

Non-linear Dimensionality Reduction on Extracellular Waveforms Reveals Physiological, Functional, and Laminar Diversity in Premotor Cortex

Eric Kenji Lee^a, Hymavathy Balasubramanian^e, Alexandra Tsolias^b, Stephanie Anakwe^f, Maria Medalla^b, Krishna V. Shenoy^{g,i,h,j,k,l}, Chandramouli Chandrasekaran^{a,b,c,d,*}

^a*Department of Psychological and Brain Sciences, Boston University, Boston, MA, USA*

^b*Department of Anatomy & Neurobiology, Boston University School of Medicine, Boston, MA, USA*

^c*Center for Systems Neuroscience, Boston University, Boston, MA, USA*

^d*Department of Biomedical Engineering, Boston University, Boston, MA, USA*

^e*Bernstein Center for Computational Neuroscience, Technische Universität, Berlin, Germany*

^f*Undergraduate Program in Neuroscience, Boston University, Boston, MA, USA*

^g*Department of Electrical Engineering, Stanford University, CA, USA*

^h*Department of Neurobiology, Stanford University, CA, USA*

ⁱ*Department of Bioengineering, Stanford University, CA, USA*

^j*Wu Tsai Neurosciences Institute, Stanford University, CA, USA*

^k*Bio-X Institute, Stanford University, CA, USA*

^l*Howard Hughes Medical Institute, Stanford University, CA, USA*

Abstract

Cortical circuits involved in decision-making are thought to contain a large number of cell types—each with different physiological, functional, and laminar distribution properties—that coordinate to produce behavior. Current *in vivo* methods rely on clustering of specified features, such as trough to peak duration of extracellular spikes, to identify putative cell types, but these can only capture a small amount of variation. Here, we develop a new method (*WaveMAP*) that combines non-linear dimensionality reduction with graph clustering to identify putative cell types. We apply *WaveMAP* to extracellular waveforms recorded from dorsal premotor cortex of macaque monkeys performing a decision-making task. Using *WaveMAP*, we robustly establish eight waveform clusters and show that these clusters recapitulate previously identified narrow- and broad-spiking types while also revealing undocumented diversity within these sub types. The eight clusters exhibited distinct laminar distributions, characteristic firing rate patterns, and decision-related dynamics.

*Corresponding Author

Email address: cchandr1@bu.edu (Chandramouli Chandrasekaran)

By revealing additional cell type diversity, *WaveMAP* facilitates a more nuanced understanding of how the dynamics of cell types unfolds across cortical layers during decision-making.

Keywords:

dimensionality reduction, cell types, extracellular waveform, unsupervised clustering

1 **Significance (120 words)**

2 How different cell types sculpt activity patterns in brain areas associated with decision-making
3 is a fundamentally unresolved problem in neuroscience. In monkeys, and other species where
4 transgenic access is not yet possible, identifying physiological types *in vivo* relies on only a few
5 discrete user-specified features of extracellular waveforms to identify cell types. Here, we show
6 that non-linear dimensionality reduction with graph clustering applied to the entire extracellular
7 waveform can delineate many different putative cell types and does so in an interpretable manner.
8 We show that this method reveals previously undocumented physiological, functional, and laminar
9 diversity in the dorsal premotor cortex of monkeys, a key brain area implicated in decision-making.

10 **Introduction**

11 The processes involved in decision-making, such as deliberation on sensory evidence and the
12 preparation and execution of motor actions, are thought to emerge from the coordinated dynamics
13 within and between cortical layers^{1,2}, cell types³⁻⁶ and brain areas^{7,8}. A large body of research has
14 described differences in decision-related dynamics across brain areas⁹⁻¹² and a smaller set of studies
15 has provided insight into layer-dependent dynamics during decision-making^{1,2,13,14}. However, we
16 currently do not understand how decision-related dynamics emerge across putative cell types.
17 Here, we address this open question by first developing a new method, *WaveMAP*, that combines
18 non-linear dimensionality reduction and graph-based clustering. We apply this technique to
19 extracellular waveforms to identify putative cell classes and examine their physiological, functional,
20 and laminar distribution properties.

21 In mice, and to some extent in rats, transgenic tools allow the *in vivo* detection of particular cell
22 types^{3,5} while *in vivo* studies in primates are largely restricted to using features of the extracellular
23 action potential (EAP) such as shape and cell firing rate (FR). Early *in vivo* monkey work¹⁵
24 introduced the importance of EAP features—such as spike duration and action potential (AP)
25 width—in identifying cell types. These experiments introduced the concept of broad- and narrow-
26 spiking neurons. Later experiments in the guinea pig¹⁶ and cat¹⁷ then established the idea that
27 these broad- and narrow-spiking waveform shapes mostly correspond to excitatory pyramidal
28 and inhibitory cells respectively and this has been assumed to also be true in primates. This
29 method of identifying cell types in mammalian cortex *in vivo* is widely used in neuroscience but
30 it is insufficient to capture the known structural and transcriptomic diversity of cell types in the
31 monkey and the mouse^{18,19}. Furthermore, recent observations in the monkey defy this simple
32 classification of broad- and narrow-spiking cells corresponding to excitatory and inhibitory cells
33 respectively. Two such examples in the primate that have resisted this idea are narrow-spiking
34 pyramidal tract neurons in deep layers of M1 (Betz cells)^{20,21} and narrow-spiking excitatory cells
35 in layer III of V1, V2 and MT²²⁻²⁵.

36 To capture a more representative diversity of cell types *in vivo*, more recent studies have in-
37 corporated additional features of EAPs (beyond AP width) such as trough to peak duration²⁶,
38 repolarization time^{27,28}, and triphasic waveform shape^{29,30}. Although these user-specified methods
39 are amenable to human intuition, they often inadequately distinguish between previously identified
40 cell types^{20,31,32}. It is also unclear how to choose these user-specified features in a principled
41 manner (i.e., one that maximizes explanatory power) as they are often highly correlated with one
42 another. This results in different studies choosing between different sets of specified features each
43 yielding different inferred cell classes^{27,33–35}—thus it is difficult to compare putative cell types
44 across literature. Some studies even conclude that there is no single set of specified features that is
45 a reliable differentiator of type^{32,36}. These issues led us to investigate techniques that don't require
46 feature specification but are designed to find patterns in complex datasets through non-linear
47 dimensionality reduction. Such methods have seen usage in diverse neuroscientific contexts such
48 as single-cell transcriptomics^{37,38}, in analyzing models of biological neural networks^{39,40}, the
49 identification of behavior^{41–43}, and in electrophysiology^{44–47}.

50 Here, in a novel technique that we term *WaveMAP*, we combine a non-linear dimensionality
51 reduction method (Universal Manifold Approximation and Projection [UMAP]⁴⁸) with graph
52 community detection (Louvain clustering⁴⁹) to understand the physiological, functional, and
53 laminar distribution diversity of cell types during decision-making. We applied *WaveMAP* to
54 extracellular waveforms collected from neurons in macaque dorsal premotor cortex (PMd) in a
55 decision-making task using laminar multi-channel probes (16 electrode "U probes"). We found
56 that *WaveMAP* significantly outperformed current approaches without need for user-specification
57 of waveform features like trough to peak duration. This data-driven approach exposed more
58 diversity in extracellular waveform shape than any single constructed spike feature in isolation or
59 in combination. Using interpretable machine learning, we also show that *WaveMAP* picks up on
60 nuanced and meaningful biological variability in waveform shape.

61 *WaveMAP* revealed three broad-spiking and five narrow-spiking waveform types that differed in
62 shape, physiological, functional, and laminar distribution properties. Although most narrow-spiking
63 cells had the high maximum firing rates typically associated with inhibitory neurons, some had
64 firing rates similar to broad-spiking neurons which are typically considered to be excitatory. The
65 time at which choice selectivity ("discrimination time") emerged for many narrow-spiking cell
66 classes was earlier than broad-spiking neuron classes—except for the narrow-spiking cells that had
67 broad-spiking like maximum firing rates. Finally, many clusters had distinct laminar distributions
68 that appear layer-dependent in a manner matching certain anatomical cell types. This clustering
69 explains variability in discrimination time over and above previously reported laminar differences¹.
70 Together, this constellation of results reveals previously undocumented relationships between
71 waveform shape, physiological, functional, and laminar distribution properties that are missed by
72 traditional approaches. Our results provide powerful new insights into how candidate cell classes
73 can be better identified and how these types coordinate with specific timing, across layers, to
74 shape decision-related dynamics.

75

76 Results

77 *Task and Behavior*

78 Two male rhesus macaques (T and O) were trained to perform a red-green reaction time decision-
79 making task (Fig. 1A). The task was to discriminate the dominant color of a central static
80 red-green checkerboard cue and to report their decision with an arm movement towards one of
81 two targets (red or green) on the left or right (Fig. 1A).

82 The timeline of the task is as follows: a trial began when the monkey touched the center target
83 and fixated on a cross above it. After a short period, two targets red and green appeared on
84 the either side of the center target (see Fig. 1B, top). The target configuration was randomized:
85 sometimes the left target was red and the right target was green or vice versa. Another short
86 viewing period later, a red-green checkerboard appeared in the center of the screen with a variable
87 mixture of red and green squares. We parameterized the variability of the checkerboard by its
88 color coherence (C). Larger coherence indicates that there is more of one color than the other (an
89 easy trial) versus a low coherence which indicates that the two colors are more equal in number
90 (a difficult trial). In signed color coherence (SC), large positive numbers indicate the presence of
91 more red squares while large negative numbers indicate more green squares. Numbers close to
92 zero (positive or negative) indicate an almost even number of either. The signed color coherence
93 provides an estimate of whether there are more red or green squares in the checkerboard whereas
94 coherence describes the difficulty of the stimulus (Fig. 1B, bottom).

95 Our monkeys demonstrated the range of behaviors typically observed in decision-making tasks:
96 monkeys made more errors and were slower for more ambiguous checkerboards compared to less
97 ambiguous completely red or green checkerboards (Fig. 1C, D). We used coherence and reaction
98 times (RT) to analyze the structure of decision-related neural activity.

99 *Recordings and Single Neuron Identification*

100 While monkeys performed this task, we recorded single neurons from the caudal aspect of dorsal
101 premotor cortex (PMd; Fig. 1E, top) using single tungsten or linear multi-contact electrodes (Plexon
102 U-Probes, 625 neurons, 490 U-probe waveforms; Fig. 1E, right) and a BlackRock Microsystems
103 Acquisition System (Cerebus). In this study, we analyzed the average extracellular action potential
104 (EAP) waveforms of these neurons. All waveforms were analyzed after being filtered by a 4th
105 order high-pass Butterworth filter (250 Hz). A 1.6 ms snippet of the waveform was recorded for
106 each spike and used in these analyses, which is a duration longer than many studies of waveform
107 shape. We restricted our analysis to well-isolated single neurons. Extracellular waveforms were
108 identified as single neurons by only accepting waveforms with minimal ISI violations ($1.5\% < 1.5$
109 ms) and also by careful online inspection of the extracellular waveforms and subsequent offline
110 spike sorting (Plexon Offline Sorter). Every effort was made online to ensure thresholds were
111 conservative so that the best quality spikes were included in the analysis.

112 We used previously reported approaches to align, average, and normalize spikes^{32,50}. Spikes were
113 aligned in time via their depolarization trough and normalized between -1 and 1 in normalized
114 amplitude units. Other common pre-processing methods, such as trough alignment and normaliza-
115 tion to trough depth, showed no differences in downstream analyses (Fig. S5B). “Positive spiking”
116 units with large positive amplitude pre-hyperpolarization spikes were dropped from the analysis

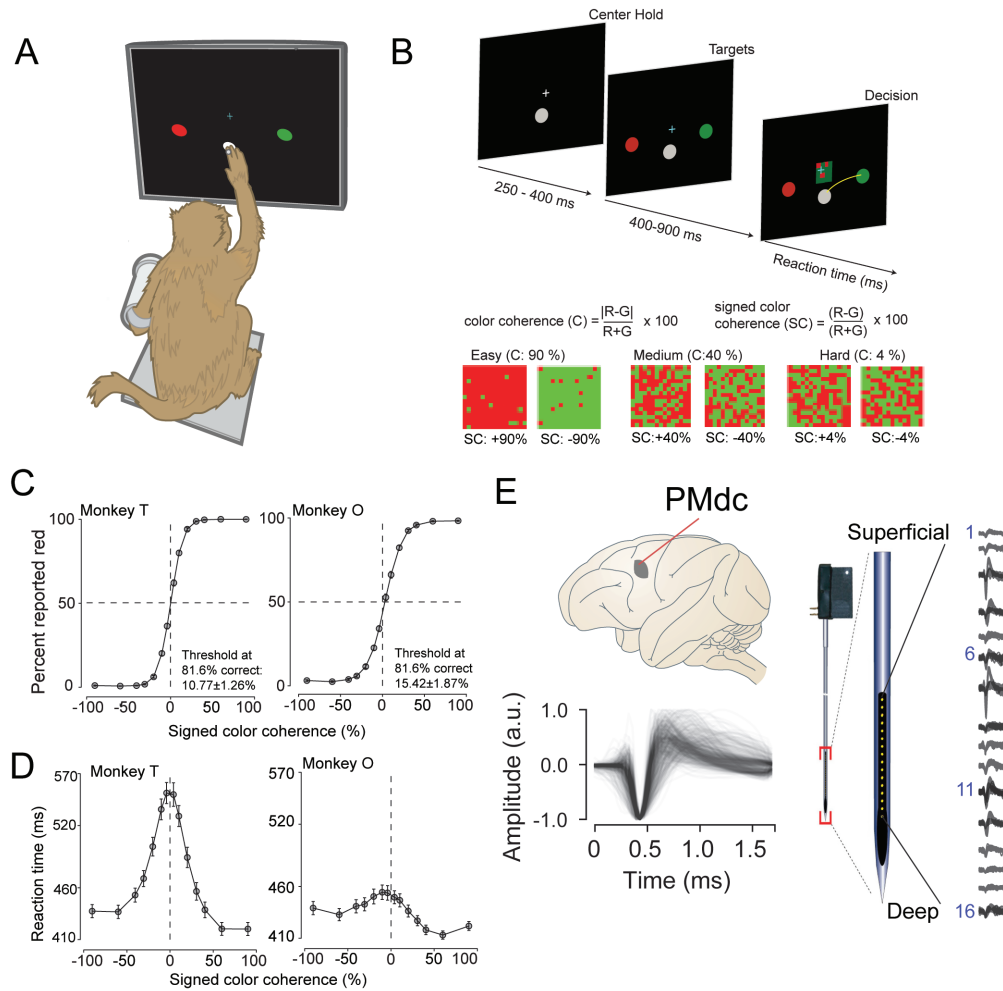


Figure 1: Recording locations, waveform shapes, techniques, task, and discrimination behavior (A) An illustration of the behavioral setup in the discrimination task. The monkey was seated with one arm free and one arm gently restrained in a plastic tube via a cloth sling. An infrared-reflecting (IR) bead was taped to the forefinger of the free hand and was used in tracking arm movements. This gave us an instantaneous readout of the hand's position and allowed us to mimic a touch screen. **(B)** A timeline of the decision-making task (top). At bottom is defined the parametrization of difficulty in the task in terms of color coherence and signed color coherence. **(C)** Average discrimination performance and **(D)** RT over sessions of the two monkeys as a function of the signed color coherence of the checkerboard cue. RT plotted here includes both correct and incorrect trials for each session and then averaged across sessions. Gray markers show measured data points along with $2 \times$ S.E.M. estimated over sessions. For many data points in (C), the error bars lie within the marker. X-axes in both (C), (D) depict the SC in %. Y-axes depict the percent responded red in (C) and RT in (D). Also shown in the inset of (C) are discrimination thresholds (mean \pm S.D. over sessions) estimated from a Weibull fit to the overall percent correct as a function of coherence. The discrimination threshold is the color coherence (CC) at which the monkey made 81.6% correct choices. 75 sessions for monkey T (128,989 trials) and 66 sessions for monkey O (108,344 trials) went into the averages. **(E)** The recording location in caudal PMd (top); normalized and aligned isolated single-unit waveforms (n = 625, 1.6 ms each, bottom); and schematic of the 16-ch Plexon U-probe (right) used during the behavioral experiment.

117 due to their association with dendrites and axons^{29,35,51}. Recordings were pooled across monkeys
 118 to increase statistical power for *WaveMAP*.

119 *Non-linear Dimensionality Reduction with Graph Clustering Reveals Robust Low-Dimensional*
120 *Structure in Extracellular Waveform Shape*

121 In *WaveMAP*, we use a three step strategy for the analysis of extracellular waveforms: We first
122 parsed intrinsic waveform structure using a non-linear dimensionality reduction method called
123 universal manifold approximation and projection (UMAP)⁴⁸. Second, we used the high-dimensional
124 graph (a fuzzy simplicial complex) found by UMAP (the approximation step, the "A" in UMAP)
125 and passed it through a graph-based community detection method, Louvain clustering,⁴⁹ to
126 delineate high-dimensional clusters. Third, we then used UMAP to project the high-dimensional
127 data into two dimensions (projection step, the "P" in UMAP). We colored the data points in this
128 projected space according to their Louvain cluster membership found in step two.

129 The time series for each normalized and trough-aligned waveform (see Fig. 1E, bottom) was
130 passed directly into the UMAP algorithm without any prior dimensionality reduction (except
131 for normalization) and without any feature specification. To identify the presence of clusters
132 in the high-dimensional graph (produced by UMAP), we applied Louvain clustering which finds
133 communities of highly inter-connected data points⁴⁹. These communities appear as color-coded
134 clusters in the UMAP space after projection (Fig. 2A). Louvain clustering was used because it
135 operates directly in high dimensions (Fig. S1). This algorithm is also fairly deterministic (after
136 fixing a random seed) and not susceptible to the variability (in local optima) induced by clustering
137 directly on data in the projected space (see Supplementary Information and Fig. S2).

138 Louvain clustering identified eight clusters in total. The UMAP algorithm then non-linearly
139 projected the waveforms into a two-dimensional space in a way that preserved local and global
140 variation. Fig. 2A shows results of this analysis. UMAP provides a clear organization without the
141 need for prior specification of important features.

142 For expository reasons and to link to prior literature^{16,52}, we loosely subdivide based on the trough
143 to peak duration these eight clusters into "narrow-spiking" and "broad-spiking" supraclusters.
144 The broad-spiking clusters had a median trough to peak duration of 0.76 ± 0.24 ms (median
145 \pm S.D.) and the narrow-spiking clusters had median trough to peak duration of 0.36 ± 0.1
146 ms (median \pm S.D.). The narrow-spiking neurons are shown in warm colors (including green)
147 at right in Fig. 2A and the broad-spiking neurons are shown in cool colors at left in the same
148 figure. The narrow-spiking supracluster was composed of five clusters with "narrow-spiking"
149 waveforms (clusters ①, ②, ③, ④, ⑤) and comprised $\sim 12\%$, $\sim 12\%$, $\sim 18\%$, $\sim 7\%$, and
150 $\sim 19\%$ ($n = 72, 78, 113, 43,$ and 116) of the total waveforms, for $\sim 68\%$ of total waveforms. The
151 broad-spiking supracluster was composed of three "broad-spiking" waveform clusters (⑥, ⑦,
152 and ⑧) comprising $\sim 13\%$, $\sim 5\%$, and $\sim 15\%$ ($n = 80, 29,$ and 94) respectively and collectively
153 $\sim 32\%$ of total waveforms.

154 To validate that *WaveMAP* finds a "real" representation of the data, we examined if a very
155 different method could pick on the same representation. We trained a gradient-boosted random
156 forest classifier on the waveform data using five-fold cross-validation with the cluster indices as
157 labels. Hyperparameters were tuned with a grid search and final parameters are shown in Table
158 S1. After training, the held-out classification accuracy averaged over clusters was 91%. Fig. 2C
159 shows the associated confusion matrix which contains accuracies for each class along the main
160 diagonal and misclassification rates on the off-diagonals. Such successful classification at high

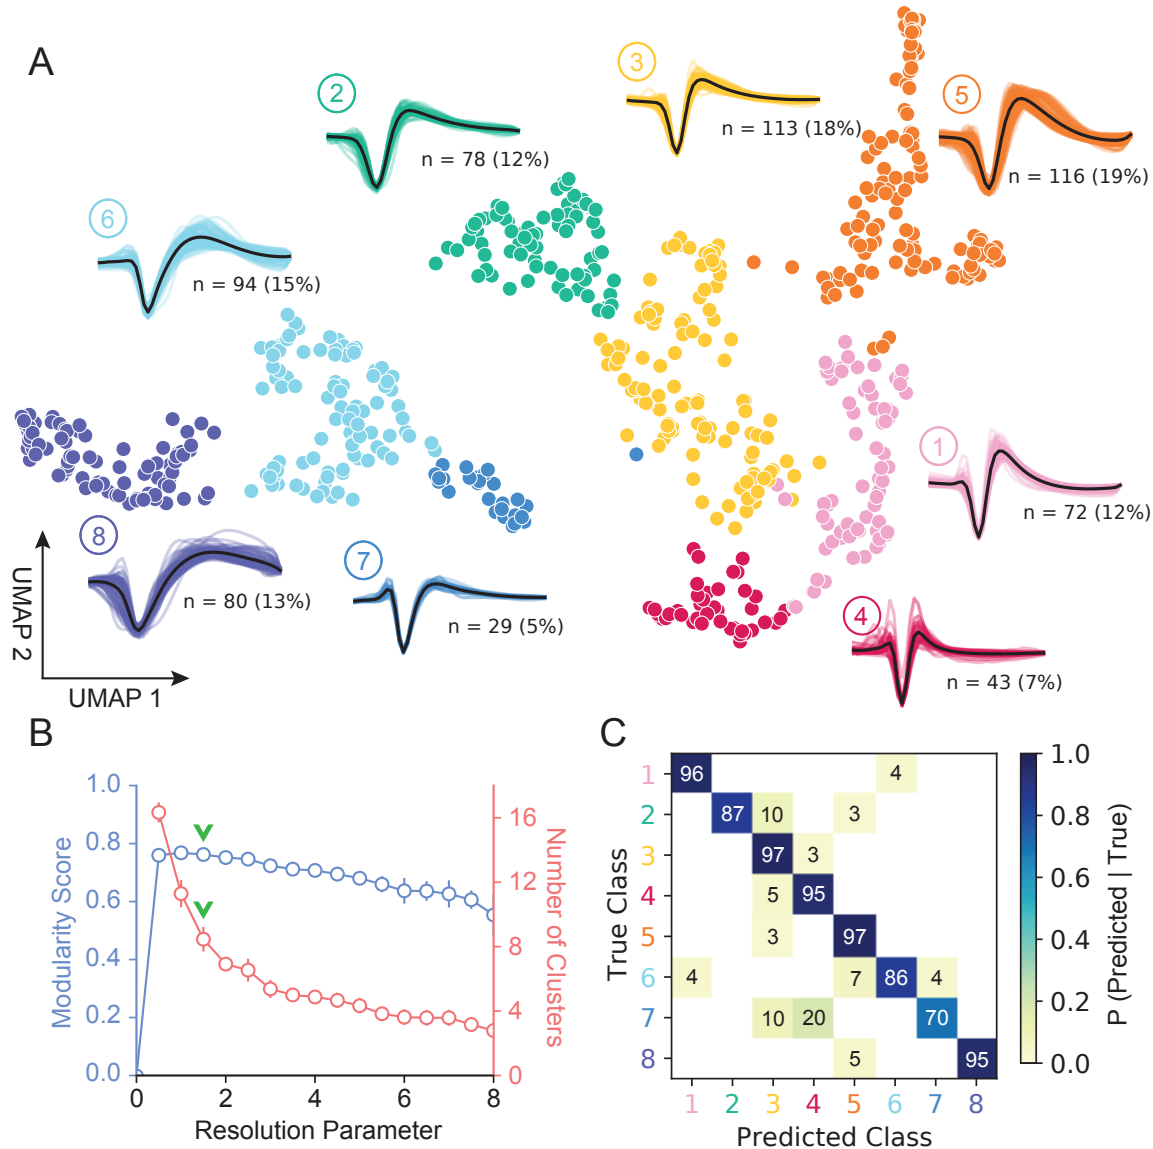


Figure 2: UMAP and Louvain clustering reveal a robust diversity of averaged single-unit waveform shapes (A) Scatter plot of normalized EAP waveforms in UMAP space colored by Louvain cluster membership. Adjacent to each numbered cluster (① through ⑧) is shown all member waveforms and the average waveform shape (in black). Each waveform is 1.6 ms. Percentages do not add to 100% due to rounding. **(B)** Louvain clustering resolution parameter versus modularity score (in blue, at left) and the number of clusters (communities) found (in red, at right). This was averaged over 25 random seeds for UMAP using the full dataset at each resolution parameter from 0 to 8 in 0.5 unit increments. Each data point is the mean \pm S.D. with many S.D. bars smaller than the marker size. Green chevrons indicate the resolution parameter of 1.5 chosen and its position along both curves. **(C)** The confusion matrix of a random forest classifier with five-fold cross-validation. The main diagonal shows the held-out classification accuracy for each cluster and the off-diagonals show the misclassification rates for each cluster to each other cluster. The average accuracy for each cluster was 91% across all clusters.

161 levels of accuracy was only possible because there were learnable clusterings of similar waveform
162 shapes in the high-dimensional space revealed by UMAP.

