

# A deep learning framework for inference of single-trial neural population activity from calcium imaging with sub-frame temporal resolution

Feng Zhu<sup>1,2</sup>, Harrison A. Grier<sup>3</sup>, Raghav Tandon<sup>1</sup>, Changjia Cai<sup>4</sup>, Andrea Giovannucci<sup>4,5,6,\*</sup>, Matthew T. Kaufman<sup>7,8,^,\*</sup>, Chethan Pandarinath<sup>1,9,10,^,\*</sup>

1. Wallace H. Coulter Department of Biomedical Engineering, Emory University and Georgia Institute of Technology, Atlanta, GA, USA
2. Neuroscience Graduate Program, Graduate Division of Biological and Biomedical Sciences, Emory University, Atlanta, GA, USA
3. Committee on Computational Neuroscience, The University of Chicago, Chicago, IL, USA
4. Joint Biomedical Engineering Department University of North Carolina at Chapel Hill and North Carolina State University. Chapel Hill, NC, USA.
5. Neuroscience Center, University of North Carolina at Chapel Hill. Chapel Hill, NC, USA.
6. Closed-Loop Engineering for Advanced Rehabilitation (CLEAR). North Carolina State University. Raleigh, NC. USA.
7. Department of Organismal Biology and Anatomy, The University of Chicago, Chicago, IL, USA
8. Neuroscience Institute, The University of Chicago, Chicago, IL, USA
9. Department of Neurosurgery, Emory University, Atlanta, GA, USA
10. Center for Machine Learning, Georgia Institute of Technology, Atlanta, GA, USA

<sup>^</sup> Authors contributed equally

\* Correspondence: chethan [at] gatech.edu, agiovann [at] email.unc.edu, mattkaufman [at] uchicago.edu

## Abstract:

In many brain areas, neural populations act as a coordinated network whose state is tied to behavior on a moment-by-moment basis and millisecond timescale. Two-photon (2p) calcium imaging is a powerful tool to probe network-scale computation, as it can measure the activity of many individual neurons, monitor multiple layers simultaneously, and sample from identified cell types. However, estimating network states and dynamics from 2p measurements has proven challenging because of noise, inherent nonlinearities, and limitations on temporal resolution. Here we describe RADICaL, a deep learning method to overcome these limitations at the population level. RADICaL extends methods that exploit dynamics in spiking activity for application to deconvolved calcium signals, whose statistics and temporal dynamics are quite distinct from electrophysiologically-recorded spikes. It incorporates a novel network training strategy that exploits the timing of 2p sampling to recover network dynamics with high temporal precision. In synthetic tests, RADICaL infers network states more accurately than previous methods, particularly for high-frequency components. In real 2p recordings from sensorimotor areas in mice performing a “water grab” task, RADICaL infers network states with close correspondence to single-trial variations in behavior, and maintains high-quality inference even when neuronal populations are substantially reduced.

## Introduction

In recent years, advances in neural recording technologies have enabled simultaneous monitoring of the activity of large neural populations<sup>1</sup>. These technologies are enabling new insights into how neural populations implement the computations necessary for motor, sensory, and cognitive processes<sup>2</sup>. However, different recording technologies impose distinct tradeoffs in the types of questions that may be asked<sup>3</sup>. Modern electrophysiology enables access to hundreds to thousands of neurons within and across brain areas with high temporal fidelity. Yet in any given area, electrophysiology is limited to a sparse sampling of relatively active, unidentified neurons (**Fig. 1a**). In contrast, two photon (2p) calcium imaging offers the ability to monitor the activity of vast populations of neurons - rapidly increasing from many tens of thousands to millions<sup>4-6</sup> - in 3-D, often with identified layers and cell types of interest<sup>7,8</sup>. Thus 2p imaging is a powerful tool for understanding how neural circuitry gives rise to function.

