

Spike Train SIMilarity Space (SSIMS): a framework for single neuron and ensemble data analysis

**Carlos E. Vargas-Irwin¹, David M. Brandman¹, Jonas B. Zimmermann¹,
John P. Donoghue¹, Michael J. Black²**

¹Department of Neuroscience, Brown University, Providence RI, USA.

²Max Planck Institute for Intelligent Systems, Tübingen, Germany.

Keywords: neural population, multi-electrode-array recordings, visualization

Abstract

Increased emphasis on circuit level activity in the brain makes it necessary to have methods to visualize and evaluate large scale ensemble activity, beyond that revealed by raster-histograms or pairwise correlations. We present a method to evaluate the relative similarity of neural spiking patterns by combining spike train distance metrics with dimensionality reduction. Spike train distance metrics provide an estimate of similarity between activity patterns at multiple temporal resolutions. Vectors of pair-wise distances are used to represent the intrinsic relationships between multiple activity patterns at the level of single units or neuronal ensembles. Dimensionality reduction is then used to project the data into concise representations suitable for clustering analysis as well as exploratory visualization. Algorithm performance and robustness are eval-

26 uated using multielectrode ensemble activity data recorded in behaving primates. We
27 demonstrate how Spike train SIMilarity Space (SSIMS) analysis captures the relation-
28 ship between goal directions for an 8-directional reaching task and successfully segre-
29 gates grasp types in a 3D grasping task in the absence of kinematic information. The
30 algorithm enables exploration of virtually any type of neural spiking (time series) data,
31 providing similarity-based clustering of neural activity states with minimal assumptions
32 about potential information encoding models.

33 **1 Introduction**

34 Examining network function at larger and larger scales is now recognized as an impor-
35 tant next step to understand key principles of brain network function and will require
36 new methods to visualize and perform statistical comparisons between activity patterns
37 observed over large sets of neurons (Alivisatos et al., 2013). Neurons often display
38 complex response properties reflecting multiple behavioral and cognitive parameters
39 (Sanes and Donoghue, 2000; Churchland et al., 2010; Rigotti et al., 2013). Character-
40 izing these complex spiking patterns and describing how information from individual
41 neurons is combined at the level of local ensembles and far-reaching networks is an
42 ongoing challenge in neuroscience.

43 Many experiments involve recording ensemble activity (often in multiple areas) un-
44 der various behavioral or cognitive conditions. Data analysis typically involves compar-
45 ing binned firing rates across conditions using standard statistical tests, or fitting neu-
46 ronal responses using models such as cosines or Gaussian distributions (Georgopoulos
47 et al., 1982; Dushanova and Donoghue, 2010; Fluet et al., 2010; Li and DiCarlo, 2010;
48 Pearce and Moran, 2012; Arimura et al., 2013). These methods often involve averag-
49 ing across repetitions of a particular behavior, or otherwise summarizing neural activ-
50 ity patterns to a level where the ensemble properties are reduced to the equivalent of
51 joint perievent histograms. This approach is prone to averaging out changes in a neu-

52 ral activity across trials. Furthermore, this level of data analysis and display becomes
53 impractical as larger ensembles of neurons are recorded simultaneously. Methods to
54 efficiently capture and display both spatial and temporal activity patterns in time series
55 data are essential to both visualize and compare large-scale activity patterns and their
56 relationship to behavior or activity in other brain areas.

57 At their core, most neural data analysis methods are interested in an assessment
58 of similarity. For instance: when an experimental condition is changed, are neuronal
59 spiking patterns similar or different, and what is the relative magnitude of the change?
60 We have formulated a novel technique that provides a quantitative measure of similarity
61 between neuronal firing patterns expressed on individual trials by either single neurons
62 or ensembles. Our approach involves the combination of two key components: spike
63 train distance metrics and dimensionality reduction.

64 Spike train metrics, as developed by Victor and Purpura, provide a measure of sim-
65 ilarity between pairs of spike trains by calculating the most direct way to transform one
66 spike train to another by inserting, deleting, or moving spikes such that both patterns
67 coincide (Victor and Purpura, 1996, 1997; Victor, 2005). Adding up a cost assigned to
68 each of these operations provides quantitative measure of the similarity between activity
69 patterns. The use of spike train metrics makes it possible to analyze long time periods
70 (on the order of seconds) while preserving structure inherent in millisecond scale spike
71 timing. Changing the cost assigned to temporal shifts offers the opportunity to examine
72 neural activity at multiple temporal resolutions.

73 Dimensionality reduction is often accomplished by model fitting, such as by fit-
74 ting tuning functions. When the model relating neural observations with the behav-
75 ior/stimulus is unknown, model-free methods such as principal component analysis can
76 be used to gain insight into the relationship. Here we employ t-Distributed Stochastic
77 Neighbor Embedding (t-SNE) (van der Maaten and Hinton, 2008) to project the high-
78 dimensional space defined by pair-wise spike train distances into a low-dimensional
79 representation which not only facilitates visualization, but also improves pattern dis-

80 crimination. This method is well suited to this type of analysis because it is based on
81 pair-wise similarity estimates and explicitly seeks to preserve the structure within local
82 neighborhoods (in this case, clusters of individual trials with similar activity patterns).

83 The proposed algorithm transforms neural data to produce a low dimensional ‘Spike
84 train SIMilarity Space’ (SSIMS) that represents the relationships between activity pat-
85 terns generated on individual trials. In the SSIMS projection, similar neural activity
86 patterns cluster together, while increasingly different activity patterns are projected fur-
87 ther apart. The degree of similarity between activity patterns of interest can be clearly
88 visualized and quantified. Furthermore, SSIMS projections can be used to evaluate the
89 similarity between training data and new samples, providing a direct basis for pattern
90 classification (decoding). The goal of this report is to describe the method, illustrate its
91 implementation, and examine the strengths and limitations of the approach.

92 We tested and validated the SSIMS algorithm using the activity of multiple single
93 neurons recorded simultaneously in primate primary motor and premotor cortex, suc-
94 cessfully separating neural activity patterns reflecting the behaviors performed in both
95 a planar center-out reaching task and a 3D reaching and grasping task. The method
96 provides a useful framework for data analysis and visualization well suited to the study
97 of large neuronal ensembles engaged in complex behaviors.

98 **2 Description of the SSIMS algorithm**

99 The goal of the SSIMS algorithm is to numerically quantify the similarity between
100 multiple neural activity patterns. We define the ‘state’ of a given ensemble of neurons
101 over a specific time period as the precise timing of each spike fired by each neuron; for
102 example, if the patterns of activity for all neurons during two different time periods can
103 be perfectly aligned, the corresponding ensemble states are considered to be identical.

104 The algorithm consists of two parts. First, pair-wise similarity estimates between
105 spike trains are obtained using the distance metric proposed by Victor and Purpura,

106 which uses a cost function to quantify the addition, deletion or temporal shifting of
107 spikes necessary to transform one spike train into another (Victor and Purpura, 1996).
108 This process results in a high-dimensional space representing pair-wise similarities be-
109 tween the sampled ensemble firing patterns (for example, a series of trials in a be-
110 havioral task). In order to facilitate statistical analysis and data visualization, the sec-
111 ond part of the algorithm refines the high-dimensional space defined in terms of these
112 pair-wise distances using the t-SNE dimensionality reduction technique developed by
113 van der Maaten and Hinton (2008). Within SSIMS projections, distances between
114 points denote the degree of similarity between the ensemble firing patterns (putative
115 network ‘states’) they represent; clustering of points that correlate with experimental
116 labels (such as behavioral conditions) allows an unbiased assessment of the relation-
117 ship between neural states within the context of the experimental variables.

118 **2.1 Measuring the similarity between two spike trains**

119 Victor and Purpura introduced cost-based metrics designed to evaluate the similarity
120 between spike trains (Victor and Purpura, 1996). A given spike train, A , can be trans-
121 formed into second spike train, B , using three basic operations: the addition of a spike,
122 the deletion of a spike, or the shifting of a spike in time. Each of these operations is
123 assigned a ‘cost’; the distance between the two spike trains is defined as the (minimum)
124 summed cost of the operations needed to transform one into the other. The cost of spike
125 insertion or deletion is set to 1, while the cost of shifting a spike in time is set to be
126 proportional to the length of time the spike is to be shifted. This last value is defined
127 using a parameter q , with the cost of shifting a spike being $q\Delta t$. Note that displacing a
128 spike by a time interval $1/q$ has a cost equivalent to deleting it. In this way, the value
129 of q is related to the temporal precision of the presumed spike code, in the sense that
130 it determines how far a spike can be moved in time while still considering it to be the
131 ‘same’ spike (that is, without having to resort to removing it). Setting $q = 0$ makes the

132 timing of a spike irrelevant, reducing all shifting costs to zero. In this case the distance
 133 function is effectively reduced to a difference in spike counts. In this way, this method
 134 can be used to probe possible values for the temporal resolution of neural data, from
 135 millisecond timing to pure rate codes.

136 **2.2 Creating a similarity space based on pair-wise distances**

137 Let us consider a set of n neurons, whose activities are simultaneously recorded over
 138 a set of m trials (with each neuron generating a spike train during each trial). Let
 139 $D_{\text{spike}}(A, B)$ denote the spike train distance metric as defined by Victor and Purpura
 140 (1996): the minimum cost of transforming spike train A into spike train B . Let $S_{i,j}$
 141 represent the spike train recorded from neuron j during the i -th trial. Let the pairwise
 142 similarity vector for spike train $S_{i,j}$ be defined as:

$$143 \quad \mathbf{d}_{\text{pw}}(S_{i,j}) = (D_{\text{spike}}(S_{i,j}, S_{1,j}), D_{\text{spike}}(S_{i,j}, S_{2,j}), \dots, D_{\text{spike}}(S_{i,j}, S_{m,j}))$$

144 Thus, each spike train from a single neuron can be mapped to a m -dimensional space
 145 by representing it as a vector of pair-wise distances to the other spike trains fired by the
 146 same neuron. An ensemble pair-wise similarity vector for trial i is formed by concate-
 147 nating the \mathbf{d}_{pw} vectors of the n neurons:

$$148 \quad \mathbf{D}_i^{\text{ensemble}} = (\mathbf{d}_{\text{pw}}(S_{i,1}), \dots, \mathbf{d}_{\text{pw}}(S_{i,n}))^T$$

149 Thus, the neural activity for each individual trial is represented by a $1 \times mn$ di-
 150 mensional vector which includes m similarity measurements for each neuron. When
 151 the vectors for each of the m trials are combined into a matrix for an ensemble of n
 152 neurons, the result is an $m \times mn$ matrix we refer to as $\mathbf{D}^{\text{ensemble}}$ which constitutes a
 153 relational embedding of the entire data set. Note that in this formulation the informa-
 154 tion obtained from a given neuron is represented in a separate subset of dimensions of

155 the matrix D^{ensemble} (instead of summing cost metrics across neurons to obtain a sin-
156 gle measure of ensemble similarity). The next part of the algorithm seeks to project
157 D^{ensemble} into a lower dimensional space.

