spike. While an error in the baseline estimate would have an effect on the detected spike power, it would have little effect on both the estimated spike width and symmetry. Temporally overlapping spikes were less likely as the temporal interval for detection was restricted the duration of the depolarization phase (i.e. 0.5 ms or less) and a linear combination fitting was not necessary in our recordings. Note that we did not capture the repolarization phase of a spike at all, however, we argue that due to smoothness constraints, the shape of the repolarization phase covaried with its symmetry, and its duration was hard to estimate due to potential drifts in baseline. Matching a large set of potential templates was computationally expensive, but also well suited to run on a GPU. Our implementation ran about twice as long as recording the data for 64 electrodes sampled at 30 kHz. Marginal histograms of shapes obtained for an example recording are shown in Figure 8 C.

Spike clusters appear as local density maxima in these histograms. To show that this is indeed the case, we sorted spikes with a widely used spike sorting algorithm (Kilosort, ***Pachitariu et al.*** (***2016***)). For that, we used a low threshold for splitting clusters in the Kilosort algorithm and extracted the shapes of the corresponding spikes from our template matching strategy. This allowed us to perform the manual step of merging clusters in an automated procedure, using the Jensen-Shannon divergence between shape histograms as a distance metric.

We obtained three dimensional histograms of shape parameters for spikes from each Kilosort cluster (Figure 8 D). We compared Kilosort clusters to clusters obtained by running the watershed algorithm on shape histograms and found a good match for high amplitude clusters (Figure 8 E). The latter clusters were (by construction) better localized in our histograms and we decided to use them instead of Kilosort clusters in the following analyses.

Possible extensions

We implemented the spike sorting for the case of single, isolated electrodes. An extension to dense arrays is beyond the scope of this article, but we will briefly discuss potential implementation issues here.

1. Linear arrays/stereotrodes: can be treated as another dimension, like the phase. This just requires one to set a spatial extent of spikes, creating spatially shifted templates. With this method, one could determine maxima at each time frame for each spatial shift, and do a recursive maximization in a second step to obtain spatially isolated maxima.

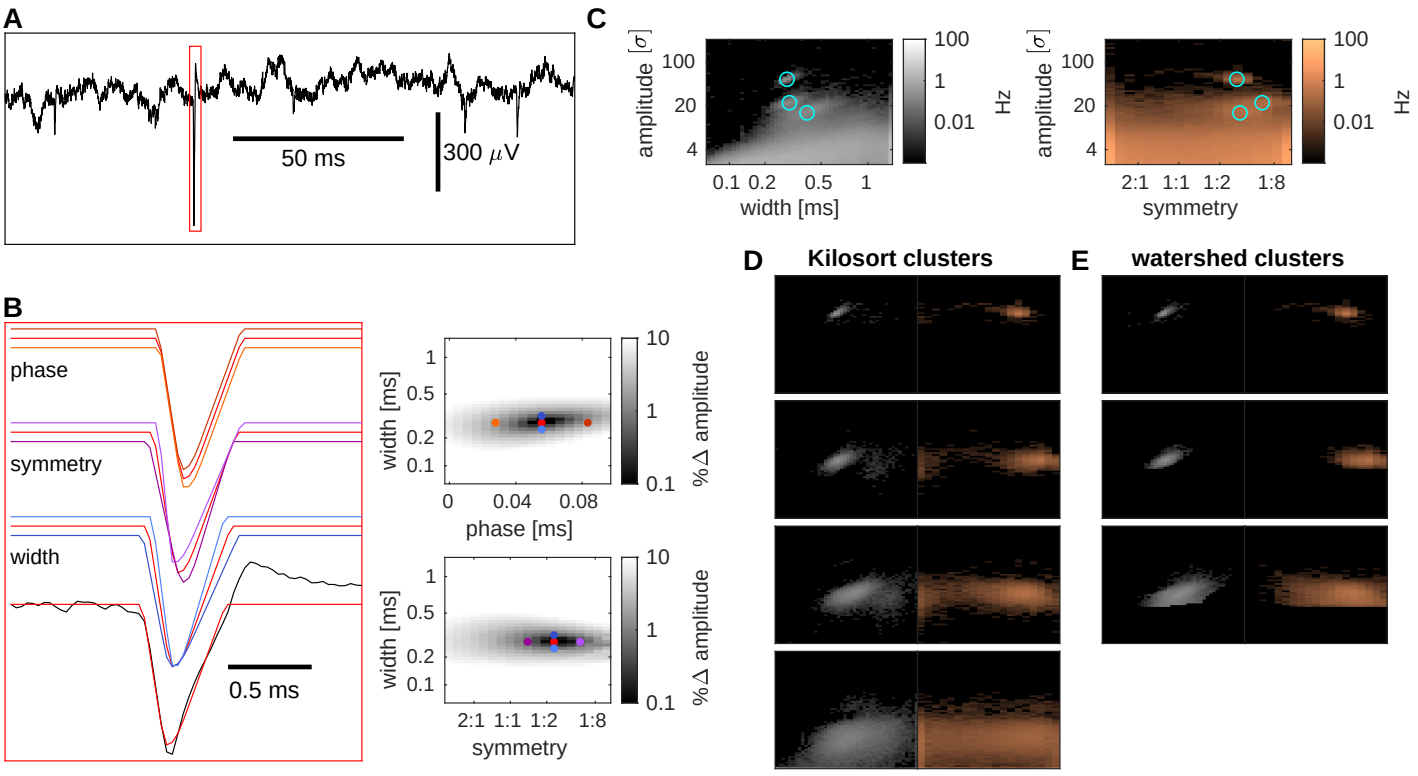

Figure 8. Spike detection and sorting. Raw voltage traces from single electrodes (**A**) are matched in a sliding window to a set of triangular, unimodal templates (examples in **B**, upper left) differing in width, symmetry and phase offset. Local maxima of template - raw trace matches in this parameter space (right plots, red dots) are then detected as putative spikes with a shape characterized by the corresponding width, symmetry and signal power (dot product of template and raw trace). (**C**) Histograms of shapes for an example electrode and recording (marginal distributions). Locations of clusters determined by a watershed algorithm are marked with cyan circles. (**D**) Shapes of events detected by Kilosort on the same electrode, grouped into clusters by an automated procedure. (**E**) Clusters determined by the watershed algorithm (corresponding to the cyan circles in (C)).

2. Spatial grids: memory constraints on the GPU will currently require chunking the array into rows of electrodes.

Our current implementation does not include a template generation and matching step, potentially resulting in suboptimal detection performance. A potential improvement, while still avoiding the baseline issue, could be to generate templates, smooth them with a kernel and generate template versions with different widths and phases by interpolation. We would need to normalize the templates to unit power and reduce positive (repolarization) parts of the templates (e.g. divide by 2), to reduce a potential baseline effect. Then we would replace the predefined templates of a given cluster (obtained from the watershed algorithm) with these templates, while keeping the other predefined templates as alternative options (for events that do not match a particular template). Next, we could rerun the detection with the modified set of templates, considering events which are best matching the inserted templates as spikes.

Cross-session merges

We computed pairwise Jensen-Shannon divergences between existing clusters from the previous 2 sessions and clusters from the current session allowing for small shifts in amplitude, width and symmetry for a penalty. Specifically, we did multiply the Jensen-Shannon divergence with the inverse of Hanning kernels with a half-width of 7 (for amplitude) and 3 (width and symmetry) bins. Each cluster from the current session was then merged with the existing cluster with the smallest Jensen-Shannon divergence if this was below a threshold of $0.3 \ln(2)$, otherwise it was labeled as a

new cluster. To allow for slow temporal drifts, the merged cluster was then assigned a shape density equal to the average of the previous and current density (resulting in effective down-weighting of earlier densities).

Motion direction tuning

Using data from all trials across sessions, we determined sensitivity maps (in spike width and amplitude space) for directional tuning for each electrode in a temporal window from 20-470 ms from stimulus onset. For that, we took spike responses with given amplitude and width, and determined the sensitivity index of direction with the maximum firing rate vs. the opposite direction. That yielded a map of sensitivity indices in dependence of spike shapes which was then thresholded, and split into connected areas exceeding the threshold. The largest of these areas was taken as a mask for the neural response of the electrode. We averaged responses for each of the 12 stimulus directions, temporally filtering with a 20 ms kernel. To see how tuning responses at a given electrode site change across sessions, we determined tuning curves for each session. Theoretically, a drift in firing rate or sensitivity could signal a change in coupling between neurons and the electrode, eventually caused by z-drift. Likewise, due to the spatial organization of area MT, a change in phase could reflect a lateral movement of the electrode.