A key tradeoff, however, is that the fluorescence transients measured via 2p imaging are a low-passed and nonlinearly-distorted transformation of the underlying spiking activity (**Fig. 1b**). Further, because neurons are serially scanned by a laser that traverses the field of view (FOV), a trade-off exists between the size of the FOV (and hence the number of neurons monitored), the sampling frequency, and the pixel size (and therefore the signal-to-noise with which each neuron is sampled). These factors together limit the fidelity with which the activity of large neuronal populations can be monitored and extracted via 2p, and thus limit our ability to link 2p activity to neural computation and behavior on fine timescales.

53 In recent years, a large amount of effort has been dedicated to improving the inference of spike trains from 2p data by  
54 detecting calcium influx events, *i.e.*, time points where single spikes, or multiple spikes in close succession, produce  
55 detectable fluorescence transients<sup>9</sup>. Ideally the spikes-to-fluorescence transformation would be invertible, such that  
56 analyzing calcium events would be equivalent to analyzing spiking activity<sup>3</sup>. However, recent benchmarks illustrate that a  
57 variety of algorithms to infer calcium events reach a similar ceiling of performance and make consistent predictions, and all  
58 achieve limited correspondence to ground truth spiking activity obtained with electrophysiology, particularly on fine  
59 timescales<sup>10,11</sup>. Comparisons of calcium imaging and electrophysiology suggest that the two methods may therefore lead to  
60 divergent scientific findings<sup>3,12–14</sup>, largely due to limitations when inferring spikes from calcium traces.

61  
62 Rather than focusing on the responses of individual neurons, an alternative approach is to characterize patterns of  
63 covariation across a neuronal population to reveal the internal state of the underlying network. These “latent variable models”,  
64 or simply “latent models”, describe each neuron’s activity as a reflection of the network’s state. For example, when applied  
65 to electrophysiological data, latent models assume that an individual neuron’s spiking is a noisy observation of a latent “firing  
66 rate”, which fluctuates in a coordinated way with the firing rates of other neurons in the population. Despite their abstract  
67 nature, the network states inferred by latent models can reveal key insights into the computations being performed by the  
68 brain areas of interest<sup>2</sup>. Inferred network states can also enhance our ability to relate neural activity to behavior. For example,  
69 one state-of-the-art deep learning method to estimate network states from electrophysiological spiking data is Latent Factor  
70 Analysis via Dynamical Systems (LFADS)<sup>15,16</sup>. In applications to data from motor, sensory, and cognitive regions, LFADS  
71 reveals rules that govern how network states progress over time and that are consistent across behavioral conditions, while  
72 also revealing tight correspondences with single-trial behavior on a millisecond timescale<sup>16,17</sup>.

73  
74 Given the success of latent models in uncovering network states from electrophysiological data, here we test whether such  
75 models can achieve accurate inference of network states from activity monitored through 2p calcium imaging. We first begin  
76 with LFADS, and evaluate network state inference using simulated 2p data in which activity reflects known, nonlinear  
77 dynamical systems, and with real 2p data from mice performing a water reaching task. LFADS uncovers network state with  
78 substantially higher accuracy than standard approaches (e.g., deconvolution plus Gaussian smoothing). We then develop  
79 and test a new approach, the Recurrent Autoencoder for Discovering Imaged Calcium Latents (RADICaL), to further improve  
80 inference over LFADS. RADICaL extends LFADS with innovations tailored specifically for 2p data. In particular, we modify  
81 the network architecture to better account for the statistics of deconvolved calcium signals, and develop a novel network  
82 training strategy that exploits the staggered timing of 2p sampling of neuronal populations to achieve subframe temporal  
83 resolution. Our new approach substantially improves inference of network states from 2p data, shown in synthetic data  
84 through accurate recovery of high-frequency features (up to 20 Hz), and in real data through improved prediction of neuronal  
85 activity, as well as prediction of single-trial variability in hand kinematics during rapid reaches (lasting 200–300 ms).  
86 Ultimately, RADICaL provides an avenue to tie precise, population-level descriptions of neural computation with the  
87 anatomical and circuit details revealed via calcium imaging.