163 *Traditional Clustering Methods with Specified features Fail to Capture Full Variability*

164 Our unsupervised approach (Fig. 2) generates a stable clustering of waveforms. However, is our
165 method better than the traditional approach of using specified features^{27,32,50,53,54}? To compare
166 how *WaveMAP* performs relative to traditional clustering methods built on specified features,
167 we applied a Gaussian mixture model (GMM) to the three-dimensional space produced by three
168 commonly used waveform features using the same dataset. In accordance with previous work, the
169 features we chose (Fig. 3A) were action potential (AP) width of the spike (width in milliseconds
170 of the full-width half minimum of the depolarization trough, a common feature in intra-cellular
171 studies of waveforms²⁰), the peak ratio which is the ratio of pre-hyperpolarization peak (A1) to
172 the post-hyperpolarization peak (A2)²⁹, and the trough to peak duration (time in ms from the
173 depolarization trough to post-hyperpolarization peak) which is the most common feature used in
174 analysis of extracellular recordings^{32,50}.

175 The GMM result applied to these three measures is shown in Fig. 3B. This method identified
176 four waveform clusters that roughly separated into broad-spiking (BS, ~33%, n = 208) which
177 were putatively excitatory, narrow-spiking (NS, ~43%, n = 269) which are putatively inhibitory,
178 broad-spiking triphasic (BST, ~9%, n = 55), and narrow-spiking triphasic (NST, ~15%, n = 93)
179 (Fig. 3B). Broad- and narrow-spiking types are well-documented on the basis of AP width and
180 trough to peak duration. Triphasic waveforms, thought to be neurons with myelinated axons or
181 neurites^{29,30,35,56,57}, contain an initial positive spike before the trough and can be identified by
182 taking the ratio of the pre-hyperpolarization peak with the post-hyperpolarization peak around
183 1.0 ms (Fig. 3A). These GMM clusters are similar to those obtained from other clusterings of
184 EAP's in macaque cortex^{27,58}. We selected four clusters by calculating the Bayesian information
185 criterion (BIC) statistic as a function of the number of clusters (Fig. 3D). For this dataset and
186 number of features, we saw no meaningful decrease in the BIC statistic beyond 4 clusters (green
187 chevron in Fig. 3C).

188 To compare the learnability of this representation with the representation provided by UMAP, we
189 trained the same random forest classifier on the waveform data (after separate hyperparameter
190 tuning, Fig. S1) but this time using the four GMM classes as target labels. After training,
191 the accuracy across all four classes averaged ~78% with no classification accuracy over 95%
192 and misclassifications between every class (Fig. 3D). The classifier trained on specified features
193 under-performed the classifier trained on the whole waveform found by *WaveMAP*. In *WaveMAP*,
194 the individual classification accuracy of most classes exceeded 95% with few misclassifications
195 between groups even though there were double the number of clusters. This result suggests that
196 the clusters based on specified features are less differentiable than *WaveMAP* clusters even when
197 a much lower cluster number is considered.

198 This deficit can be understood as an inability of the GMM to fully capture the structure of the
199 variability of clusters. If we examine the gray data point shadows (Fig. 3B), no features contain
200 clear clusters and neither do they contain Gaussian distributions which is an assumption of the
201 GMM model. Examining the marginal distributions in Fig. 3B, none of the features induce a clear
202 separability between the clusters alone or in conjunction. Furthermore, the reproducible clusters

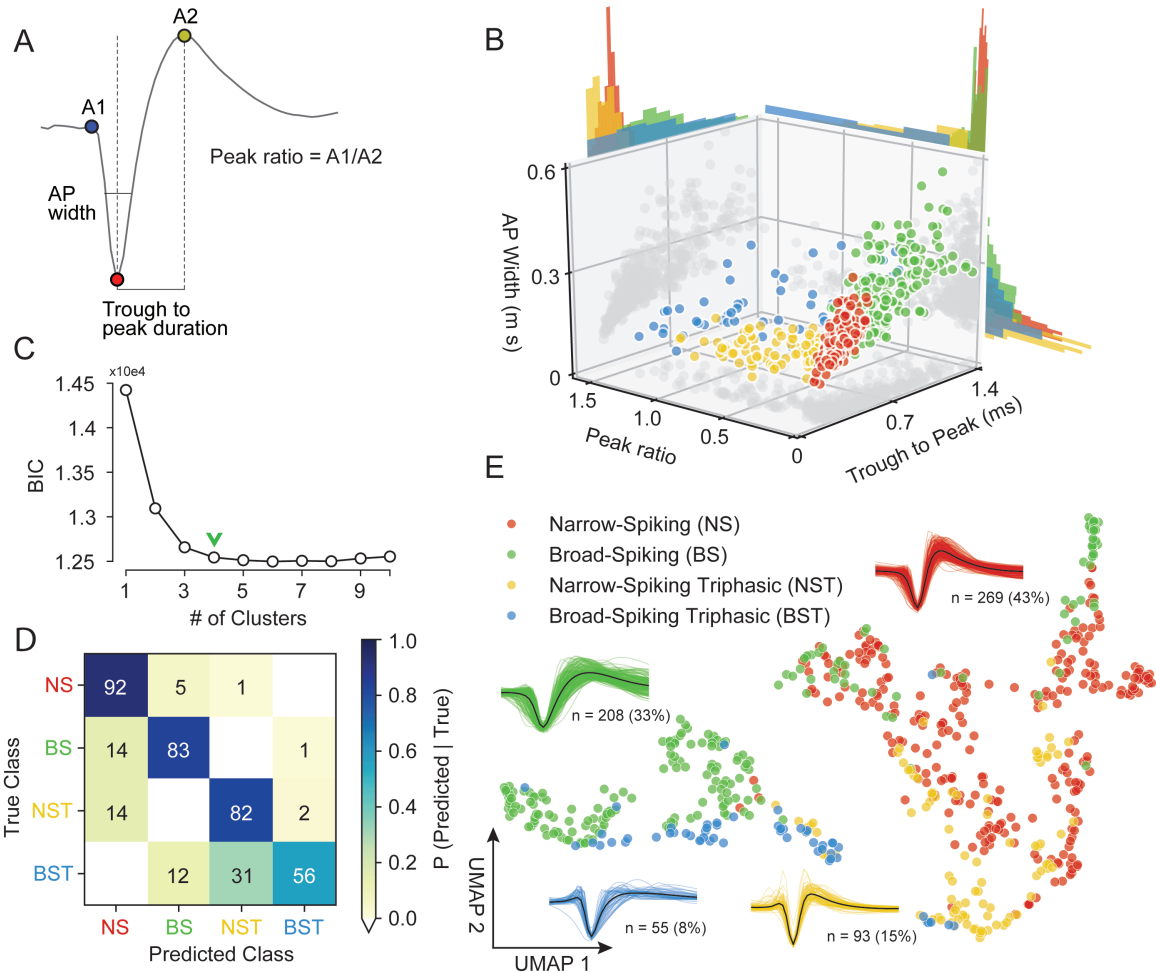


Figure 3: Gaussian mixture model clustering on specified features fails to capture the breadth of waveform diversity (A) The three EAP waveform landmarks used to generate the specified features passed to the GMM on a sample waveform. ● is the pre-hyperpolarization peak (A1); ● is the depolarization trough; and ● is the post-hyperpolarization peak (A2). **(B)** A three-dimensional scatter plot with marginal distributions of waveforms and GMM classes on the three specified features in (A). Narrow-spiking (NS) are in red; broad-spiking (BS) in green; narrow-spiking triphasic (NST) in yellow; and broad-spiking triphasic (BST) types are in blue. Trough to peak was calculated as the time between ● and ●; peak ratio was determined as the ratio between the heights of ● and ● (A1/A2); and AP width was determined as the width of the depolarization trough ● using the MLIB toolbox⁵⁵. **(C)** The optimal cluster number in the three-dimensional feature space in (B) was determined to be 4 clusters using the Bayesian information criterion (BIC)²⁷. The number of clusters was chosen to be at the “elbow” of the BIC curve (green chevron). **(D)** A confusion matrix for a gradient-boosted random classifier with five-fold cross-validation. The main diagonal contains the classification accuracy percentages across the four GMM clusters and the off-diagonal contains the misclassification rates. The average accuracy per each class was 78%. **(E)** The same scatter plot of normalized EAP waveforms in UMAP space as in Fig. 2E but now colored by GMM category.

203 found by *WaveMAP* are linearly inseparable in the feature space of the three GMM features.
 204 Labeling the data according to *WaveMAP* cluster identity yields an inseparable collection of points
 205 (Fig. S3A).

206 Note, this is not an artifact of using a lower cluster number in the GMM as opposed to the eight

207 found by *WaveMAP*. Even if the GMM is instantiated with eight clusters (Fig. S3B), a classifier
208 is still unable to learn this clustering with high accuracy (Fig. S3C) even if the waveforms found
209 by the GMM with eight clusters seem somewhat sensible (Fig. S3D). Thus, it is a deficit of the
210 constructed feature-based approach to capture the full diversity of waveforms and not a peculiarity
211 of the model parameters chosen.

212 We also investigated the representation of specified features in the projected UMAP space. We
213 color coded the waveforms in UMAP, as in Fig. S4, according to each point's feature values across
214 AP width (Fig. S4B), trough to peak duration (Fig. S4C), and peak ratio (Fig. S4D). We find that
215 *WaveMAP* implicitly captures each of these specified feature shown as a monotonic gradation of
216 values. Our method also exposes the correlation between certain specified features: the gradient
217 (direction of steepest change) between trough to peak duration and AP width points roughly in
218 the same direction thus both features are highly correlated. This correlation between features is
219 another reason why traditional approaches fail to capture the full diversity of waveform shapes.

220 To obtain a clearer picture of how *WaveMAP* captures clustering missed by specified features,
221 we color the points in UMAP space by their GMM cluster identity in Fig. 3E. Here, *WaveMAP*
222 is able to recapitulate the same structure observed by specified features as a gradation from
223 triphasic to biphasic along the UMAP-2 direction. Our technique also captures the transition from
224 broad- to narrow-spiking along the UMAP-1 direction. *WaveMAP* is also able to capture clusters
225 that occupy an intermediate identity between GMM classes. For instance, *WaveMAP* cluster ②
226 (Fig. 2A) is nearly equal parts broad- and narrow-spiking in the GMM clustering (Fig. 3E). If a
227 GMM were used, ② would be split between two classes despite it having a distinct waveform
228 shape characterized by a small pre-hyperpolarization peak, a moderate post-hyperpolarization
229 peak, and relatively constant repolarization slope.

230 *WaveMAP* interpretably recapitulates and expands upon known waveform features

231 We have established that *WaveMAP* has the ability to discover extracellular waveform clusters
232 but a common contention with such methods is that these approaches are uninterpretable. Here
233 using an interpretable machine learning approach, we show that *WaveMAP* is sensible^{59,60}. To
234 identify the features our algorithm is paying attention to, we first computed the inverse mapping
235 of the UMAP transform to probe the projected space in a systematic way. Second, we leveraged
236 the random forest classifier in Fig. 2C and used SHapley Additive exPlanations (SHAP values) to
237 reveal what waveform features are implicitly used to differentiate clusters⁶¹.

238 To quantify the differences between Louvain clusters, we applied a grid of “test points” to the
239 UMAP projected space (Fig. 4A, top) and inverted the transform at each location; each of these
240 test points is a coordinate on a grid (black x's) and shows the waveform associated with every point
241 in the projected space (Fig. 4A, bottom). On the bottom is shown the waveform that corresponds
242 to each point in UMAP space color-coded to the nearest cluster or to gray if there were no nearby
243 clusters. As UMAP-1 increases, there is a smooth transition in the sign of the inflection of the
244 repolarization slope (the second derivative) from negative to positive (slow to fast repolarization
245 rate). That is, the post-hyperpolarization peak becomes more sharp as we increase in the UMAP-1
246 direction. As UMAP-2 increases, we see a widening of the post-hyperpolarization slope distinct
247 from the change in its inflection (UMAP-1). These two UMAP dimensions recapitulate the
248 known importance of hyperpolarization properties in clustering waveforms. Both hyperpolarization

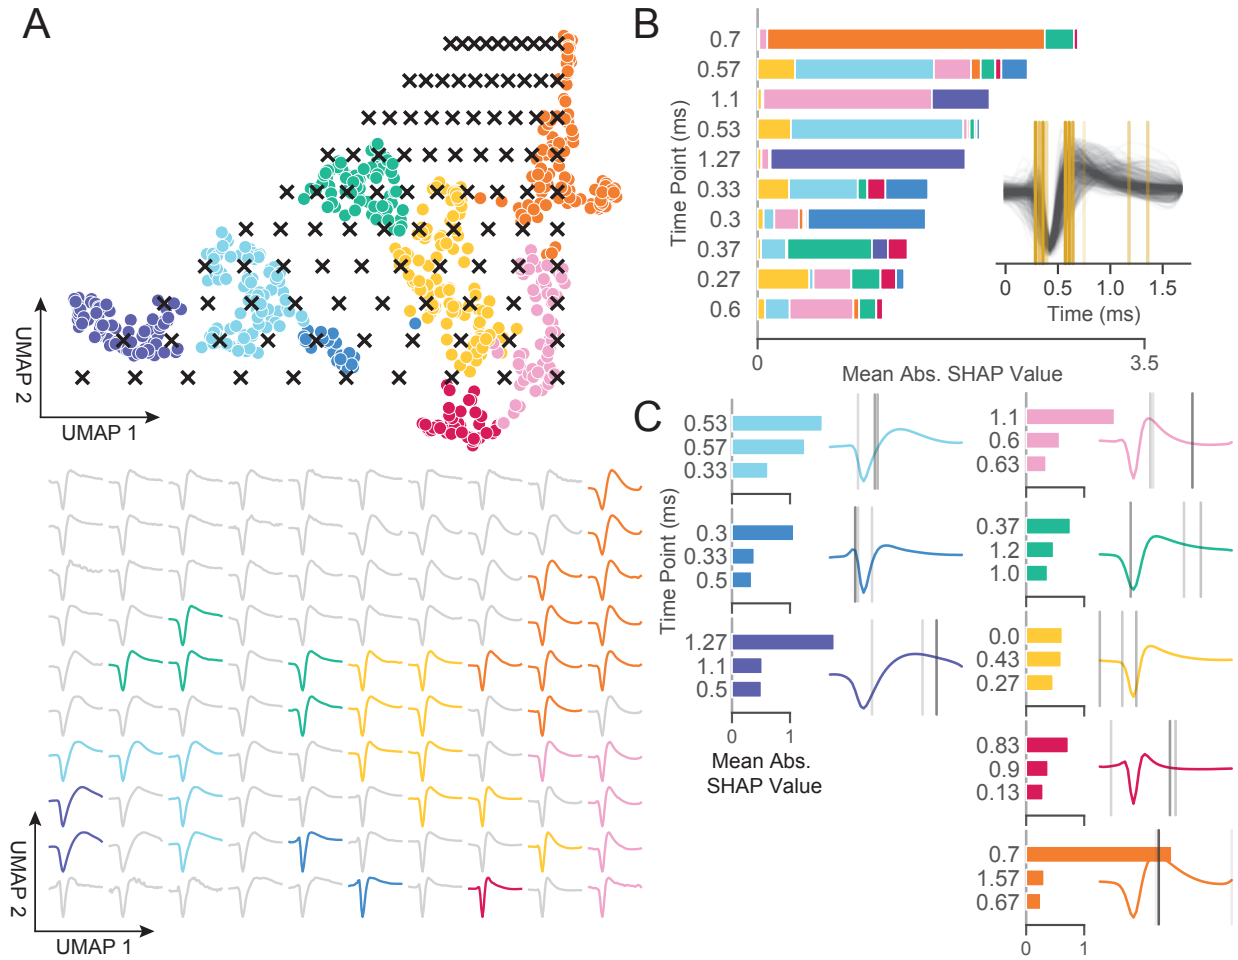


Figure 4: *WaveMAP* provides interpretable representations that both validate and extend known and unknown features importances (A) *WaveMAP* applied to the EAP's as in Fig. 2A overlaid with a grid of test points (black x's, top). At bottom, the inverse UMAP transform is used to show the predicted waveform at each test point. These are plotted and assigned the color of the nearest cluster or in gray if no cluster is nearby. Note that there exists instability in the waveform shape (see waveforms at corners) as test points leave the learned manifold. (B) The mean absolute SHAP values for ten time points along all waveforms subdivided according to the SHAP values contributed by each *WaveMAP* cluster. These SHAP values were informed by a random forest classifier trained on the waveforms with the *WaveMAP* clusters as labels. In the inset, all waveforms are shown and in goldenrod are shown the time points for which the SHAP values are shown on the left. Each vertical line is such that the most opaque line contains the greatest summed mean absolute SHAP value across *WaveMAP* clusters; the least opaque, the smallest summed mean absolute SHAP value. (C) Each averaged Louvain waveform cluster is shown with the three time points containing the greatest SHAP values for each cluster individually. As before, the SHAP value at each time point is proportional to the opacity of the gray vertical line but now also shown as a bar graph; this bar graph shows the relative SHAP values for the three locations specified by the gray lines.

249 rate (proportional to trough to peak width) and hyperpolarization slope inflection (proportional
250 to repolarization time) are separate but highly informative properties^{26,27}. Furthermore, since
251 repolarization rate and post-hyperpolarization width associate with different UMAP dimensions,
252 this implies that these two processes are somewhat independent factors shaping the waveform and
253 that potassium channel dynamics—which principally govern repolarization rates—might play an
254 important role in differentiating cell types via waveform shape. Thus, *WaveMAP* not only finds
255 an interpretable and smoothly varying low-dimensional space (a “manifold”) it also offers possibly
256 yielding biological insights, in this case how cell types might differ according to channel protein
257 expression and dynamics.

258 In Fig. 4B, we made use of SHAP values to identify which aspects of waveform shape the random
259 forest classifier utilizes in assigning what waveform to which cluster^{61,62}. SHAP values build off
260 of the game theoretic quantity of Shapley values^{63,64} which poses that each feature (point in
261 time along the waveform) is of variable importance in influencing the classifier to decide whether
262 the data point belongs to a specific class or not. Operationally, SHAP values are calculated
263 by examining the change in classifier performance as each feature is obscured (the waveform’s
264 amplitude at each time point in this case), one-by-one⁶¹. Fig. 4B shows the top-10 time points
265 (in terms of total SHAP value, i.e. importance in explaining a cluster) and their location. It is
266 important to note that not every time point is equally informative for distinguishing every cluster
267 individually and thus each bar is subdivided into the mean absolute SHAP value contribution of
268 the eight constituent waveform classes. For instance, the 0.7 ms location is highly informative for
269 cluster ⑤ and the 0.3 ms point is highly informative for cluster ⑦.

270 In the inset is shown all waveforms along with each of the top ten time points (in goldenrod)
271 with higher SHAP value shown with more opacity. The time points with highest SHAP value
272 tend to cluster around two different locations giving us an intuition for which locations are most
273 informative for telling apart the Louvain clusters. For instance, the 0.5 to 0.65 ms region contains
274 high variability amongst waveforms and is important in separating out broad- from narrow-spiking
275 clusters. This region roughly contains the post-hyperpolarization peak which is a feature of known
276 importance and incorporated into nearly every study of EAP waveform shape (see Table 2 in
277 ref.²⁰). Similarly, SHAP values implicate the region around 0.3 ms to 0.4 ms as time points
278 that are also of importance and these correspond to the pre-hyperpolarization peak which is
279 notably able to partition out triphasic waveforms²⁹. Importance is also placed on the location
280 at 0.6 ms corresponding to the inflection point which is similarly noted as being informative^{27,28}.
281 These methods also implicate other regions of interest that have not been previously noted in
282 the literature to the best of our knowledge: two other locations are highlighted farther along the
283 waveform at 1.1 and 1.27 ms and are important for differentiating ⑧ and ① from the other
284 waveforms. This result suggests that using only up to 1.0 ms or less of the waveform may obscure
285 diversity.

286 In Fig. 4C, we show the three locations that are most informative for delineating a specific
287 cluster; these appear as gray lines with their opacity proportional to their importance. These
288 individually-informative features often do align with those identified as globally-informative but do
289 so with cluster-specific weights. Put another way, not every time point is equally informative for
290 identifying waveforms individually and these “most informative” parts of each waveform don’t
291 always perfectly align with globally-informative features. In summary, *WaveMAP* independently

292 and sensibly arrived at a more nuanced incorporation of the very same features identified in previous
 293 work—and several novel ones—using a completely unsupervised framework which obviated the
 294 need to specify waveform features. In the second half of the paper, we investigate whether these
 295 clusters have distinct physiological (in terms of firing rate), functional, and laminar distribution
 296 properties which would give evidence that they connect to real cell types.

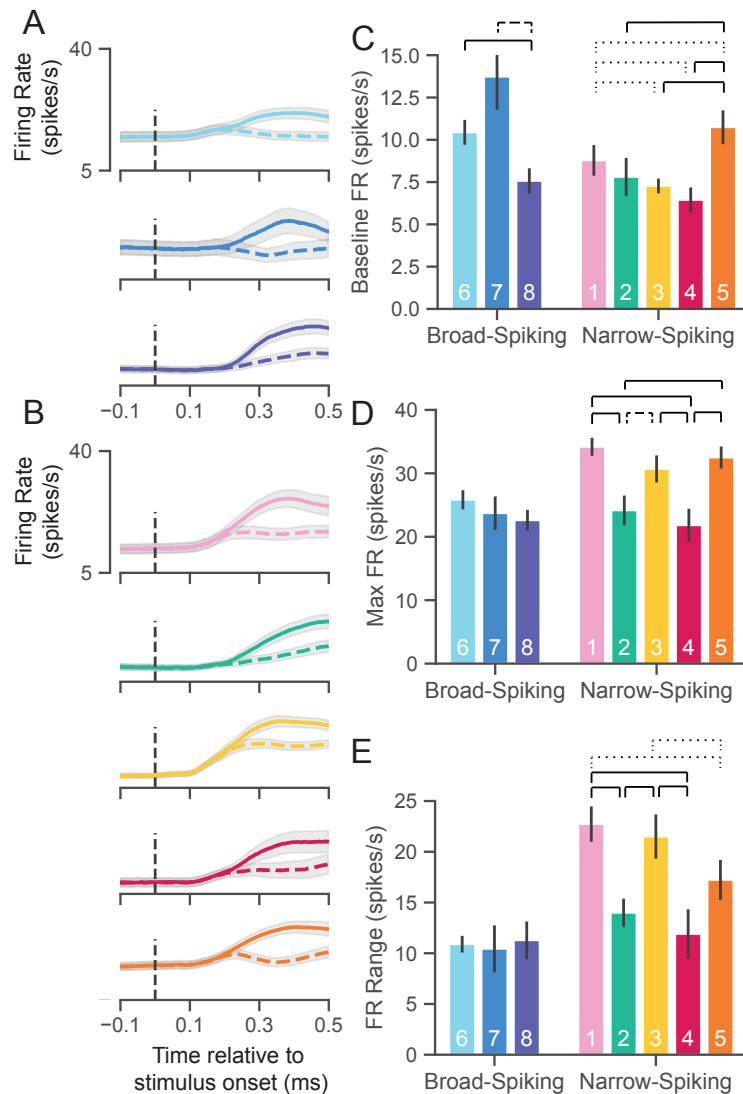


Figure 5: UMAP clusters exhibit distinct physiological properties (A) Stimulus-aligned trial-averaged firing rate activity in PMd for broad-spiking *WaveMAP* clusters. The traces shown are separated into trials for PREF direction reaches (solid lines) and NONPREF direction reaches (dashed lines) and across the corresponding *WaveMAP* clusters. Shaded regions correspond to bootstrapped standard error of the mean. **(B)** The same plots as in (A) but for narrow-spiking *WaveMAP* clusters. **(C)** Baseline median firing rates with S.E.M. for the neurons in the eight different classes. Baselines were calculated as the average firing rate from the first 200 ms of recording before the stimulus appeared. **(D)** Maximum median firing rates with S.E.M. for the neurons in the eight different clusters. This was calculated as the median maximum firing rate for each neuron across the entire trial. **(E)** FR range with S.E.M. calculated as the median difference, per neuron, between its baseline and max FR. $p < 0.05$; ---- $p < 0.01$; — $p < 0.005$; Mann-Whitney *U* test, FDR adjusted.

297 *WaveMAP clusters have distinct physiological properties*

298 A defining aspect of cell types is that they vary in their physiology and especially firing rate
299 properties^{15,16,52,65–67}. However, these neuronal characterizations via waveform *ex vivo* are not
300 always conserved when the same waveform types are observed *in vivo* during behavior^{68,69}. To
301 connect our waveform clusters to physiological cell types *in vivo*, beyond waveform shape, we
302 identified each cluster's firing rate properties. We performed several analyses using the firing rate
303 (FR) in spikes per second (spikes/s) for each cluster during the decision-making task described in
304 Fig. 1.

305 The trial-averaged FRs are aligned to stimulus onset (stim-aligned) and separated into preferred
306 (PREF, solid trace) or non-preferred (NONPREF, dashed trace) reach direction trials. This is shown
307 for both broad- (Fig. 5A) and narrow-spiking (Fig. 5B) clusters. A neuron's preferred direction
308 (right or left) was determined as the reach direction in which it had a higher FR on average in the
309 100 ms time period before movement onset.

310 To further quantify the FR differences between clusters we calculated three properties of the FR
311 response to stimulus: baseline firing rate, max firing rate, and FR range.