158 **2.3 Dimensionality reduction with t-SNE**

159 As we will show later, it is possible to create low dimensional representations based on
160 neural ensemble pairwise similarity data that increase the accuracy of pattern classifi-
161 cation, preserving nearest-neighbor relationships without information loss. The SSIMS
162 method uses the t-SNE algorithm, which is particularly well suited to our approach
163 because it explicitly models the local neighborhood around each point using pair-wise
164 similarity measures (van der Maaten and Hinton, 2008). The general intuition for the
165 algorithm is as follows: given a particular data point in a high dimensional space, one
166 is interested in picking another point that is similar; that is, another point that is in
167 the same ‘local neighborhood’. However, instead of deterministically picking a single
168 closest point, one selects the local neighbor in a stochastic manner, according to a prob-
169 ability (making the probability of selecting points that are close together high, and those
170 that are very far apart low). The set of resulting conditional probabilities (given point A ,
171 what is the likelihood that point B is a local neighbor?) effectively represents similarity
172 between data points. The local neighborhoods around each point are modeled as t-
173 distributions. Rather than using a fixed value for the width of the distribution (σ) across
174 the entire space, the algorithm uses multiple values of σ determined by the data density
175 in the local neighborhood around each point. The span of each of these local neigh-
176 borhoods is determined by the ‘perplexity’ parameter setting of the algorithm, which
177 determines effective number of points to include. Note that if a given dataset contains
178 a dense cluster and a sparse cluster, the size of the local neighborhoods in the sparse
179 cluster will be larger than those in the dense cluster. This dynamic adaptation of local
180 neighborhood size serves to mitigate the ‘crowding problem’, which arises when at-

181 tempting to separate clusters with different densities using a single fixed neighborhood
182 size (which potentially leads to over-sampling the dense cluster or under-sampling the
183 sparse one). Probability distributions describing local neighborhoods are modeled us-
184 ing pair-wise distances, which can be evaluated regardless of the dimensionality of the
185 space. It is therefore possible to compare the similarity of the local neighborhoods for
186 high and low dimensional versions of a given dataset. By minimizing the difference
187 between the two sets of conditional probabilities, the local neighborhood structure is
188 preserved in the low-dimensional mapping.

189 In order to reduce computational complexity, we perform a preliminary round of di-
190 mensionality reduction using principal component analysis (PCA) to project the $\mathbf{D}^{\text{ensemble}}$
191 matrix into a 100-dimensional space. The t-SNE algorithm then refines the resulting lin-
192 ear transform by minimizing the Kullback-Leibler divergence between local neighbor-
193 hood probability functions for this starting point and progressively lower dimensional
194 spaces via gradient descent. Using the terminology from the previous section, the fi-
195 nal output of the t-SNE algorithm is a $mn \times d$ matrix (the t-SNE transform), which
196 projects the $m \times mn$ $\mathbf{D}^{\text{ensemble}}$ matrix into the desired d dimensional space (where n is
197 the number of neurons and m is the number of spike trains).

198 **2.4 Software, Hardware, and processing time**

199 Calculations were performed using MatLab on a Mac workstation with a 2.93 GHz
200 quad-core Intel Xeon processor and 12 GB of RAM. Using this hardware, producing
201 a two-dimensional representation of neural activity for ~100 trials based on the firing
202 patterns of ~100 neurons over one second took, on average, five seconds (including
203 the processing time required to calculate all pair-wise distances between spike trains
204 starting from a list of spike timestamps for each neuron). The source code used for
205 data analysis will be made freely available for non-commercial use at the Donoghue
206 lab website. The algorithm could be modified for near real-time discrete classification

207 in the following manner: First, a training dataset with exemplars in each desired cate-
208 gory would be collected. After calculating all pair-wise distances the t-SNE transform
209 would be calculated as described above (taking only a few seconds after data has been
210 collected). The resulting SSIMS space would provide a relational reference frame to
211 interpret new incoming data. Note that once the t-SNE transform is calculated, pro-
212 jecting new data samples into the resulting SSIMS representation would only take a
213 fraction of the time since the gradient descent part of t-SNE is no longer required. It
214 would still be necessary to calculate pair wise distances for new data samples, but this
215 would involve only m operations per neuron in order to project a new trial into the
216 original SSIMS representation (as opposed to the $n \times m^2$ operations needed to gener-
217 ate the initial embedding). Furthermore, pair-wise distance calculation is well suited
218 to parallel computing and could be further optimized using multi-threading or special-
219 ized hardware. Parallel streams could also be used to independently update the t-SNE
220 transform incorporating new data, providing updated SSIMS embeddings on demand.
221 Overall, the limiting factor on processing time would be the duration of the time win-
222 dow to be analyzed, which would depend on the precise nature of the spiking patterns
223 being classified. The results presented in the following sections suggest that an 8-way
224 classification with $>95\%$ accuracy could be accomplished in under one second.

225 **3 SSIMS algorithm validation using primate cortical en-** 226 **semble activity**

227 Performance of the algorithm was evaluated using cortical ensemble activity recorded
228 in rhesus macaques (*Macaca mulatta*) using 96 channel chronically implanted micro-
229 electrode arrays. Details of the implantation procedure are described in Suner et al.
230 (2005) and Barrese et al. (2013). All procedures were approved by the Brown Univer-
231 sity Institutional Animal Care and Use Committee. Two datasets were used to illustrate

232 the implementation of the method and its properties. The first consisted of neural data
233 recorded in primary motor cortex (MI) from a monkey performing a planar center-out
234 reaching task. The second dataset consisted of neural data recorded in ventral premotor
235 cortex (PMv) from a monkey performing a naturalistic reaching and grasping task that
236 involved intercepting and holding moving objects in a 3D workspace.

237 **3.1 Electrophysiological Recording**

238 During each recording session, signals from up to 96 electrodes were amplified (gain
239 5000), bandpass filtered between 0.3 kHz and 7.5 kHz, and recorded digitally at 30 kHz
240 per channel using a Cerebus acquisition system (Blackrock Microsystems, Salt Lake
241 City, UT). Waveforms were defined in 1.6 ms data windows starting 0.33 ms before the
242 voltage crossed a threshold of at least -4.5 times the channel root mean square variance.
243 These waveforms were then sorted using a density clustering algorithm (Vargas-Irwin
244 and Donoghue, 2007), the results of which were reviewed using Offline Sorter (Plexon,
245 Dallas TX) to eliminate any putative units with multiunit signals (defined by interspike
246 intervals (ISI) < 1 ms) or signal to noise ratios (SNR) less than 1.5.

247 **3.2 Center-out (COUT) task**

248 One monkey was operantly trained to move a cursor that matched the monkey's hand
249 location to targets projected onto a horizontal reflective surface in front of the monkey.
250 The monkey sat in a primate chair with the right arm placed on individualized, cush-
251 ioned arm troughs secured to links of a two-joint exoskeletal robotic arm (KINARM
252 system; BKIN technologies, Kingston ON, Canada; Scott, 1999) underneath an image
253 projection surface that reflected a computer monitor display. The shoulder joint was
254 abducted 85° so that shoulder and elbow movements were made in an approximately
255 horizontal plane. The shoulder and elbow joint angles were digitized at 500 Hz by the
256 motor encoders at the joints of the robotic arm. The x and y positions of the hand were

257 computed using the standard forward kinematic equations and sampled at 200 Hz. For
258 more details on the experimental setup using the KINARM exoskeleton, refer to Rao
259 and Donoghue (2014). Neural data was simultaneously recorded from a chronically
260 implanted microelectrode array in the upper limb area of primary motor cortex. To ini-
261 tiate a trial, the monkey was trained to acquire a target in the center of the workspace.
262 A visual cue was used to signal movement direction during an instructed delay (with
263 duration 1 – 1.6 s) to one of eight radially distributed targets on a screen. At the end of
264 the instructed delay period, the central target was extinguished, instructing the monkey
265 to reach towards the previously cued target. Movement onset was defined as the time
266 when the cursor left the central target. The trajectories for each of the eight movement
267 directions are shown in Figure 1.

268 **3.3 Center-out task: Single neuron properties**

269 We first validated the algorithm by generating SSIMS projections for individual neurons
270 over a time window of one second starting 100 ms before movement onset (using $q =$
271 10, such that $1/q = 100$ ms, SSIMS dimensionality = 2 and t-SNE perplexity = 30).
272 Figure 2 shows two samples of single-neuron SSIMS projections, as well as traditional
273 raster plots. While the raster plots clearly convey the changes in the mean firing rate
274 averaged across trials, it is difficult to discriminate the variability in the firing patterns
275 for each movement direction.

276 The SSIMS plot represents the spike train for each trial as a single point. This rep-
277 resentation shows that the firing patterns for the neuron in Figure 2A are more tightly
278 clustered for the 315° direction (representing a greater degree of similarity). Further-
279 more, the figure reveals that the firing patterns are most similar between 315° and 270°
280 reaches. It is also possible to identify individual 0° trials where this neuron fires in
281 a manner very similar to 315° trials. Note that, in this case, the direction presenting
282 the most tightly clustered firing pattern is not the direction of with the highest firing

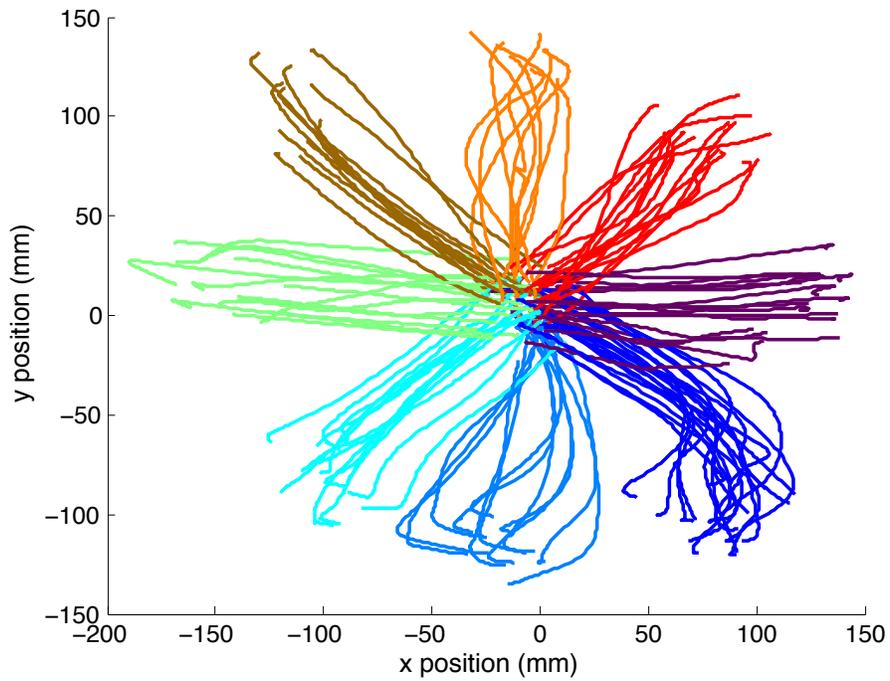


Figure 1: *Center out task kinematics*. The trajectories show the position of the tip of the index finger as the monkey performs a center-out motion to 8 peripheral targets (labeled from 0 to 315°). The trajectories shown were taken from a 1-second time window starting 100 ms before movement onset (corresponding to the main time period used for neural data analysis).