Cluster survival

Spike shapes were very similar for a large fraction of clusters. It could be that clusters only appeared to last across sessions, but in fact represented multiple different clusters that just happened to have matching shapes. Therefore we wanted to test for a bias in longevity for units with common spike shapes. We computed histograms of amplitudes, widths, symmetry and volume of shape clusters, and the average of these quantities for each better-isolated unit across sessions. We then ranked units according to the local density of shape clusters. A lot of short-lived units had low rank, but this may be a result of detection of low-amplitude units or an increased noise level in some of the recordings. Hence, we determined percentiles (in steps of 10) of the ranks of units surviving for at least 5 sessions. For all units with ranks smaller than a given percentile, we then estimated the conditional probability that a unit was lost in the subsequent session after having survived at least until that session (N). With l_i denoting the measured lifetimes of units, and Θ the Heaviside step function, that probability estimate was

$$\hat{p_N} = 1 - \frac{\sum_i (\mathrm{l_i} - \mathrm{N} - 1)\Theta(\mathrm{l_i} - \mathrm{N} - 1)}{\sum_i (\mathrm{l_i} - \mathrm{N})\Theta(l_i - N)}. \quad (1)$$

It assumes that after the N-th session, unit losses are described by a Poisson process with a fixed rate.

Receptive fields

Firing rate responses were averaged across sessions and smoothed using a 41 ms Hanning kernel. Maximum responses were obtained for each stimulus condition and visualized. The receptive field was then determined as the region where the spatially interpolated response exceeded a threshold of twice the interquartile range above the median across conditions. Data were insufficient for estimating the size of the receptive field for individual sessions. To visualize the cross-session variation of receptive field locations, we assumed periodic boundary conditions and calculated the circular mean eccentricity and direction (colored circles in Figure 7 A, G). Temporal firing responses of individual sessions (Figure 7 D, J) were smoothed using an 18 ms Hanning kernel.

Acknowledgments

This work was supported by the US BRAIN Initiative (U01 NS094330) to ACH, the University of Texas at Austin (College of Natural Sciences Catalyst Award) to ACH, the National Institute on Drug Abuse (T32 DA018926) to AJL and HCC and the National Eye Institute (T32 EY021462) to AJL.

Competing interests

The authors declare that no competing interests exist.

References

Andersen RA, Bracewell RM, Barash S, Gnadt JW, Fogassi L. Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. Journal of Neuroscience. 1990 Apr; 10(4):1176–1196. https://www.jneurosci.org/content/10/4/1176, doi: 10.1523/JNEUROSCI.10-04-01176.1990, publisher: Society for Neuroscience Section: Articles.

Brainard DH. The Psychophysics Toolbox. Spatial Vision. 1997 Jan; 10(4):433–436. https://brill.com/view/journals/sv/10/4/article-p433_15.xml, doi: 10.1163/156856897X00357, publisher: Brill Section: Spatial Vision.

Camuñas-Mesa LA, Quiroga RQ. A Detailed and Fast Model of Extracellular Recordings. Neural Computation. 2013 Mar; 25(5):1191–1212. https://doi.org/10.1162/NECO_a_00433, doi: 10.1162/NECO_a_00433.

Chung JE, Joo HR, Fan JL, Liu DF, Barnett AH, Chen S, Geaghan-Breiner C, Karlsson MP, Karlsson M, Lee KY, Liang H, Magland JF, Pebbles JA, Tooker AC, Greengard LF, Tolosa VM, Frank LM. High-Density, Long-Lasting, and Multi-region Electrophysiological Recordings Using Polymer Electrode Arrays. Neuron. 2019 Jan; 101(1):21–31.e5. http://www.sciencedirect.com/science/article/pii/S0896627318309930, doi: 10.1016/j.neuron.2018.11.002.

Debnath S, Prins NW, Pohlmeyer E, Mylavarapu R, Geng S, Sanchez JC, Prasad A. Long-term stability of neural signals from microwire arrays implanted in common marmoset motor cortex and striatum. Biomedical Physics & Engineering Express. 2018 Aug; 4(5):055025. https://doi.org/10.1088%2F2057-1976%2Faada67, doi: 10.1088/2057-1976/aada67.

Dhawale AK, Poddar R, Wolff SB, Normand VA, Kopelowitz E, Ölveczky BP. Automated long-term recording and analysis of neural activity in behaving animals. eLife. 2017 Sep; 6:e27702. https://doi.org/10.7554/eLife.27702, doi: 10.7554/eLife.27702.

Dickey AS, Suminski A, Amit Y, Hatsopoulos NG. Single-Unit Stability Using Chronically Implanted Multielectrode Arrays. Journal of Neurophysiology. 2009 Aug; 102(2):1331–1339. https://journals.physiology.org/doi/full/10.1152/jn.90920.2008, doi: 10.1152/jn.90920.2008.

Eastman KM, Huk AC. PLDAPS: A Hardware Architecture and Software Toolbox for Neurophysiology Requiring Complex Visual Stimuli and Online Behavioral Control. Frontiers in Neuroinformatics. 2012; 6. https://www.frontiersin.org/articles/10.3389/fninf.2012.00001/full, doi: 10.3389/fninf.2012.00001.

Fraser GW, Schwartz AB. Recording from the same neurons chronically in motor cortex. Journal of Neurophysiology. 2011 Dec; 107(7):1970–1978. https://journals.physiology.org/doi/full/10.1152/jn.01012.2010, doi: 10.1152/jn.01012.2010.

Fu TM, Hong G, Viveros RD, Zhou T, Lieber CM. Highly scalable multichannel mesh electronics for stable chronic brain electrophysiology. Proceedings of the National Academy of Sciences. 2017 Nov; 114(47):E10046–E10055. https://www.pnas.org/content/114/47/E10046, doi: 10.1073/pnas.1717695114.

Fu TM, Hong G, Zhou T, Schuhmann TG, Viveros RD, Lieber CM. Stable long-term chronic brain mapping at the single-neuron level. Nature Methods. 2016 Oct; 13(10):875–882. https://www.nature.com/articles/nmeth.3969/, doi: 10.1038/nmeth.3969.

Huk A, Bonnen K, He BJ. Beyond Trial-Based Paradigms: Continuous Behavior, Ongoing Neural Activity, and Natural Stimuli. Journal of Neuroscience. 2018 Aug; 38(35):7551–7558. https://www.jneurosci.org/content/38/35/7551, doi: 10.1523/JNEUROSCI.1920-17.2018, publisher: Society for Neuroscience Section: TechSights.

Jackson A, Fetz EE. Compact Movable Microwire Array for Long-Term Chronic Unit Recording in Cerebral Cortex of Primates. Journal of Neurophysiology. 2007 Nov; 98(5):3109–3118. https://journals.physiology.org/doi/full/10.1152/jn.00569.2007, doi: 10.1152/jn.00569.2007.

Jun JJ, Mitelut C, Lai C, Gratiy SL, Anastassiou CA, Harris TD. Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction. bioRxiv. 2017 Jan; p. 101030. https://www.biorxiv.org/content/10.1101/101030v2, doi: 10.1101/101030, publisher: Cold Spring Harbor Laboratory Section: New Results.

768 **Jun JJ**, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, Lee AK, Anastassiou CA, Andrei A, Aydın
769 C, Barbic M, Blanche TJ, Bonin V, Couto J, Dutta B, Gratiy SL, Gutnisky DA, Häusser M, Karsh B, Ledochow-
770 itsch P, et al. Fully integrated silicon probes for high-density recording of neural activity. Nature. 2017 Nov;
771 551(7679):232–236. https://www.nature.com/articles/nature24636, doi: 10.1038/nature24636.

772 **Kaneko H**, Tamura H, Suzuki SS. Tracking Spike-Amplitude Changes to Improve the Quality of Multi-
773 neuronal Data Analysis. IEEE Transactions on Biomedical Engineering. 2007 Feb; 54(2):262–272. doi:
774 10.1109/TBME.2006.886934.