312 **Baseline FR:** Cell types are thought to demonstrate different baseline firing rates. We estimated
313 baseline firing rate (baseline FR, Fig. 5C) as the median FR across the 200 ms time period before
314 the appearance of the red-green checkerboard and during the hold period after targets appeared
315 for the broad (Fig. 5A), and narrow-spiking clusters (Fig. 5B). The broad-spiking clusters showed
316 significant differences in baseline FR when pooled and compared against the pooled narrow-spiking
317 clusters ($p = 0.0028$, Mann-Whitney U test). Similar patterns were observed in another study
318 of narrow vs. broad spiking neurons in PMd during an instructed delay task⁵³. We also found
319 that not all broad-spiking neurons had low baseline firing rates and not all narrow-spiking neurons
320 had high firing rates. The broad-spiking clusters ⑥ and ⑦ were not significantly different but
321 both differed significantly from ⑧ in that their baseline FR was much higher (10.3 ± 0.7 and
322 13.2 ± 1.9 spikes/s vs. 7.6 ± 0.75 spikes/s [median \pm bootstrap S.E.]; $p = 0.0052$, $p = 0.0029$
323 respectively, Mann-Whitney U test, FDR adjusted). The narrow-spiking clusters (Fig. 5B, right)
324 ②, ③, and ④ had relatively low median baseline FRs (7.5 ± 1.1 , 7.4 ± 0.4 , 6.5 ± 0.7 spikes/s,
325 median \pm bootstrap S.E.) and were not significantly different from one another but all were
326 significantly different from ① and ⑤ ($p = 0.04$, $p = 2.8e-4$, $p = 2.8e-7$, $p = 4.9e-5$ respectively,
327 Mann-Whitney U test, FDR adjusted; see Fig. 5C).

328 **Maximum FR:** A second important property of cell types is their maximal firing rate^{15,16,65}. We
329 estimated the maximum FR for a cluster as the median of the maximum FR of neurons in the
330 cluster in a 1200 ms period aligned to movement onset (800 ms before and 400 ms after movement
331 onset; Fig. 5D). In addition to significant differences in baseline FR, pooled broad- vs. pooled
332 narrow-spiking neurons showed a significant difference in max FR ($p = 1.60e-5$, Mann-Whitney
333 U test). Broad-spiking clusters were fairly homogeneous with low median max FR (24.3 ± 1.0 ,
334 median \pm bootstrap S.E.) and no significant differences between distributions. In contrast, there
335 was significant heterogeneity in the FR's of narrow-spiking neurons: three clusters (①, ③, and ⑤
336) had uniformly higher max FR (33.1 ± 1.1 , median \pm bootstrap S.E.) while two others (② and
337 ④) were uniformly lower in max FR (23.0 ± 1.4 , median \pm bootstrap S.E.) and were comparable
338 to the broad-spiking clusters. Nearly each of the higher max FR narrow-spiking clusters were

339 significantly different than each of the lower max FR clusters (all pairwise relationships $p < 0.001$
340 except ② to ③ which was $p = 0.007$, Mann-Whitney U test, FDR adjusted).

341 **FR Range:** Many neurons, especially inhibitory types, display a sharp increase in firing rates
342 during behavior relative to their baseline. In addition, their FR's span a wide range during
343 behavior^{1,32,53,70,71}. To examine this change over the course of a trial, we took the median
344 difference across trials between the max FR and baseline FR per neuron to calculate the FR range.
345 We again found the group difference between pooled broad- and pooled narrow-spiking clusters to
346 be significant ($p = 0.0002$, Mann-Whitney U test). Each broad-spiking cluster (⑥, ⑦, and ⑧)
347 had a median increase of around 10.8 spikes/s (10.8 ± 0.8 , 10.7 ± 2.3 , and 10.9 ± 1.9 spikes/s
348 respectively, median \pm bootstrap S.E.) and each was nearly identical in FR range differing by
349 less than 0.2 spikes/s. In contrast, the narrow-spiking clusters differed greatly in their FR range
350 showing, in the same manner as max FR, with significant differences between the clusters ①, ③,
351 and ⑤ having a large range (20.3 ± 1.1 spikes/s, median \pm bootstrap S.E.) and the clusters ②
352 and ④ having a relatively smaller range (13.4 ± 1.3 spikes/s, median \pm bootstrap S.E.). This
353 demonstrates that some narrow-spiking clusters, in addition to having high baseline firing rates,
354 highly modulated their firing rates over the course of a behavioral trial.

355 Such physiological heterogeneity in narrow-spiking cells has been noted before in their FR
356 properties^{26,28,72} and in some cases attributed to different subclasses of a single inhibitory class^{73,74}.
357 Other work also strongly suggests that narrow-spiking cells contain excitatory neurons with distinct
358 FR properties which may contribute to this diversity^{20,24}.

359 Furthermore, if *WaveMAP* has truly arrived at a closer delineation of underlying cell types compared
360 to previous methods, it should produce a "better" clustering of physiological properties beyond
361 just a better clustering of waveform shape. To address this issue, we calculate the same firing
362 rate traces and physiological properties as in Fig. 5 but with the GMM clusters (Fig. S6). While
363 the FR traces maintain the same trends (BS does not increase its FR prior to the split into PREF
364 and NONPREF while NS does; compare to *WaveMAP* broad-spiking vs. narrow-spiking clusters
365 respectively), much of the significant differences between clusters is lost across all physiological
366 measures even though less groups are compared (Fig. S6B, C, and D). We also quantitatively reify
367 these differences by calculating the effect sizes (Cohen's f^2) across the *WaveMAP* and GMM
368 clusterings with a one-way ANOVA. The effect size was larger for *WaveMAP* vs. GMM clustering
369 respectively for every physiological property: baseline firing rate (0.070 vs. 0.013), maximum
370 firing rate (0.035 vs. 0.011), and firing rate range (0.055 vs. 0.034).

371 *UMAP clusters have distinct decision-related dynamics*

372 Our analysis in the previous section showed that there is considerable heterogeneity in their physio-
373 logical properties. Are these putative cell types also functionally different? Prior literature suggests
374 and argues that neuronal cell types have distinct functional roles during cortical computation
375 with precise timing. For instance, studies of macaque premotor⁷⁵, interior temporal (IT)⁷⁶, and
376 frontal eye field (FEF)⁹ areas show differences in decision-related functional properties: between
377 broad- and narrow-spiking neurons, narrow-spiking neurons exhibit choice-selectivity earlier than
378 broad-spiking neurons. In the mouse, specific aspects of behavior are directly linked with inhibitory
379 cell types^{3,4}. This would suggest that inhibitory cell types in the monkey are also implicated in
380 behavior although this has yet to be directly assessed. Here we examine the functional properties

381 of each cluster based on two inferred statistics: choice-related dynamics and discrimination time.

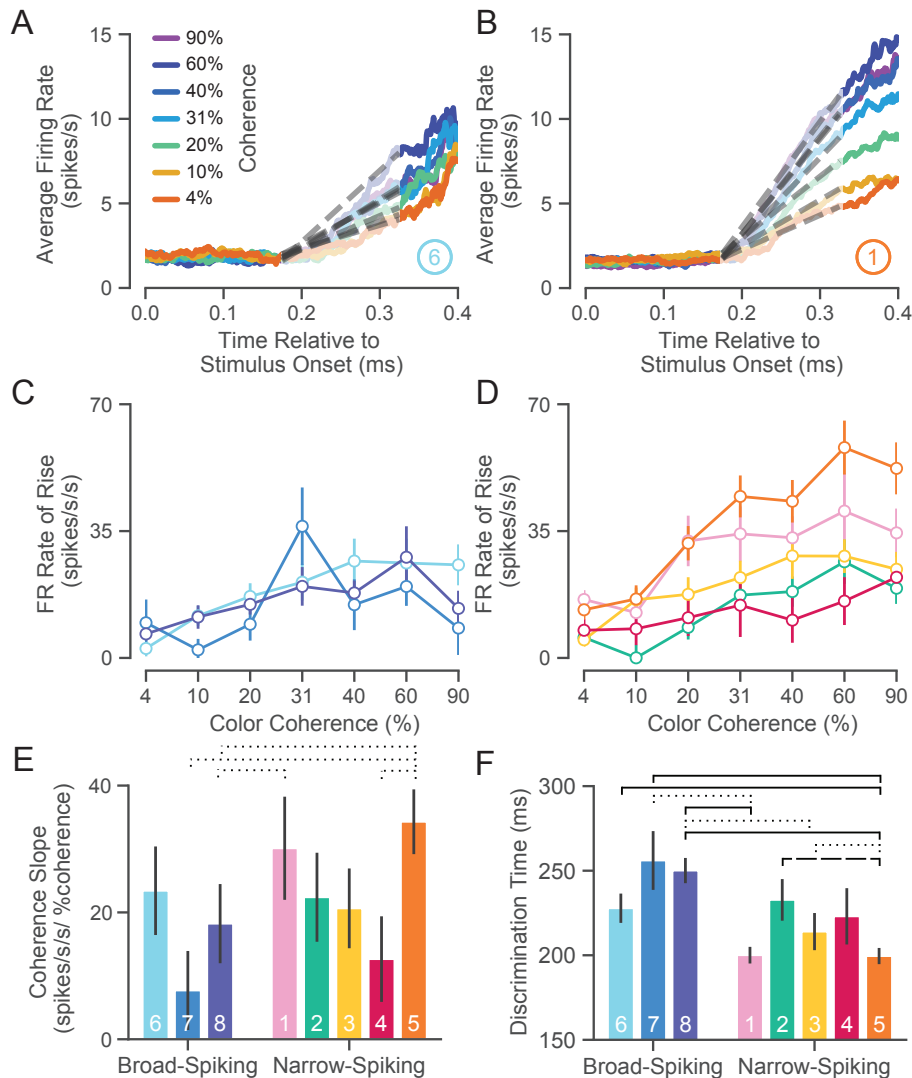


Figure 6: UMAP clusters exhibit distinct functional properties (A) Average FR over time for ⑥ across different coherences. (B) Average FR over time for ① across different coherences. (C) FR rate of rise vs. coherence for broad- and (D) narrow-spiking clusters. (E) Bootstrapped median coherence slope is shown with the bootstrapped standard error of the median for each cluster on a per-neuron basis. Coherence slope is a linear regression of the cluster-specific lines in the previous plots Fig. 6C and D. (F) Bootstrapped median discrimination time for each cluster with bootstrapped standard error of the median. Discrimination time was calculated as the earliest time at which the choice-selective signal could be differentiated from the FR or the hold period¹. $p < 0.05$; ---- $p < 0.01$; — $p < 0.005$; Mann-Whitney U test, FDR adjusted.

382 **Choice-related Dynamics:** The first property we assessed for these *WaveMAP* clusters was the
 383 dynamics of the choice-selective signal. The neural prediction made by computational models of
 384 decision-making (for neurons that covary with an evolving decision) is the build-up of average
 385 neural activity in favor of a choice is faster for easier compared to harder color coherences^{1,9,11}.
 386 Build-up activity is measured by analyzing the rate of change of choice-selective activity vs.
 387 time. We therefore examined the differences in averaged stimulus-aligned choice-selectivity signals

388 (defined as $|\text{left} - \text{right}|$) for different checkerboard color coherences for each cluster.

389 In Fig. 6A and B, we show average choice-selectivity signals across seven coherence levels for an
390 example broad- (⑥) and narrow-spiking cluster (①). For ⑥ (Fig. 6A), easier stimuli (higher
391 coherence) only led to modest increases in the rate at which the choice selectivity signal increases.
392 In contrast, ① (Fig. 6B) shows faster rates for the choice-selective signal as a function of
393 coherence. We summarized these effects by measuring the rate of change for the choice-selective
394 signal between 175 and 325 ms for stimulus-aligned trials in each coherence condition (dashed lines
395 in Fig. 6A,B). This rate of rise for the choice-selective signal (spikes/s/s) vs. coherence is shown
396 for broad- (Fig. 6C) and narrow-spiking (Fig. 6D) clusters. The broad-spiking clusters demonstrate
397 fairly similar coherence-dependent changes with each cluster being somewhat indistinguishable and
398 only demonstrating a modest increase with respect to coherence. In contrast, the narrow-spiking
399 clusters show a diversity of responses with ① and ⑤ demonstrating a stronger dependence of
400 choice-related dynamics on coherence compared to the other three narrow-spiking clusters which
401 were more similar in response to broad-spiking neurons.

402 We further summarized these plots by measuring the dependence of the rate of rise of the
403 choice-selective signal as a function of coherence measured as the slope of a linear regression
404 performed on the rate of rise vs. coherence for each cluster (Fig. 6E). The coherence slope for
405 broad-spiking clusters was moderate and similar to ②, ③, and ④ while the coherence slope for
406 ① and ⑤ was steeper. Consistent with Fig. 6C,D, the choice selective signal for ① and ⑤
407 showed the strongest dependence on stimulus coherence.

408 **Discrimination time:** The second property that we calculated was the discrimination time for
409 clusters which is defined as the first time in which the choice-selective signal (again defined as
410 $|\text{left} - \text{right}|$) departed from the FR of the hold period. We calculated the discrimination
411 time on a neuron-by-neuron basis by computing the first time point in which the difference in FR
412 for the two choices was significantly different from baseline using a bootstrap test (at least 25
413 successive time points significantly different from baseline FR corrected for multiple comparisons¹).
414 Discrimination time for broad-spiking clusters (255 ± 94 ms, median \pm bootstrap S.E.) was
415 significantly later than narrow-spiking clusters (224 ± 89 ms, $p < 0.005$, median \pm bootstrap
416 S.E., Mann-Whitney *U* test). Clusters ① and ⑤, with the highest max FRs (34.0 ± 1.4 and
417 33.0 ± 1.8 spikes/s, median \pm S.E.) and most strongly modulated by coherence, had the fastest
418 discrimination times as well (200.0 ± 4.9 and 198.5 ± 4.9 ms, median \pm S.E.).

419 Together the analysis of choice-related dynamics and discrimination time showed that there is
420 considerable heterogeneity in the properties of narrow-spiking neuron types. Not all narrow-spiking
421 neurons are faster than broad-spiking neurons and choice-selectivity signals have similar dynamics
422 for many broad-spiking and narrow-spiking neurons. ① and ⑤ have the fastest discrimination
423 times and strongest choice dynamics. In contrast, the broad-spiking neurons have uniformly slower
424 discrimination times and weaker choice-related dynamics.

425 *WaveMAP clusters contain distinct laminar distributions*

426 In addition to having certain physiological properties and functional roles, numerous studies have
427 shown that cell types across phylogeny, verified by single-cell transcriptomics, are defined by
428 distinct patterns of laminar distribution in cortex^{18,77}. Here we examined the laminar distributions
429 of *WaveMAP* clusters and compared them to laminar distributions of GMM clusters. The number

430 of waveforms from each cluster was counted at each of sixteen U-probe channels separately. These
431 channels were equidistantly spaced every 0.15 mm between 0.0 to 2.4 mm. This spanned the
432 entirety of PMd which is 2.5 mm in depth from the pial surface to white matter⁷⁸. However,
433 making absolute statements about layers is difficult with these measurements because of errors
434 in aligning superficial electrodes with layer I across different days. This could lead to shifts in
435 estimates of absolute depth; up to 0.15 mm (the distance between the first and second electrode)
436 of variability is induced in the alignment process (see Methods). However, relative comparisons are
437 likely better preserved. Thus, we use relative comparisons to describe laminar differences between
438 distributions and in comparison to anatomical counts in fixed tissue in later sections.

439 Above each column of Fig. 7A, B are the laminar distributions for all waveforms in the associated
440 supracluster (in gray); below these are the laminar distributions for each supracluster's constituent
441 clusters. On the right (Fig. 7C), we show the distribution of all waveforms collected at top in
442 gray with each GMM cluster's distribution shown individually below.

443 The overall narrow- and broad-spiking populations did not differ significantly according to their
444 distribution ($p = 0.24$, Kolmogorov-Smirnov test). The broad-spiking supracluster of neurons (⑥,
445 ⑦, and ⑧) are generally thought to contain cortical excitatory pyramidal neurons enriched in middle
446 to deep layers^{16,79}. Consistent with this view, we found these broad-spiking clusters (Fig. 7A) were
447 generally centered around middle to deep layers with broad distributions and were not significantly
448 distinguishable in laminarity (all comparisons $p > 0.05$, two-sample Kolmogorov-Smirnov test,
449 FDR adjusted).

450 In contrast, narrow-spiking clusters (Fig. 7B) were distinctly varied in their distribution such that
451 almost every cluster had a unique laminar distribution. Cluster ① contained a broad distribution.
452 It was significantly different in laminar distribution from clusters ② and ④ ($p=0.002$ and $p =$
453 0.013 respectively, two-sample Kolmogorov-Smirnov, FDR adjusted).

454 Cluster ② showed a strongly localized concentration of neurons at a depth of 1.1 ± 0.33 mm
455 (mean \pm S.D.). It was significantly different from almost all other narrow-spiking clusters ($p =$
456 0.002 , $p = 1e-5$, $p = 0.010$ for ①, ④, and ⑤ respectively; two-sample Kolmogorov-Smirnov
457 test, FDR adjusted). Similarly, cluster ③ also showed a strongly localized laminar distribution
458 but was situated more superficially than ② with a heavier tail (1.0 ± 0.6 mm, mean \pm S.D.).
459 This cluster was more superficial than ②.

460 Cluster ④ was restricted uniquely deep in its cortical distribution (1.70 ± 0.44 , mean \pm
461 S.D.). These neurons had a strongly triphasic waveform shape characterized by a large pre-
462 hyperpolarization peak. These waveforms have been implicated as arising from myelinated
463 excitatory pyramidal cells²⁹ especially dense in this caudal region of PMd⁸⁰.

464 The last cluster, ⑤, like ① was characterized by a broad distribution across cortical depths
465 unique among narrow-spiking neurons and was centered around a depth of 1.3 ± 0.65 mm (mean
466 \pm S.D.) and present in all layers⁷⁸.

467 Such laminar differences were not observed when we used GMM clustering. Laminar distributions
468 for BS, BST, NS, and NST did not significantly differ from each other (Fig. 7C; BS vs. BST had
469 $p = 0.067$, all other relationships $p > 0.2$; two-sample Kolmogorov-Smirnov test, FDR adjusted).
470 Each GMM cluster also exhibited broad distributions across cortex which is at odds with our

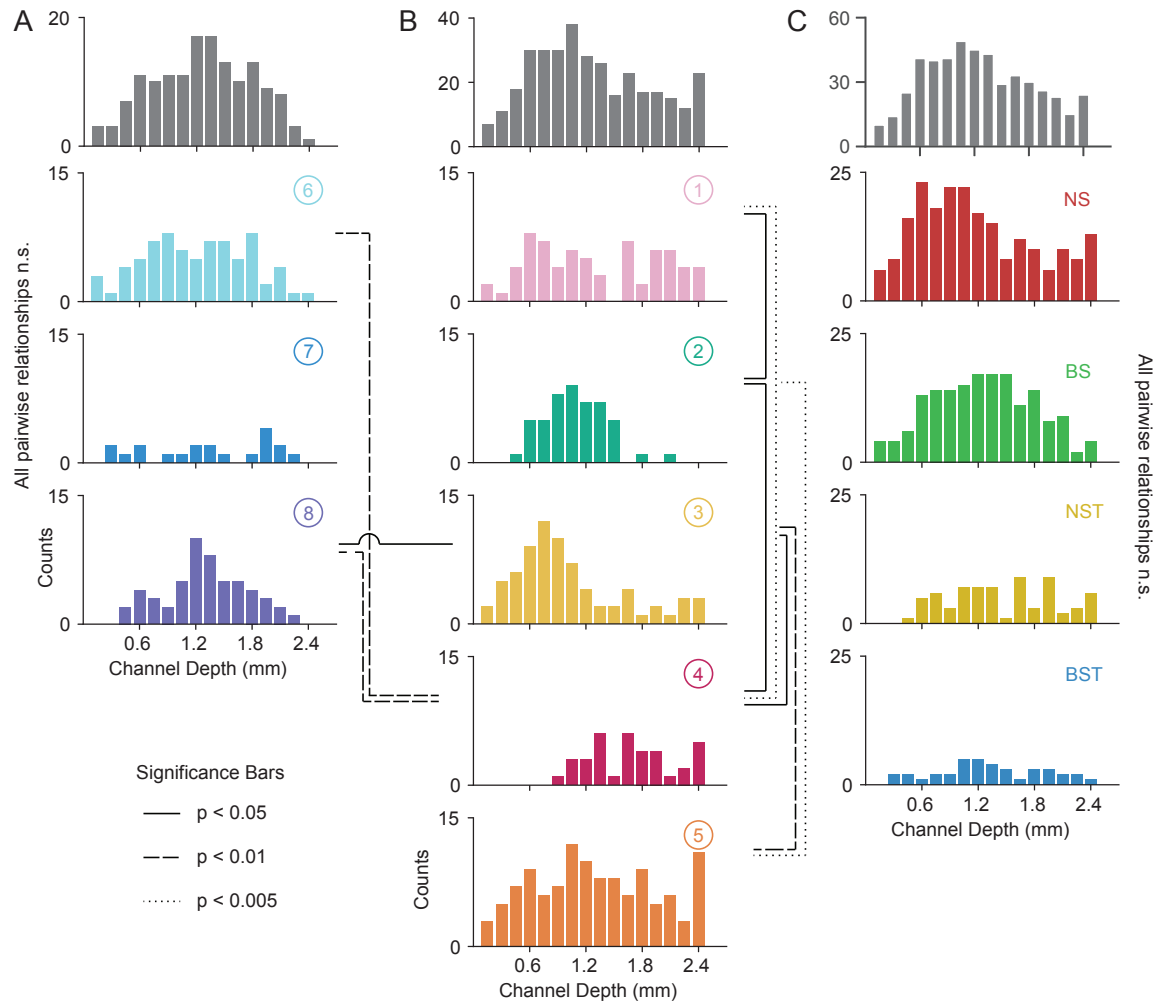


Figure 7: Laminar distribution of WaveMAP waveform clusters (A, B) The overall histogram for the broad- and narrow-spiking waveform clusters are shown at top across cortical depths on the left and right respectively (in gray); below are shown histograms for their constituent WaveMAP clusters. These waveforms are shown sorted by the cortical depth at which they were recorded from the (0.0 mm (pial surface) to 2.4 mm in 0.15 mm increments). Broad-spiking clusters were generally centered around middle layers and were less distinct in their differences in laminar distribution. Narrow-spiking clusters are shown on the right and were varied in their distribution with almost every cluster significantly varying in laminar distribution from every other. **(C)** Depth histograms for all waveforms collected (top, in gray) and every GMM cluster (below). p < 0.05; - - - p < 0.01; — p < 0.005; two-sample Kolmogorov-Smirnov Test, FDR adjusted.

471 understanding of cell types using histology.

472 *Some narrow-spiking WaveMAP cluster laminar distributions align with inhibitory subtypes*

473 We have shown that WaveMAP clusters have more distinct laminarity than GMM clusters. If
 474 WaveMAP clusters are consistent with cell type, we should expect their distributions to be relatively
 475 consistent with distributions from certain anatomical types visualized via immunohistochemistry
 476 (IHC). An especially well-studied set of non-overlapping anatomical inhibitory neuron types in the
 477 monkey are parvalbumin-, calretinin-, and calbindin-positive GABAergic interneurons (PV⁺, CR⁺,
 478 and CB⁺ respectively)⁸¹. Using IHC, we examined tissue from macaque rostral PMd stained for

479 each of these three interneuron types. We then conducted stereological counting of each type
480 averaged across six exemplars to quantify cell type distribution across cortical layers (see Fig. 8A
481 and B)⁸² and compared it to the distributions in Fig. 7.

482 Both CB⁺ and CR⁺ cells (Fig. 8C and D respectively) exhibited a similarly restricted superficial
483 distribution most closely resembling ③. In addition, CR⁺ and CB⁺ cells are known to have very
484 similar physiological properties and spike shape⁸³. An alternative possibility is that one of CR⁺ or
485 CB⁺ might correspond to ② and the other to ③ but this is less likely given their nearly identical
486 histological distributions (Fig. 8C and D) and similar physiology⁸³.

487 In contrast, *WaveMAP* cluster ①, had laminar properties consistent with PV⁺ neurons (Fig. 8B):
488 both were concentrated superficially but proliferated into middle layers (Fig. 8E). In addition, there
489 were striking physiological and functional similarities between ① and PV⁺ cells. In particular, both
490 ① and PV⁺ cells have low baseline firing rate, early responses to stimuli and robust modulation
491 of firing rate. This feature of PV⁺ cells was elicited through optogenetic studies in the mouse
492 that show PV⁺ neurons in M1 have low baseline firing rate, early responsiveness to stimuli, high
493 peak firing rate, and short discrimination time⁴. Cluster ⑤ also had similar properties to ① and
494 could also correspond to PV⁺ cells.

495 Together, these results from immunohistochemistry suggest that the narrow-spiking clusters
496 identified from *WaveMAP* potentially map on to different inhibitory types.

497 *Heterogeneity in decision-related activity emerges from both cell type and layer*

498 Our final analysis examines whether these *WaveMAP* clusters can explain some of the heterogeneity
499 observed in decision-making responses in PMd above previous methods¹. Heterogeneity in decision-
500 related activity can emerge from two sources: laminar location that presumably reflects inputs
501 and outputs or variability that can emerge from different cell types within each layer. To examine
502 whether *WaveMAP* clusters and neuron depth captured the same or different aspects of the
503 microcircuit, we regressed discrimination time on both separately and together. To compare these,
504 we examined the change in variance explained (adjusted R^2). We then compared this against the
505 GMM clusters with depth information to show that *WaveMAP* better explains the heterogeneity
506 of decision-related responses.

507 We previously showed that some of the variability in decision-related responses is explained by the
508 layer from which the neurons are recorded¹. However, as our results here show, there is considerable
509 variability within cortical layers. We therefore examined if adding cell type information in addition
510 to cortical depth improved our explained variability in decision-related activity. Consistent with
511 previous work, we found that cortical depth explains some variability in discrimination time (1.7%).
512 We next examined if the *WaveMAP* clusters identified explain variability in discrimination time: a
513 categorical regression between *WaveMAP* clusters and discrimination time, explained a much larger
514 6.6% of variance. Including both cortical depth and cluster identity in the regression explained
515 7.3% of variance in discrimination time.