## 88 89 **Results**

### 90 **Leveraging population dynamics to infer network states from 2p imaging data**

91 Dynamical systems models such as LFADS rely on two key principles to infer network states from neural population activity.  
92 First, simultaneously recorded neurons exhibit coordinated spatial patterns of activation that reflect the state of the  
93 network<sup>18,19</sup>. Due to this coordination, network states might be reliably estimated even if the measurement of individual  
94 neurons’ activity is unreliable. Second, these coordinated spatial patterns evolve over time based on consistent rules  
95 (dynamics)<sup>2,20</sup>. Thus, while it may be challenging to accurately estimate the network’s state based on activity at a single time  
96 point, knowledge of the network’s dynamics provides further information to help constrain network state estimates using data  
97 from multiple time points.


98  
99 To apply these principles to improve inference from 2p data, we extended LFADS to produce RADICaL (**Fig. 1c**). Both  
100 LFADS and RADICaL model neural population dynamics using recurrent neural networks (RNNs) in a sequential  
101 autoencoder configuration (details in *Methods*, and in previous work<sup>15,16</sup>). This configuration is built on the assumption that  
102 the network states underlying neural population activity can be approximated by an input-driven dynamical system, and that  
103 observed activity is a noisy observation of the state of the dynamical system. The dynamical system itself is modeled by an

104 RNN (the ‘generator’). For any given trial, the time-varying network states can be captured by three pieces of information:  
105 the initial state of the dynamical system (trial-specific), the dynamical rules that govern state evolution (shared across trials),  
106 and any external inputs that may affect the dynamics (trial-specific). The states of the generator are linearly mapped onto a  
107 latent space to produce a ‘factors’ representation, which is then transformed to produce the time-varying output for each  
108 neuron (detailed below). The model has a variety of hyperparameters that control training and prevent overfitting, whose  
109 optimal settings are not known *a priori*. To ensure that these hyperparameters were optimized properly for each dataset, we  
110 built RADICaL on top of a powerful, large-scale hyperparameter optimization framework we recently developed known as  
111 AutoLFADS<sup>17,21</sup>.

### 112 113 **Novel features of RADICaL**

114 RADICaL incorporates two major innovations over LFADS and AutoLFADS. First, we modified RADICaL’s observation model  
115 to better account for the statistics of deconvolved events. In LFADS, discrete spike count data are modeled as samples from  
116 an underlying time-varying Poisson process for each neuron. However, deconvolving 2p calcium signals results in  
117 continuous-valued, time-varying events, with limited correspondence to the actual spike times<sup>10</sup>. In RADICaL, deconvolved  
118 events are therefore modeled as samples from a time-varying zero-inflated gamma (ZIG) distribution, which has been shown  
119 to be more appropriate for calcium data<sup>22</sup>, and whose parameters are taken as the output of the generator RNN (**Fig. 1c**;  
120 details in *Methods*). We then define the network state at any given time point as a vector containing the inferred (i.e., de-  
121 noised) event rates of all neurons, where the de-noised event rate is taken as the mean of each neuron’s inferred ZIG  
122 distribution at each time point. The de-noised event rates are latent variables that are tied to the underlying network state at  
123 each time point. We note that in RADICaL, the generator RNN must produce multiple parameters that control the shape of  
124 the ZIG distributions, which may result in activity at the level of the generator and factors that does not directly correspond  
125 to the biological network’s activity. To avoid this complication, rather than using the factors as an estimate of the biological  
126 network’s state, we used the de-noised event rates. Doing so for both RADICaL and AutoLFADS allowed us to compare  
127 methods as directly as possible.