775 **Karumbaiah L**, Saxena T, Carlson D, Patil K, Patkar R, Gaupp EA, Betancur M, Stanley GB, Carin L, Bel-
776 lamkonda RV. Relationship between intracortical electrode design and chronic recording function. Biomate-
777 rials. 2013 Nov; 34(33):8061–8074. http://www.sciencedirect.com/science/article/pii/S0142961213007989, doi:
778 10.1016/j.biomaterials.2013.07.016.

779 **Kloosterman F**, Layton SP, Chen Z, Wilson MA. Bayesian decoding using unsorted spikes in the rat hippocam-
780 pus. Journal of Neurophysiology. 2013 Oct; 111(1):217–227. https://journals.physiology.org/doi/full/10.1152/
781 jn.01046.2012, doi: 10.1152/jn.01046.2012.

782 **Knöll J**, Pillow JW, Huk AC. Lawful tracking of visual motion in humans, macaques, and marmosets in a natu-
783 ralistic, continuous, and untrained behavioral context. Proceedings of the National Academy of Sciences of
784 the United States of America. 2018; 115(44):E10486–E10494. doi: 10.1073/pnas.1807192115.

785 **Linderman MD**, Gilja V, Santhanam G, Afshar A, Ryu S, Meng TH, Shenoy KV. Neural Recording Stability of
786 Chronic Electrode Arrays in Freely Behaving Primates. In: *2006 International Conference of the IEEE Engineering*
787 *in Medicine and Biology Society*; 2006. p. 4387–4391. doi: 10.1109/IEMBS.2006.260814, iSSN: 1557-170X.

788 **Lindén H**, Tetzlaff T, Potjans TC, Pettersen KH, Grün S, Diesmann M, Einevoll GT. Modeling the spatial reach of
789 the LFP. Neuron. 2011; 72(5):859–872. doi: 10.1016/j.neuron.2011.11.006.

790 **Nicolelis MAL**, Dimitrov D, Carmena JM, Crist R, Lehew G, Kralik JD, Wise SP. Chronic, multisite, multielec-
791 trode recordings in macaque monkeys. Proceedings of the National Academy of Sciences. 2003 Sep;
792 100(19):11041–11046. https://www.pnas.org/content/100/19/11041, doi: 10.1073/pnas.1934665100, pub-
793 lisher: National Academy of Sciences Section: Biological Sciences.

794 **Okun M**, Lak A, Carandini M, Harris KD. Long Term Recordings with Immobile Silicon Probes in the Mouse Cor-
795 tex. PLoS ONE. 2016 Mar; 11(3). https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4784879/, doi: 10.1371/jour-
796 nal.pone.0151180.

797 **Pachitariu M**, Steinmetz N, Kadir S, Carandini M, D HK. Kilosort: realtime spike-sorting for extracellular elec-
798 trophysiology with hundreds of channels. bioRxiv. 2016 Jun; p. 061481. https://www.biorxiv.org/content/10.
799 1101/061481v1, doi: 10.1101/061481.

800 **Paxinos G**, Watson C, Petrides M, Rosa M, Tokuno H. The Marmoset Brain in Stereotaxic Coordinates. El-
801 sevier Academic Press; 2012. https://espace.curtin.edu.au/handle/20.500.11937/40725, accepted: 2017-01-
802 30T14:45:05Z.

803 **Pedreira C**, Martinez J, Ison MJ, Quian Quiroga R. How many neurons can we see with current spike sorting
804 algorithms? Journal of Neuroscience Methods. 2012 Oct; 211(1):58–65. http://www.sciencedirect.com/science/
805 article/pii/S0165027012002749, doi: 10.1016/j.jneumeth.2012.07.010.

806 **Prodanov D**, Delbeke J. Mechanical and Biological Interactions of Implants with the Brain and Their Impact on
807 Implant Design. Frontiers in Neuroscience. 2016; 10. https://www.frontiersin.org/articles/10.3389/fnins.2016.
808 00011/full, doi: 10.3389/fnins.2016.00011, publisher: Frontiers.

809 **Quian Quiroga R**. What is the real shape of extracellular spikes? Journal of Neuroscience Meth-
810 ods. 2009 Feb; 177(1):194–198. http://www.sciencedirect.com/science/article/pii/S0165027008005797, doi:
811 10.1016/j.jneumeth.2008.09.033.

812 **Rennaker RL**, Street S, Ruyle AM, Sloan AM. A comparison of chronic multi-channel cortical implantation tech-
813 niques: manual versus mechanical insertion. Journal of Neuroscience Methods. 2005 Mar; 142(2):169–176.
814 http://www.sciencedirect.com/science/article/pii/S016502700400305X, doi: 10.1016/j.jneumeth.2004.08.009.

815 **Sedaghat-Nejad E**, Herzfeld DJ, Hage P, Karbasi K, Palin T, Wang X, Shadmehr R. Behavioral training of mar-
816 mosets and electrophysiological recording from the cerebellum. Journal of Neurophysiology. 2019 Aug;
817 122(4):1502–1517. https://journals.physiology.org/doi/full/10.1152/jn.00389.2019, doi: 10.1152/jn.00389.2019.

Siegle JH, López AC, Patel YA, Abramov K, Ohayon S, Voigts J. Open Ephys: an open-source, plugin-based platform for multichannel electrophysiology. Journal of Neural Engineering. 2017 Jun; 14(4):045003. https://doi.org/10.1088%2F1741-2552%2Faa5eea, doi: 10.1088/1741-2552/aa5eea.

Todorova S, Sadtler P, Batista A, Chase S, Ventura V. To sort or not to sort: the impact of spike-sorting on neural decoding performance. Journal of Neural Engineering. 2014 Aug; 11(5):056005. https://doi.org/10.1088%2F1741-2560%2F11%2F5%2F056005, doi: 10.1088/1741-2560/11/5/056005.

Trautmann EM, Stavisky SD, Lahiri S, Ames KC, Kaufman MT, O'Shea DJ, Vyas S, Sun X, Ryu SI, Ganguli S, Shenoy KV. Accurate Estimation of Neural Population Dynamics without Spike Sorting. Neuron. 2019 Jul; 103(2):292–308.e4. http://www.sciencedirect.com/science/article/pii/S0896627319304283, doi: 10.1016/j.neuron.2019.05.003.

Vaidya M, Dickey A, Best MD, Coles J, Balasubramanian K, Suminski AJ, Hatsopoulos NG. Ultra-long term stability of single units using chronically implanted multielectrode arrays. In: *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*; 2014. p. 4872–4875. doi: 10.1109/EMBC.2014.6944715, iSSN: 1558-4615.