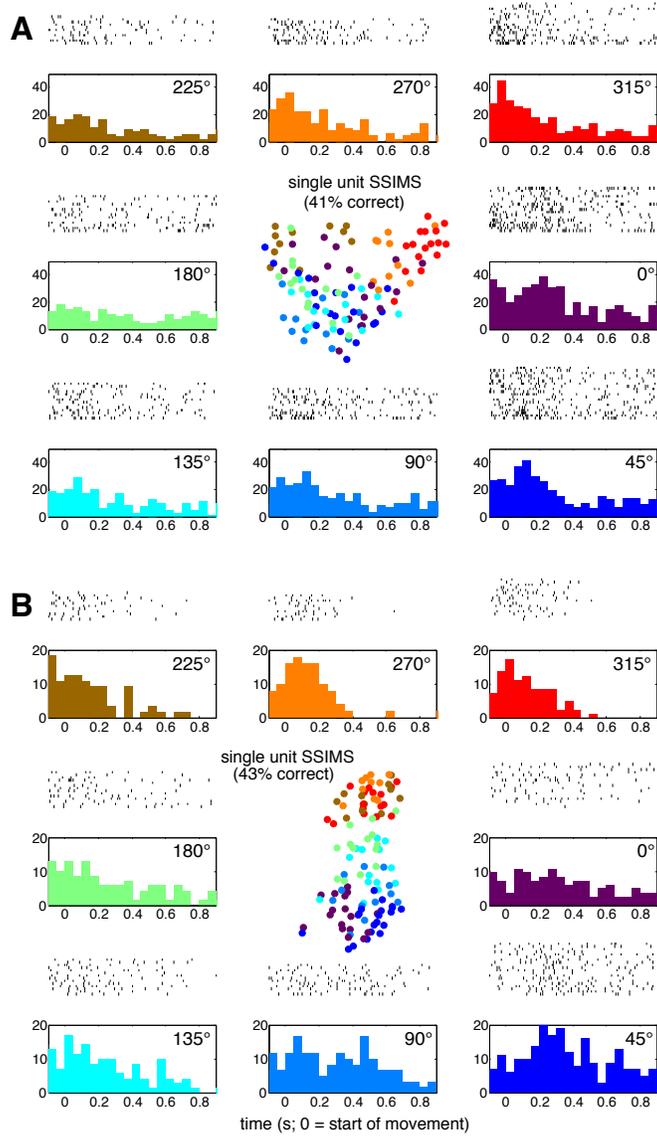


Figure 2: *Single neuron SSIMS in the Center-out task.* *A.* The outer plots show traditional raster-histograms (50 ms bins) for each of the 8 movement directions (radially arranged to represent their relative position on the workspace as shown in Fig. 1). The central plot shows the SSIMS representation for the same data. Each trial shown in the raster plots corresponds to a single point in the SSIMS representation. Color coding is used to match SSIMS points with the corresponding movement directions. A KNN classifier operating on the SSIMS representation of this single unit was capable of correctly predicting the direction of 41% of the trials (see main text for details). *B.* Similar comparison with a second neuron.

283 rate (0°), which would be labeled as the ‘preferred direction’ if firing rates were pa-
284 rameterized with a standard cosine fit. Also note that the most tightly clustered pattern
285 does not correspond to the direction with lowest firing rate, as might be expected if a
286 Poisson noise model is assumed. The neuron shown in Figure 2B is also difficult to de-
287 scribe in terms of standard models, since the timing of the peak in firing rate appears to
288 change as a function of direction. The preferred direction for this neuron would there-
289 fore change as a function of time if it were evaluated using short time windows. The
290 SSIMS algorithm is able to display spiking patterns over a time frame encompassing
291 the entire movement. The resulting plot clearly shows that the greatest difference in
292 spiking patterns exists between 225° , 270° , and 315° reaches compared to 0° and 45° ,
293 with the remaining directions roughly in the middle. This layout reflects the relation-
294 ships between the neural activity patterns observed across reach directions that would
295 be difficult to capture using standard tuning functions.

296 We tested for significant direction-related clustering at the level of single neurons
297 by comparing the distribution of SSIMS distances within and between directions using
298 a Kruskal-Wallis test. Neurons were identified as being directionally selective when the
299 median SSIMS distance was smaller between trials in the same direction compared to
300 trials in different directions. A 10D SSIMS projection was used for this operation, to
301 encompass high dimensional features not visible in 2D projections. Eighty-three out of
302 103 recorded neurons ($\sim 81\%$) were determined to be directionally selective using this
303 method (Kruskall-Wallis $p < 0.001$). For comparison, a Kruskall-Wallis test performed
304 directly on the firing rates for the same time period only produced p values < 0.001 for
305 70% of the neurons.

306 The magnitude of directional selectivity for individual neurons was evaluated us-
307 ing a nearest neighbor (NN) classifier implemented using leave-one-out cross valida-
308 tion. Each trial was classified based on the direction of the nearest neighbor in the 10D
309 SSIMS projection. The percent of correctly classified trials was used as a measure of
310 directional information for a given neuron. The distribution of average single-neuron

311 classification results is shown in Figure 3A. These values were used to rank the neurons
312 from most to least informative.

313 Ensemble decoding was performed using two different strategies: neurons were
314 added to the decoding ensemble from most to least informative (providing an approxi-
315 mate upper bound for classification) or in the reverse order (to generate an approximate
316 lower bound). Classification accuracy (using a KNN classifier with $k = 1$, implemented
317 with leave-one-out cross validation) is shown as a function of ensemble size for both
318 curves in Figure 3B. Figure 3C–F displays the relationship between ensemble activity
319 in each of the 8 movement directions as neurons are progressively added. Although
320 classification was performed in a 10-dimensional space, the SSIMS algorithm was used
321 to project the data down to two dimensions for ease of visualization (classification using
322 2 or 3D SSIMS produced similar results on average, but with greater variability). Note
323 that when the entire ensemble is used, the shape of the clusters matches the directions of
324 movement, generating a circular pattern where clusters are arranged from 0 to 315 de-
325 grees. This structure emerges solely from the relationship between the firing patterns,
326 since clustering is performed without any information about the movement direction
327 associated with each trial. Color coding is added after the fact for visualization; this
328 information about the task is not utilized by the SSIMS algorithm.

329 **3.4 Free Reach-to-Grasp (FRG) task**

330 In the Free Reach-to-Grasp (FRG) task, monkeys were required to intercept and hold
331 objects swinging at the end of a string (Figure 4A). After successfully holding an object
332 for one second, they received a juice reward and were required to release the object to
333 initiate a new trial. The objects were presented at different positions and speeds. Three
334 different objects were used (one at a time) in order to elicit different grasping strate-
335 gies. The first object was a vertical plate 10 cm high by 7 cm wide by 0.3 cm thick.
336 The second object was a vertical 18 cm long cylinder with a 2.5 cm diameter. The third

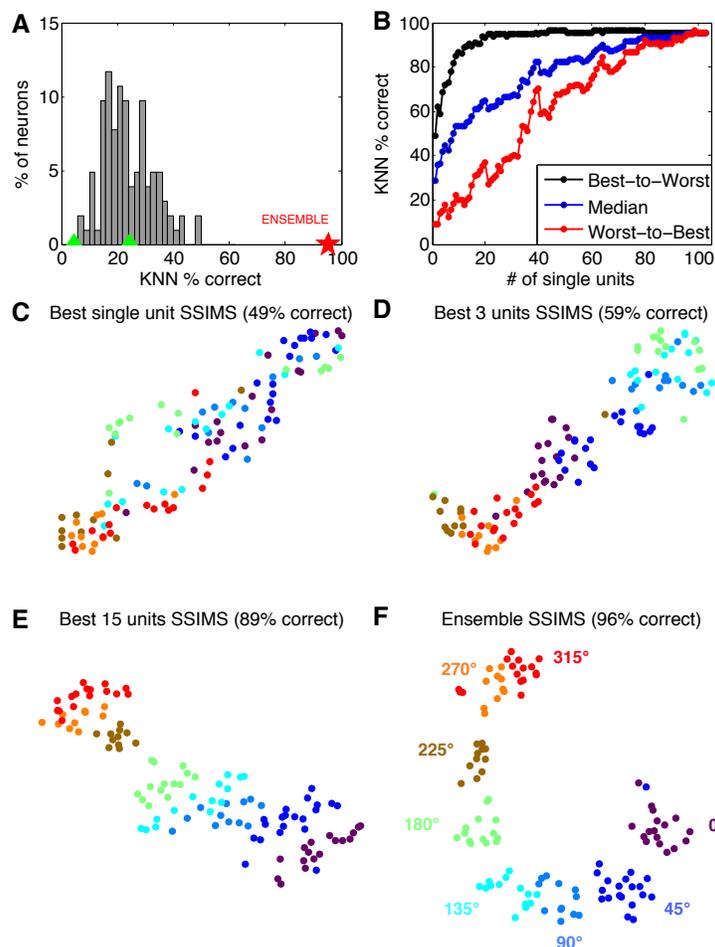


Figure 3: *Center out task: From single neurons to ensembles.* *A.* Single neuron performance in 8-direction classification (10D SSIMS, NN classification using data from individual neurons separately). Classification accuracy using the combined data from all neurons is highlighted with a red star for comparison. Green triangles denote the 95% confidence interval of the chance distribution (calculated over 10,000 random shuffles of the trial labels). *B.* Classification performance as a function of ensemble size (10D SSIMS). Neurons were ranked according to single-unit NN results and added to the decoding ensemble from best to worst (black) or worst to best (red). The median value between these two extremes is shown in blue, representing the expected trend for randomly chosen neurons. *C–F.* SSIMS projections for various ensemble sizes (2D SSIMS). Color coding denotes reach direction using the same conventions as figure 2 (directions are also highlighted in panel F).