128  
129 Second, we developed a novel neural network training strategy, selective backpropagation through time (SBTT), that  
130 leverages the precise sampling times of individual neurons to enable recovery of high-frequency network dynamics. Since  
131 standard multiphoton acquisition systems rely on point-by-point raster scanning of a laser beam to acquire frames, it is  
132 possible to increase temporal precision of individual neurons’ sample times beyond the timing of individual frames, by  
133 exploiting the relationship between scanning position and accumulated scanning time within the frame (**Fig. 1d**). To leverage  
134 this information to improve inference of high-frequency network dynamics on single trials, we reframe the underlying  
135 interpolation problem as a missing data problem: we treat low-sampling rate data from each neuron as high-sampling rate  
136 data at the level of the population. In this framing, each neuron is effectively sampled sparsely in time, *i.e.*, the majority of  
137 time points for each neuron do not contain valid data (**Fig. 1e**). Such sparsely sampled data creates a challenge when  
138 training the underlying neural network: briefly, neural networks are trained by adjusting their parameters (weights), and  
139 performing this adjustment requires evaluating the gradient of a cost function with respect to weights. SBTT allows us to  
140 compute this gradient using only the valid data, and ignore the many missing samples (**Fig. 1f**; see *Methods*). Because this  
141 feature only affects how we compute the gradient and update the weights, the network still infers event rates for every neuron  
142 at every time point, regardless of whether samples exist at that time point or not. This allows the trained network to accept  
143 sparsely-sampled observations as input, and produce high-temporal resolution event rate estimates at its output.



**Figure 1 | Improving inference of network states from 2p imaging.** (a) Calcium imaging offers the ability to monitor the activity of hundreds or thousands of neurons, in 3-D, often with cell types of interest and layers identified. In contrast, electrophysiology sparsely samples the neurons in the vicinity of a recording electrode, and may be biased toward neurons with high firing rates. (b) 2p fluorescence transients are a low-passed and lossy transformation of the underlying spiking activity. Spike inference methods may provide a reasonable estimate of neurons' activity on coarse timescales (*left*), but yield poor estimates on fine timescales (*right*; data from ref. <sup>23</sup>). (c) RADiCaL uses a recurrent neural network-based generative model to infer network states - *i.e.*, de-noised event rates for the population of neurons - and assumes a time-varying ZIG observation model. (d) *Top*: in 2p imaging, the laser's serial scanning results in different neurons being sampled at different times within the frame. *Bottom*: individual neurons' sampling times are known with sub-frame precision (colors) but are typically analyzed with whole-frame precision (grey). (e) Sub-frame binning precisely captures individual neurons' sampling times but results in neuron-time points without data. The numbers in the table indicate the deconvolved event in each frame. (f) SBTT is a novel network training method for sparsely sampled data that prevents unsampled time-neuron data points from affecting the gradient computation.


## RADiCaL uncovers high-frequency features from simulated data

144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155  
156  
157  
158

159 We first tested RADiCaL using simulated 2p data, which provides a valuable tool for quantifying performance because the  
 160 underlying network states are known and are parameterizable. We hypothesized that the new features of RADiCaL would  
 161 allow it to infer high-frequency features with greater accuracy than standard approaches, such as Gaussian-smoothing the  
 162 deconvolved events (“s-deconv”) or the simulated fluorescence traces themselves (“s-sim-fluor”), or state-of-the-art tools for  
 163 electrophysiology analysis, such as AutoLFADS. We generated synthetic spike trains by simulating a population of neurons  
 164 whose firing rates were linked to the state of a Lorenz system<sup>15,24</sup> (detailed in *Methods* and **Supp. Fig. 1a**). We ran the  
 165 Lorenz system at various speeds, allowing us to investigate the effects of temporal frequency on the quality of network state  
 166 recovery achieved by different methods. In the 3-dimensional Lorenz system, the Z dimension contains the highest-  
 167 frequency content (**Supp. Fig. 1b**). Here we denote the frequency of each Lorenz simulation by the peak frequency of the  
 168 power spectrum of its Z dimension (**Supp. Fig. 1c**).