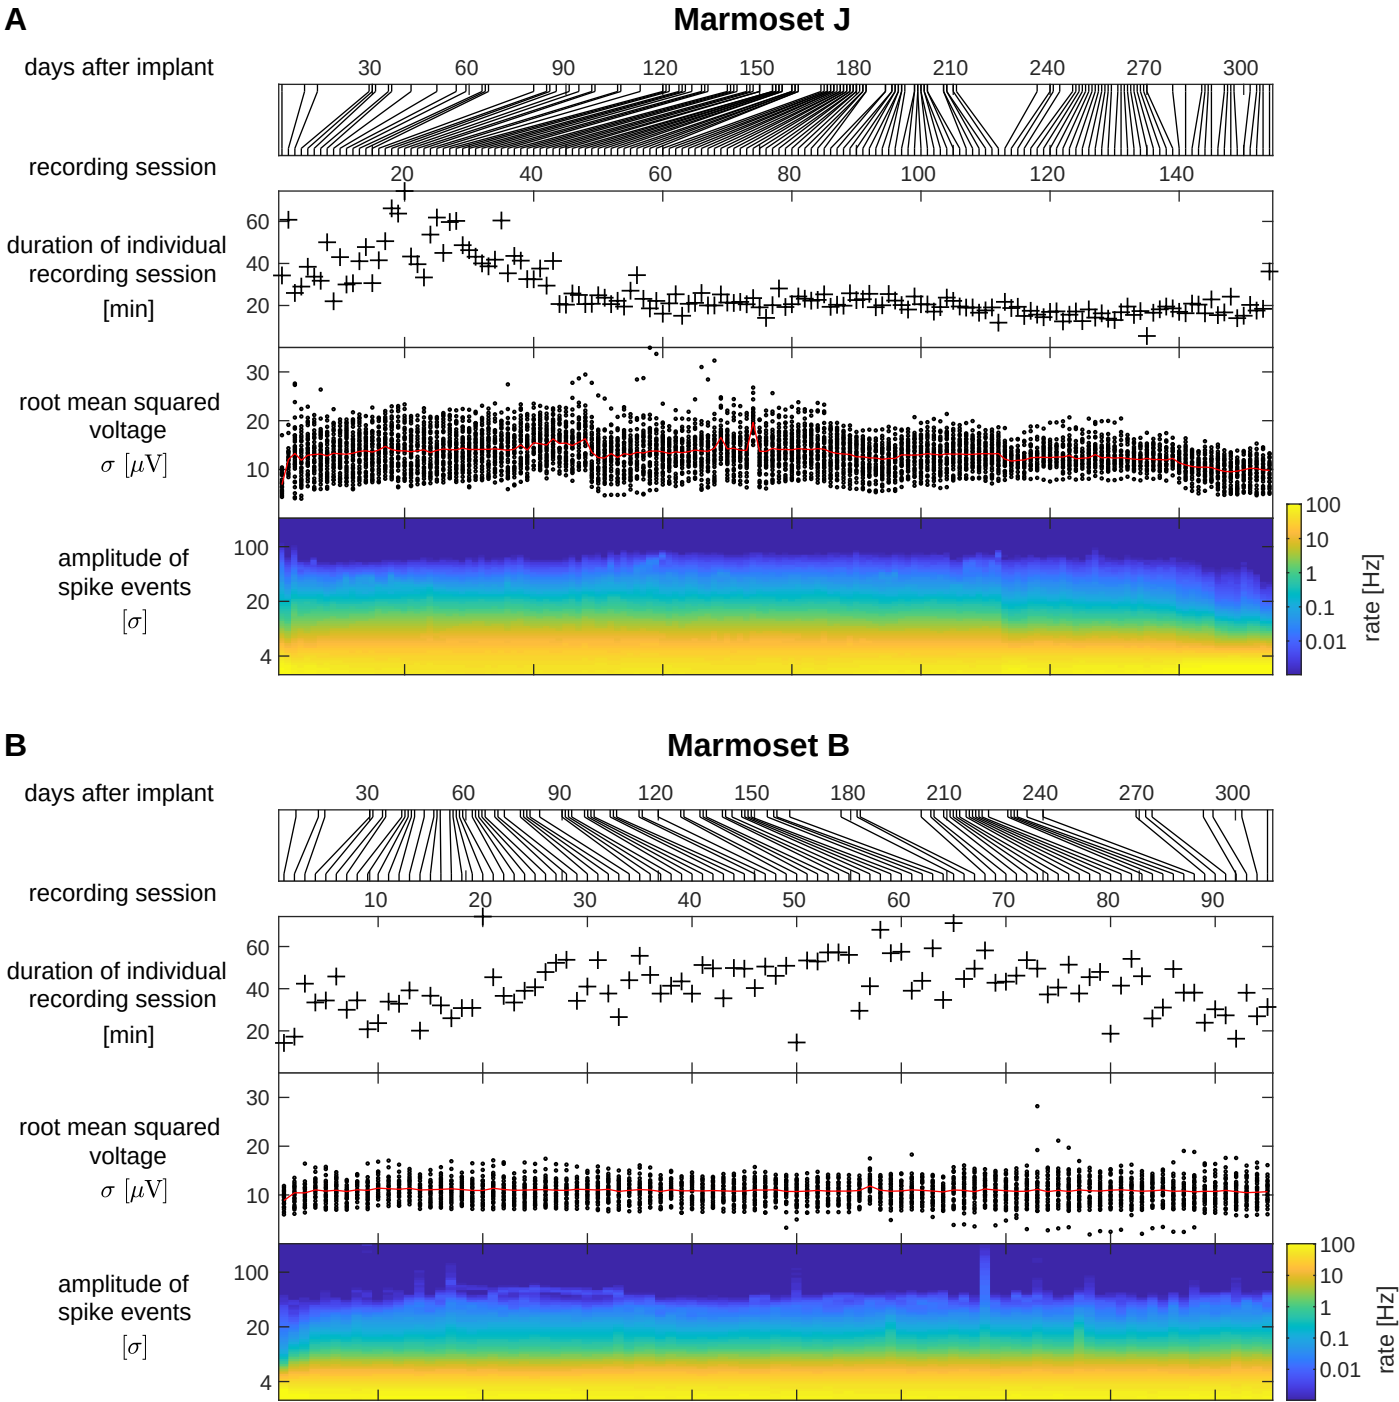
A
Marmoset J
days after implant
30 60 90 120 150 180 210 240 270 300
recording session
20 40 60 80 100 120 140
duration of individual recording session [min]
60 40 20
root mean squared voltage σ [μV]
30 20 10
amplitude of spike events [σ]
100 20 4
100 10 1 0.1 0.01
rate [Hz]
B
Marmoset B
days after implant
30 60 90 120 150 180 210 240 270 300
recording session
10 20 30 40 50 60 70 80 90
duration of individual recording session [min]
60 40 20
root mean squared voltage σ [μV]
30 20 10
amplitude of spike events [σ]
100 20 4
100 10 1 0.1 0.01
rate [Hz]