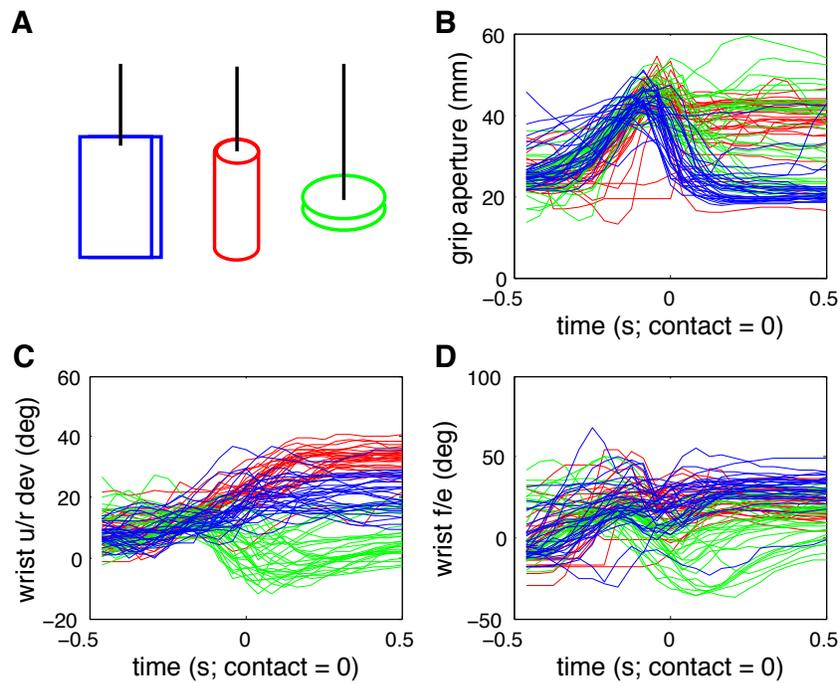


Figure 4: *Free Reach-to-Grasp task kinematics*. A. Diagram of the target objects (not to scale). Each one was presented at the end of a string moving through points in the workspace. B–D. Hand kinematics measured using optical motion capture spanning one second centered on object contact. Color coding matches object color in panel A (blue = vertical plate, red = cylinder, green = disk). Grip aperture was measured as the distance between markers placed on the distal-most joints of the index and thumb. Wrist u/r dev = ulnar/radial deviation; f/e = flexion/extension.

337 object was a horizontal disk 7.5 cm in diameter and 0.3 cm thick. The monkey’s move-
338 ments were measured using an optical motion capture system (Vicon Motion Systems
339 Ltd. UK) to track reflective markers attached to the skin as described in Vargas-Irwin
340 et al. (2010). For this dataset we measured grip aperture (the distance between mark-
341 ers placed on the distal interphalangeal joint of the index finger and thumb) as well as
342 wrist flexion/extension and ulnar/radial deviation subsampled at 24 Hz (Figure 4B–D).
343 Object contact was detected using capacitative switches built into the objects.

344 **3.5 Free Reach-to-Grasp task: Single neuron properties**

345 Spike trains, one second in duration, were recorded from PMv and centered on each
346 successful object contact event (where the grip was maintained for at least one second).
347 Neural activity and kinematics were collected for a total of 90 trials (30 with each
348 object). SSIMS projections for classification were derived from the neural data using
349 $q = 10$, such that $1/q = 100$ ms, SSIMS dimensionality = 10 and t-SNE perplexity
350 = 30.

351 Single unit properties were tested using the same strategy employed in the center-out
352 task. We tested for significant grasping-related clustering by comparing the distribution
353 of SSIMS distances within and between categories using a Kruskal-Wallis test. Neu-
354 rons were identified as being object selective when the median SSIMS distance was
355 smaller between trials with the same object compared to trials with different objects.
356 Forty-seven out of 126 recorded neurons (~37%) were determined to be selective us-
357 ing this method (Kruskall-Wallis $p < 0.001$). For comparison, a Kruskall-Wallis test
358 performed directly on the firing rates for the same time period only produced p values
359 < 0.001 for 19% of the neurons. As with the center-out data, the magnitude of direc-
360 tional selectivity for individual neurons was evaluated using a nearest neighbor (NN)
361 classifier implemented using leave-one-out cross validation. Single-unit classification
362 results are summarized in Figure 5A. These values were used to rank the neurons from

363 most to least informative. Classification accuracy (using a NN classifier) is shown as a
364 function of ensemble size in Figure 5B (for neurons added from best to worst, or in the
365 inverse order).

366 Figure 5C–F displays the relationship between ensemble activity patterns associated
367 with the three objects as neurons are progressively added to the ensemble (for ease of
368 visualization 2D SSIMS projections are shown). The target object clearly emerges as
369 the dominant feature in the SSIMS projections; this can be seen in the post-hoc color
370 coding. Note that this result does not imply that other kinds of information – such as
371 hand position – are not represented in the neural data. With greater numbers of neurons
372 cluster separation and classification performance gradually increase. A NN classifier
373 (implemented with leave-one-out cross validation) applied to the full ensemble SSIMS
374 projections correctly identified the target object in ~96% of the trials, exceeding re-
375 sults obtained using a similar classifier applied directly on all kinematic measurements
376 shown in Figure 4 spanning the same time duration (89% correct). Measuring addi-
377 tional kinematics and or dynamics could potentially narrow the gap between neural and
378 kinematic classification. However, our results demonstrate that the SSIMS algorithm
379 is capable of capturing grasp-related activity patterns with fidelity on par with detailed
380 kinematic measurements. The method can successfully discriminate activity patterns
381 in complex tasks involving many interacting degrees of freedom, and is therefore a po-
382 tentially useful tool for the analysis of high-dimensional motor, sensory, or cognitive
383 neural responses.

384 **3.6 Comparison with other methods**

385 The SSIMS algorithm combines spike train similarity metrics with t-SNE in order to
386 generate low-dimensional representations of neural spiking data. It is possible to gen-
387 erate similar outputs by combining different pre-processing and dimensionality reduc-
388 tion techniques. In order to examine the contributions different approaches, we tested

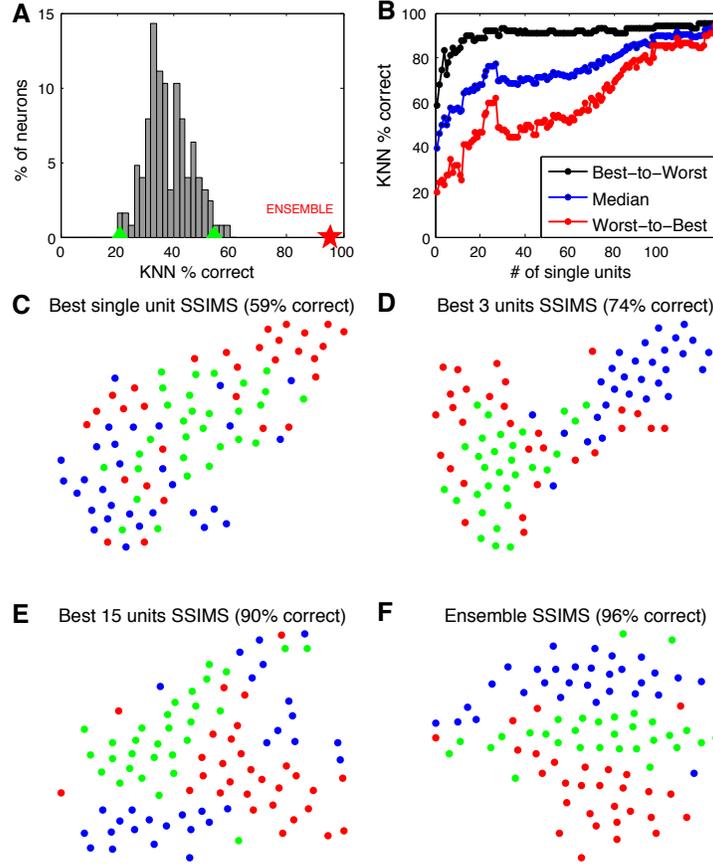


Figure 5: *FRG task: from single neurons to ensembles*. *A*. Single neuron performance in 3-object classification (10D SSIMS, NN classification using data from individual neurons separately). Classification accuracy using the combined data from all neurons is highlighted with a red star for comparison. Green triangles denote the 95% confidence interval of the chance distribution (calculated over 10,000 random shuffles of the trial labels). *B*. Classification performance as a function of ensemble size (10D SSIMS). Neurons were ranked according to single-unit NN results and added to the decoding ensemble from best to worst (black) or worst to best (red). The median value between these two extremes is shown in blue, representing the expected trend for randomly chosen neurons *C–F*. SSIMS projections for various ensemble sizes (2D SSIMS). Color denotes the object being grasped (blue = vertical plate, red = cylinder, green = disk).

389 two pre-processing methods with three dimensionality reduction algorithms. The pre-
390 processing methods analyzed were spike train similarity metrics (SIM) and binned spike
391 counts (SC), while the dimensionality reduction algorithms were t-SNE, multidimen-
392 sional scaling (MDS), and principal component analysis (PCA). Each combination was
393 evaluated using a NN classifier (as described in previous sections) for both the COUT
394 and FRG task data. Each pre-processing method was evaluated at two temporal accu-
395 racy settings (100msec bins, equivalent to $1/q = 100\text{msec}$, and 10msec bins, equivalent
396 to $1/q = 10\text{msec}$). In all comparisons one second of neural data was used. Each di-
397 mensionality reduction algorithm was used to generate a 10D space (well-suited for
398 classification) as well as a 2D space (for ease of visualization). Additionally, we ran
399 the NN classifier on data without the benefit of dimensionality reduction as a baseline
400 comparison. Results are summarized in Table 1.

401 Across all of the comparisons evaluated, methods using spike counts produced, on
402 average, 67% correct classification (s.dev = 20), while methods based on spike train
403 similarity averaged 80%. Methods including PCA averaged 65% (s.dev = 24), while
404 the average for MDS was 74% (s.dev = 20), and the average for t-SNE was 83% (s.dev
405 = 17). For any given task, dimensionality, and temporal accuracy, the combination of
406 techniques used in the SSIMS algorithm (SIM + t-SNE) produced the highest accuracy
407 observed, with the exception of COUT, 2D, and 100msec, where it was 1% below t-SNE
408 + spike counts.