169  
 170 We used the synthetic spike trains to generate realistic noisy fluorescence signals consistent with GCAMP6f (detailed in  
 171 *Methods* and **Supp. Fig. 2**). To recreate the variability in sampling times due to 2p laser scanning, fluorescence traces were  
 172 simulated at 100 Hz and then sub-sampled at 33.3 Hz, with offsets in each neuron’s sampling times consistent with spatial  
 173 distributions across a simulated FOV. We then deconvolved the generated fluorescence signals to extract events<sup>25,26</sup>.  
 174 Because RADiCaL uses SBTT, it could be applied directly to the deconvolved events with offset sampling times. In contrast,  
 175 for both AutoLFADS and s-deconv, deconvolved events for all neurons were treated as all having the same sampling times  
 176 (i.e., consistent with the frame times), as is standard in 2p imaging (detailed in *Methods*).

177  
 178 Despite the distortion introduced by the fluorescence simulation and deconvolution process, RADiCaL was able to infer  
 179 event rates that closely resembled the true underlying rates (**Fig. 2a**). To assess whether each method accurately inferred  
 180 the time-varying state of the Lorenz system, we mapped the representations from the different approaches - i.e., the event  
 181 rates inferred by RADiCaL or AutoLFADS, the smoothed deconvolved events, and the smoothed simulated fluorescence  
 182 traces - onto the true underlying Lorenz states using ridge regression. We then quantified performance using the coefficient  
 183 of determination ( $R^2$ ), which quantifies the fraction of the variance of the true latent variables captured by the estimates.  
 184 **Figure 2b** shows the Lorenz Z dimension for example trials from three Lorenz speeds, as well as the recovered values for  
 185 three of the methods. RADiCaL inferred latent states with high fidelity ( $R^2 > 0.8$ ) up to 15 Hz, and significantly outperformed  
 186 other methods across a range of frequencies (**Fig. 2c**; performance for the X and Y dimensions is shown in **Supp. Fig. 3**;  
 187  $p < 0.05$  for all frequencies and dimensions, paired, one-sided t-Test, detailed in *Methods*).



188  
 189 **Figure 2 | Application of RADiCaL to synthetic data.** (a) Example firing rates and spiking activity from a Lorenz system simulated at 7  
 190 Hz, deconvolved calcium events (inputs to RADiCaL), and the corresponding rates and factors inferred by RADiCaL. Simulation  
 191 parameters were tuned so that the performance in inferring spikes using OASIS matched previous benchmarks<sup>10</sup> (see *Methods*). (b) True

192 and inferred Lorenz latent states ( $Z$  dimension) for a single example trial from Lorenz systems simulated at three different frequencies.  
193 Black: true. Colored: inferred. (c) Performance in estimating the Lorenz  $Z$  dimension as a function of simulation frequency was quantified  
194 by variance explained ( $R^2$ ) for all 4 methods.  
195

196 These synthetic results provide an important proof-of-principle that RADICaL can infer high-frequency features of the network  
197 activity underlying 2p signals, which is readily validated when ground truth is known. However, it is important to acknowledge  
198 limitations of the simulation process that might constrain the generality of these results when applying to real data. In  
199 particular, the parameter space is very large, especially considering the variety of calcium indicators, protein expression  
200 patterns, imaging settings, cell types and firing rate patterns. An exhaustive search of this parameter space is infeasible,  
201 and we chose parameters so that the resulting signal-to-noise regime produced similar correlations between real and inferred  
202 spike trains as those typically observed<sup>10</sup> (see *Methods*). However, it is difficult to know whether the results from any  
203 particular choice of simulation parameters (or a variety of choices) can be extrapolated to real experimental conditions. Thus,  
204 we next benchmarked performance on real data to demonstrate RADICaL's utility in the real world.  
205


### 206 **RADICaL improves network state inference in data from a mouse water grab task**

207 We next tested RADICaL on 2p recordings from mice performing a forelimb water grab task (**Fig. 3a, top**). We analyzed data  
208 from four experiments: two mice, and two sessions from each mouse in which different brain areas were imaged (M1, S1).  
209 Our task is a variant of the water-reaching task of Galiñanes & Huber<sup>27</sup>. In each trial, the mouse was cued by the pitch of an  
210 auditory tone to reach to a left or right spout and retrieve a droplet of water with its right forepaw (**Fig. 3a, bottom**; see  
211 *Methods*). The forepaw position was tracked at 150 frames per second with DeepLabCut<sup>6</sup> for 420-560 trials per experiment.  
212 To test whether each method could reveal structure in the neural activity at finer resolution than left vs. right reaches, we  
213 divided trials from each condition into subgroups based on forepaw height during the reach (**Fig. 3a, top right**; see *Methods*).  
214 Two-photon calcium imaging from GCaMP6f transgenic mice was performed at 31 Hz, with 430-510 neurons within the FOV  
215 in each experiment (**Fig. 3b**).  
216