516 Repeating this analysis for GMM clusters, we found that GMM clusters regressed against dis-
517 crimination time only explained 3.3% of variance and the inclusion of both GMM cluster and
518 cortical depth only explained 4.6% of variance. Thus, we find that *WaveMAP* explains a much
519 larger variance explained relative to depth alone or with traditional methods. This demonstrates

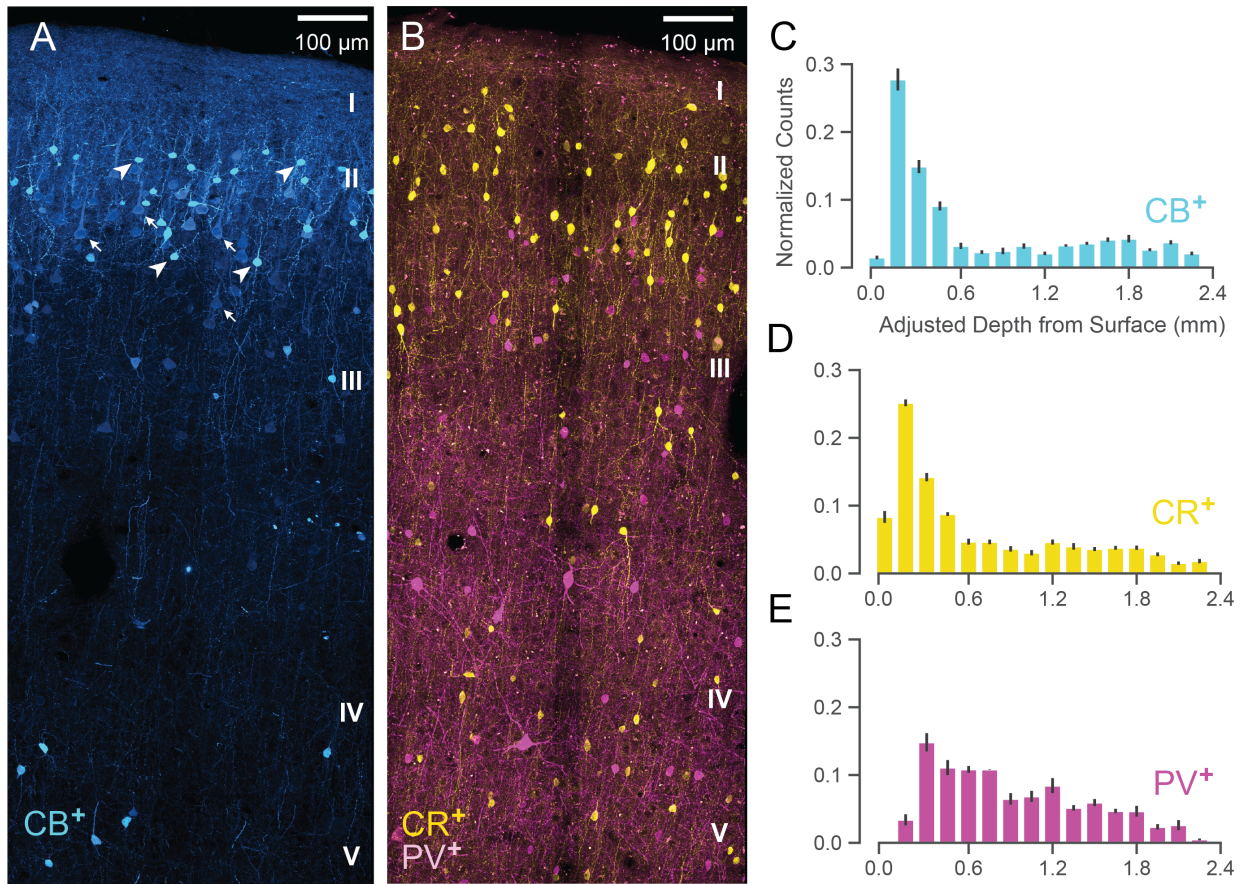


Figure 8: Anatomical labeling of three inhibitory interneuron types in PMd: (A) Maximum intensity projection of immunohistological staining of rostral PMd CB⁺ interneurons in blue. Note the many weakly-positive excitatory pyramidal neurons (arrows) in contrast to the strongly-positive interneurons (arrowheads). Only the interneurons were considered in stereological counting. In addition, only around first 1.5 mm of tissue is shown (top of layer V) but the full tissue area was counted down to 2.4 mm. Layer IV exists as a thin layer in this area. Layer divisions were estimated based on depth and referencing Arikuni *et al.*⁷⁸. **(B)** Maximum intensity projection of immunohistological staining of PMd CR⁺ and PV⁺ interneurons in yellow and fuschia respectively. The same depth of tissue and layer delineations were used as in (A). **(C, D, E)** Stereological manual counts⁸² (mean ± S.D.) of CB⁺, CR⁺, PV⁺ cells in PMd respectively. Counts were collected from six specimens and normalized at each depth.

520 that *WaveMAP* clusters come closer to cell types than previous efforts and are not artifacts of
521 layer-dependent decision-related inputs. That is, both the cortical layer in which a cell type is
522 found as well *WaveMAP* cluster membership contributes to the variability in decision-related
523 responses. Furthermore, they outperform GMM clusters as regressors of a functional property
524 associated with cell types. These results further highlight the power of *WaveMAP* to separate out
525 putative cell types and help better understand decision-making circuits.

526 Discussion

527 Our goal in this study was to further understand the relationship between waveform shape and the
528 physiological, functional, and laminar distribution diversity of cell populations in dorsal premotor
529 cortex during perceptual decision-making. Our approach was to develop a new method *WaveMAP*
530 that combines a recently developed non-linear dimensionality reduction technique (UMAP) with
531 graph clustering (Louvain community detection) to uncover hidden diversity in extracellular
532 waveforms. We found this approach not only replicated previous studies by distinguishing between
533 narrow- and broad-spiking neurons but did so in a way that 1) revealed additional diversity, and
534 2) obviated the need to examine particular waveform features. In this way, our results demonstrate
535 how traditional feature-based methods obscure biological detail, which is more faithfully revealed
536 by our *WaveMAP* method. Furthermore, through interpretable machine learning, we show our
537 approach not only leverages many of the features already established as important in the literature
538 but expands upon them in a more nuanced manner—all with minimal supervision or stipulation of
539 priors. Given the successes of standard UMAP across a variety of data domains, we envision that
540 *WaveMAP* will provide improvements in these areas with its additions of graph-based clustering and
541 interpretable machine learning. These diverse applications include computational ethology^{41,42,84},
542 analyzing multi-scale population structure⁸⁵, and even metascientific analyses of the literature⁸⁶.

543 *Advantages of WaveMAP over traditional methods*

544 At the core of *WaveMAP* is UMAP which has important advantages over other non-linear
545 dimensionality reduction methods that have been applied in this context, such as t-distributed
546 stochastic neighborhood embedding (t-SNE)⁴⁴. While providing visualization, these methods are
547 difficult to cluster upon because they return a different mapping on every initialization. UMAP, in
548 contrast, is a manifold learning method that projects efficiently into arbitrary dimensionalities while
549 also returning an invertible transform. That is, we can pass new data into the projected space
550 without having to recompute the mapping. This property provides two advantages over other
551 dimensionality reduction approaches: First, it allows exploration of any region of the manifold.
552 This yields an intuitive understanding of how UMAP non-linearly transforms the data which might
553 be related to underlying biological phenomena. Thus, UMAP allows *WaveMAP* to go beyond
554 a “discriminative model” typical of other clustering techniques and function as a “generative
555 model” with which to make predictions. Second, it enables cross-validation of a classifier trained
556 on cluster labels. This is impossible with methods that don’t return an invertible transform. To
557 cross-validate unsupervised methods, unprocessed test data must be passed into the transform
558 computed *only* on training data⁸⁷. This is only possible if a proper transform is admitted by the
559 method of dimensionality reduction as in UMAP.

560 A final advantage of UMAP is that it inherently allows for not just unsupervised but supervised and
561 semi-supervised learning whereas other methods do not⁸⁸. This key difference enables “transductive

562 inference” which is making predictions on unlabeled test points based upon information gleaned
563 from labeled training points. This opens up a diverse number of novel applications in neuroscience
564 through informing the manifold learning process with biological ground truths (in what is called
565 “metric learning”)^{89,90}. Experimentalists could theoretically pass biological ground truths (e.g., cell
566 type information derived using opto- or pharmaco-genetic methods *in vivo* or transcriptomics post-
567 experiment) to *WaveMAP* as training labels. This would in effect “teach” *WaveMAP* to produce
568 a manifold that more closely hews to true underlying diversity. For instance, if experimentalists
569 “opto-tag” neurons of a particular cell type^{44,91–94}, this information can be passed along with the
570 extracellular waveform to *WaveMAP* which would, in a semi-supervised manner, learn manifolds
571 better aligned to biological truth. A learned manifold could be used in future experiments to
572 identify cell types in real-time without opto-tagging. This could be done by projecting the averaged
573 waveforms found within an experiment into the learned *WaveMAP* manifold. This method would
574 be especially useful in a scenario in which the number of electrodes exceeds the number of
575 channels available to record from simultaneously and not all cell types are of equal interest to
576 record (e.g., Neuropixels probes which have 960 electrodes but simultaneously record from only
577 384 channels^{95,96}). We believe this is a rich area that can be explored in future work.

578 We use a fully-unsupervised method for separating and clustering waveform classes associated with
579 distinct laminar distributions and functional properties in a decision-making task. One concern
580 with fully unsupervised methods is that the features used for separation are unclear. However,
581 by applying an interpretable machine learning method^{61,63}, we showed that our unsupervised
582 methods utilized many of the same waveform features derived by hand in previous work but
583 did so in a single unifying framework. Our interpretable machine learning approach shows how
584 each waveform feature delineates certain waveform clusters at the expense of others and—more
585 importantly—shows how they can be optimally combined to reveal the full diversity of waveform
586 shapes.

587 Our novel approach of using non-linear dimensionality reduction with graph clustering on the
588 population of extracellular action potentials compared to specified waveform features has parallels
589 with the evolution of new approaches for the analysis of neuronal firing rates in relevant brain
590 areas^{97–101}. Classically, the approach to analyzing firing rates involved in cognition was to develop
591 simple metrics that separated neurons recorded in relevant brain areas. For instance, tuning
592 is used to separate neurons in the motor¹⁰² and visual cortex¹⁰³. Similarly, visuomotor indices
593 that categorize neurons along a visual to motor continuum are used to understand firing rates
594 during various tasks in the frontal eye fields¹⁰⁴ and premotor cortex¹. However, these specified
595 features quash other aspects of a firing rate profile in favor of focusing on only a few other
596 aspects. New approaches to analyze firing rates use dimensionality reduction techniques such as
597 principal component analysis^{97,98,105}, tensor component analysis¹⁰⁶, demixed principal component
598 analysis¹⁰⁷, targeted dimensionality reduction⁹⁹, and autoencoder neural networks¹⁰⁸. These
599 methods have provided insight into heterogeneous neural activity patterns in many brain areas
600 without the need for specified features like tuning or a visuomotor index. Our study adopts a similar
601 approach for analysis of extracellular waveforms by examining whole populations and strongly
602 suggests that non-linear dimensionality reduction methods applied to the entire extracellular
603 waveform are better than using hand-derived waveform features such as trough to peak duration,
604 repolarization time, spike width and other metrics. This progression also follows similar trends in
605 machine learning towards data-driven approaches.

606 *Waveform cluster shapes are unlikely to arise from electrode placement*

607 It is a possibility is that the diversity of waveforms we observe is just an artifact of electrode
608 placement relative to the site of discharge. This supposes that waveform shape changes with
609 with respect to the distance between the neuron and electrode. This is unlikely because both *in*
610 *vitro* studies⁵⁶ and computational simulations¹⁰⁹ show distance from the soma induces changes
611 in amplitude but not shape. We controlled for this variation in amplitude by normalizing spikes
612 during preprocessing.

613 It is another possibility that different neuronal structures present different waveform shapes. Put
614 another way, the diversity we see could be due to recording from different morphological structures
615 (dendrites, soma, or axons) rather than different cell types. While it is true that there are some cell
616 structures associated with different waveform shapes (such as triphasic waveforms near neurites,
617 especially axons^{29,30,35,56}), highly-controlled *in vitro* studies show that a large majority of EAP's
618 are from somata (86%)⁵⁶. In concordance with this, we only observed one cluster (④, 6% of all
619 EAP's) with a triphasic shape. These waveforms were only found in deep layers where myelination
620 is prevalent. Thus we believe that almost all of our waveforms come from somata, with the
621 possible exclusion of ④.

622 It is also unlikely that the waveform diversity seen is due to the location of the electrode relative
623 to the neuron given the observed differences in firing rate (Fig. 5), functional role (Fig. 6), and
624 laminarity (Fig. 7) associated with each *WaveMAP* cluster.

625 Given that electrode location has little effect on waveform shape, we might then ask what about
626 a neuron's waveform shape, in terms of cellular physiology, is captured by *WaveMAP*? We
627 propose that the space found by UMAP-1 and UMAP-2 sensibly covaries according to documented
628 properties of K⁺ ion channel dynamics. As UMAP-1 increases, we observe a smooth transition of
629 the inflection of the repolarization slope from negative to positive (slow to fast repolarization rate).
630 Said differently, the post-hyperpolarization peak becomes sharper as we increase in the UMAP-1
631 direction. These observations are consistent with the same gradual change in intracellular AP
632 repolarization slope facilitated by the kinetics of the fast voltage-gated Kv3 potassium-channel in an
633 activity-dependent manner¹¹⁰. These channels are necessary for sustained high-frequency firing¹¹¹.
634 In the UMAP-2 direction, there is a smooth decrease in the width of the post-hyperpolarization
635 peak and this direction roughly traverses from broad- to narrow-spiking to triphasic waveforms.
636 This gradual change too has been noted as being associated with the kinetics of the Kv3 potassium-
637 channel: blocking this channel in a dose-dependent manner with tetraethylammonium induces a
638 gradual widening of post-hyperpolarization peak width^{112,113}. Both of these changes in intracellular
639 waveform shape likely have a strong effect on the shape of extracellular waveforms¹¹⁴.

640 *Better parcellation of waveform variability leads to biological insight*

641 Our results show a greater proportion of narrow- (putatively inhibitory) vs. broad-spiking (putatively
642 excitatory) neurons (69% vs. 31% respectively); this appears inconsistent with anatomical studies.
643 These studies demonstrate, through direct labeling of cell type, that in the macaque cortex,
644 65-80% of neurons are excitatory while 20-35% are inhibitory. We are not the only study to report
645 this puzzling result: Kaufman and colleagues⁵³ note a proportion of "intermediary" cells that
646 could be classified as either narrow- or broad-spiking. Onorato and colleagues²⁴ also report greater
647 numbers of narrow-spiking compared to broad-spiking neurons in monkey V1. Thus, care must be

648 taken when attempting links between spike waveform metrics and cell type. A resolution to this
649 discrepancy is to rethink equating narrow-spiking to inhibitory cells and broad-spiking to excitatory
650 cells. In the monkey, not all inhibitory neurons have narrow spikes^{58,115} and not all excitatory
651 neurons have broad spikes^{20,21,24}. Anatomical studies show that a substantial number of excitatory
652 neurons in the monkey motor and visual cortices express the Kv3.1b potassium channel which
653 is known to confer neurons with the ability to produce action potentials of narrow spike width
654 and high firing rate^{22,25,116,117}. Furthermore, *in vivo* studies show narrow-spiking neurons can be
655 excitatory in motor^{20,21} and visual cortices^{22,24,25,116} of the macaque.

656 We therefore believe prior studies have underexplored the diversity of classes accessed by their
657 physiological recordings. Histograms of peak width (and other specified features) across literature
658 are often not cleanly bimodal^{27,31,53}, and the relative proportions of narrow vs. broad is often
659 dependent on the cutoff chosen for narrow- vs. broad-spiking neurons which span a wide range²⁰.
660 Analyses like ours which look at entire waveforms—rather than a few specified features—extract
661 this diversity from extracellular recordings whereas specified features mix waveform classes.

662 We also find that many narrow-spiking subtypes in PMd signal choice earlier than broad-spiking
663 neurons in our decision-making task (Fig. 6F). These observations are consistent with another
664 study of PMd in monkeys in reach target selection and movement production⁷⁵. In this study,
665 narrow-spiking neurons signaled the selected target 25 ms earlier than broad-spiking neurons. Our
666 results are also consistent with other studies of narrow- vs. broad-spiking neurons in the frontal
667 eye fields (FEF)⁹ and inferior temporal area (IT)⁷⁶ during decision-making. In these studies,
668 narrow-spiking neurons had higher firing rates before movement onset compared to broad-spiking
669 neurons—a result consistent with our observations for some “narrow-spiking” PMd neurons. Our
670 analyses recapitulate these results and provide additional insights into how different narrow-spiking
671 cell types correlate with decisions. We reproduce the result that narrow-spiking cells, as a whole,
672 have a lower discrimination time than broad-spiking cells but in addition we show that certain
673 narrow-spiking cells respond as slowly as broad-spiking cells (② and ④; Fig. 6F). This lends
674 further evidence to our theory that ② and ④ are likely narrow-spiking excitatory cells.

675 We observed further heterogeneity in narrow-spiking cells with different clusters occupying different
676 discrete areas of cortex. One of our narrow-spiking clusters (cluster ②) was localized to more
677 superficial layers (Fig. 7B) and had functional properties—low baseline firing rate and longer
678 discrimination times—which are thought to be more closely aligned to properties of excitatory
679 neurons. Furthermore, it is known that there are large populations of narrow-spiking excitatory
680 cells in other areas (PMd is as of yet unstudied in this manner)²¹. Similarly, another narrow-spiking
681 cluster (④) exhibited physiological and functional properties to ② (all comparisons not significant
682 in Fig. 5C, D, and E or Fig. 6E and F) but with a distinct laminar distribution Fig. 7B and highly
683 triphasic waveform shape Fig. S7B. In contrast to ② which concentrated in layer III, ④ was
684 restricted to deep layers which are known to contain large pyramidal cells, often corticospinal.
685 This is also where large-amplitude triphasic waveforms appear most often^{29,30,35}. We suspect that
686 cluster ② is similar to the one reported in a recent study of the primary visual cortex²⁴ while ④
687 likely contains axons or initial segments of deep narrow-spiking pyramidal cells^{20,21,117}.

688 *Neuropixels and optogenetics can provide better insight into candidate cell classes in the monkey*

689 Our recordings here were performed with 16 channel U-probes and provided reasonable estimates of

690 laminar organization for these different putative cell classes. Use of high-density electrophysiological
691 methods providing higher electrode counts perpendicular to the cortical surface would provide
692 further insight into the laminar organization of different cell types⁹⁶. High-density recording would
693 allow us to perform *WaveMAP* in an additional dimension (across multiple electrodes) to increase
694 confidence in identified cell classes and localization of signal to somata⁴⁴. Sensitive electrodes
695 providing spatial access to neural activity⁹⁶ can also improve our understanding of how these cell
696 classes are organized both parallel and perpendicular to cortical surface^{118,119}. This would allow
697 for the identification of “me-types” through electromorphology^{37,45}.

698 Another powerful tool that has been leveraged in the study of cell types during behavior is
699 optogenetics^{3,5,6}. Although in its infancy relative to its use in the mouse, optogenetics in monkeys
700 offers direct interrogation of cell types. Future studies that combine opto-tagging in monkeys with
701 simultaneous application of *WaveMAP* and other waveform classification methods will allow us to
702 more precisely link our putative cell classes to function¹²⁰ and do so *in vivo*. NHP optogenetics is
703 slowly advancing and efforts in many research groups around the world are producing new methods
704 for *in vivo* optogenetics¹²¹. We expect future experiments using the promising new mDlx¹²² and
705 h56d¹²³ promoter sequences to selectively opto-tag inhibitory neurons or PV⁺ neurons directly¹²⁴
706 will greatly benefit validation of these derived cell classes.

707 **Author Contributions**

708 CC trained monkeys and recorded in PMd using multi-contact electrodes under the mentorship of
709 KVS. EKL developed the UMAP and Louvain Clustering method (*WaveMAP*) with interpretable
710 machine learning under the mentorship of CC. HB, EKL, and CC expanded the previous GMM
711 approach for clustering based on waveform features. EKL and CC analyzed *WaveMAP* clusters
712 for their properties. AT and MM performed immunohistochemistry experiments and imaging. SA,
713 EKL, and AT performed stereological cell counting. EKL and CC wrote initial drafts of the paper.
714 All authors refined further drafts contributing analyses, insights, and writing.

715 **Acknowledgments**

716 We thank M. Noichl and L. McInnes for their suggestions on graph community detection and
717 UMAP visualization. We also thank Dr. Kathleen Rockland, Dr. Jennifer Luebke and Dr.
718 Jonathan Kao for detailed comments on the previous versions of the manuscript.

719 CC was supported by a NIH/NINDS R00 award R00NS092972, the Moorman-Simon Interdis-
720 ciplinary Career Development Professorship from Boston University, the Whitehall foundation,
721 and the Young Investigator Award from the Brain and Behavior Research Foundation. MM was
722 supported by a NIH-NIMH K99/R00MH101234 grant and a NIH-NIMH R01MH116008 grant.
723 KVS was supported by the following awards: NIH Director's Pioneer Award 8DP1HD075623,
724 NIDCD R01-DC014034, NIDCD U01-DC017844, NINDS UH2-NS095548, NINDS UO1-NS098968,
725 DARPA-BTO 'REPAIR' Award N66001-10-C-2010, DARPA-BTO 'NeuroFAST' award W911NF-
726 14-2-0013, Simons Foundation Collaboration on the Global Brain awards 325380 and 543045,
727 Office of Naval Research award N000141812158, Larry and Pamela Garlick, Wu Tsai Neuro-
728 sciences Institute at Stanford, the Hong Seh and Vivian W. M. Lim endowed professorship and
729 the Howard Hughes Medical Institute. The funders had no role in study design, data collection
730 and interpretation, or the decision to submit the work for publication.

731 **Declaration of interests**

732 K.V.S. consults for Neuralink Corp. and CTRL-Labs Inc. (part of Facebook Reality Labs) and is
733 on the scientific advisory boards of MIND-X Inc., Inscopix Inc. and Heal Inc. All other authors
734 have no competing interests. These companies provided no funding and had no role in study
735 design, data collection and interpretation, or the decision to submit the work for publication.

736 **Methods (3,000 words)**

737 *Subjects and Surgery*

738 Our experiments were conducted using two adult male macaque monkeys (*Macaca mulatta*;
739 monkey T, 7 years, 14 kg; O, 11 years, 15.5 kg) that were trained to reach to visual targets for a
740 juice reward. Our Monkeys were housed in a social vivarium with a normal day/night cycle. All
741 experimental protocols were approved by the Stanford University Institutional Animal Care and Use
742 Committee (IACUC). After initial training to come out of the cage and sit comfortably in a chair,
743 monkeys underwent sterile surgery for implantation of head restraint holders (Crist Instruments,
744 cylindrical head holder) and standard recording cylinders (Crist Instruments, Hagerstown, MD).

745 We placed our cylinders over caudal PMd (+16, 15 stereotaxic coordinates) and surface normal to
746 the cortex. We covered the skull within the cylinder with a thin layer of dental acrylic/PALACOS
747 bone cement.

748 *Apparatus*

749 Monkeys sat in a customized chair (Crist Instruments, Snyder Chair) with the head restrained via
750 the surgical implant. The arm not used for reaching was loosely restrained using a tube and a
751 cloth sling. Experiments were controlled and data were collected under a custom computer control
752 system (xPC target and Psychtoolbox-3¹²⁵). Visual stimuli were displayed on an Acer HN2741
753 computer screen placed approximately 30 cm from the monkey and a photodetector (Thorlabs
754 PD360A) was used to record the onset of the visual stimulus at a 1 ms resolution. Every session,
755 we taped a small infrared reflective bead (11.5 mm, NDI Digital passive spheres) 1 cm from the
756 tip of the middle digit of the right hand (left hand, monkey O). The position of this bead was
757 tracked optically in the infrared (60 Hz, 0.35 mm root mean square accuracy; Polaris system;
758 Northern Digital).

759 Eye position was tracked with an overhead infrared camera (estimated accuracy of 1°, Iscan,
760 Burlington, MA). To get a stable eye image for the overhead infrared camera, an infrared dichroic
761 mirror was positioned at a 45° angle (facing upward) immediately in front of the nose. This mirror
762 reflected the image of the eye in the infrared range while letting visible light pass through. A visor
763 placed around the chair prevented the monkey from touching the infrared mirror, the juice tube,
764 or bringing the bead to his mouth.