409 Overall, similarity metrics tended to outperform spike counts and produce represen-
410 tations which were more stable across different dimensionality settings. The largest dif-
411 ferences between dimensionality reduction algorithms were observed in the 2D spaces,
412 where t-SNE was clearly superior. For 10D spaces the performance of different al-
413 gorithms was relatively similar (especially when using spike train similarity as a pre-
414 processing step). This pattern suggests that even for cases where discrete classification
415 accuracy for MDS and t-SNE is roughly equivalent, t-SNE consistently produces more
416 informative 2D plots for visualization purposes. Samples of 2D plots produced using

417 different methods are shown in Figs. 6 and 7. Note that PCA fails to capture the circular
418 arrangement of targets in the COUT task (Fig. 6). This pattern is revealed by MDS, but
419 the clusters tend to be more diffuse than those obtained using t-SNE. The differences
420 are more pronounced for the FRG task, where only the full SSIMS algorithm shows a
421 clear recognizable pattern in 2D (Fig. 7).

422 **3.7 Effects of Parameter Setting on SSIMS algorithm performance**

423 We tested the performance of the SSIMS algorithm under a range of parameter settings
424 spanning a range of spike train durations, temporal offsets, dimensionality, temporal
425 resolution (q values), and perplexity. Algorithm performance was evaluated based on
426 classification accuracy of either reaching direction in the COUT task or target object for
427 grasping in the FRG task. In both cases, a nearest neighbor classifier with leave-one-out
428 cross validation was applied as previously described.

429 For both of the tasks examined, accurate pattern classification (greater than 85%
430 correct) was observed for a wide range of time windows (Figure 8). For the COUT task,
431 the most informative time period for direction classification was around the time of start
432 of movement. In the FRG task, the most informative period for grip classification was
433 roughly 500 ms before contact with the object, coinciding with the transport phase that
434 includes hand pre-shaping. The duration of the time window analyzed had a relatively
435 small effect on performance. During the most informative time periods, time windows
436 of as short as 200 ms were sufficient for accurate classification. Extending the time
437 window by an order of magnitude (up to 2 s) did not adversely affect performance.
438 These results show that the SSIMS method is suitable for exploring neural data at a
439 broad range of time scales.

440 We also examined the effect of SSIMS dimensionality and temporal accuracy (q
441 value) on classification performance. For this part of the analysis, we selected fixed 1-
442 second time windows coinciding with highly informative periods in each task: starting

		RAW	RAW	PCA	PCA	MDS	MDS	t-SNE	SSIMS	
		COUNTS	SIM.	COUNTS	SIM.	COUNTS	SIM.	COUNTS		
A	FRG (100ms)	79%	89%	10D	76%	91%	73%	91%	83%	96%*
				2D	40%	29%	42%	58%	70%	87%*
	COJT (100ms)	91%	96%*	10D	96%*	95%	95%	95%	96%*	96%*
				2D	60%	57%	86%	86%	98%*	97%
B	FRG (10ms)	40%	91%	10D	57%	90%	51%	91%	62%	92%*
				2D	36%	42%	42%	50%	44%	88%*
	COJT (10ms)	19%	92%	10D	84%	95%	56%	95%	54%	96%*
				2D	55%	39%	89%	86%	68%	96%*

Table 1: *Pairing neural data pre-processing and dimensionality reduction strategies*
Classification results obtained using a NN classifier on data processed using different combinations of algorithms. Column headings denote the dimensionality reduction algorithm: principal component analysis (PCA), multidimensional scaling (MDS), t-distributed stochastic neighbor embedding (t-SNE), or 'RAW' when no dimensionality reduction was performed. Each column heading also lists the data pre-processing method: spike counts (COUNTS), or spike train similarity metrics (SIM). The highest classification values for each task, dimensionality, and temporal accuracy setting (rows) are highlighted. A. Results for temporal accuracy of 100msec ($1/q = 100\text{msec}$ for SIM, bin size = 100msec. for COUNTS) B. Results for temporal accuracy of 10msec. In 7 out of 8 combinations of dataset, temporal accuracy setting, and dimensionality (table rows) the SSIMS algorithm (t-SNE + SIM) outperformed or matched the classification accuracy obtained using any of the other methods evaluated.

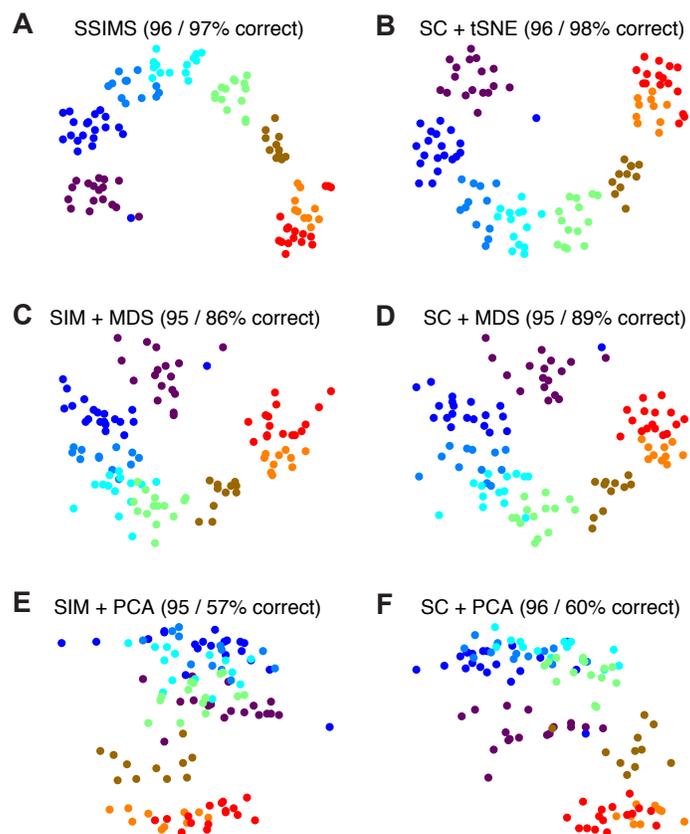


Figure 6: *Neural data visualization: COUT task*. Top row shows results using tSNE for the dimensionality reduction step (A,B), middle row represents MDS (C,D) and bottom row PCA (E,F). Left column shows results for methods using spike train similarity as a pre-processing step (A,C,E), right column shows results for methods based on spike counts (B,D,F).

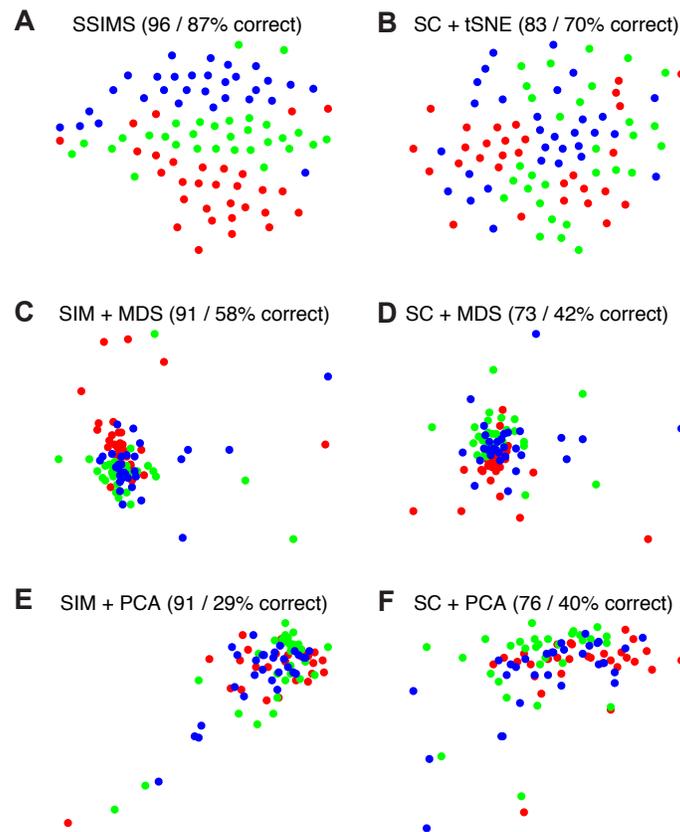


Figure 7: *Neural data visualization: FRG task*. Top row shows results using tSNE for the dimensionality reduction step (A,B), middle row represents MDS (C,D) and bottom row PCA (E,F). Left column shows results for methods using spike train similarity as a pre-processing step (A,C,E), right column shows results for methods based on spike counts (B,D,F).

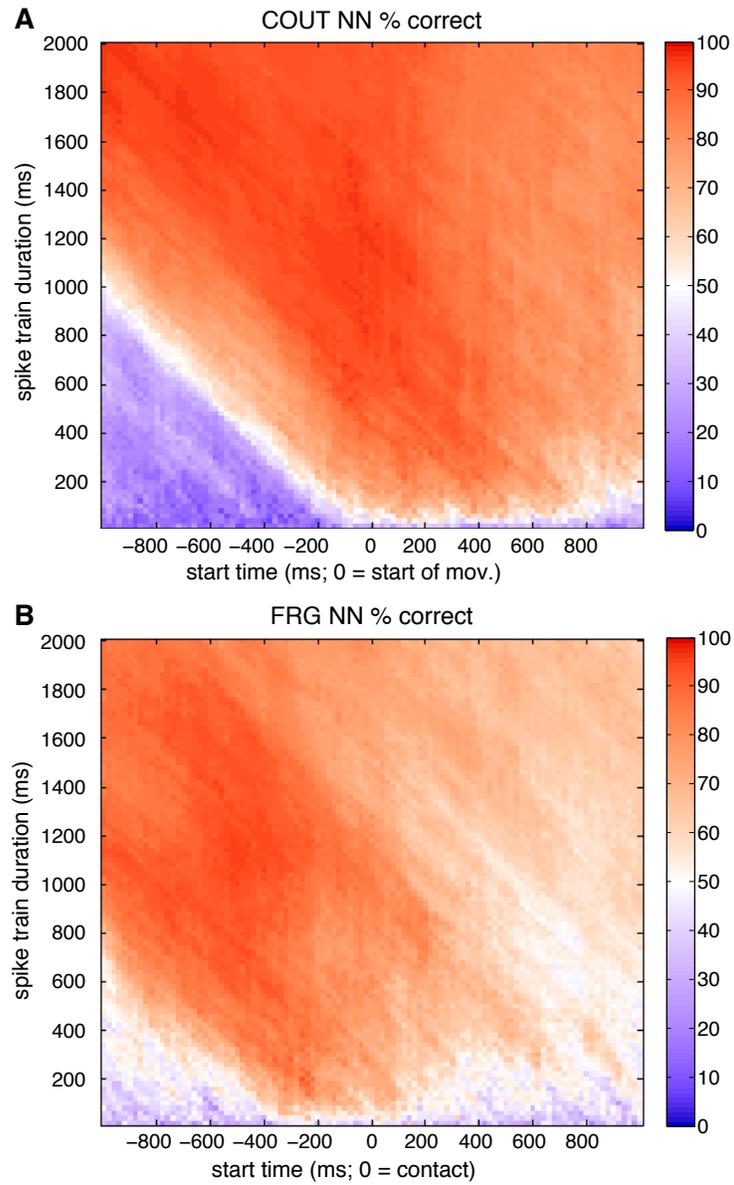


Figure 8: *Effect of spike train duration and temporal offset on SSIMS.* A. Effects of temporal offset and spike train duration on COUT direction classification. The abscissa is the start time for the window used to generate the SSIMS projection (centered around start of movement; negative values are before the onset of movement). The ordinate varies the length of the time window. These results were obtained holding $q = 10$ (corresponding to a temporal precision of 0.1 s), perplexity = 30, and SSIMS dimensionality = 10. B. Effects of temporal offset and spike train duration on FRG grip classification (same conventions as A).