217 With real datasets, a key challenge when benchmarking latent variable inference is the lack of ground truth data for  
218 comparison. A useful first-order assessment is whether the event rates inferred for individual trials match the empirical peri-  
219 stimulus time histograms (PSTHs), *i.e.*, the rates computed by averaging noisy single-trial data across trials with similar  
220 behavioral characteristics<sup>16,17</sup>. While this approach obscures meaningful across-trial variability, it provides a 'de-noised'  
221 estimate that is useful for coarse performance quantification and comparisons. To compute empirical PSTHs, we averaged  
222 the smoothed deconvolved events (s-deconv rates) across trials within each subgroup.  
223

224 We found that RADICaL-inferred event rates recapitulated features of individual neurons' activity that were apparent in the  
225 empirical PSTHs, both when averaging across trials, but also on individual trials (**Fig. 3c**). Importantly, RADICaL is an  
226 unsupervised method, meaning that it was not provided any behavioral information, such as whether the mouse reached to  
227 the left or right on a given trial, or which subgroup a trial fell into. Yet the single-trial event rates inferred by RADICaL showed  
228 clear separation not only between left and right reach conditions, but also between subgroups of trials within each condition.  
229 This separation was not clear with the single-trial s-deconv rates. We quantified the correspondence between the single-trial  
230 inferred event rates and the empirical PSTHs via Pearson's correlation coefficient ( $r$ , see *Methods*). RADICaL single-trial  
231 event rates showed substantially higher correlation with the empirical PSTHs than s-deconv rates (**Fig. 3d**) or those inferred  
232 by AutoLFADS (**Supp. Fig. 4**). Importantly, these improvements were not limited to a handful of neurons, but instead were  
233 broadly distributed across the population.  
234

235 We next tested whether the population activity inferred by RADICaL also showed meaningful structure on individual trials.  
236 We produced low-dimensional visualizations of the population's activity by applying principal component analysis (PCA) to  
237 the RADICaL-inferred or s-deconv event rates after log-transforming and trial-averaging, and then projected the single-trial  
238 event rates (also log-transformed) into the subspace formed by the top three PCs. The low-D trajectories computed from the  
239 RADICaL-inferred rates showed consistent, clear single-trial structure that corresponded to behavioral conditions and  
240 subgroups for all four experiments (**Fig. 3e, top row**; **Supp. Fig. 5, top row**), despite RADICaL receiving no direct information  
241 about which trials belonged to which condition. In comparison, low-D trajectories computed from the s-deconv rates showed  
242 noisy single-trial structure with little correspondence to behavioral subgroups (**Fig. 3e, bottom row**; **Supp. Fig. 5, bottom**  
243 **row**).



**Figure 3 | Application of RADiCaL to real two-photon calcium imaging of a water grab task.** (a) Task. *Top left*: Mouse performing the water grab task. Pink trace shows paw centroid trajectory. *Bottom*: Event sequence/task timing. RT: reaction time. ITI: inter-trial interval. *Top right*: Individual reaches colored by subgroup identity. (b) *Top*: an example field of view (FOV), identified neurons colored randomly. *Bottom left*:  $dF/F$  from a single trial for 5 example neurons. *Bottom right*: Allen Atlas M1/S1 brain regions imaged. (c) Comparison of trial-averaged (left) and single-trial (right) rates for 8 individual neurons for two different brain areas (left vs. right) and two different mice (top half vs. bottom half) for s-deconv and RADiCaL (alternating rows). *Left*: each trace represents a different reach subgroup (4 in total) with error bars indicating s.e.m. *Right*: each trace represents an individual trial (same color scheme as trial-averaged panels). *Odd rows*: s-deconv event rates (Gaussian kernel: 40 ms s.d.). *Even rows*: RADiCaL-inferred event rates. Horizontal scale bar represents 200 ms. Vertical scale bar denotes event rate (a.u.). Vertical dashed line denotes lift onset time. (d) Performance of RADiCaL and s-deconv in capturing the empirical PSTHs on single trials. Correlation coefficient  $r$  was computed between the inferred single-trial event rates and empirical PSTHs. Each point represents an individual neuron. (e) Single-trial neural trajectories derived from RADiCaL