765 *Behavioral training*

766 Our animals were trained using the following operant conditioning protocol. First, the animal
767 was rewarded for arm movements toward the screen and learnt to take pieces of fruit on the
768 screen. Once the animal acquired the association between reaching and reward, the animal was
769 conditioned to reach and touch a target for a juice reward. The position, as well as the color of this
770 target, was then randomized as the monkey learned to touch targets of various colors at different
771 locations on the screen. We then used a design in which the monkey first held the central hold for
772 a brief period, and then a checkerboard cue, which was nearly 100% red or 100% green, appeared
773 for 400–600 ms and finally the two targets appeared. The monkey received a reward for making a
774 reach to the color of the target that matched the checkerboard cue. Two-target “Decision” blocks
775 were interleaved with single target blocks to reinforce the association between checkerboard color
776 and the correct target. After two weeks of training with this interleaved paradigm, the animal
777 reliably reached to the target matching the color of the central checkerboard cue. We switched
778 the paradigm around by adopting a design in which the targets appeared before the checkerboard
779 cue onset. We initially trained on holding periods (where the monkeys view targets) from 300 to
780 1800 ms. We trained the animal to maintain the hold on the center until the checkerboard cue
781 appeared by providing small amounts of juice at rough time intervals. When the animal reliably
782 avoided breaking central hold during the holding period, we stopped providing small amounts of
783 juice. After the animal learned to stay still during the target viewing period, we introduced more
784 difficult checkerboard cues to the animal while reducing the maximal holding period to 900 ms.
785 We then trained the animal to discriminate the checkerboard as accurately and as fast as possible
786 while discouraging impulsivity by adopting timeouts.

787 *Electrophysiological recordings*

788 Known response properties of PMd and M1, the neural responses to muscle palpation, and
789 stereotactic coordinates were used as our guides for electrophysiological recordings. Our chambers
790 were placed surface normal to cortex to align with the skull of the monkey and recordings were
791 performed perpendicular to the surface of the brain. Recordings were made anterior to the central
792 sulcus, lateral to the spur of the arcuate sulcus, and lateral to the precentral dimple. For both
793 monkeys, we were able to identify the upper and lower arm representation by repeated palpation
794 at a large number of sites to identify muscle groups associated with the sites. Recordings were
795 performed in the PMd and M1 contralateral to the arm used by the monkey. Monkey T used his
796 right arm (O used his left arm) to perform tasks.

797 A subset of the electrophysiological recordings were performed using traditional single electrode
798 recording techniques. Briefly, we made small burr holes through the PALACOS/acrylic using
799 handheld drills. We then used a Narishige drive with a blunt guide tube placed in firm contact
800 with the dura. Recordings were obtained using FHC electrodes to penetrate the overlying dura.
801 Every effort was made to isolate single units during the recordings with FHC electrodes by online
802 monitoring and seeking out well-isolated signals.

803 We performed linear multi-contact electrode (U-probe) recordings in the same manner as sin-
804 gle electrode recordings with some minor modifications. We used a slightly sharpened guide
805 tube to provide more purchase on dura. We also periodically scraped away, under ketamine-
806 dexmetomidine anesthesia, any overlying tissue on the dura. Both these modifications greatly
807 facilitated penetration of the U-probe. We typically penetrated the brain at very slow rates (~2–5
808 $\mu\text{m/s}$). Once we felt we had a reasonable sample population of neurons, potentially spanning
809 different cortical layers, we stopped and waited for 45–60 min for the neuronal responses to
810 stabilize. The experiments then progressed as usual. We used 180- μm thick 16-electrode U-probes
811 with an inter-electrode spacing of 150 μm ; electrode contacts were 100 $\text{K}\Omega$ in impedance.

812 We attempted to minimize the variability in U-probe placement on a session-by-session basis.
813 Our approach was to place the U-probe so that the most superficial electrodes (electrodes 1,
814 2 on the 16 channel probe) were in layer I and able to record multi-unit spiking activity. Any
815 further movement of the electrode upwards resulted in the spiking activity disappearing and a
816 change in the overall activity pattern of the electrode (suppression of overall LFP amplitudes).
817 Similarly, driving the electrodes deeper resulted in multiphasic extracellular waveforms and also a
818 change in auditory markers which were characterized by decreases in overall signal intensity and
819 frequency content. Both markers suggested that the electrode entered white matter. Recording
820 yields and electrode placement were in general much better in monkey T (average of ~16 units
821 per session) than monkey O (average of ~9 units per session). We utilized these physiological
822 markers as a guide to place electrodes and thus minimize variability in electrode placement on a
823 session-by-session basis. Importantly, the variability in placement would act against our findings
824 of depth-related differences shown in [Fig. 7](#).

825 *Preprocessing of single-unit recordings*

826 We obtained 996 extracellularly recorded single units (778 units recorded with the U-probe) from
827 PMd across two monkeys (450 from Monkey O and 546 from Monkey T). Of these, we identified
828 801 units whose ISI violations (refractory period ≤ 1.5 ms) $\leq 1.5\%$ ¹. Our waveforms were filtered

829 with a 4th-order 250 Hz high-pass Butterworth filter. The waveforms for each of the units were
830 extracted for a duration of 1.6 ms with a pre-trough period of 0.4 ms, sampled at 30 kHz.

831 *Alignment and normalization of waveforms*

832 In order to calculate the mean waveform for each single unit, we upsampled individual waveforms
833 calculated over different trials by a factor of 10 and aligned them based on the method proposed
834 in³². For each waveform, we calculated its upswing slope (slope between trough to peak) and the
835 downswing slope (slope to the trough) and re-aligned to the midpoint of the slope that exceeded
836 the other by a factor of 1.5. Following this alignment, we chose the best set of waveforms for
837 calculating the mean as those that satisfied the criteria (1) less the 2 standard deviations (S.D.)
838 from the mean at each point and (2) average deviation from the mean across time is less than
839 0.4³². The final set of waveforms for each unit was averaged and downsampled to 48 time points.
840 Upon visual inspection, we then identified 761 units (625 single units with 490 U-probe recorded
841 units) whose average waveforms qualified the criteria of exhibiting a minimum of two phases
842 with a first trough trajectory and the remaining waveforms, unless stated otherwise here, were
843 removed from the analysis. We excluded positive-spiking waveforms because of their association
844 with axons³⁵. Finally, we normalized the waveforms by dividing the extreme value of the amplitude
845 such that the maximum deviation is ± 1 unit⁵⁰.

846 It is important to note that the preprocessing we use, individual mean subtraction and ± 1 unit
847 normalization, operates independently of the data. Using another commonly used preprocessing
848 normalization, normalization to trough depth, we obtained extremely similar results. We found
849 ± 1 unit trough to peak normalization had virtually the same number of clusters as normalization
850 to trough (8.29 ± 0.84 vs. 8.16 ± 0.65 clusters, mean \pm S.D.; Fig. S5A and C). Furthermore,
851 both normalizations picked out the same structure (Fig. S5B and D; the normalization to trough
852 did have a 9th cluster splitting off of ⑤ but this was something also seen with ± 1 unit trough to
853 peak normalization in certain data subsets as well.

854 *WaveMAP: UMAP and Louvain clustering*

855 The normalized extracellular waveforms were passed into the Python package `umap 0.4.0rc3`⁴⁸
856 with the parameters shown in S1. The `n_neighbors` value was increased to 20 to induce more
857 emphasis on global structure. UMAP utilizes a stochastic k -nearest neighbor search to establish
858 the graph and stochastic gradient descent to arrive at the embedding thus it produces similar but
859 different embeddings in the projected space. For reproducibility reasons, the `random_state` was
860 fixed in the algorithm and in `numpy`. The choice of random seed only impacted the projection and
861 not the clustering (Fig. S2A). From here, the graph provided by `umap.graph_` was passed into
862 the Louvain community detection algorithm to generate the clustering seen in Fig. 2A. For details
863 of the UMAP algorithm, see Supplementary Information.

864 Graph networks are often hierarchical and it has been recommended that the Louvain resolution
865 parameter be chosen to elicit the phenomenon of interest^{126,127}. To select the resolution parameter
866 t , we chose a value that best balanced a maximization of modularity score (a measure of the
867 ratio between connections within a cluster vs. incoming from outside of it; see Supplementary
868 Information) while still returning an interpretable number of clusters given our data size. We
869 selected a resolution parameter at the “elbow” of the number of clusters and at the top of
870 the modularity score plateau of maximum values (yellow marker on Fig. 2B). Choosing a lower

871 resolution parameter resulted in communities of size lower than 20 units which precluded statistical
872 comparison against other communities [Fig. S2A](#). These scores were calculated over 25 random
873 UMAP instantiations of the full dataset. For algorithmic details of Louvain clustering, see
874 [Supplementary Information](#).

875 To validate that our parameter selection was stable and produced the same number of clusters
876 reliably, we used a bootstrap and applied the *WaveMAP* procedure to random subsets of the full
877 data set [Fig. S2B](#). We obtained 100 random samples from 10% to 100% of the full data set in
878 10% increments while simultaneously choosing a different random seed for the UMAP algorithm
879 each time. Thus the variances due to sampling and random instantiation are compounded and
880 shown together in [Fig. S2B](#).

881 Ensemble clustering for graphs (ECG)^{128,129} was used to validate the clusters found in [Fig. 2A](#)
882 (see [Fig. S2C](#)). We added the algorithm ([https://github.com/ftheberge/Ensemble-Clustering-
883 for-Graphs](https://github.com/ftheberge/Ensemble-Clustering-for-Graphs)) into the `python-igraph` package¹³⁰ and passed UMAP graphs into it directly. We
884 set the number of partitions k to be 10 to produce the plot in [Fig. S2C](#). This algorithm uses k
885 different randomized instantiations of the clusters in the graph followed by one round of Louvain
886 clustering ([Fig. S1B](#). Each of these k (called level-1 partitions since one round of Louvain was
887 performed) graphs are then combined as a single graph such that when edges co-occur between
888 nodes in one of the k graphs, it is more heavily weighted. This ensembling of several graphs via
889 the weight function $W_{\mathcal{P}}$ (see [Supplemental Methods](#)) yields the final ECG graph.

890 *Gradient-boosted random forest classifier*

891 We then trained a gradient boosted random forest classifier in `xgboost 1.0.2`¹³¹ with five-fold
892 cross-validation on the *WaveMAP* cluster labels in [Fig. 2A](#). A 30% - 70% test-train split was used.
893 Optimal hyperparameters were obtained after a grid search using `scikit-learn's GridSearchCV`
894 function with parameters in [Table S1](#). The default objective function `binary:logistic` was also
895 used. The percent accuracy for each cluster against all others is plotted as a confusion matrix in
896 [Fig. 2C](#).

897 The same procedure was used when training on the GMM labels found in [Fig. 3D](#) and for the eight
898 cluster GMM labels in [Fig. S3B](#). Each of these classifiers also separately underwent hyperparameter
899 tuning using `scikit-learn's GridSearchCV` function as well with final hyperparameters shown
900 in [S1](#).

901 It is important to note that cross-validation was done after the cluster labels were generated which
902 results in data leakage^{87,132} hurting out-of-dataset performance and thus classifier performance
903 only demonstrates UMAP's ability to sensibly separate waveforms within-dataset relative to
904 traditional GMM methods ([Fig. 3D](#)). [Fig. 2C](#) is not meant to show how such a classifier would
905 perform out-of-dataset.

906 *Specified waveform shape features*

907 To compute specified features for each normalized waveforms ([Fig. 3A](#)), we first up-sampled
908 the waveforms from 48 to 480 time points using a cubic spline interpolation method. We then
909 used this up-sampled waveform to compute three separate features: trough to peak duration, AP
910 width, and peak ratio. Trough to peak is the time from the bottom of the depolarization trough
911 (global minimum) to the post-hyperpolarization peak (subsequent local maximum). AP width was

912 calculated as the width of the depolarization trough at the full-width half-minimum point. Both
913 these measures were found using the `mwave` function from the MLIB 1.7.0.0 toolbox. Peak ratio
914 was the ratio of heights (above baseline) between the pre-hyperpolarization (maximum before
915 trough) and the post-hyperpolarization peak (maximum after trough).

916 *Gaussian mixture model clustering*

917 Using the specified feature values (trough to peak, AP width, and peak ratio), the normalized
918 waveforms were clustered in the three-dimensional feature space using a Gaussian mixture model
919 (GMM) with hard-assignment (each data point belongs to one cluster) through MATLAB's
920 `fitgmdist` function across 50 replicates (Fig. 3B). Each replicate is a different random instantiation
921 of the GMM algorithm and the model with the largest log likelihood is chosen.

The Bayesian information criterion (BIC) was used to determine the optimal cluster number and is defined as

$$\text{BIC} = -2 \ln P(X|\theta) + K \ln(n) \quad (1)$$

922 where the first term $-2 \ln P(X|\theta)$ is the negative log likelihood function i.e. the conditional
923 probability of observing the sample X given a vector of parameters θ . In the particular case of
924 GMM, the function $P(X)$ is a sum of multivariate Gaussians. The second term $K \ln(n)$ is a
925 penalty on model complexity (a regularization) thus capturing the idea that “simple is better”
926 and ultimately constraining the number of Gaussians used to fit the data. It is a function of the
927 number of free model parameters K and the size of the dataset n .

Assuming we have N_f features and N_c clusters we can calculate K using the following framework. For each Gaussian, the total number of parameters is N_f means and $N_f(N_f + 1)/2$ covariance parameters. Another free parameter that is learned is the weight for each Gaussian that sums up to 1, leaving us with $N_c - 1$ unique weights. Thus the K which is the effective number of parameters is,

$$K = N_c \left(N_f + \frac{N_f(N_f + 1)}{2} \right) + N_c - 1 \quad (2)$$

928 The “best” model in a BIC-sense will have the set of parameters θ maximizing the likelihood
929 function (thus minimizing the negative log likelihood) for a given model or model family—a
930 number of multivariate Gaussians in a three-dimensional feature space in this case. To arrive at
931 the parameters best approximating the Gaussian distribution giving rise to the data (Maximum
932 Likelihood Estimation or MLE), the Expectation-Maximization (EM) algorithm was used. The
933 optimal cluster number was selected as the lowest number of clusters between 1 and 10 at which
934 the change in BIC was minimal (at the “elbow” in Fig. 3C).

935 *Interpretable machine learning: UMAP inverse transform and SHAP*

936 To facilitate interpretability, we used the invertibility of the UMAP transform (which itself is based
937 on Delauney triangulation) to generate test waveforms tiling the manifold of the projected space
938 Fig. 4A. 100 evenly-spaced test coordinates were generated spanning the manifold and passed
939 backwards through the UMAP transform using `umap`'s built-in `inverse_transform` function.
940 The waveform generated at each test point is shown color-coded to the nearest cluster color or in
941 gray if the distance exceeds 0.5 units in UMAP space.

942 Using the package `shap` (<https://github.com/slundberg/shap>), SHAP values were calculated for
943 the classifier trained on *WaveMAP* identified clusters. The trained XGBoost model was passed
944 directly into the model-agnostic `shap.TreeExplainer`⁶² which then calculated SHAP values for
945 all waveform time points (features). TreeExplainer assigned SHAP values for every time point
946 class-by-class and these were used to generate the class-specific SHAP plots (Fig. 4C). The SHAP
947 values for each time point, across classes, was summed to generate the overall SHAP values for
948 each time point (Fig. 4B).

949 *Choice-selective signal*

950 We use an approach developed by Meister and Huk to estimate the choice-selective signal¹³³. We
951 chose such an approach because decision-related activity of PMd neurons is not simply increases
952 and decreases in firing rate and often shows considerable temporal modulation. We estimated
953 for each neuron a choice-selective signal on a time point-by-time point basis as absolute value of
954 the firing rate difference between left and right choice trials ($|\text{left} - \text{right}|$) or equivalently
955 $|\text{PREF} - \text{NONPREF}|$. We use this choice-selective signal to understand choice-related dynamics and
956 estimate discrimination time.

957 *Discrimination time*

958 We identified the discrimination time, that is the time at which the neuron demonstrated significant
959 choice selectivity, on a neuron-by-neuron basis. We compared the choice-selective signal at each
960 point to the 95th percentile of the bootstrap estimates of baseline choice-selective signal (i.e.,
961 before checkerboard stimulus onset). We enforced the condition that the choice-selective signal
962 should be significantly different from the baseline for at least 25 ms after this first identified time
963 to be included as an estimate of a time of significant choice selectivity for that neuron. Using
964 longer windows provided very similar results.

965 *Choice-related dynamics*

966 To understand the dynamics of the choice-selectivity signal as a function of the unsigned checker-
967 board coherence, we performed the following analysis. As described above, we first estimated the
968 choice-selectivity signal in spikes/s for each neuron and each checkerboard coherence as shown
969 for example in Fig. 6A,B. We then estimated the slope of this choice-selectivity signal in the
970 175-325 millisecond period after checkerboard onset. Repeating this analysis for each unsigned
971 checkerboard coherence provided us with an estimate of the rate of change of the choice selectivity
972 signal (η) for all the coherences is spikes/s/s. Averaging over neurons for each cluster provided us
973 with the graphs in Fig. 6C,D. We then estimated the dependence of η on unsigned coherence by
974 regressing η and coherence to estimate how strong choice-selectivity signals in a particular cluster
975 are modulated by the checkerboard input and are summarized in Fig. 6E.

976 *Experimental subjects (anatomical data)*

977 Archived tissues were harvested from six young rhesus macaques of both sexes (9 ± 1.13 years,
978 *Macaca mulatta*). These subjects were close in age to the macaques used in the main study and
979 were part of part of a larger program of studies on aging and cognition led by Dr. Douglas Rosene.
980 These monkeys were obtained from the Yerkes National Primate Center and housed individually in
981 the Laboratory Animal Science Center at the Boston University School of Medicine; these facilities
982 are fully accredited by the Association for Assessment and Accreditation of Laboratory Animal

983 Care (AAALAC). Research was conducted in strict accordance with the guidelines of the National
984 Institutes of Health's Guide for the Care and Use of Laboratory Animals and Public Health Service
985 Policy on the Humane Care and Use of Laboratory Animals.

986 *Perfusion and fixation*

987 All brain tissue for histological studies was fixed and harvested using our well-established two-stage
988 perfusion protocol as described¹³⁴. After sedation with ketamine hydrochloride (10 mg/ml)
989 and deep anesthetization with sodium pentobarbital (to effect, 15 mg/kg i.v.), monkeys were
990 perfused through the ascending aorta with ice-cold Krebs–Henseleit buffer containing (in mM): 6.4
991 Na₂HPO₄, 1.4 NaH₂PO₄, 137.0 NaCl, 2.7 KCl, and 1.0 MgCl₂ at pH 7.4 (Sigma-Aldrich) followed
992 by fixation with 4% paraformaldehyde in 0.1M phosphate buffer (PB, pH 7.4, 37 °C). The fixed
993 brain sample was blocked, *in situ*, in the coronal plane, removed from the skull and cryoprotected
994 in a series of glycerol solutions, and flash frozen in 70 °C isopentane¹³⁵. The brain was cut on
995 a freezing microtome in the coronal plane at 30 µm and were systematically collected into 10
996 matched series and stored in cryoprotectant (15% glycerol, in 0.1M PB, pH 7.4) at -80 °C¹³⁶.

997 *Immunohistochemistry*

998 To assess the laminar distribution of interneurons we batch processed 30 µm coronal sections
999 through the rostral dorsal premotor cortex area (PMdr) from six specimens. Sections were
1000 immunolabelled for inhibitory neuronal subtypes based on their expression of calcium binding
1001 proteins, calbindin (CB), calretinin (CR), and parvalbumin (PV), which label non-overlapping
1002 populations in primates⁸¹. Free floating sections were first rinsed (3 × 10 min, 4 °C) in 0.01M
1003 phosphate-buffered saline (PBS) and incubated in 50 mM glycine for 1 hr at 4 °C. Sections were
1004 then rinsed in 0.01M PBS (3 × 10 min, 4 °C), and antigen retrieval was performed with 10 mM
1005 sodium citrate (pH 8.5) in a 60-70 °C water bath for 20 min. Sections were then rinsed in 0.01M
1006 PBS (3 × 10 min, 4 °C) and incubated in pre-block (0.01M PBS, 5% bovine serum albumin
1007 [BSA], 5% normal donkey serum [NDS], 0.2% Triton X-100) to reduce any non-specific binding
1008 of secondary antibodies. Primary antibodies (Fig. S2) were diluted in 0.1 M PB, 0.2% acetylated
1009 BSA (BSA-c), 1% NDS, 0.1% Triton X-100. To increase the penetration of the antibody, two
1010 microwave incubation sessions (2 × 10 min at 150 watts) using the Pelco Biowave Pro (Ted
1011 Pella), followed by a 48 hour incubation at 4 °C with gentle agitation. After rinsing (3 × 10
1012 min) in 0.01M PBS at 4 °C, sections were co-incubated with secondary antibodies diluted in
1013 incubation buffer (see S2), microwaved 2 × 10 min at 150 W, and placed at 4 °C for 24 hours
1014 with gentle agitation. Sections were then rinsed (3 × 10min) in 0.1M PB, mounted onto slides
1015 and coverslipped with prolong anti-fade gold mounting medium (ThermoFisher) and cured at
1016 room temperature in the dark.

1017 *Confocal microscopy and visualization of immunofluorescent labeling*

1018 Immunofluorescent labeling was imaged using a laser-scanning confocal microscope (Leica SPE)
1019 using 488 and 561 nm diode lasers. For each coronal section, two sets of tile scan images of
1020 a cortical column, ~200 µm wide and spanning, from pia to the white matter boundary, were
1021 obtained in the PMdr. This corresponded to the area 6FR in cytoarchitectural maps^{80,137,138} and
1022 area F7 in several functional maps^{139,140}. The two columns were spaced 200 µm apart. All images
1023 were acquired using a plain apochromat 40x/1.3 NA oil-immersion objective at a resolution of
1024 0.268 × 0.268 × 0.5 µm voxel size. The resulting image stacks were deconvolved and converted to
1025 8-bit images using AutoQuant (Media Cybernetics) to improve the signal to noise ratio¹³⁴.

1026 *Stereological cell counting*

1027 Due to its demonstrated ability in producing minimally-biased results, 3D stereologic cell counting⁸²
1028 was utilized to count parvalbumin- (PV⁺), calretinin- (CR⁺) and calbindin-positive (CB⁺) cells.
1029 Using the CellCounter plugin in Fiji¹⁴¹ on each image stack after maximum intensity projection,
1030 the inhibitory cells were counted slice by slice, recognized by their round shape (as opposed to
1031 pyramids), lack of apical dendrite, and relatively high uniform intensity. Cells at the bottom slice
1032 of each image stack and touching the left image border were excluded to avoid double-counting.

1033 *Statistics*

1034 All statistical tests (Kolmogorov-Smirnov, Kruskal-Wallis, and Mann-Whitney *U*) were conducted
1035 using the package `scipy.stats`¹⁴². Multiple comparisons were corrected for using false detection-
1036 rate adjusted p-values (Benjamini-Hochberg); this was done using `scipy.stats.multitest`
1037 and `scikit-posthocs`¹⁴³. Ordinary least squares regressions were conducted in the package
1038 `statsmodels`¹⁴⁴. Bootstrapped standard errors of the median were calculated by taking 5,000
1039 random samples with replacement (a bootstrap) of a dataset and then the standard deviation
1040 of each bootstrap was taken. Effect sizes were given as adjusted R^2 values or Cohen's f^2 (of
1041 a one-way ANOVA) using `statsmodels.formula.api.ols` and `statsmodels.stats.oneway`
1042 respectively.

1043 **References**

- 1044 1. Chandrasekaran C, Peixoto D, Newsome WT, Shenoy KV (2017) Laminar differences in
1045 decision-related neural activity in dorsal premotor cortex. *Nature Communications* 8(1):614.
- 1046 2. Finn ES, Huber L, Jangraw DC, Molfese PJ, Bandettini PA (2019) Layer-dependent activity
1047 in human prefrontal cortex during working memory. *Nature Neuroscience* 22(10):1687–1695.
- 1048 3. Pinto L, Dan Y (2015) Cell-type-specific activity in prefrontal cortex during goal-directed
1049 behavior. *Neuron* 87(2):437–450.
- 1050 4. Estebanez L, Hoffmann D, Voigt BC, Poulet JFA (2017) Parvalbumin-Expressing GABAergic
1051 Neurons in Primary Motor Cortex Signal Reaching. *Cell Reports* 20(2):308–318.
- 1052 5. Lui JH, et al. (2020) Differential encoding in prefrontal cortex projection neuron classes
1053 across cognitive tasks. *Cell*.
- 1054 6. Kvitsiani D, et al. (2013) Distinct behavioural and network correlates of two interneuron
1055 types in prefrontal cortex. *Nature* 498(7454):363–366.
- 1056 7. Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annual review of*
1057 *neuroscience* 30.
- 1058 8. Cisek P (2012) Making decisions through a distributed consensus. *Current opinion in*
1059 *neurobiology* 22(6):927–936.
- 1060 9. Ding L, Gold JI (2012) Neural correlates of perceptual decision making before, during, and
1061 after decision commitment in monkey frontal eye field. *Cerebral Cortex* 22(5):1052–1067.
- 1062 10. Thura D, Cisek P (2014) Deliberation and commitment in the premotor and primary motor
1063 cortex during dynamic decision making. *Neuron* 81(6):1401–1416.
- 1064 11. Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a
1065 combined visual discrimination reaction time task. *Journal of neuroscience* 22(21):9475–9489.
- 1066 12. Hanks TD, et al. (2015) Distinct relationships of parietal and prefrontal cortices to evidence
1067 accumulation. *Nature* 520(7546):220–223.