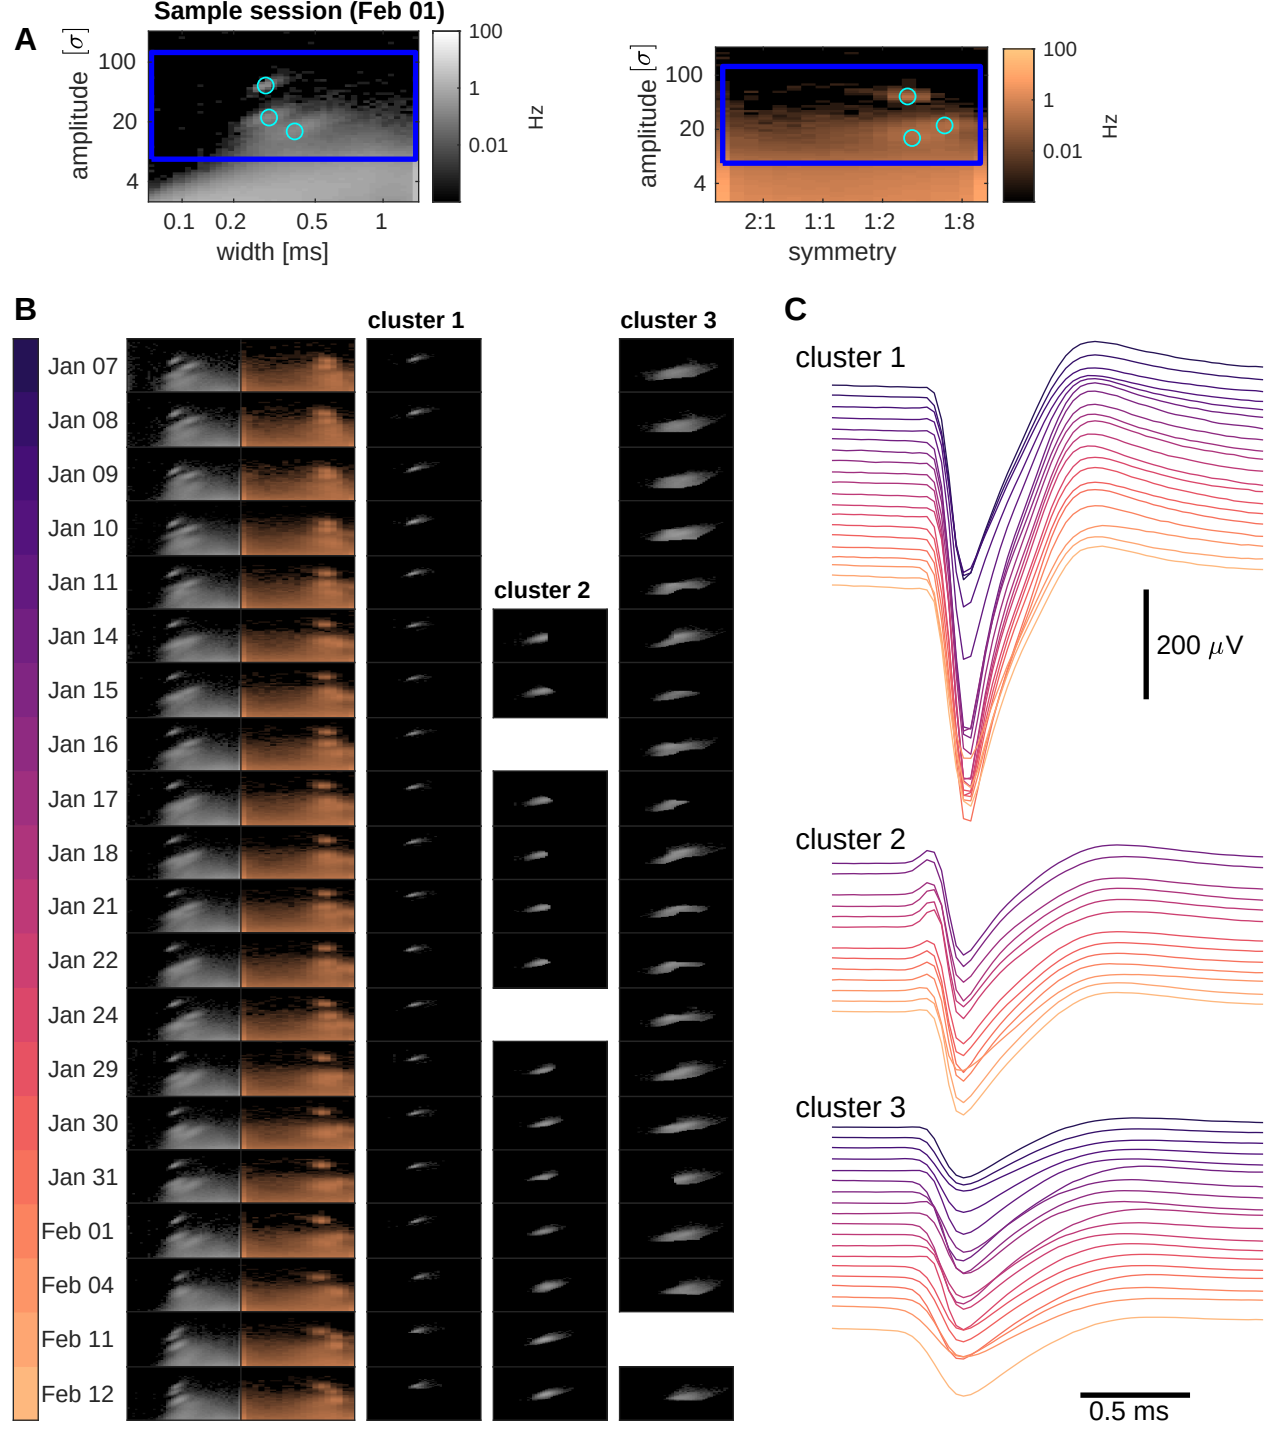
A
Sample session (Feb 01)
amplitude [σ]
100
20
4
0.1
0.2
0.5
1
width [ms]
100
1
0.01
Hz
amplitude [σ]
100
20
4
2:1
1:1
1:2
1:8
symmetry
100
1
0.01
Hz
B
cluster 1
cluster 2
cluster 3
Jan 07
Jan 08
Jan 09
Jan 10
Jan 11
Jan 14
Jan 15
Jan 16
Jan 17
Jan 18
Jan 21
Jan 22
Jan 24
Jan 29
Jan 30
Jan 31
Feb 01
Feb 04
Feb 11
Feb 12
C
cluster 1
200 μV
cluster 2
cluster 3
0.5 ms

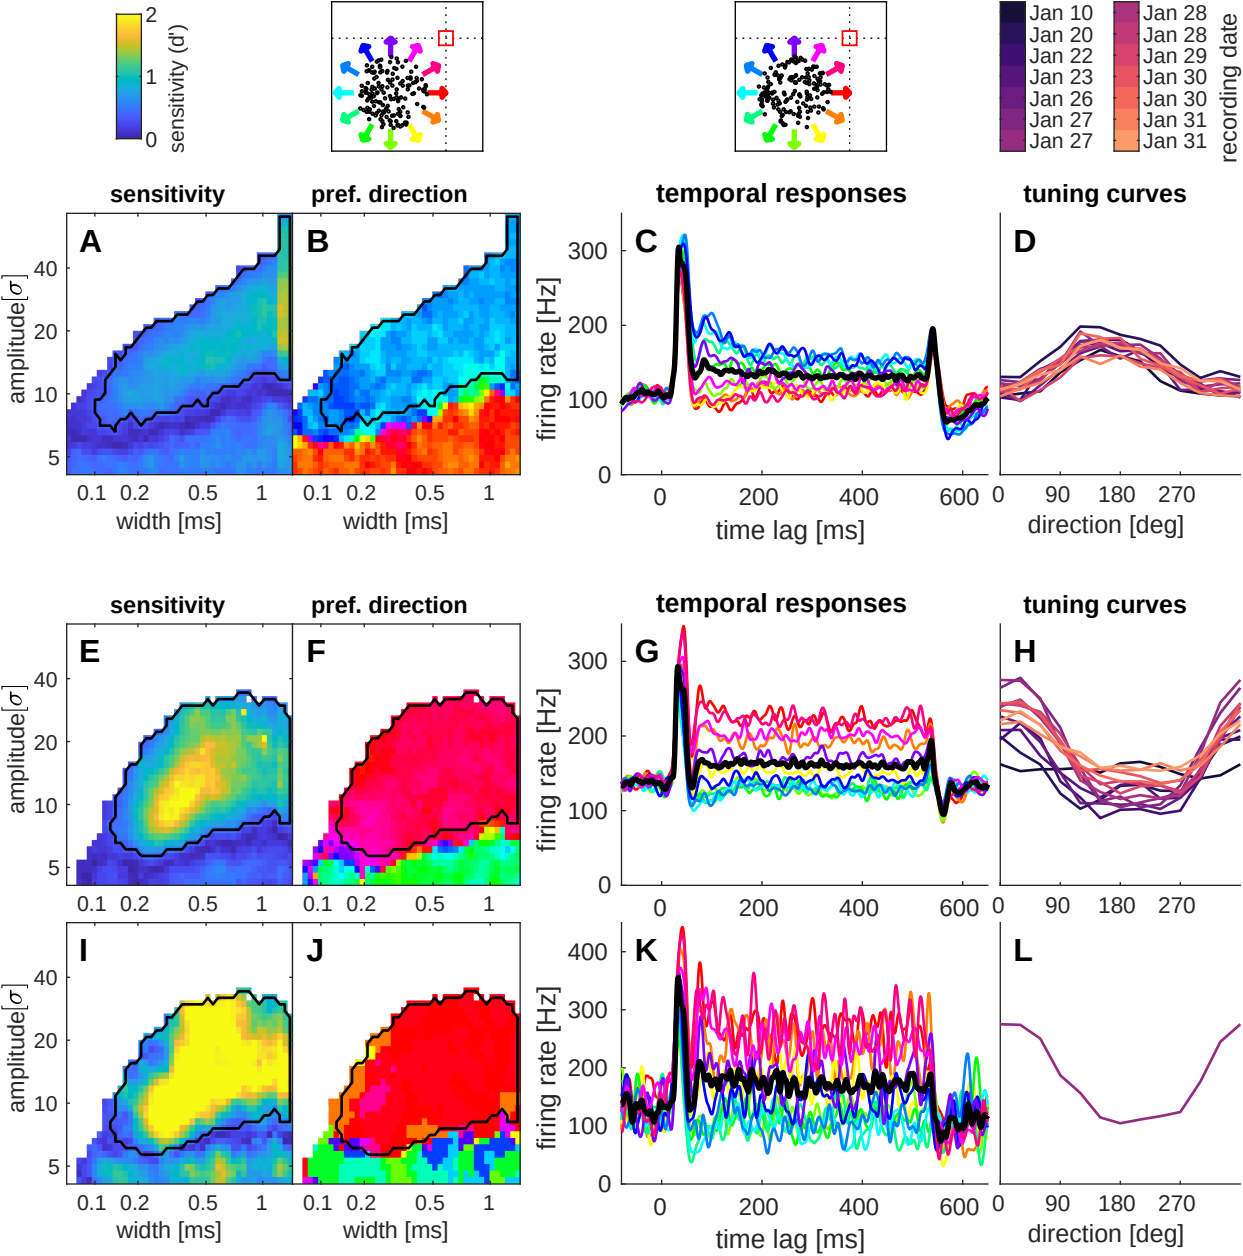
sensitivity (d')
2
1
0
Jan 10
Jan 20
Jan 22
Jan 23
Jan 26
Jan 27
Jan 27
Jan 28
Jan 28
Jan 29
Jan 30
Jan 30
Jan 31
Jan 31
recording date
sensitivity
pref. direction
temporal responses
tuning curves
A
B
C
D
E
F
G
H
I
J
K
L
amplitude[σ]
width [ms]
firing rate [Hz]
time lag [ms]
direction [deg]