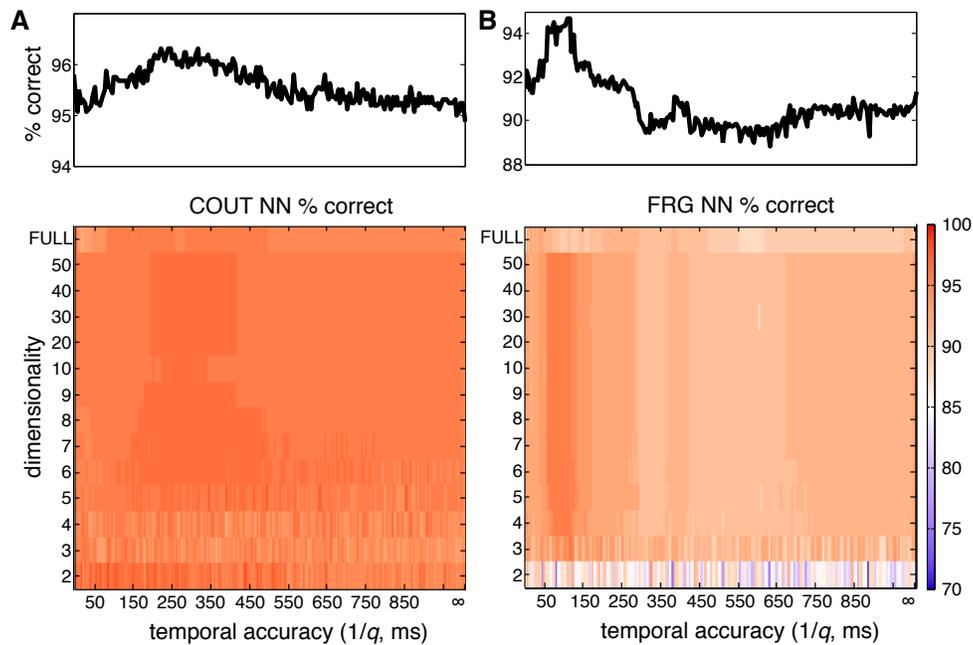


Figure 9: *Effect of dimensionality and q on SSIMS. A.* Effect of q and dimensionality on direction classification in the COUT task. In the ‘FULL’ dimensionality condition classification was performed directly on the pairwise distance matrices without applying t-SNE. Infinite temporal resolution corresponds to setting $q = 0$ (pure rate code). The following parameters were held constant: window start time = -0.1 s, spike train duration = 1 s. The marginal distribution averaging percent correctly classified trials across dimensions is shown above each plot. *B.* Effect of q and dimensionality on grip classification in the FRG task. Same conventions as A. Window start time = -0.5 s, spike train duration = 1 s.

443 100 ms before movement onset for COUT and 500 ms before object contact for FRG.
 444 While holding spike train duration and temporal offset constant, we examined classi-
 445 fication performance as a function of dimensionality and q (Figure 9). For both of the
 446 tasks, dimensionality reduction did not have an adverse effect on classification, suggest-
 447 ing that the low dimensional spaces successfully characterize the patterns present in the
 448 original high dimensional pair-wise similarity matrix. In the COUT task, 2 dimensions

449 were sufficient for accurate decoding, while in the FRG task performance was more
450 stable with 3 or more dimensions. We explicitly tested clustering without the benefit of
451 dimensionality reduction (labeled as ‘FULL’ dimensionality in Figure 9); for both tasks
452 a modest but consistent increase in classification was observed when dimensionality
453 reduction was applied (more pronounced for the FRG task). Adjusting the temporal
454 resolution of the algorithm (q value) produced different effects in the two tasks exam-
455 ined. Recall that q determines the cost of shifting, such that a shift of more than $1/q$
456 has a cost equivalent to removing a spike and inserting a new one. This cutoff deter-
457 mines when the algorithm treats spikes as temporally shifted versions of each other,
458 rather than unrelated events. Changing the value of q had a relatively small effect on
459 classification accuracy for the COUT task. However, there was a gradual trend towards
460 better classification for temporal accuracy values of 250 ms. The FRG task displayed
461 a clearer effect of temporal resolution, with a consistent increase in classification accu-
462 racy for $1/q$ values around 100 ms. Overall, incorporating spike timing provided better
463 performance than assuming a pure rate code (setting $q = 0$). This finding demonstrated
464 the advantage of incorporating spike timing information rather than only spike counts.

465 We also tested the effect of varying the perplexity setting in t-SNE (which deter-
466 mines the effective number of neighbors for each point). Algorithm performance did
467 not vary for perplexity values between 1 and 50 (data not shown).

468 **4 Algorithm validation using synthetic data**

469 In the two data sets analyzed, classification accuracy showed systematic variation as a
470 function of the q settings in the SSIMS algorithm. However, the true degree of tem-
471 poral accuracy for the behaviors examined is not known. In order to test whether the
472 SSIMS algorithm is sensitive to the temporal resolution of spiking patterns, we con-
473 ducted additional tests using synthetic spike trains with predetermined degrees of tem-
474 poral precision. Artificial data was generated based on eight one-second spike trains



Figure 10: *Synthetic spike trains*. Eight different spike trains recorded in primary motor cortex served as templates for synthetic data generation. For each synthetic dataset each spike train was jittered and randomly subsampled removing between 0 and 20% of the spikes. Samples of spike trains jittered by ± 1 , 10, and 100 ms are shown in *A*, *B*, and *C*, respectively.

475 recorded from a sample neuron recorded in the COUT data set (one spike train for each
 476 movement direction). In order to simulate a stochastic response, synthetic spike trains
 477 were generated by applying a random jitter to each recorded spike train (drawn from a
 478 uniform distribution) and then removing a percentage of the spikes (chosen randomly
 479 between 0 and 20%). The magnitude of the introduced jitter was used as a model for
 480 the temporal accuracy of the neural code. Fifty-one synthetic datasets were generated
 481 with jitter values ranging from 1 to 500 ms. Each dataset included 20 samples for each
 482 of the eight directions. Sample spike trains with varying levels of jitter are shown in
 483 (Figure 10).

484 Each synthetic dataset (representing neural codes with varying degrees of temporal

485 consistency) was evaluated in separate runs of the SSIMS algorithm using values of the
486 q parameter ranging from 0 to 1000, resulting in values of $1/q$ ranging from 1000 (ef-
487 fectively infinite) to 1 ms. For all tests performed, the algorithm yielded above chance
488 classification (with a minimum of 40%, significantly above the expected chance value
489 of 12.5% for eight categories). For jitter values of up to 100 ms, the peak in classifi-
490 cation as a function of $1/q$ closely matched the true temporal accuracy (jitter) of the
491 synthetic data (Figure 11A). This observation shows that the SSIMS algorithm can be
492 used to detect precise temporal patterns in spiking data and estimate their precision. As
493 temporal codes progressively deteriorate (at higher jitter values), classification accuracy
494 becomes less sensitive to the q parameter setting (Figure 11B). These findings suggest
495 that optimization of q is not critical for rate-based codes, but can become an important
496 factor in the discrimination of activity patterns where information is contained in the
497 timing of individual spikes.

498 **5 Discussion**

499 Although neuronal spiking patterns contain large amounts of information, parameteriz-
500 ing the outputs of individual neurons is challenging, since their activity often reflects
501 complex interactions of multiple (often unknown) variables and noise, leading to trial-
502 by-trial variation that is difficult to characterize. Furthermore, the response properties of
503 individual neurons are not stationary, but instead are subject to rapid context-dependent
504 changes (Donoghue et al., 1990; Sanes et al., 1992; Hepp-Reymond et al., 1999; Moore
505 et al., 1999; Li et al., 2001; Tolias et al., 2005; Stokes et al., 2013). Limiting data anal-
506 ysis to sub-populations of neurons that can be described using relatively simple models
507 may severely distort conclusions drawn from an experiment and disregard important re-
508 lationships that emerge at large scales. With technological advancements allowing for
509 the simultaneous recording of ensembles approaching thousands of neurons, address-
510 ing these challenges is becoming increasingly important (Grewe et al., 2010; Ahrens

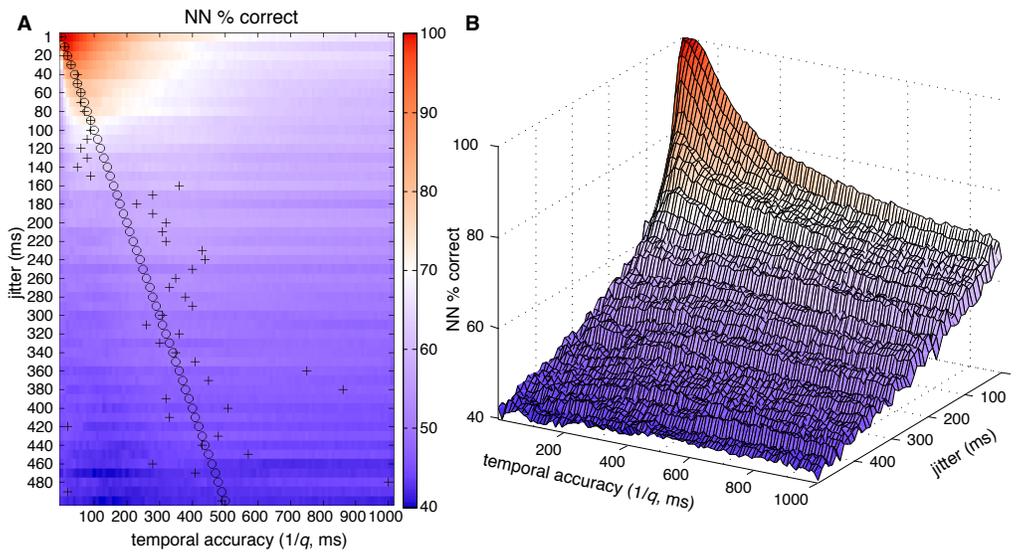


Figure 11: *Estimating the temporal accuracy of neural codes.* *A.* Classification accuracy (using a nearest neighbor classifier implemented using leave-one-out cross validation) is plotted as a function of the jitter used to generate the synthetic data (y -axis) and the q value setting for the SSIMS algorithm (x -axis). Classification results shown are the average value obtained across 20 iterations of synthetic data generation. For each synthetic dataset (row) the jitter value is highlighted by a circle. Similarly, the value of $1/q$ yielding the highest NN classification is highlighted with a '+' sign. *B.* 3D projection of the data presented in panel A. This view highlights the large effects of q parameter settings on classification of spiking patterns with high temporal accuracy (small jitter). The same variation in q has a much less pronounced effect on low accuracy temporal codes (with hundreds of milliseconds of jitter).