- 1068 13. Chandrasekaran C, Bray IE, Shenoy KV (2019) Frequency shifts and depth dependence of
1069 premotor beta band activity during perceptual decision-making. *The Journal of neuroscience*
1070 : the official journal of the Society for Neuroscience 39:1420–1435.
- 1071 14. Bastos AM, Loonis R, Kornblith S, Lundqvist M, Miller EK (2018) Laminar recordings
1072 in frontal cortex suggest distinct layers for maintenance and control of working memory.
1073 *Proceedings of the National Academy of Sciences* 115(5):1117–1122.
- 1074 15. Mountcastle VB, Talbot WH, Sakata H, Hyvärinen J (1969) Cortical neuronal mechanisms
1075 in flutter-vibration studied in unanesthetized monkeys. neuronal periodicity and frequency
1076 discrimination. *Journal of neurophysiology* 32(3):452–484.
- 1077 16. McCormick DA, Connors BW, Lighthall JW, Prince DA (1985) Comparative electrophysiology
1078 of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of Neurophysiology*
1079 54(4):782–806.
- 1080 17. Azouz R, Gray CM, Nowak LG, McCormick DA (1997) Physiological properties of inhibitory
1081 interneurons in cat striate cortex. *Cerebral Cortex* 7(6):534–545.
- 1082 18. Hodge RD, et al. (2019) Conserved cell types with divergent features in human versus mouse
1083 cortex. *Nature* pp. 1–8.
- 1084 19. Krienen FM, et al. (2020) Innovations present in the primate interneuron repertoire. *Nature*
1085 pp. 1–8.
- 1086 20. Vigneswaran G, Kraskov A, Lemon RN (2011) Large Identified Pyramidal Cells in Macaque
1087 Motor and Premotor Cortex Exhibit “Thin Spikes”: Implications for Cell Type Classification.
1088 *The Journal of Neuroscience* 31(40):14235–14242.
- 1089 21. Soares D, et al. (2017) Expression of Kv3.1b potassium channel is widespread in macaque
1090 motor cortex pyramidal cells: A histological comparison between rat and macaque. *Journal*
1091 *of Comparative Neurology* 525(9):2164–2174.
- 1092 22. Constantinople CM, Disney AA, Maffie J, Rudy B, Hawken MJ (2009) Quantitative analysis
1093 of neurons with Kv3 potassium channel subunits, Kv3.1b and Kv3.2, in macaque primary
1094 visual cortex. *Journal of Comparative Neurology* 516(4):291–311.
- 1095 23. Amatrudo JM, et al. (2012) Influence of Highly Distinctive Structural Properties on the
1096 Excitability of Pyramidal Neurons in Monkey Visual and Prefrontal Cortices. *The Journal of*
1097 *Neuroscience* 32(40):13644–13660.
- 1098 24. Onorato I, et al. (2019) A Distinct Class of Bursting Neurons with Strong Gamma Synchron-
1099 ization and Stimulus Selectivity in Monkey V1. *Neuron*.
- 1100 25. Kelly JG, García-Marín V, Rudy B, Hawken MJ (2018) Densities and Laminar Distributions of
1101 Kv3.1b-, PV-, GABA-, and SMI-32-Immunoreactive Neurons in Macaque Area V1. *Cerebral*
1102 *Cortex* 29(5):1921–1937.
- 1103 26. Ardid S, et al. (2015) Mapping of functionally characterized cell classes onto canonical circuit
1104 operations in primate prefrontal cortex. *The Journal of Neuroscience* 35(7):2975–2991.
- 1105 27. Trainito C, Nicolai Cv, Miller EK, Siegel M (2019) Extracellular Spike Waveform Dissociates
1106 Four Functionally Distinct Cell Classes in Primate Cortex. *Current Biology*.
- 1107 28. Boroujeni KB, Tiesinga P, Womelsdorf T (2020) Interneuron specific gamma synchronization
1108 encodes uncertain cues and prediction errors in lateral prefrontal and anterior cingulate
1109 cortex. *bioRxiv*.
- 1110 29. Barry JM (2015) Axonal activity in vivo: technical considerations and implications for the
1111 exploration of neural circuits in freely moving animals. *Frontiers in Neuroscience* 9:153.
- 1112 30. Robbins AA, Fox SE, Holmes GL, Scott RC, Barry JM (2013) Short duration waveforms

- 1113 recorded extracellularly from freely moving rats are representative of axonal activity. *Frontiers*
1114 *in Neural Circuits* 7:181.
- 1115 31. Krimer LS, et al. (2005) Cluster Analysis–Based Physiological Classification and Morphological
1116 Properties of Inhibitory Neurons in Layers 2–3 of Monkey Dorsolateral Prefrontal Cortex.
1117 *Journal of Neurophysiology* 94(5):3009–3022.
- 1118 32. Kaufman MT, Churchland MM, Shenoy KV (2013) The roles of monkey M1 neuron classes
1119 in movement preparation and execution. *Journal of neurophysiology* 110(4):817–25.
- 1120 33. Viskontas IV, Ekstrom AD, Wilson CL, Fried I (2007) Characterizing interneuron and
1121 pyramidal cells in the human medial temporal lobe in vivo using extracellular recordings.
1122 *Hippocampus* 17(1):49–57.
- 1123 34. Katai S, et al. (2010) Classification of extracellularly recorded neurons by their discharge
1124 patterns and their correlates with intracellularly identified neuronal types in the frontal cortex
1125 of behaving monkeys. *European Journal of Neuroscience* 31(7):1322–1338.
- 1126 35. Sun SH, et al. (2021) Analysis of extracellular spike waveforms and associated receptive
1127 fields of neurons in cat primary visual cortex. *The Journal of Physiology*.
- 1128 36. Weir K, Blanquie O, Kilb W, Luhmann HJ, Sinning A (2015) Comparison of spike parameters
1129 from optically identified GABAergic and glutamatergic neurons in sparse cortical cultures.
1130 *Frontiers in Cellular Neuroscience* 8:460.
- 1131 37. Tasic B, et al. (2018) Shared and distinct transcriptomic cell types across neocortical areas.
1132 *Nature* 563(7729):72–78.
- 1133 38. Becht E, et al. (2018) Dimensionality reduction for visualizing single-cell data using UMAP.
1134 *Nature Biotechnology* 37(1):38–44.
- 1135 39. Maheswaranathan N, Williams AH, Golub MD, Ganguli S, Sussillo D (2019) Universality
1136 and individuality in neural dynamics across large populations of recurrent networks.
- 1137 40. Kleinman M, Chandrasekaran C, Kao JC (2019) Recurrent neural network models of multi-
1138 area computation underlying decision-making. *bioRxiv*.
- 1139 41. Bala PC, et al. (2020) OpenMonkeyStudio: Automated Markerless Pose Estimation in Freely
1140 Moving Macaques. *bioRxiv* p. 2020.01.31.928861.
- 1141 42. Hsu AI, Yttri EA (2020) B-soid: An open source unsupervised algorithm for discovery of
1142 spontaneous behaviors. *bioRxiv*.
- 1143 43. Dolensek N, Gehrlach DA, Klein AS, Gogolla N (2020) Facial expressions of emotion states
1144 and their neuronal correlates in mice. *Science (New York, N.Y.)* 368(6486):89–94.
- 1145 44. Jia X, et al. (2019) High-density extracellular probes reveal dendritic backpropagation and
1146 facilitate neuron classification. *Journal of Neurophysiology* 121(5):1831–1847.
- 1147 45. Gouwens NW, et al. (2020) Integrated Morphoelectric and Transcriptomic Classification of
1148 Cortical GABAergic Cells. *Cell* 183(4):935–953.e19.
- 1149 46. Klempíř O, et al. (2020) Application of spike sorting algorithm to neuronal signals originated
1150 from boron doped diamond micro-electrode arrays. *Physiological Research* 69(3).
- 1151 47. Markanday A, et al. (2020) Using deep neural networks to detect complex spikes of cerebellar
1152 Purkinje cells. *Journal of Neurophysiology* 123(6):2217–2234.
- 1153 48. McInnes L, Healy J, Melville J (2018) Umap: Uniform manifold approximation and projection
1154 for dimension reduction. *arXiv*.
- 1155 49. Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in
1156 large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008(10):P10008.
- 1157 50. Snyder AC, Morais MJ, Smith MA (2016) Dynamics of excitatory and inhibitory networks are

- 1158 differentially altered by selective attention. *Journal of Neurophysiology* 116(4):1807–1820.
- 1159 51. Gold C, Girardin CC, Martin KAC, Koch C (2009) High-Amplitude Positive Spikes Recorded
1160 Extracellularly in Cat Visual Cortex. *Journal of Neurophysiology* 102(6):3340–3351.
- 1161 52. Connors BW, Gutnick MJ, Prince DA (1982) Electrophysiological properties of neocortical
1162 neurons in vitro. *Journal of Neurophysiology* 48(6):1302–1320.
- 1163 53. Kaufman MT, et al. (2010) Roles of monkey premotor neuron classes in movement preparation
1164 and execution. *Journal of neurophysiology* 104(2):799–810.
- 1165 54. Mitchell JF, Sundberg KA, Reynolds JH (2007) Differential attention-dependent response
1166 modulation across cell classes in macaque visual area v4. *Neuron* 55(1):131–141.
- 1167 55. Stuttgart M (2020) Mlib - toolbox for analyzing spike data.
- 1168 56. Deligkaris K, Bullmann T, Frey U (2016) Extracellularly Recorded Somatic and Neuritic Signal
1169 Shapes and Classification Algorithms for High-Density Microelectrode Array Electrophysiology.
1170 *Frontiers in Neuroscience* 10:421.
- 1171 57. Bakkum DJ, et al. (2013) Tracking axonal action potential propagation on a high-density
1172 microelectrode array across hundreds of sites. *Nature Communications* 4(1):2181.
- 1173 58. Gur M, Beylin A, Snodderly DM (1999) Physiological Properties of Macaque V1 Neurons
1174 are Correlated With Extracellular Spike Amplitude, Duration, and Polarity. *Journal of
1175 Neurophysiology* 82(3):1451–1464.
- 1176 59. Molnar C (2020) *Interpretable Machine Learning: A Guide for Making Black Box Models
1177 Explainable*. (LULU COM).
- 1178 60. Azodi CB, Tang J, Shiu SH (2020) Opening the Black Box: Interpretable Machine Learning
1179 for Geneticists. *Trends in Genetics* 36(6):442–455.
- 1180 61. Lundberg S, Lee SI (2017) A Unified Approach to Interpreting Model Predictions. *arXiv*.
- 1181 62. Lundberg SM, et al. (2020) From local explanations to global understanding with explainable
1182 AI for trees. *Nature Machine Intelligence* 2(1):56–67.
- 1183 63. Shapley LS (1988) *A value for n-person games*, ed. Roth AE. (Cambridge University Press),
1184 p. 31–40.
- 1185 64. Štrumbelj E, Kononenko I (2014) Explaining prediction models and individual predictions
1186 with feature contributions. *Knowledge and Information Systems* 41(3):647–665.
- 1187 65. Connors BW, Gutnick MJ (1990) Intrinsic firing patterns of diverse neocortical neurons.
1188 *Trends in Neurosciences* 13(3):99–104.
- 1189 66. Nowak LG, Azouz R, Sanchez-Vives MV, Gray CM, McCormick DA (2003) Electrophysio-
1190 logical Classes of Cat Primary Visual Cortical Neurons In Vivo as Revealed by Quantitative
1191 Analyses. *Journal of Neurophysiology* 89(3):1541–1566.
- 1192 67. Contreras D (2004) Electrophysiological classes of neocortical neurons. *Neural Networks*
1193 17(5-6):633–646.
- 1194 68. Steriade M (2004) Neocortical cell classes are flexible entities. *Nature Reviews Neuroscience*
1195 5(2):121–134.
- 1196 69. Steriade M, Timofeev I, Dürmüller N, Grenier F (1998) Dynamic Properties of Corticothalamic
1197 Neurons and Local Cortical Interneurons Generating Fast Rhythmic (30–40 Hz) Spike Bursts.
1198 *Journal of Neurophysiology* 79(1):483–490.
- 1199 70. Johnston K, DeSouza JFX, Everling S (2009) Monkey Prefrontal Cortical Pyramidal and
1200 Putative Interneurons Exhibit Differential Patterns of Activity Between Prosaccade and
1201 Antisaccade Tasks. *Journal of Neuroscience* 29(17):5516–5524.
- 1202 71. Hussar CR, Pasternak T (2009) Flexibility of Sensory Representations in Prefrontal Cortex

- 1203 Depends on Cell Type. *Neuron* 64(5):730–743.
- 1204 72. Quirk MC, Sosulski DL, Feierstein CE, Uchida N, Mainen ZF (2009) A defined network of
1205 fast-spiking interneurons in orbitofrontal cortex: responses to behavioral contingencies and
1206 ketamine administration. *Frontiers in Systems Neuroscience* 3:13.
- 1207 73. Povysheva NV, Zaitsev AV, Gonzalez-Burgos G, Lewis DA (2013) Electrophysiological
1208 Heterogeneity of Fast-Spiking Interneurons: Chandelier versus Basket Cells. *PLoS ONE*
1209 8(8):e70553.
- 1210 74. Zaitsev AV, et al. (2009) Interneuron Diversity in Layers 2–3 of Monkey Prefrontal Cortex.
1211 *Cerebral Cortex* 19(7):1597–1615.
- 1212 75. Song JH, McPeck RM (2010) Roles of Narrow- and Broad-Spiking Dorsal Premotor Area
1213 Neurons in Reach Target Selection and Movement Production. *Journal of Neurophysiology*
1214 103(4):2124–2138.
- 1215 76. Mruczek REB, Sheinberg DL (2012) Stimulus selectivity and response latency in puta-
1216 tive inhibitory and excitatory neurons of the primate inferior temporal cortex. *Journal of*
1217 *Neurophysiology* 108(10):2725–2736.
- 1218 77. Tosches MA, et al. (2018) Evolution of pallium, hippocampus, and cortical cell types revealed
1219 by single-cell transcriptomics in reptiles. *Science* 360(6391):881–888.
- 1220 78. Arikuni T, Watanabe K, Kubota K (1988) Connections of area 8 with area 6 in the brain of
1221 the macaque monkey. *Journal of Comparative Neurology* 277(1):21–40.
- 1222 79. Nandy AS, Nassi JJ, Reynolds JH (2017) Laminar Organization of Attentional Modulation
1223 in Macaque Visual Area V4. *Neuron* 93(1):235–246.
- 1224 80. Barbas H, Pandya DN (1987) Architecture and frontal cortical connections of the premotor
1225 cortex (area 6) in the rhesus monkey. *The Journal of Comparative Neurology* 256(2):211–228.
- 1226 81. DeFelipe J (1997) Types of neurons, synaptic connections and chemical characteristics
1227 of cells immunoreactive for calbindin-D28K, parvalbumin and calretinin in the neocortex.
1228 *Journal of Chemical Neuroanatomy* 14(1):1–19.
- 1229 82. Schmitz C, et al. (2014) Current automated 3D cell detection methods are not a suitable
1230 replacement for manual stereologic cell counting. *Frontiers in Neuroanatomy* 8(27):1–14.
- 1231 83. Zaitsev A, et al. (2005) Localization of Calcium-binding Proteins in Physiologically and Mor-
1232 phologically Characterized Interneurons of Monkey Dorsolateral Prefrontal Cortex. *Cerebral*
1233 *Cortex* 15(8):1178–1186.
- 1234 84. Ali M, Jones MW, Xie X, Williams M (2019) TimeCluster: dimension reduction applied to
1235 temporal data for visual analytics. *The Visual Computer* 35(6-8):1013–1026.
- 1236 85. Diaz-Papkovich A, Anderson-Trocmé L, Ben-Eghan C, Gravel S (2019) UMAP reveals cryptic
1237 population structure and phenotype heterogeneity in large genomic cohorts. *PLoS Genetics*
1238 15(11):e1008432.
- 1239 86. Noichl M (2019) Modeling the structure of recent philosophy. *Synthese* pp. 1–12.
- 1240 87. Moscovich A, Rosset S (2019) On the cross-validation bias due to unsupervised pre-processing.
1241 *arXiv*.
- 1242 88. Sainburg T, McInnes L, Gentner TQ (2020) Parametric UMAP: learning embeddings with
1243 deep neural networks for representation and semi-supervised learning. *arXiv*.
- 1244 89. Bellet A, Habrard A, Sebban M (2013) A Survey on Metric Learning for Feature Vectors
1245 and Structured Data. *arXiv*.
- 1246 90. Yang L, Jin R (2006) Distance metric learning: A comprehensive survey.
- 1247 91. Roux L, Stark E, Sjulson L, Buzsáki G (2014) In vivo optogenetic identification and manipu-

- 1248 lation of GABAergic interneuron subtypes. *Current opinion in neurobiology* 26:88–95.
- 1249 92. Deubner J, Coulon P, Diester I (2019) Optogenetic approaches to study the mammalian
1250 brain. *Current Opinion in Structural Biology* 57:157–163.
- 1251 93. Cohen JY, Haesler S, Vong L, Lowell BB, Uchida N (2012) Neuron-type-specific signals for
1252 reward and punishment in the ventral tegmental area. *Nature* 482(7383):85–88.
- 1253 94. Hangya B, Ranade SP, Lorenc M, Kepecs A (2015) Central Cholinergic Neurons Are Rapidly
1254 Recruited by Reinforcement Feedback. *Cell* 162(5):1155–1168.
- 1255 95. Trautmann EM, et al. (2019) Accurate Estimation of Neural Population Dynamics without
1256 Spike Sorting. *Neuron*.
- 1257 96. Jun JJ, et al. (2017) Fully integrated silicon probes for high-density recording of neural
1258 activity. *Nature* 551(7679):232–236.
- 1259 97. Shenoy KV, Sahani M, Churchland MM (2013) Cortical control of arm movements: a
1260 dynamical systems perspective. *Annual review of neuroscience* 36.
- 1261 98. Churchland MM, et al. (2012) Neural population dynamics during reaching. *Nature*
1262 487(7405):51–56.
- 1263 99. Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by
1264 recurrent dynamics in prefrontal cortex. *nature* 503(7474):78–84.
- 1265 100. Remington ED, Narain D, Hosseini EA, Jazayeri M (2018) Flexible sensorimotor computations
1266 through rapid reconfiguration of cortical dynamics. *Neuron* 98(5):1005 – 1019.e5.
- 1267 101. Wang J, Narain D, Hosseini EA, Jazayeri M (2017) Flexible timing by temporal scaling of
1268 cortical responses. *Nature Neuroscience* 21(1):102–110.
- 1269 102. Georgopoulos A, Schwartz A, Kettner R (1986) Neuronal population coding of movement
1270 direction. *Science* 233(4771):1416–1419.
- 1271 103. Hubel DH, Wiesel TN (1959) Receptive fields of single neurones in the cat's striate cortex.
1272 *The Journal of physiology* 148(3):574.
- 1273 104. Bruce CJ, Goldberg ME (1985) Primate frontal eye fields. i. single neurons discharging before
1274 saccades. *Journal of neurophysiology* 53(3):603–635.
- 1275 105. Cunningham JP, Byron MY (2014) Dimensionality reduction for large-scale neural recordings.
1276 *Nature neuroscience* 17(11):1500–1509.
- 1277 106. Williams AH, et al. (2018) Unsupervised Discovery of Demixed, Low-Dimensional Neu-
1278 ral Dynamics across Multiple Timescales through Tensor Component Analysis. *Neuron*
1279 98(6):1099–1115.e8.
- 1280 107. Kobak D, et al. (2016) Demixed principal component analysis of neural population data.
1281 *Elife* 5:e10989.
- 1282 108. Pandarinath C, et al. (2018) Inferring single-trial neural population dynamics using sequential
1283 auto-encoders. *Nature methods* 15(10):805–815.
- 1284 109. Lindén H, et al. (2014) LFPy: a tool for biophysical simulation of extracellular potentials
1285 generated by detailed model neurons. *Frontiers in Neuroinformatics* 7.
- 1286 110. Kaczmarek LK, Zhang Y (2017) Kv3 Channels: Enablers of Rapid Firing, Neurotransmitter
1287 Release, and Neuronal Endurance. *Physiological Reviews* 97(4):1431–1468.
- 1288 111. Ding S, Matta SG, Zhou FM (2011) Kv3-Like Potassium Channels Are Required for Sustained
1289 High-Frequency Firing in Basal Ganglia Output Neurons. *Journal of Neurophysiology*
1290 105(2):554–570.
- 1291 112. Erisir A, Lau D, Rudy B, Leonard CS (1999) Function of Specific K⁺ Channels in Sustained
1292 High-Frequency Firing of Fast-Spiking Neocortical Interneurons. *Journal of Neurophysiology*

- 1293 82(5):2476–2489.
- 1294 113. Bean BP (2007) The action potential in mammalian central neurons. *Nature Reviews*
1295 *Neuroscience* 8(6):451–465.
- 1296 114. Henze DA, et al. (2000) Intracellular Features Predicted by Extracellular Recordings in the
1297 Hippocampus In Vivo. *Journal of Neurophysiology* 84(1):390–400.
- 1298 115. Torres-Gomez S, et al. (2020) Changes in the Proportion of Inhibitory Interneuron Types
1299 from Sensory to Executive Areas of the Primate Neocortex: Implications for the Origins of
1300 Working Memory Representations. *Cerebral cortex (New York, N.Y. : 1991)*.
- 1301 116. Kelly JG, Hawken MJ (2020) GABAergic and non-GABAergic subpopulations of Kv3.1b-
1302 expressing neurons in macaque V2 and MT: laminar distributions and proportion of total
1303 neuronal population. *Brain Structure and Function* 225(3):1135–1152.
- 1304 117. Ichinohe N, et al. (2004) A voltage-gated potassium channel, Kv3.1b, is expressed by
1305 a subpopulation of large pyramidal neurons in layer 5 of the macaque monkey cortex.
1306 *Neuroscience* 129(1):179–185.
- 1307 118. Saleh MS, et al. (2019) Cmu array: A 3d nano-printed, customizable ultra-high-density
1308 microelectrode array platform. *bioRxiv*.
- 1309 119. Mosher CP, et al. (2020) Cellular classes in the human brain revealed in-vivo by heartbeat-
1310 related modulation of the extracellular action potential waveform. *Cell Reports* 30(10):3536
1311 – 3551.e6.
- 1312 120. Courtin J, et al. (2014) Prefrontal parvalbumin interneurons shape neuronal activity to drive
1313 fear expression. *Nature* 505(7481):92–96.
- 1314 121. Tremblay S, et al. (2020) An Open Resource for Non-human Primate Optogenetics. *Neuron*.
- 1315 122. De A, El-Shamayleh Y, Horwitz GD (2020) Fast and reversible neural inactivation in macaque
1316 cortex by optogenetic stimulation of gabaergic neurons. *Elife* 9:e52658.
- 1317 123. Mehta P, et al. (2019) Functional access to neuron subclasses in rodent and primate forebrain.
1318 *Cell Reports* 26(10):2818 – 2832.e8.
- 1319 124. Vormstein-Schneider D, et al. (2020) Viral manipulation of functionally distinct interneurons
1320 in mice, non-human primates and humans. *Nature Neuroscience* pp. 1–8.
- 1321 125. Kleiner M, et al. (2007) What's new in psychtoolbox-3. *Perception* 36(14):1–16.
- 1322 126. Porter MA, Onnela JP, Mucha PJ (2009) Communities in Networks. *arXiv*.
- 1323 127. Lambiotte R, Delvenne JC, Barahona M (2008) Laplacian dynamics and multiscale modular
1324 structure in networks. *arXiv preprint arXiv:0812.1770*.
- 1325 128. Poulin V, Théberge F (2018) Ensemble Clustering for Graphs. *arXiv*.
- 1326 129. Poulin V, Théberge F (2019) Ensemble clustering for graphs: comparisons and applications.
1327 *Applied Network Science* 4(1):51.
- 1328 130. Csardi G, Nepusz T (2006) The igraph software package for complex network research.
1329 *InterJournal Complex Systems*:1695.
- 1330 131. Chen T, Guestrin C (2016) Xgboost: A scalable tree boosting system. arxiv 2016. *arXiv*
1331 *preprint arXiv:1603.02754*.
- 1332 132. Kaufman S, Rosset S, Perlich C (2011) Leakage in data mining: formulation, detection, and
1333 avoidance in *In Proceedings of the 2011 conference on Knowledge Discovery in Data Mining*.
1334 pp. 556–563.
- 1335 133. Meister MLR, Hennig JA, Huk AC (2013) Signal multiplexing and single-neuron computations
1336 in lateral intraparietal area during decision-making. *Journal of Neuroscience* 33(6):2254–2267.
- 1337 134. Medalla M, Luebke JI (2015) Diversity of Glutamatergic Synaptic Strength in Lateral

- 1338 Prefrontal versus Primary Visual Cortices in the Rhesus Monkey. *The Journal of Neuroscience*
1339 35(1):112–127.
- 1340 135. Rosene DL, Roy NJ, Davis BJ (1986) A cryoprotection method that facilitates cutting
1341 frozen sections of whole monkey brains for histological and histochemical processing without
1342 freezing artifact. *The journal of histochemistry and cytochemistry : official journal of the*
1343 *Histochemistry Society* 34(10):1301–1315.
- 1344 136. Estrada LI, et al. (2016) Evaluation of Long-Term Cryostorage of Brain Tissue Sections for
1345 Quantitative Histochemistry. *Journal of Histochemistry & Cytochemistry* 65(3):153–171.
- 1346 137. Morecraft R, Cipolloni P, Stilwell [U+2010] Morecraft K, Gedney M, Pandya D (2004) Cy-
1347 toarchitecture and cortical connections of the posterior cingulate and adjacent somatosensory
1348 fields in the rhesus monkey. *Journal of Comparative Neurology* 469(1):37–69.
- 1349 138. Morecraft RJ, et al. (2019) Terminal organization of the corticospinal projection from the
1350 lateral premotor cortex to the cervical enlargement (C5–T1) in rhesus monkey. *Journal of*
1351 *Comparative Neurology* 527(16):2761–2789.
- 1352 139. Matelli M, Luppino G (1996) Thalamic input to mesial and superior area 6 in the macaque
1353 monkey. *Journal of Comparative Neurology* 372(1):59–87.
- 1354 140. Rizzolatti G, Luppino G, Matelli M (1998) The organization of the cortical motor system:
1355 new concepts. *Electroencephalography and Clinical Neurophysiology* 106(4):283–296.
- 1356 141. Shindelin J, et al. (2012) Fiji: An Open-source Platform for Biological-image Analysis. *Nature*
1357 *Methods* 9(7):676—682.
- 1358 142. Virtanen P, et al. (2020) SciPy 1.0: Fundamental Algorithms for Scientific Computing in
1359 Python. *Nature Methods* 17:261–272.
- 1360 143. Terpilowski M (2019) scikit-posthocs: Pairwise multiple comparison tests in python. *The*
1361 *Journal of Open Source Software* 4(36):1169.
- 1362 144. Seabold S, Perktold J (2010) statsmodels: Econometric and statistical modeling with python
1363 in *9th Python in Science Conference*.
- 1364 145. Verleysen M, Francois D, Simon G, Wertz V (2003) Artificial Neural Nets Problem Solving
1365 Methods, 7th International Work-Conference on Artificial and Natural Neural Networks,
1366 IWANN2003 Maó, Menorca, Spain, June 3–6, 2003 Proceedings, Part II. pp. 105–112.
- 1367 146. Beyer KS, Goldstein J, Ramakrishnan R, Shaft U (1999) When is "nearest neighbor"
1368 meaningful? in *Proceedings of the 7th International Conference on Database Theory,*
1369 *ICDT '99*. (Springer-Verlag, Berlin, Heidelberg), p. 217–235.
- 1370 147. McInnes L (2019) Umap uniform manifold approximation and projection for dimension
1371 reduction.
- 1372 148. Tenenbaum JB, Silva Vd, Langford JC (2000) A Global Geometric Framework for Nonlinear
1373 Dimensionality Reduction. *Science* 290(5500):2319–2323.
- 1374 149. Maaten Lvd, Hinton G (2008) Visualizing data using t-sne. *Journal of machine learning*
1375 *research* 9(Nov):2579–2605.
- 1376 150. Lee JA, Verleysen M, eds. (2007) *Nonlinear Dimensionality Reduction*. (Springer New York).
- 1377 151. Spivak DI (2009) Metric realization of fuzzy simplicial sets.
- 1378 152. Dong W, Moses C, Li K (2011) Efficient k-nearest neighbor graph construction for generic
1379 similarity measures in *Proceedings of the 20th international conference on World wide web*.
1380 pp. 577–586.
- 1381 153. Belkin M, Niyogi P (2002) Laplacian eigenmaps and spectral techniques for embedding and
1382 clustering in *Advances in Neural Information Processing Systems 14*, eds. Dietterich TG,

- 1383 Becker S, Ghahramani Z. (MIT Press), pp. 585–591.
- 1384 154. Bengio Y, Vincent P, Paiement J (2003) Spectral clustering and kernel pca are learning
1385 eigenfunctions.
- 1386 155. Kobak D, Linderman GC (2019) UMAP does not preserve global structure any better than
1387 t-SNE when using the same initialization. *bioRxiv* p. 2019.12.19.877522.
- 1388 156. Newman MEJ, Girvan M (2004) Finding and evaluating community structure in networks.
1389 *Phys. Rev. E* 69(2):026113.
- 1390 157. Lambiotte R (2007) Finding communities at different resolutions in large networks.