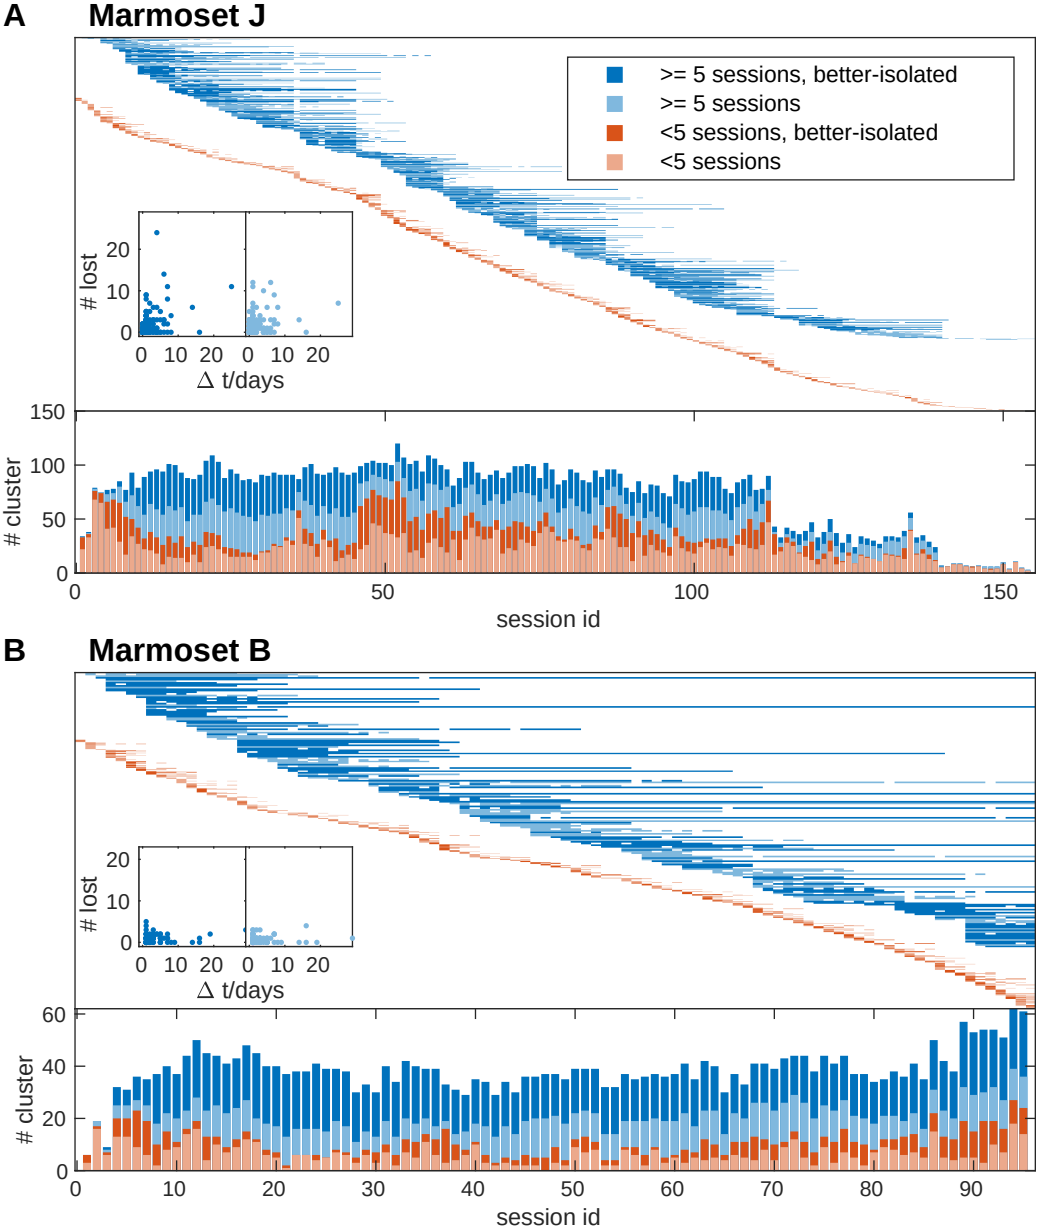
A
Marmoset J
>= 5 sessions, better-isolated
>= 5 sessions
<5 sessions, better-isolated
<5 sessions
lost
Δ t/days
cluster
session id
B
Marmoset B
lost
Δ t/days
cluster
session id

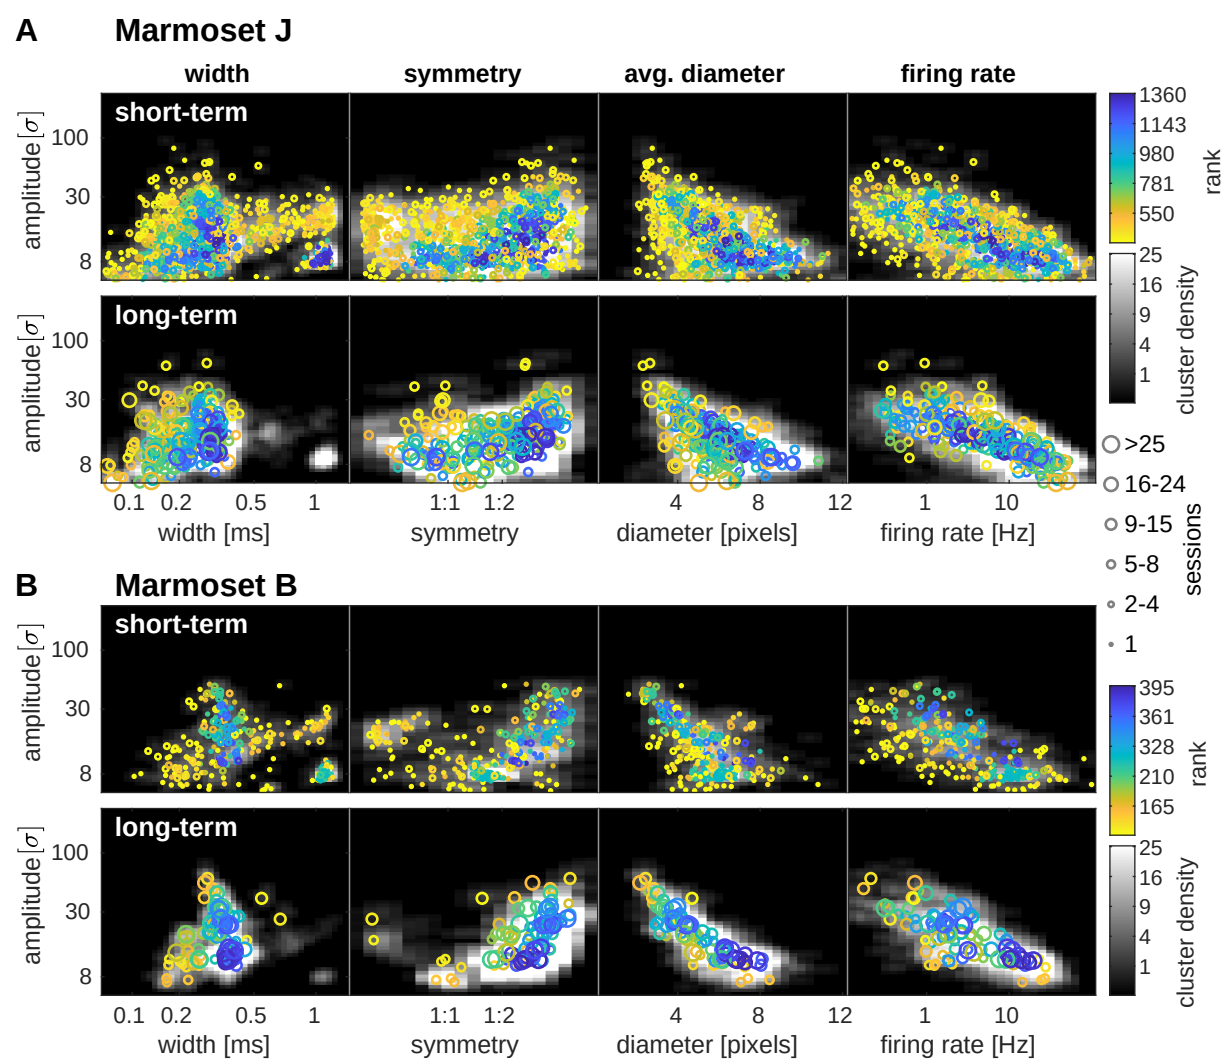
A
Marmoset J
width
symmetry
avg. diameter
firing rate
short-term
long-term
amplitude [σ]
100
30
8
0.1
0.2
0.5
1
1:1
1:2
4
8
12
1
10
width [ms]
symmetry
diameter [pixels]
firing rate [Hz]
1360
1143
980
781
550
rank
25
16
9
4
1
cluster density
>25
16-24
9-15
5-8
2-4
1
sessions
B
Marmoset B
395
361
328
210
165