244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255


256  
257  
258  
259  
260

rates (top row) and s-deconv rates (bottom row) for two experiments (*left*: Mouse2 M1; *right*: Mouse1 S1), colored by subgroups. Each trajectory is an individual trial, plotting from 200 ms before to 400 ms after lift onset. Lift onset times are indicated by the dots in the same colors with the trajectories. Grey dots indicate 200 ms prior to lift onset time. Neural trajectories from additional experiments are shown in **Supp. Fig. 5**.

### RADICaL captures dynamics that improve hand kinematics prediction

261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276

We next tested whether the RADICaL-inferred event rates were closely linked to behavior by decoding forepaw positions and velocities from the inferred event rates using cross-validated ridge regression (**Fig. 4a**). Decoding using RADICaL-inferred rates significantly outperformed results from s-deconv rates, or from the AutoLFADS-inferred rates (**Fig. 4b**; position: average  $R^2$  of 0.90 across all experiments, versus 0.72 and 0.83 for s-deconv and AutoLFADS, respectively; velocity: average  $R^2$  of 0.59 across the mice/areas, versus 0.32 and 0.45 for s-deconv and AutoLFADS, respectively;  $p < 0.05$  for position and velocity for all individual experiments, paired, one-sided t-test, detailed in *Methods*). Importantly, the performance advantage was not achieved by simply predicting the mean event rates for all trials of a given condition: RADICaL also outperformed AutoLFADS and s-deconv in decoding the kinematic residuals (i.e., the single-trial deviations from the mean; **Supp. Fig. 7**). To assess how these decoding improvements were distributed as a function of frequency, we computed the coherence between the true and decoded positions and velocities for each method (**Fig. 4c**). RADICaL predictions showed higher coherence with behavior than predictions from s-deconv or AutoLFADS across a wide range of frequencies, and the difference in coherence between RADICaL and AutoLFADS widened (especially for position) at higher frequencies (5-15 Hz). Notably, decoding was improved due to both innovations in RADICaL (i.e., modeling events with a ZIG distribution, and SBTT), and the combination of the two innovations significantly improved performance over each innovation alone (**Supp. Fig. 8**).



277  
278  
279  
280  
281  
282


**Figure 4 | RADICaL improves prediction of behavior.** (a) Decoding hand kinematics using ridge regression. Each column shows an example mouse/area. Row 1: true hand positions trajectories, colored by subgroups. Rows 2–4: predicted hand positions using ridge regression applied to the event rates inferred by RADICaL or AutoLFADS, or s-deconv rates (Gaussian kernel: 40 ms s.d.). Hand positions from additional experiments are shown in **Supp. Fig. 6**. (b) Decoding accuracy was quantified by measuring variance explained ( $R^2$ ) between the true and decoded position (top) and velocity (bottom) across all trials across each of the 4 datasets (2 mice for M1, denoted



283 by squares, and 2 mice for S1, denoted by triangles), for all 3 techniques. (c) Quality of reconstructing the kinematics across frequencies  
284 was quantified by measuring coherence between the true and decoded position (top) and velocity (bottom) for individual trials across all  
285 4 datasets, for all 3 techniques.  
286