1391 **Supplementary Information**

1392 *Clustering in high-dimensionality and the curse of dimensionality*

1393 Clustering in high-dimensions is a difficult problem. In particular, “concentration of measure”
1394 results in Euclidean distances, used by k-Means clustering, becoming meaningless as a measure of
1395 distance. Specifically, as dimensionality increases, the difference between distances of randomly
1396 chosen points all converge to the same constant distance^{145,146}. UMAP counters this by assuming
1397 “connectedness” between each point and its nearest neighbors. This results in the local unit distance
1398 metric around each data point stretching or contracting¹⁴⁷ (see Supplementary Information). In
1399 addition, traditionally used clustering methods like a Gaussian mixture model (GMM) or k -Means
1400 do not consistently capture these clusters because they assume normality in cluster shape. This
1401 is violated by the perturbations introduced by the non-linearity of the UMAP projection even if
1402 distributions are Gaussian in the original high-dimensional space (which they are also likely not
1403 either).

1404 *UMAP dimensionality reduction*

1405 UMAP is among the class of non-linear dimensionality reductions known as manifold learning
1406 algorithms which also includes other well-known methods in neuroscience such as Isomap¹⁴⁸ and
1407 t-SNE¹⁴⁹ (see¹⁵⁰ for a review of methods). Key to this algorithm is the presumption that although
1408 data may not be uniformly spaced in the ambient space, it is uniform on some low-dimensional
1409 manifold. It is also assumed that the underlying manifold is locally connected (i.e. doesn't
1410 have any breaks or isolated points). This leads to the conclusion that the underlying notion of
1411 distance (Riemannian metric) changes in each region of the manifold: the notion of a unit distance
1412 “stretches” in areas of sparser density and “shortens” in areas of higher density. This is formalized
1413 beginning with defining how a local Riemannian metric g should be constructed by **Lemma 1**
1414 of⁴⁸:

1415 **Lemma 1:** Let \mathcal{M} be a Riemannian manifold equipped with a metric g in ambient
1416 space \mathbb{R}^n . Let $p \in \mathcal{M}$ be a point in this space. If g is locally constant about the point
1417 p in an open neighborhood U such that g is a constant diagonal matrix in diagonal
1418 coordinates, then in a ball $B \subseteq U$ centered at p with a volume $\frac{\pi^{n/2}}{\Gamma(n/2+1)}$ with respect
1419 to g , the geodesic distance from p to any point $q \in B$ is $\frac{1}{r}d_{\mathbb{R}^n}(p, q)$, where r is the
1420 radius of the ball in the ambient space and $d_{\mathbb{R}^n}$ is the existing metric on the ambient
1421 space.

1422 Using this definition of g , each ball B of fixed volume should contain the same number of data
1423 points in X regardless of where on the manifold B is located. This also implies that a ball centered
1424 on data point x_i should contain the k -nearest neighbors of x_i in a fixed volume no matter which
1425 $x_i \in X$ is chosen. Thus the geodesic distance around each data point is normalized by its distance
1426 to its k^{th} nearest neighbor.

To compensate for the impact of certain nearest neighbors in the ball lying much further than those closer by, the normalizing distances are transformed by the exponential function,

$$\sum_{j=1}^k \exp\left(\frac{-|x_i - x_{i_j}|}{r_i}\right) = \log_2(k)$$

1427 To unite these disparate metric spaces (each data point has a possibly unique local notion of
 1428 distance), category theory is used to convert this representation to that of a fuzzy simplicial set
 1429 via an adjunction that will not be defined here (see **Section 3** of⁴⁸ and the Nerve Theorem). In
 1430 this way, the topological structure of the data can be equivalently represented as a metric space or
 1431 a set of fuzzy simplices. One large benefit of this construction is that while normalized distances
 1432 in high-dimensional spaces suffer from the Curse of Dimensionality in the form of concentration
 1433 of measure, normalized nearest neighbor distances do not¹⁴⁷. The end result of this process is
 1434 an approximation of the topology of the manifold by fuzzy simplicial sets in the form of a Čech
 1435 complex.

With this fuzzy topological representation, the low-dimensional representation can be found through
 an optimization procedure that minimizes the cross-entropy of fuzzy simplicial sets containing the
 same objects. Given two fuzzy sets with the same members A and separate membership strength
 functions μ and ν (of Spivak's characteristic form¹⁵¹), the cross-entropy $C((A, \mu), (A, \nu))$ is
 defined as,

$$C((A, \mu), (A, \nu)) \triangleq \sum_{a \in A} \left(\mu(a) \log \frac{\mu(a)}{\nu(a)} + (1 - \mu(a)) \log \frac{1 - \mu(a)}{1 - \nu(a)} \right).$$

1436 The first term $\mu(a) \log \frac{\mu(a)}{\nu(a)}$ captures the attractive force minimised if short edges in high-dimension
 1437 correspond to short edges in low-dimension and $(1 - \mu(a)) \log \frac{1 - \mu(a)}{1 - \nu(a)}$ is the repulsive forces that
 1438 are minimised if long edges in high-dimension correspond to long edges in low-dimension.

1439 From a computational perspective, this whole UMAP process proceeds in two steps: construction
 1440 of a k -nearest neighbor graph and layout of the graph into a low-dimensional manifold. Note
 1441 that after the first step, the k -nearest neighbor graph is passed to Louvain community detection
 1442 and thus the clustering is not dependent on the embedding of the graph, just on its construction
 1443 and associated UMAP parameters such as `n_neighbors` and `metric` but not `min_dist`. The
 1444 embedding *is* however used for visualization and interpretability.

Graph construction: Given a set of data points $X = \{x_1, \dots, x_N\}$ and a metric d ,
 the construction of an undirected weighted k -nearest neighbor graph (captured by
 an adjacency matrix capturing the connection weights between nodes) is conducted
 using a nearest neighbor descent algorithm¹⁵². For each data point $x_i \in X$ and fixed
 nearest neighbor hyperparameter k , we have the set $\{x_{i_1} \dots x_{i_k}\}$ the set of k -nearest
 neighbors of x_i under the local Riemannian metric d . We define ρ_i and σ_i such that,

$$\rho_i = \min \{d(x_i, x_{i_j}) | 1 \leq j \leq k, d(x_i, x_{i_j}) > 0\}$$

and setting σ_i such that,

$$\log_2(k) = \sum_{j=1}^k \exp \left(\frac{-\max(0, d(x_i, x_{i_j}) - \rho_i)}{\sigma_i} \right).$$

The weighted graph $\bar{G} = (V, E, w)$ is defined in terms of the vertices V , edges

$E = \{(x_i, x_{i_j}) | 1 \leq j \leq k, 1 \leq i \leq N\}$, and weight function w as,

$$w((x_i, x_{i_j})) = \sum_{j=1}^k \exp\left(\frac{-\max(0, d(x_i, x_{i_j}) - \rho_i)}{\sigma_i}\right)$$

. If A is the weighted adjacency matrix of \overline{G} , we can get the undirected weighted graph B by the relationship,

$$B = A + A^\top - A \circ A^\top$$

1445

where \circ is the Hadamard product.

Graph layout: The UMAP algorithm finds a low-dimensional projection (manifold) of the high-dimensional data by a force directed layout of the constructed graph. Before this is done though, the graph is spectrally embedded to aid in consistency and convergence of the algorithm through initialization¹⁵³⁻¹⁵⁵. The symmetric normalized Laplacian L of the graph is calculated for the 1-skeleton of the weighted graph which is analogous to the Laplace-Beltrami operator (divergence of the gradient, $\Delta f = \nabla^2 f$) on a manifold. If D is the degree matrix (a diagonal matrix containing the degree of each vertex) of the adjacency matrix A , we compute the Laplacian matrix as,

$$L = D^{\frac{1}{2}}(D - A)D^{\frac{1}{2}}$$

with associated eigenvectors y and eigenvalues λ ,

$$Ly = \lambda Dy.$$

After the spectral embedding of the graph with Laplacian eigenmaps, the force directed graph layout iteratively applies attractive and repulsive forces on the edges and vertices. The attractive force between two vertices i and j at coordinates \mathbf{y}_i and \mathbf{y}_j with tunable hyperparameters a and b and is determined by,

$$\frac{-2ab \|\mathbf{y}_i - \mathbf{y}_j\|_2^{2(b-1)}}{1 + \|\mathbf{y}_i - \mathbf{y}_j\|_2^2} w((x_i, x_j)) (\mathbf{y}_i - \mathbf{y}_j)$$

and the repulsive forces with hyper-parameter $\epsilon = 0.001$ to prevent division by zero,

$$\frac{b}{(\epsilon + \|\mathbf{y}_i - \mathbf{y}_j\|_2^2)(1 + \|\mathbf{y}_i - \mathbf{y}_j\|_2^2)} \left(1 - w((x_i, x_j))\right) (\mathbf{y}_i - \mathbf{y}_j).$$

1446

This optimization procedure is then completed using stochastic gradient descent to arrive at the final embedding.

1447

1448 *Louvain method for community detection*

The Louvain method for community detection (here called clustering)⁴⁹ operates on a weighted network and locates highly-interconnected nodes called a community. This “connectedness” is measured by their modularity Q with added resolution parameter t ^{156,157} (taking real values between -1 and 1 inclusive) defined as,

$$Q_t = (1 - t) + \frac{1}{2m} \sum_{i,j} \left[tA_{i,j} - \frac{k_i k_j}{2m} \right] \delta(c_i, c_j)$$

1449 where t is a parameter controlling the “characteristic scale” of the communities¹²⁷. The larger
1450 the resolution parameter, the fewer the number of communities and the larger their size. Smaller
1451 values of resolution parameter results in more communities smaller in size. Note that when $t = 1$,
1452 the simplified definition of modularity is given as appears in⁴⁹. $A_{i,j}$ is an adjacency matrix with
1453 the weights of the edges between the nodes indexed by i and j . $k_i = \sum_j A_{i,j}$ is the sum of
1454 weights of the edges connected to the node i . c_i is the community to which the node i belongs
1455 to and the function $\delta(c_i, c_j)$ is an indicator function that is 1 if $c_i = c_j$ and 0 otherwise. The
1456 value $m = \frac{1}{2} \sum_{i,j} A_{i,j}$ which is the sum of all the weights of all the edges in the network. This
1457 equation also serves as an objective function for the iterative procedure in Blondel *et al* 2008⁴⁹
1458 which proceeds in two steps,

Modularity Optimization: Each node is assigned to its own community (in the initial step) and then each node is moved into a community with a random neighbor and the change in modularity is examined. The equation for this change in modularity is,

$$\Delta Q_t = \left[\frac{\sum_{in} + k_{i,in}}{2m} - \left(\frac{\sum_{tot} + k_i^2}{2m} \right) \right] - \left[\frac{\sum_{in}}{2m} - \left(\frac{\sum_{tot}}{2m} \right) - \left(\frac{k_i}{2m} \right)^2 \right].$$

1459 where \sum_{in} is the sum of all the weights inside the community that the node i is moving
1460 into. \sum_{tot} is the sum of all the weights of the edges to nodes in the community i
1461 is moving into. k_i is the sum of all the weighted edges incident on i . $k_{i,in}$ is the sum of
1462 the weights from the edges of i to nodes in the cluster.

1463 Once the change in modularity is examined for a node before and after joining a
1464 neighboring cluster, the neighbor joins (or stays with) the community with the largest
1465 positive increase in modularity; if no increase can be found, the node remains a part
1466 of its current community. Once there can be found no increase in modularity for any
1467 points, the algorithm proceeds to the second step.

1468 **Community Aggregation:** Every node in each community in the previous step is
1469 then collapsed into a single node and their edges combined to form a new graph. The
1470 process is then repeated from the previous step.

1471 The graph produced by UMAP was passed into this Louvain clustering algorithm⁴⁹ using `cylouvain`
1472 0.2.3 with parameters in Table S1. This clustering method requires no prior specification of
1473 the number of clusters that should be present but its number does depend on the resolution
1474 parameter. To choose this parameter, a sensitivity analysis was conducted across various values of

1475 the resolution parameter with the number of communities and total modularity compared Fig. 2B.
 1476 Each waveform was then plotted in UMAP space and color-coded to its associated cluster label
 1477 found by Louvain Clustering Fig. 2A.

1478 *Ensemble clustering for graphs (ECG)*

1479 ECG is a consensus clustering method for graphs and was used to validate the Louvain clustering
 1480 algorithm. ECG consists of two steps: generation and integration.

1481 **Generation:** This step instantiates the ECG algorithm by using Louvain clustering
 1482 to produce a randomized set of k level-1 partitions $\mathcal{P} \in \{\mathcal{P}_\infty, \dots, \mathcal{P}_\parallel\}$. The ran-
 1483 domization comes from the randomization of vertices in the initial step of Louvain
 1484 clustering.

Integration: Once the k randomized level-1 Louvain partitions are obtained, Louvain
 is run on a weighted version of the initial graph $G = (V, E)$. These weights $W_{\mathcal{P}}(u, v)$
 are obtained via co-association of edges $e = (u, v) \in E$. These weights are defined
 as,

$$W_{\mathcal{P}}(u, v) \triangleq \begin{cases} w_* + (1 - w_*) \cdot \left(\frac{\sum_{i=1}^k v_{\mathcal{P}_i}(u, v)}{k} \right) & , \text{ if } (u, v) \text{ is in the 2-core of } G \\ w_* & , \text{ otherwise} \end{cases}$$

1485 where we have the minimum ECG weight $0 < w_* < 1$ and the co-occurrence of edges
 1486 u and v as $v_{P_i} = \sum_{j=1}^{l_i} \mathbb{1}_{C_i^j}(u) \cdot \mathbb{1}_{C_i^j}(v)$ where $\mathbb{1}_{C_i^j}(u)$ is an indicator function of if
 1487 the edge u occurs in the cluster of P_i or not.

1488 With this function, ECG combines these level-1 Louvain partitions as a single weighted graph
 1489 which serves as the result.

1490 *SHapley Additive exPlanations (SHAP)*

SHAP values build off of Shapley values⁶³ and provides interpretability to machine learning models
 by computing the contributions of each feature towards the overall model. These explanations of
 machine learning models are models in and of themselves and are referred to as “additive feature
 attribution methods”. These explanation models use simplified inputs x' which are mapped to
 original inputs through a function $x = h_x(x')$ and try to ensure that $g(z') \approx f(h_x(z'))$ whenever
 $z' \approx x'$ where f is the machine learning model and g is the explanation model. This yields the
 additive form which is a linear combination of binary variables:

$$g(z') = \phi_0 + \sum_{i=1}^M \phi_i z'_i$$

1491 where $z' \in \{0, 1\}^M$ is the binary value specifying the inclusion or exclusion of a number of
 1492 simplified feature inputs M ; $\phi_i \in \mathbb{R}$ is the effect of each feature.

1493 Work in⁶¹ devises such a model satisfying three important properties within this framework:

Local accuracy/efficiency: the explanation's features with their effects ϕ_i must sum for each feature i to the output $f(x)$.

$$f(x) = \phi_0(f) + \sum_{i=1}^M \phi_i(f, x)$$

1494 where $\phi_0(f) = \mathbb{E}[f(z)] = f_x(\emptyset)$

Consistency/Monotonicity: If a model changes so that the effect of a feature increases or stays the same regardless of other inputs, that input's attribution should not decrease. For any two models f and f' if,

$$f'_x(S) - f'_x(S \setminus i) \geq f_x(S) - f_x(S \setminus i)$$

1495 where $S \in \mathcal{F}$ are subsets of all features and $S \setminus i$ is the setting of feature i to zero (or
1496 some background reference value intended to be negligible) then $\phi_i(f', x) \geq \phi_i(f, x)$.

Missingness: This is the idea that features with no effect on f_x should have no assigned impact ϕ_i . This is expressed as,

$$f_x(S \cup i) = f_x(S)$$

1497 for all subsets of features $S \in \mathcal{F}$, then $\phi_i(f, x) = 0$.

The authors prove that the only possible additive feature attribution method that satisfies these three criteria is SHAP whose values are computed as the following,

$$\phi_i(f, x) = \sum_{R \in \mathbb{R}} \frac{1}{M!} \left[f_x(P_i^R \cup i) - f_x(P_i^R) \right]$$

1498 where R is the set of all feature orderings and P_i^R is the set of all features that come before the
1499 i^{th} one in ordering R and M is the number of input features.

1500 Extending SHAP values to tree classifiers, the authors create `shap.TreeExplainer`⁶² to calculate
1501 SHAP values by using path-dependent feature perturbation to yield the plots in Fig. 4B and C.

1502 *WaveMAP is stable over hyperparameter choice and random seed*

1503 **Hyperparameter Choice:** Louvain clustering⁴⁹ requires the specification of a resolution hyperparameter, t , which controls the “characteristic scale” by which network “communities” are
1504 identified; the larger this parameter, the fewer the number of clusters (communities) detected and
1505 vice versa¹⁵⁷. We selected a resolution parameter that balanced a maximal modularity score with
1506 a number of clusters selected at the “elbow” of the curve of number of clusters vs. resolution
1507 parameter. Modularity (the “connectedness” of a community, see Methods) is a summed pairwise
1508 measure between nodes in a cluster. It is defined as the difference between the weights of
1509 the edges within a cluster and the edges incoming from any other node outside of the cluster.
1510 Maximizing this value over the whole graph finds communities (sub-graphs) with high amounts
1511 of inter-connectivity. We chose t to be 1.5 which resulted in an average modularity score of
1512 0.761 ± 0.004 (mean \pm S.D.) and an average of 8.29 ± 0.84 (mean \pm S.D.) clusters across 25
1513

1514 random UMAP seeds (Fig. 2B). In this manner, we found a set of waveform clusters the balanced
1515 the diversity found by UMAP and statistical interpretability.

1516 **Robustness of solutions:** To show this hierarchy of clustering resolutions along the curve in
1517 Fig. 2C and to demonstrate *WaveMAP*'s robustness to random seed initialization, we plotted three
1518 different plots for several different resolution parameters in Fig. S2A. Each random seed produced
1519 the same clustering with only slight perturbations of scale and rotation. To validate these clusters
1520 were reliable and not an artifact of our particular data sample, we counted the number of clusters
1521 from 100 randomly permuted subsets of the full dataset at varying proportions (from 10% to
1522 100% in 10% increments) and also set each with a different UMAP random seed (Fig. S2B). As
1523 the data portion increased, we found that the number of clusters increased then tapered off to
1524 around eight clusters at ~60% of the full dataset. At 100% of the dataset, we found 8.5 ± 0.61
1525 (mean \pm S.D.) clusters with residual variability induced by the UMAP random seed. This analysis
1526 is reassuring and demonstrates that we have an adequate number of waveforms to describe the
1527 diversity of cell types in monkey PMd.

1528 **Stability:** Louvain clustering is sometimes unstable. That is, results from successive runs on the
1529 same data can show considerable variation on some datasets¹²⁹. To test whether these eight
1530 clusters consistently contained the same constituent data points run-to-run, we used ensemble
1531 clustering for graphs (ECG)^{128,129}. ECG generates k randomized level-1 (one round of Louvain
1532 clustering, Fig. S1B) partitions and combines together their graph structure via the co-occurrence
1533 of edges between nodes across partitionings. Hence the "ensemble" in the name (also called
1534 "consensus clustering"). Performing ECG with $k=10, 100$ times on UMAP graphs with different
1535 random seeds produced an average of 8.87 ± 0.74 (mean \pm S.D.) clusters which was similar to
1536 that found by Louvain clustering with resolution parameter set to 1.5. In addition, the runs of ECG
1537 that yielded eight clusters had an almost exact structure to that produced by *WaveMAP* (compare
1538 Fig. S2C to Fig. 2A). The runs of ECG with more than eight clusters contained small clusters ($n \leq$
1539 20) splitting off from ⑤ which were too small to allow us to make rigorous conclusions statistically.
1540 We therefore chose the more conservative eight cluster solution that balanced maximizing cluster
1541 number while ensuring adequate cluster sizes.