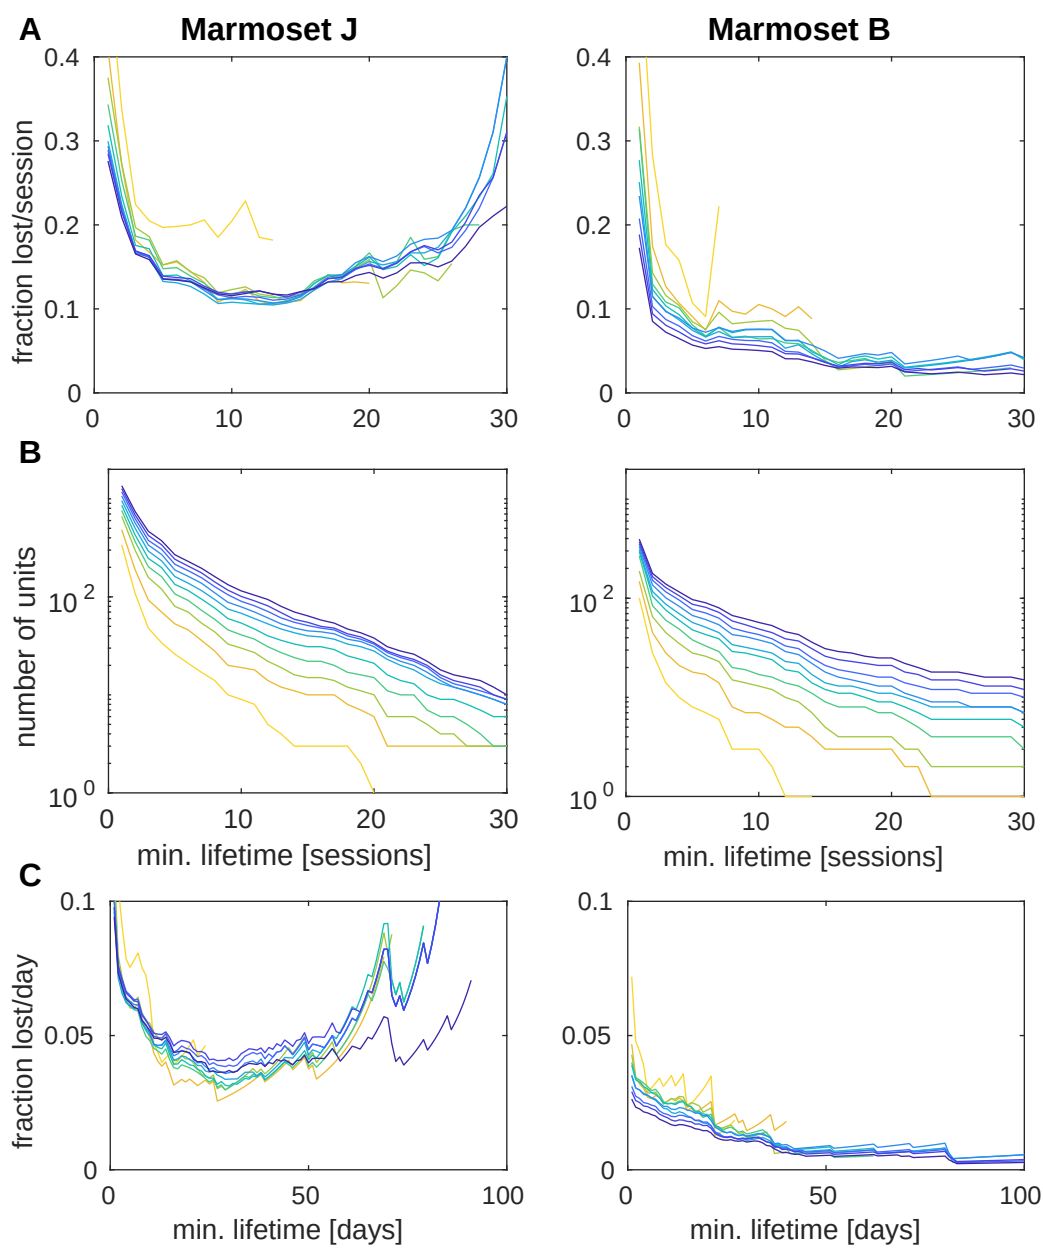
A
Marmoset J
Marmoset B
fraction lost/session
0.4
0.3
0.2
0.1
0
0
10
20
30
B
number of units
10^2
10^0
min. lifetime [sessions]
C
fraction lost/day
0.1
0.05
0
50
100
min. lifetime [days]