511 et al., 2012, 2013). The SSIMS algorithm allows the direct comparison of neuronal
512 firing patterns with minimal assumptions regarding the specific nature of neural encod-
513 ing of the underlying behavioral task or stimulus presentation. Using a similarity-based
514 methodology circumvents the problems of over and under-parameterization: in effect,
515 the templates used to evaluate spiking activity are supplied by the neuronal data. This
516 relational approach is solely based on the intrinsic properties of neural activity, and
517 does not require a direct mapping between neuronal firing patterns and extrinsic vari-
518 ables (measured in the external world). As highlighted in a recent review by Lehky et
519 al., intrinsic, unlabeled, relational, approaches to neural data analysis provide robust,
520 physiologically plausible encoding models (Lehky et al., 2013). Our results demon-
521 strate the flexibility of intrinsic coding implemented in the SSIMS framework. We
522 were able to apply an almost identical analysis (differing only in the number of cate-
523 gories to discriminate) for neural activity elicited in very different behavioral contexts
524 without having to adjust any parameters relating firing patterns to extrinsic variables.
525 Avoiding the need for ‘extrinsic labeling’ is one of the main features that makes this
526 kind of model appealing from a biological standpoint (Lehky et al., 2013).

527 Our results also demonstrate how accurate movement decoding (of either reach di-
528 rection and grip type) can be achieved by applying relatively simple algorithms (such as
529 nearest neighbor classifiers) to SSIMS representations. The algorithm can successfully
530 discriminate between ensemble spiking patterns associated with a planar 8-directional
531 reaching task, accurately reflecting the relationships between reach directions (Fig-
532 ure 3). SSIMS projections can also be used to separate three different grasping strate-
533 gies used in the Free Reach-to-Grasp task, despite the higher number of degrees of
534 freedom engaged (Figure 5). In both tasks, stable cluster separation was achieved over
535 a broad range of physiologically relevant parameter settings (Figures 8, 9). Classifica-
536 tion accuracy was consistently improved by the application of dimensionality reduction
537 as well as the inclusion of spike timing information (Figure 9). This finding highlights
538 the advantages of the two core techniques that form the basis of the SSIMS algorithm.

539 Evaluating neural data under various parameter settings can potentially reveal features
540 related to the inherent dimensionality as well as spike timing precision. In the COUT
541 task, optimal pattern classification was observed with temporal accuracy settings of ap-
542 proximately 250 ms; whereas in the FRG task, classification peaked for $1/q$ values of
543 approximately 100 ms. Although these observations suggest a greater degree of tempo-
544 ral accuracy for spiking during grasping than reaching behaviors, it must be stressed that
545 the values represent only 2 datasets collected from different animals. Further research
546 involving the comparison of multiple subjects engaged in both tasks would be required
547 to explore this hypothesis. Although pursuing this inquiry is outside the scope of the
548 current manuscript, this finding shows how the application of the SSIMS method can be
549 used to fuel data-driven hypothesis generation. The SSIMS algorithm provides outputs
550 that can be conveniently visualized and quantitatively evaluated. Visual examination of
551 the ensemble SSIMS plots makes it easy to fine-tune algorithm performance: for exam-
552 ple, given the overlap between the categories in Figure 3F, we could reasonably expect
553 100% correct classification for a 4-directional decoder. Of course, this prediction as-
554 sumes that the properties of the data being recorded are stable over time, an ongoing
555 challenge for on-line neural control (Barrese et al., 2013). SSIMS visualization may
556 also prove useful in this respect, providing and intuitive display of the trial by trial vari-
557 ation of single-unit or ensemble neural activity patterns which would make it easier to
558 detect and address variations in decoder performance. This kind of application may be
559 a valuable tool for the challenge of developing reliable neuromotor prosthetics.

560 **Relationship to existing neural dimensionality reduction algorithms**

561 Evaluating the information content of neuronal ensembles using machine-learning meth-
562 ods for classification and decoding is a widely used strategy. This approach often in-
563 cludes an implicit element of dimensionality reduction: for example, estimating the 2D
564 position of the arm using a Kalman filter (Wu et al., 2006) is a dimensionality reduction

565 operation guided by kinematic parameters. Other algorithms such as population vector
566 decoding (Georgopoulos et al., 1986) can also be viewed as a kinematic-dependent su-
567 pervised form of dimensionality reduction (since preferred directions must be assigned
568 beforehand). Methods like these require parametrization of neural data with respect
569 to an externally measured covariate. By contrast, relational, intrinsic decoding meth-
570 ods such as SSIMS perform dimensionality reduction in an unsupervised way, with no
571 reference to continuous kinematic variables (Lehky et al., 2013).

572 Non-supervised dimensionality reduction techniques based on principal component
573 analysis (PCA) have also been successfully used to produce concise representations of
574 neural ensemble activity without *a priori* knowledge of external variables (Churchland
575 et al., 2007, 2010, 2012; Mante et al., 2013). This approach has revealed structured
576 transitions from movement preparation to execution not evident using traditional analy-
577 sis methods focusing on single-unit changes in firing rate. Several studies have applied
578 relational encoding methods using multidimensional scaling (MDS) to examine corti-
579 cal ensemble activity in the primate visual system (Young and Yamane, 1992; Rolls
580 and Tovee, 1995; Op de Beeck et al., 2001; Kayaert et al., 2005; Kiani et al., 2007;
581 Lehky and Sereno, 2007). Murata and colleagues have also employed similar methods
582 to examine grasp-related encoding in area AIP (Murata et al., 2000). These studies
583 have successfully generated low-dimensional spaces representing relational coding of
584 different objects and grip strategies.

585 One key difference between the SSIMS algorithm and other methods is the combi-
586 nation of dimensionality reduction with spike train similarity metrics. Instead of repre-
587 senting neuronal activity in terms of firing rates (either binned, or smoothed using a ker-
588 nel function) the SSIMS algorithm applies dimensionality reduction to sets of pair-wise
589 distances between spike trains, allowing for retention of millisecond-level spike timing
590 information. Although it is still necessary to specify a time window, the precise timing
591 of each spike is taken into account; it is therefore possible to examine relatively large
592 time windows without sacrificing temporal resolution. Previous work on spike train

593 metrics revealed no net benefits from the application of dimensionality reduction, aside
594 from convenient visualization (Victor and Purpura, 1997). Our method differs from pre-
595 vious applications in terms of how information from individual neurons is combined.
596 Instead of collapsing ensemble similarity measures by shifting spikes between neurons,
597 our approach keeps information from each neuron segregated until the dimensionality
598 reduction step. Our choice of dimensionality reduction algorithm (t-SNE, as described
599 in van der Maaten and Hinton, 2008) also differs from traditional approaches by using
600 dynamic density estimation to minimize the differences between local neighborhoods
601 in the high and low dimensional spaces.

602 We directly compared the SSIMS algorithm to methods using MDS or PCA imple-
603 mented on data represented in terms of spike counts as well as spike train similarity
604 metrics. Our results show an increase in the accuracy of pattern recognition associated
605 with both components of the SSIMS algorithm (Figs. 6 and 7, Table 1). The combi-
606 nation of spike train similarity with t-SNE allow the SSIMS algorithm to effectively
607 use dimensionality reduction to enhance pattern recognition, improving performance
608 compared to the alternative methods tested.

609 **Limitations and future work**

610 The main application for the SSIMS method is the comparison of discrete experimental
611 conditions with the goal of clustering similar activity patterns. SSIMS coordinates are
612 determined by the relative similarity of the activity patterns analyzed. It is therefore not
613 possible to directly map SSIMS projections generated from different ensembles into the
614 same space (for example, from different subjects, or different brain areas). However,
615 normalized clustering statistics (for example the ratio between within and between-
616 cluster distances) could be used to compare SSIMS representations from different en-
617 sembles. Decoding results (such as the nearest-neighbor classifier demonstrated here)
618 can also be used to quantify and compare the separation between activity patterns from

619 different sources.

620 Although the SSIMS method provides useful visualization and quantification of the
621 main trends present in the data, it should not be regarded as a comprehensive repre-
622 sentation of all the information contained in a given set of neural activity patterns. For
623 example, time-varying continuous variables may fail to produce clear clusters unless
624 there are underlying repeating motifs centered around the time epochs of interest. Fur-
625 thermore, while low dimensional representations may reveal the principal organizing
626 patterns for a dataset, more subtle trends may not be evident without taking into ac-
627 count higher dimensional spaces. Note that while this may hinder visualization, the
628 statistical techniques described for cluster evaluation can be used to determine the opti-
629 mal dimensionality to discriminate patterns in a given task.

630 For the current implementation of the algorithm, it is necessary to align spike trains
631 using an external reference event, which inevitably introduces temporal jitter related
632 to the sensor and detection system used. Metrics based on inter-spike intervals could
633 help mitigate possible misalignments (Victor and Purpura, 1996). Future versions of
634 the algorithm may also refine spike train alignment using other biological signals, such
635 as local field potentials (for example, in addition to comparing the timing of spikes, it
636 may be useful to compare their phase alignment with respect to ongoing oscillations
637 at specific frequencies). The current metric also lacks an explicit model of potential
638 interactions between different neurons. Incorporating similarity between pairs or neu-
639 rons, or measures of synchrony between them could potentially expand the sensitivity
640 of the algorithm. Tracking the evolution of SSIMS cluster statistics using sliding time
641 windows, will also be also possible to see how particular activity patterns converge or
642 diverge over time, providing insight into ensemble dynamics.

643 The t-SNE algorithm is well suited for the separation and classification of neural
644 activity patterns based on pair-wise similarity metrics because of the emphasis it places
645 on comparisons among neighboring points. However, the dynamic density estimation
646 used to define local neighborhoods can potentially have a normalization effect on the

647 variance of individual clusters. Therefore, if the goal of the analysis is to estimate
648 the inherent variability of neural responses in different conditions, it may be better to
649 perform the comparison using other dimensionality reduction methods (or foregoing
650 dimensionality reduction altogether).

651 Note that to demonstrate the application of SSIMS for classification we used a sim-
652 ple NN method. NN, however, is not a part of the main SSIMS algorithm. Of course,
653 more sophisticated classifiers could be applied to the SSIMS output, likely providing
654 further improvements in decoding accuracy.