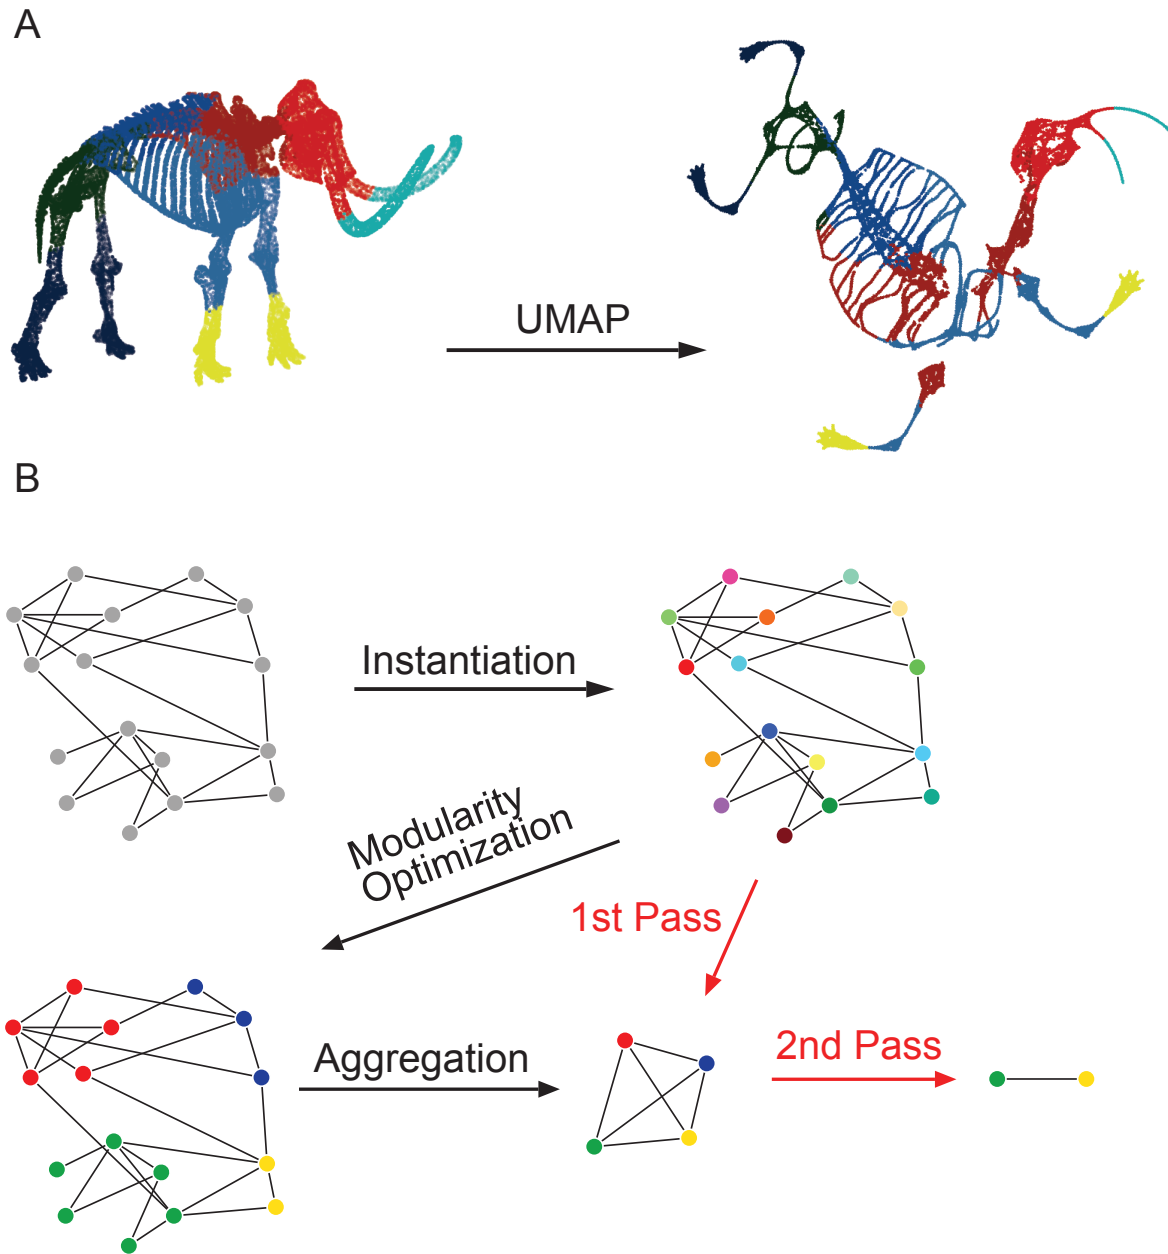

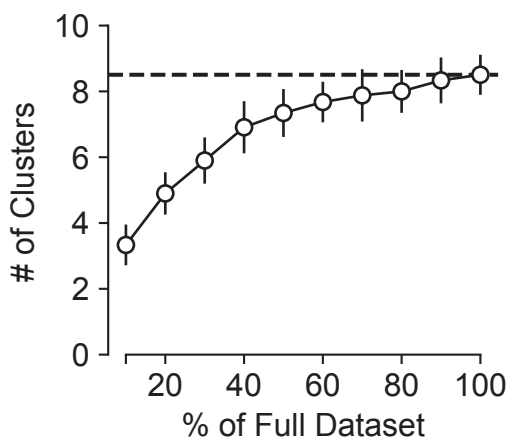


Figure S1: Diagrams of UMAP and Louvain Community Detection: (A) A demonstration of UMAP projection on a 3D point cloud skeleton of a woolly mammoth. Local and global structures are incorporated as evident in the preservation of the individual bone shapes and the sensible spatial relationships between the body parts. Idea from Maximilian Noichl (<https://github.com/MNoichl/UMAP-examples-mammoth->) and mammoth skeleton from the Smithsonian Institute's Smithsonian 3D (<https://3d.si.edu/>). (B) The Louvain community detection algorithm is applied to weighted undirected graphs and proceeds in three steps which are said to be one "pass" of the algorithm: (1) each node is assigned to its own cluster; (2) each node is randomly moved into a neighboring cluster and if modularity increases, it becomes a member of that cluster; (3) once modularity no longer increases, each cluster is collapsed into one node. This process repeats for multiple passes until a stopping condition is reached.

A



B



C




Figure S2: Stability analysis of *WaveMAP*: (A) *WaveMAP* instantiated with several different UMAP random seeds and also Louvain resolution parameters. (B) The mean \pm S.D. number of clusters produced by *WaveMAP* across various proportions of the full dataset. (C) An eight cluster result of ECG applied to the UMAP graph.

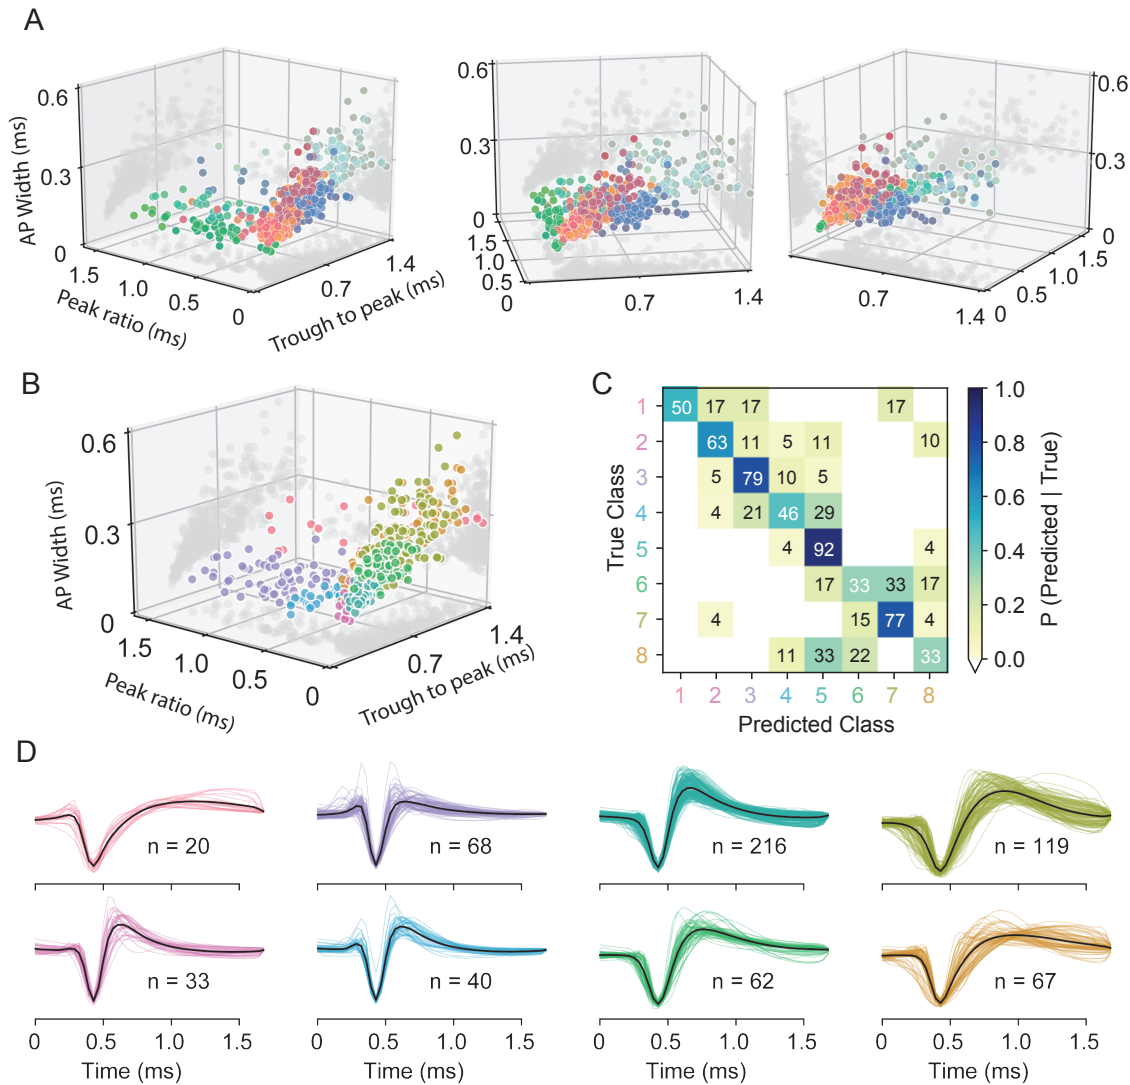


Figure S3: Comparison of GMM and UMAP in the constructed feature space: (A) Three views of the eight *WaveMAP* clusters shown in the constructed feature space. The clusters maintain some structure but are largely mixed and linearly inseparable. **(B)** A GMM instantiated with eight clusters in the constructed feature space. **(C)** Confusion matrix for a random forest classifier with the same hyperparameters as the one trained on four GMM classes S1. Numbers listed are in percent accuracy on the main diagonal and misclassification rate percentage on the off-diagonals against held-out data. **(D)** Each cluster of waveforms in the eight class GMM.

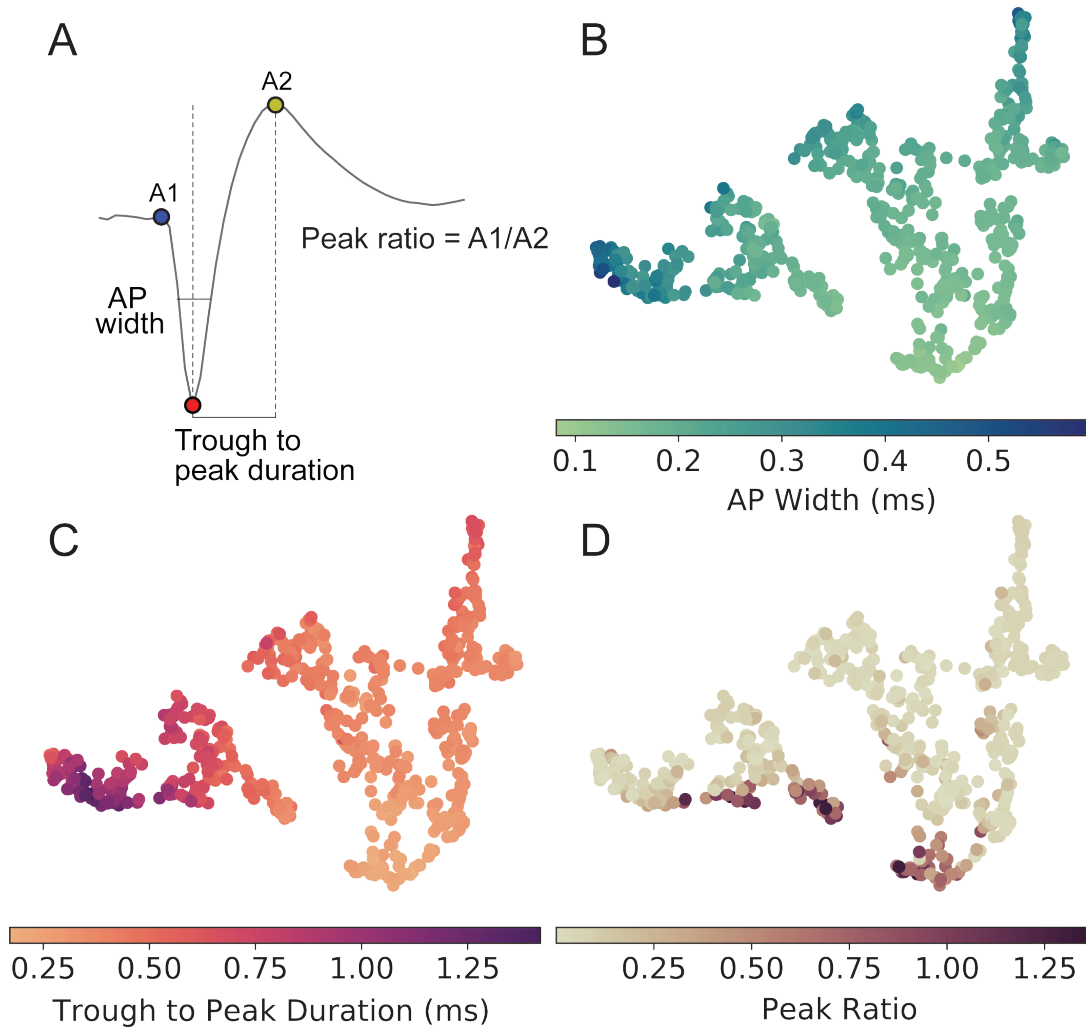


Figure S4: *WaveMAP* implicitly captures waveform features without the need for specification: (A) Three waveform shape features used in traditional clustering approaches. The three EAP waveform landmarks used to generate the specified features passed to the GMM on a sample waveform. ● is the pre-hyperpolarization peak (A1); ● is the depolarization trough; and ● is the post-hyperpolarization peak (A2). AP width is the distance in time between the falling and rising phase of the depolarization trough at its full-width half minimum. The trough to peak duration is the distance between the minimum of the depolarization trough and the peak of the post-hyperpolarization peak. The peak ratio is the height (above zero) of the pre-hyperpolarization peak over the height (again, above zero) of the post-hyperpolarization peak. The same diagram as in Fig. 3A but repeated here. **(B, C, D)** The waveform data points in the projected UMAP space and color coded according to their AP width, trough to peak duration, and peak ratio respectively.

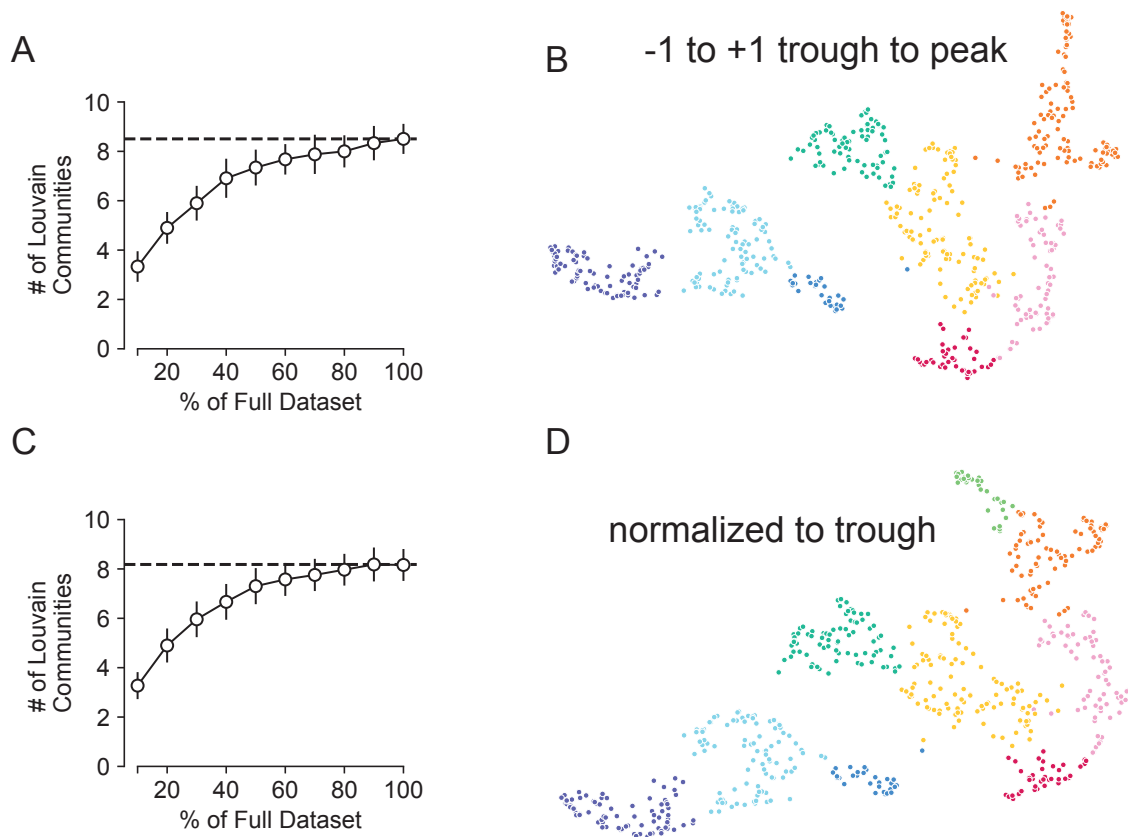


Figure S5: Different normalizations have little effect on WaveMAP structure: (A) As in Fig. S2B, the number of Louvain communities found across various random subsets and random seeds. The mean number of clusters shown on the full dataset with a dashed line. (B) the WaveMAP clusters on waveforms with ± 1 unit trough to peak normalization (used in the paper). (C) The same random subsetting and random seed strategy in (C) applied to waveform data normalized to trough depth. (D) WaveMAP clusters applied to waveform data normalized to trough depth.

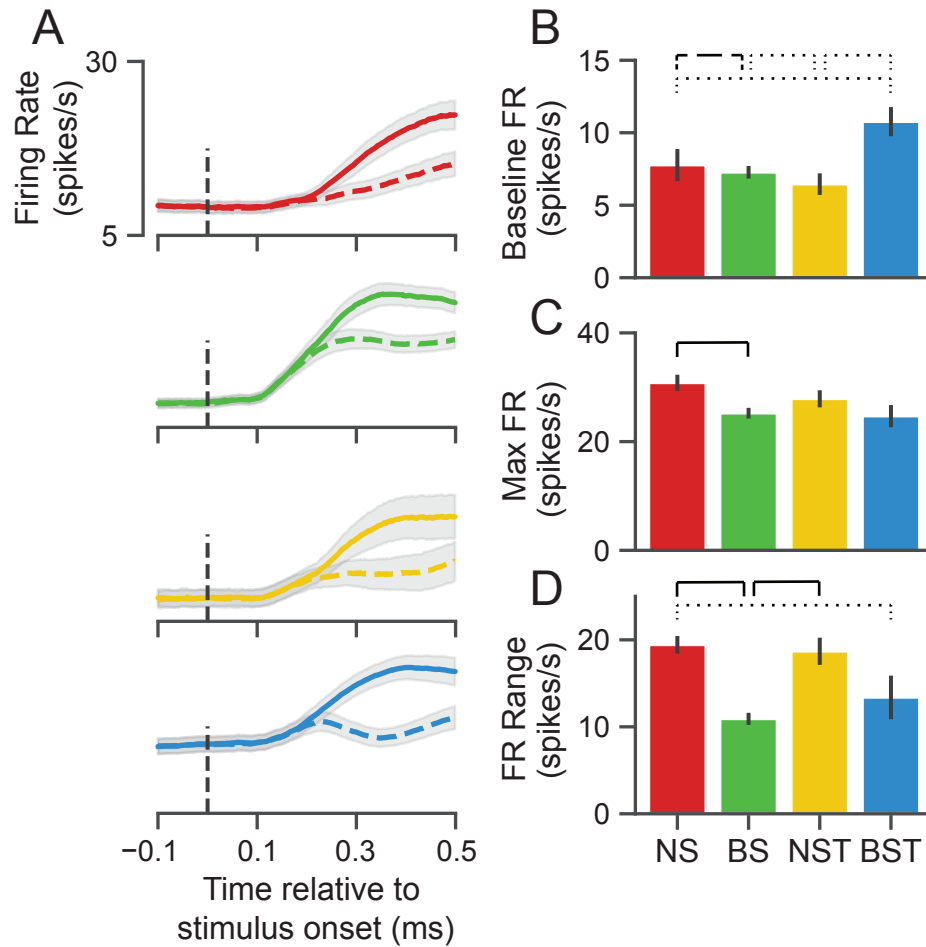


Figure S6: GMM clusters are less physiologically distinguishable than *WaveMAP* clusters: **(A)** Stimulus-aligned trial-averaged firing rate activity in PMd for GMM clusters. As in Fig. 5, the traces are separated into PREF and NONPREF trials with solid and dashed lines respectively. Shaded regions correspond to bootstrapped standard error of the mean. **(B)** Baseline median firing rates for the four GMM clusters. Baselines were calculated as the average firing rate during the first 200 ms of the trial. $p < 0.05$; ---- $p < 0.01$; — $p < 0.005$; Mann-Whitney U test, FDR adjusted.

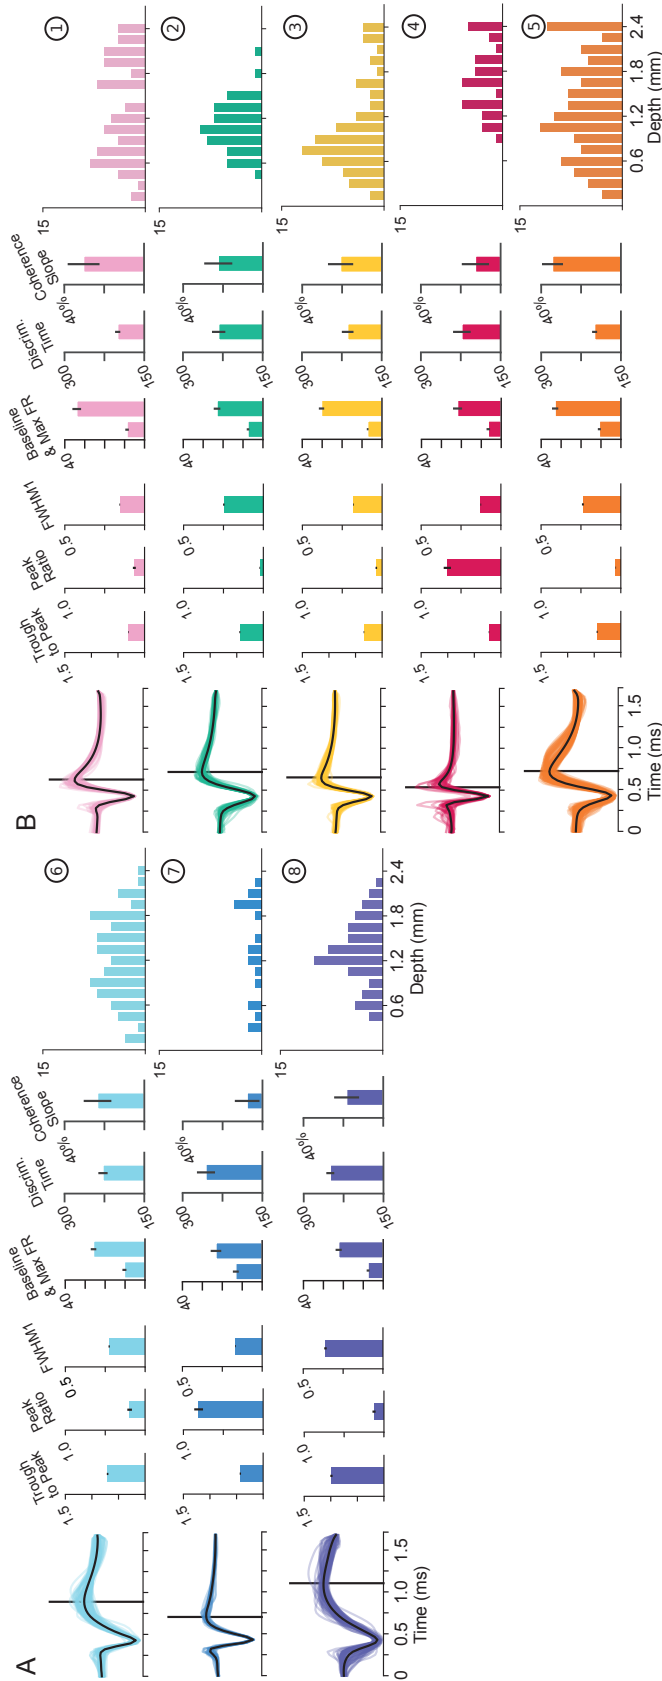


Figure S7: Detailed summary of each UMAP cluster and features: (A, B) A detailed summary of broad- (A) and narrow-spiking (B) cluster waveform shapes, physiological measures, and laminar distribution. Each waveform shape is shown at left with the average waveform shown as a black trace. The average post-hyperpolarization peak position is shown with a black line. The three waveform features used in the GMM classification (Fig. 3A) are shown in the middle as the mean \pm S.E. The baseline and max FR for each cluster are subsequently shown in spikes/s (median \pm bootstrap S.E.). Functional properties, discrimination time and coherence slope, are shown in milliseconds and spikes/s/% coherence (both shown in median \pm bootstrap S.E.). Laminar distributions are also shown with each column in the histogram being the number of each waveform found at each channel location. Channels are spaced every 0.15 mm apart from 0.0 to 2.4 mm.

Function	Function Name	Parameters	Value
UMAP Algorithm (Python)	umap.UMAP	n_neighbors	30
		min_dist	0.1
		random_state	42
		metric	'euclidean'
Louvain Clustering (Python)	cylouvain.best_partition	resolution	1.5
UMAP Random Forest (Python)	xgboost.XGBClassifier	max_depth	4
		min_child_weight	2.5
		n_estimators	100
		learning_rate	0.3
		objective	'binary:logistic'
		rand_state	42
GMM Random Forest (Python)	xgboost.XGBClassifier	max_depth	10
		min_child_weight	2.5
		n_estimators	110
		learning_rate	0.05
		objective	'binary:logistic'
		seed	42
8-Class GMM Random Forest (Python)	xgboost.XGBClassifier	max_depth	4
		min_child_weight	2.5
		n_estimators	100
		learning_rate	0.3
		objective	'binary:logistic'
		seed	42
Gaussian Mixture Model (MATLAB)	fitgmdist	k	4
		start	'randsample'
		replicates	50
		statset('MaxIter')	200

Table S1: Non-default model hyperparameters used

Primary Antibody	Host	Dilution	Vendor	Catalog	Secondary Antibody Conjugate	Host-Antigen	Dilution	Vendor	Catalog
Calbindin D-28k (CB)	Rabbit	1:2000	Swant	CB38	Alexa 546	Donkey anti-Rabbit	1:200	ThermoFisher	A10040
Calretinin D-28k (CR)	Rabbit	1:2000	Swant	7697	Alexa 546	Donkey anti-Rabbit	1:200	ThermoFisher	A10040
Parvalbumin (PV)	Guinea Pig	1:2000	Swant	GP72	Alexa 488	Donkey anti-Guinea Pig	1:200	Jackson	706-545-148

Table S2: Primary and secondary antibodies used in immunohistochemical staining