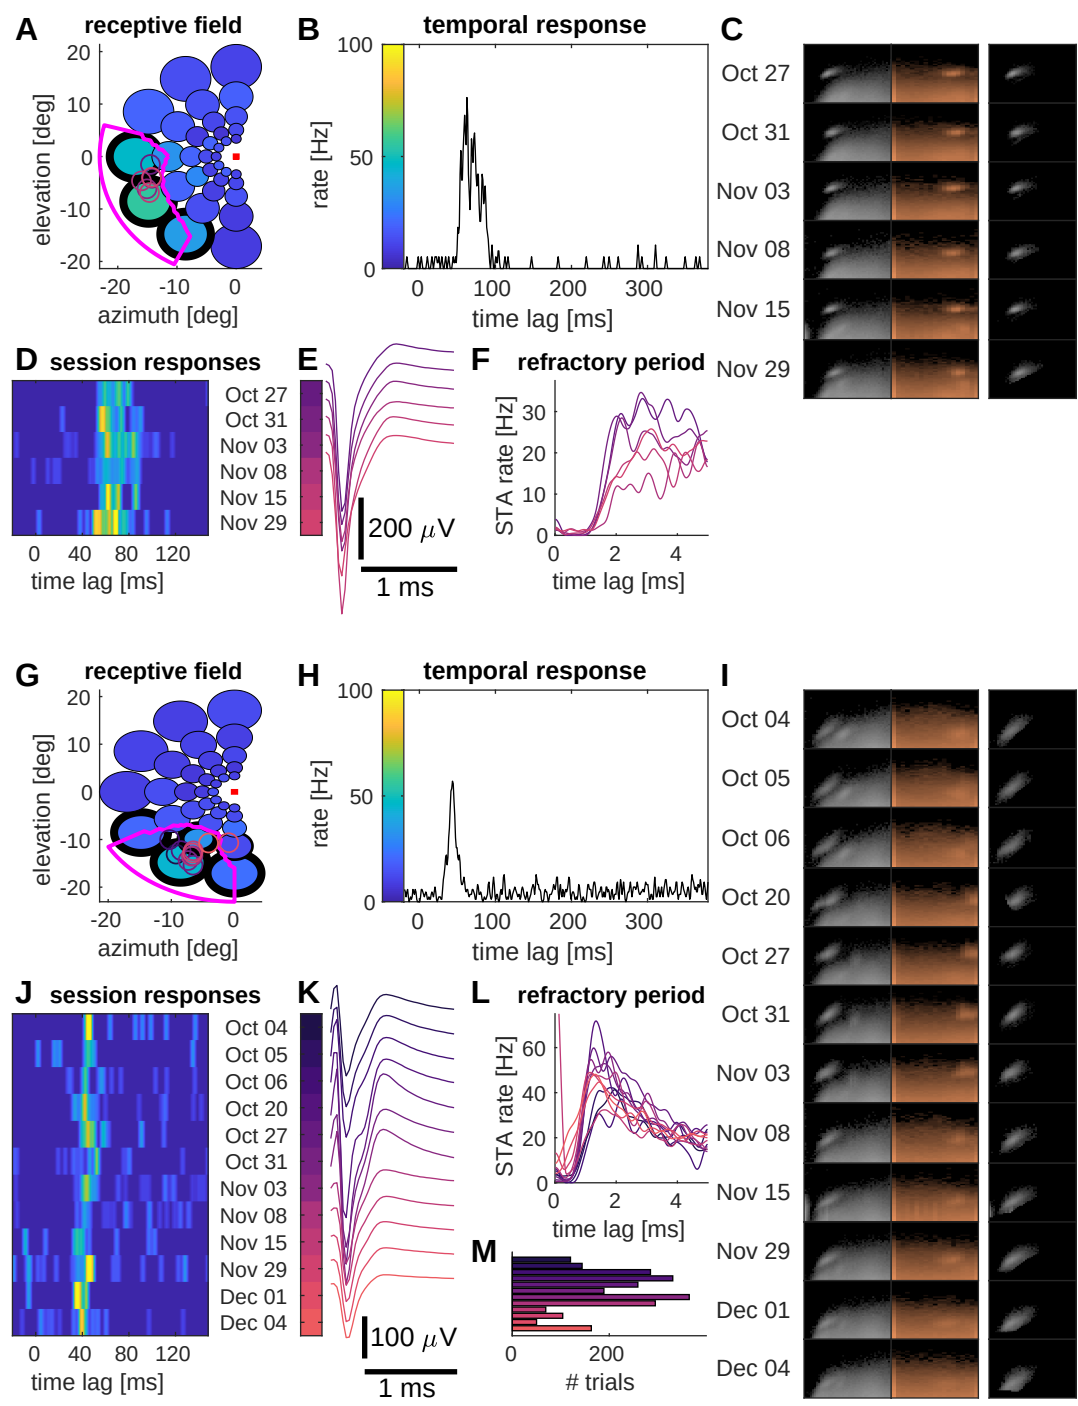
A
receptive field
elevation [deg]
azimuth [deg]
B
temporal response
rate [Hz]
time lag [ms]
C
Oct 27
Oct 31
Nov 03
Nov 08
Nov 15
Nov 29
D
session responses
time lag [ms]
E
200 μV
1 ms
F
refractory period
STA rate [Hz]
time lag [ms]
G
receptive field
elevation [deg]
azimuth [deg]
H
temporal response
rate [Hz]
time lag [ms]
I
Oct 04
Oct 05
Oct 06
Oct 20
Oct 27
Oct 31
Nov 03
Nov 08
Nov 15
Nov 29
Dec 01
Dec 04
J
session responses
time lag [ms]
K
100 μV
1 ms
L
refractory period
STA rate [Hz]
time lag [ms]
M
trials

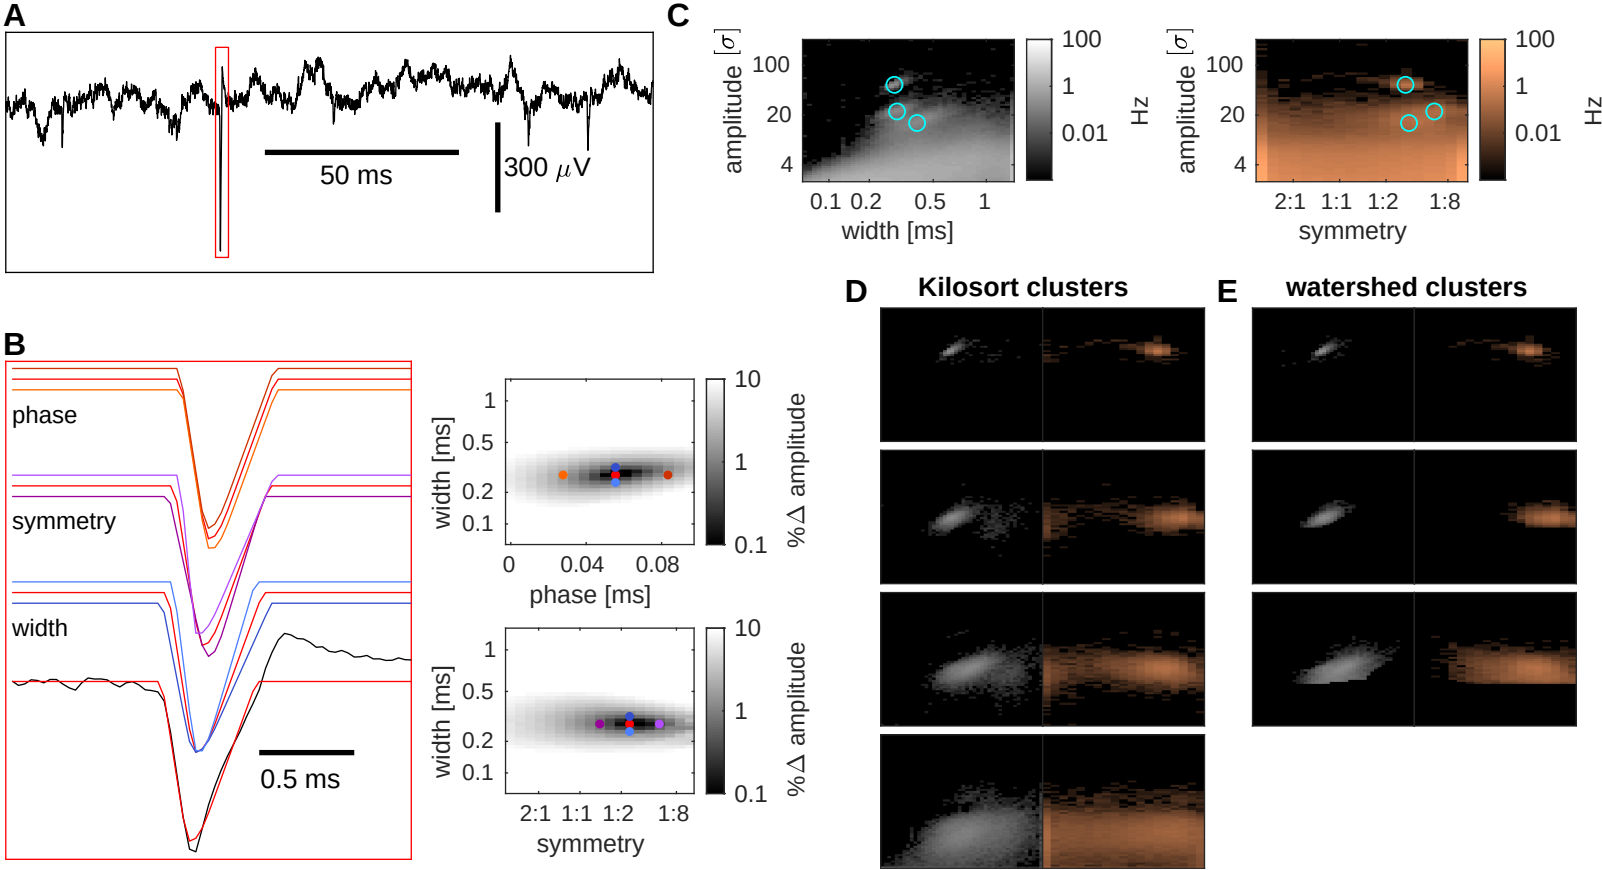
A
50 ms
300 μV
B
phase
symmetry
width
0.5 ms
width [ms]
phase [ms]
%Δ amplitude
symmetry
C
amplitude [σ]
width [ms]
Hz
symmetry
D
Kilosort clusters
E
watershed clusters