655 **Conclusion**

656 Understanding the relationship between patterns of activity emerging in large scale neu-
657 ral recordings is a key step in understanding principles of biological information pro-
658 cessing. The SSIMS algorithm provides a widely applicable framework for neural data
659 analysis allowing both straightforward visualization of of an arbitrary number of simul-
660 taneously recorded spike trains and a way to perform precise statistical comparisons
661 between activity patterns. By combining spike train metrics that capture precise spike
662 timing and a dimensionality reduction technique based on pair-wise similarity, we have
663 demonstrated that SSIMS is an effective analytical tool in two dramatically different
664 non-human primate experimental paradigms.

665 The techniques described can be employed beyond the motor domain, providing a
666 way to quantify the relationship between perceptual or cognitive states where kinemat-
667 ics do not provide an intuitive topography. Additionally applying unsupervised cluster-
668 ing algorithms (such as *k*-means) to SSIMS data could reveal clusters of similar neural
669 activity patterns without any *a priori* knowledge of the behavioral context. Using these
670 tools, it may be possible to automatically identify recurring network states as well as the
671 transitions between them, providing an intuitive framework to represent the high level
672 flow of neural computation.

673 **Acknowledgements**

674 We wish to thank Naveen Rao and Lachlan Franquemont for overseeing data collection,
675 Corey Triebwasser for his assistance with animal training, and John Murphy for his help
676 with the design and fabrication of the experimental setup. Research was supported by
677 VA-Rehab RD, NINDS-Javits (NS25074), DARPA (N66001-10-C-2010) and the Katie
678 Samson Foundation.

679 **References**

680 Ahrens, M. B., Li, J. M., Orger, M. B., Robson, D. N., Schier, A. F., Engert, F., and
681 Portugues, R. (2012). Brain-wide neuronal dynamics during motor adaptation in
682 zebrafish. *Nature*, 485(7399):471–7.

683 Ahrens, M. B., Orger, M. B., Robson, D. N., Li, J. M., and Keller, P. J. (2013). Whole-
684 brain functional imaging at cellular resolution using light-sheet microscopy. *Nat.*
685 *Methods*, 10(5):413–20.

686 Alivisatos AP, Chun M, Church GM, Deisseroth K, Donoghue JP, Greenspan RJ,
687 McEuen PL, Roukes ML, Sejnowski TJ, Weiss PS, Yuste R. (2013). The brain activ-
688 ity map *Science*, 339(6125):1284-5.

689 Arimura, N., Nakayama, Y., Yamagata, T., Tanji, J., and Hoshi, E. (2013). Involvement
690 of the globus pallidus in behavioral goal determination and action specification. *J.*
691 *Neurosci.*, 33(34):13639–53.

692 Barrese, J. C., Rao, N., Paroo, K., Triebwasser, C., Vargas-Irwin, C., Franquemont,
693 L., and Donoghue, J. P. (2013). Failure mode analysis of silicon-based intracortical
694 microelectrode arrays in non-human primates. *J. Neural Eng.*, 10(6):066014.

695 Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P.,

696 Ryu, S. I., and Shenoy, K. V. (2012). Neural population dynamics during reaching.
697 *Nature*, 487(7405):51–6.

698 Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., and Shenoy, K. V.
699 (2010). Cortical preparatory activity: representation of movement or first cog in a
700 dynamical machine? *Neuron*, 68(3):387–400.

701 Churchland, M. M., Yu, B. M., Sahani, M., and Shenoy, K. V. (2007). Techniques for
702 extracting single-trial activity patterns from large-scale neural recordings. *Curr Opin*
703 *Neurobiol*, 17(5):609–18.

704 Donoghue, J. P., Suner, S., and Sanes, J. N. (1990). Dynamic organization of primary
705 motor cortex output to target muscles in adult rats. ii. rapid reorganization following
706 motor nerve lesions. *Exp. Brain Res.*, 79(3):492–503.

707 Dushanova, J. and Donoghue, J. (2010). Neurons in primary motor cortex engaged
708 during action observation. *Eur. J. Neurosci.*, 31(2):386–98.

709 Fluet, M.-C., Baumann, M. A., and Scherberger, H. (2010). Context-specific
710 grasp movement representation in macaque ventral premotor cortex. *J. Neurosci.*,
711 30(45):15175–84.

712 Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., and Massey, J. T. (1982). On the re-
713 lations between the direction of two-dimensional arm movements and cell discharge
714 in primate motor cortex. *J. Neurosci.*, 2(11):1527–37.

715 Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986). Neuronal population
716 coding of movement direction. *Science*, 233(4771):1416–9.

717 Grewe, B. F., Langer, D., Kasper, H., Kampa, B. M., and Helmchen, F. (2010).
718 High-speed in vivo calcium imaging reveals neuronal network activity with near-
719 millisecond precision. *Nat. Methods*, 7(5):399–405.

720 Hepp-Reymond, M.-C., Kirkpatrick-Tanner, M., Gabernet, L., Qi, H. X., and Weber, B.
721 (1999). Context-dependent force coding in motor and premotor cortical areas. *Exp.*
722 *Brain Res.*, 128(1-2):123–33.

723 Kayaert, G., Biederman, I., and Vogels, R. (2005). Representation of regular and irreg-
724 ular shapes in macaque inferotemporal cortex. *Cereb. Cortex*, 15(9):1308–21.

725 Kiani, R., Esteky, H., Mirpour, K., and Tanaka, K. (2007). Object category structure
726 in response patterns of neuronal population in monkey inferior temporal cortex. *J.*
727 *Neurophysiol.*, 97(6):4296–309.

728 Lehky, S. R. and Sereno, A. B. (2007). Comparison of shape encoding in primate dorsal
729 and ventral visual pathways. *J. Neurophysiol.*, 97(1):307–19.

730 Lehky, S. R., Sereno, M. E., and Sereno, A. B. (2013). Population coding and the label-
731 ing problem: extrinsic versus intrinsic representations. *Neural Comput.*, 25(9):2235–
732 64.

733 Li, C.-S. R., Padoa-Schioppa, C., and Bizzi, E. (2001). Neuronal correlates of motor
734 performance and motor learning in the primary motor cortex of monkeys adapting to
735 an external force field. *Neuron*, 30(2):593–607.

736 Li, N. and DiCarlo, J. J. (2010). Unsupervised natural visual experience rapidly
737 reshapes size-invariant object representation in inferior temporal cortex. *Neuron*,
738 67(6):1062–75.

739 Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. (2013). Context-dependent
740 computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84.

741 Moore, C. I., Nelson, S. B., and Sur, M. (1999). Dynamics of neuronal processing in
742 rat somatosensory cortex. *Trends. Neurosci.*, 22(11):513–20.

743 Murata, A., Gallese, V., Luppino, G., Kaseda, M., and Sakata, H. (2000). Selectivity for
744 the shape, size, and orientation of objects for grasping in neurons of monkey parietal
745 area aip. *J. Neurophysiol.*, 83(5):2580–601.

746 Op de Beeck, H., Wagemans, J., and Vogels, R. (2001). Inferotemporal neurons
747 represent low-dimensional configurations of parameterized shapes. *Nat. Neurosci.*,
748 4(12):1244–52.

749 Pearce, T. M. and Moran, D. W. (2012). Strategy-dependent encoding of planned arm
750 movements in the dorsal premotor cortex. *Science*, 337(6097):984–8.

751 Rao, N. G. and Donoghue, J. P. (2014). Cue to action processing in motor cortex
752 populations. *J Neurophysiol*, 111(2):441-53.

753 Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., and Fusi,
754 S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*,
755 497(7451):585–90.

756 Rolls, E. T. and Tovee, M. J. (1995). Sparseness of the neuronal representation of
757 stimuli in the primate temporal visual cortex. *J. Neurophysiol.*, 73(2):713–26.

758 Sanes, J. N. and Donoghue, J. P. (2000). Plasticity and primary motor cortex. *Annu.*
759 *Rev. Neurosci.*, 23:393–415.

760 Sanes, J. N., Wang, J., and Donoghue, J. P. (1992). Immediate and delayed changes
761 of rat motor cortical output representation with new forelimb configurations. *Cereb.*
762 *Cortex*, 2(2):141–52.

763 Scott, S. H. (1999). Apparatus for measuring and perturbing shoulder and elbow joint
764 positions and torques during reaching. *J. Neurosci. Methods.*, 89(2):119–27.

765 Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., and Duncan, J. (2013).
766 Dynamic coding for cognitive control in prefrontal cortex. *Neuron*, 78(2):364–75.

767 Suner, S., Fellows, M. R., Vargas-Irwin, C., Nakata, G. K., and Donoghue, J. P. (2005).
768 Reliability of signals from a chronically implanted, silicon-based electrode array in
769 non-human primate primary motor cortex. *IEEE Trans. Neural. Syst. Rehabil. Eng.*,
770 13(4):524–41.

771 Tolias, A. S., Keliris, G. A., Smirnakis, S. M., and Logothetis, N. K. (2005). Neurons in
772 macaque area v4 acquire directional tuning after adaptation to motion stimuli. *Nat.*
773 *Neurosci.*, 8(5):591–3.

774 van der Maaten, L. J. P. and Hinton, G. E. (2008). Visualizing high-dimensional data
775 using t-SNE. *Journal of Machine Learning Research*, (9):2579–2605.

776 Vargas-Irwin, C. and Donoghue, J. P. (2007). Automated spike sorting using density
777 grid contour clustering and subtractive waveform decomposition. *J. Neurosci. Meth-*
778 *ods.*, 164(1):1–18.

779 Vargas-Irwin, C. E., Shakhnarovich, G., Yadollahpour, P., Mislow, J. M. K., Black,
780 M. J., and Donoghue, J. P. (2010). Decoding complete reach and grasp actions from
781 local primary motor cortex populations. *J. Neurosci.*, 30(29):9659–69.

782 Victor, J. D. (2005). Spike train metrics. *Curr Opin Neurobiol*, 15(5):585–92.

783 Victor, J. D. and Purpura, K. P. (1996). Nature and precision of temporal coding in
784 visual cortex: a metric-space analysis. *J. Neurophysiol.*, 76(2):1310–26.

785 Victor, J. D. and Purpura, K. P. (1997). Metric-space analysis of spike trains: theory,
786 algorithms and application. *Network: Computation in Neural Systems*, 8(2):127–164.

787 Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., and Black, M. J. (2006). Bayesian
788 population decoding of motor cortical activity using a kalman filter. *Neural Comput.*,
789 18(1):80–118.

790 Young, M. P. and Yamane, S. (1992). Sparse population coding of faces in the infer-
791 otemporal cortex. *Science*, 256(5061):1327–31.