### 287 RADICaL retains high decoding performance when reducing the number of neurons used in the model

288 In previous demonstrations on electrophysiological spiking data, LFADS maintained accurate performance in reconstructing  
289 single-trial neural activity and decoding even when reducing the number of sampled neurons<sup>16</sup>. Enabling the same  
290 capabilities for 2p imaging could help mitigate the effects of the tradeoffs in sampling frequency or signal-to-noise that occur  
291 when laser scanning over large FOVs. To evaluate whether this holds for RaDICAL, we performed a neuron-downsampling  
292 experiment where we gradually reduced the number of neurons used in training RADICaL or AutoLFADS (**Fig. 5a**). RADICaL  
293 retained relatively high decoding performance as the population size was reduced (**Fig. 5b**; data from Mouse2 M1). Decoding  
294 performance declined gradually, with a steeper slope for velocity. Notably, however, performance when only 121/439  
295 neurons were used for training RADICaL was similar to that of AutoLFADS - and higher than for s-deconv - even when those  
296 methods were applied to the full population of 439 neurons. Note that this analysis represents a lower bound on performance:  
297 for this proof-of-concept, we simply artificially excluded from our analysis data collected as the laser scanned outside the  
298 restricted FOVs, which resulted in substantial time periods that lacked data entirely (e.g., 2/3 of the total sampling time for  
299 the smallest FOV considered). In a real application, those time periods that were artificially excluded could instead be used  
300 to monitor other brain areas or layers, or to monitor the same neurons with higher sampling rates, either of which might be  
301 expected to provide additional information. These results provide an avenue to retain information by scanning smaller areas  
302 when capturing multiple layers or regions, opening opportunities to study interesting questions such as communication  
303 between layers or interactions between regions (see *Discussion*).



304  
305 **Figure 5 | RADICaL retains high decoding performance in a neuron downsampling experiment.** (a) The area selected to include  
306 was gradually shrunk to the center of the FOV to reduce the number of neurons included in training RADICaL or AutoLFADS. (b) Decoding  
307 performance measured using variance explained ( $R^2$ ) as a function of the number of neurons used in each technique (top: Position;  
308 bottom: Velocity). Data from Mouse2 M1.  
309

## 310 Discussion

311 2p imaging is a widely-used method for interrogating neural circuits, with the potential to monitor vast volumes of neurons  
312 and provide new circuit insights that elude electrophysiology. To date, however, it has proven challenging to precisely infer  
313 network states from imaging data, due in large part to the inherent noise, indicator dynamics, and low temporal resolution  
314 associated with 2p imaging. RADICaL bridges this gap. RADICaL is tailored specifically for 2p imaging, with a noise

emissions model that is appropriate for deconvolved calcium events, and a novel network training strategy (SBTT) that takes advantage of the specifics of 2p laser scanning to achieve substantially higher temporal resolution. Through synthetic tests, we demonstrated that RADICaL accurately infers network states and substantially outperforms alternate approaches in uncovering high-frequency fluctuations. Then, through careful validation on real 2p data, we demonstrated that RADICaL infers network states that are closely linked to single-trial behavioral variability, even on fast timescales. Finally, we demonstrated that RADICaL maintains high-quality inference of network states even as the neural population size is reduced substantially.

The ability to de-noise neural activity on single trials is especially valuable. First, de-noising improves the ability to decode behavioral information from neural activity, allowing subtle relationships between neural activity and behavior to be revealed (**Fig. 4**). Second, de-noising may enable the field to move away from experimental paradigms that evoke the stereotyped behaviors that are needed to facilitate trial-averaging of neural data. This support for reduced stereotypy could allow greater insight in experiments with animals such as mouse and marmoset, where powerful experimental tools are available but highly repeatable behaviors are challenging to achieve. A move away from trial-averaging could also enable better interpretability of more complex or naturalistic behaviors<sup>17,28-31</sup>. Third, this de-noising capability will enable greater insight into processes that fundamentally differ from trial to trial, such as learning from errors<sup>32,33</sup>, variation in internal states such as arousal<sup>34,35</sup>, or paradigms in which tuning to uninstructed movements contaminates measurement of the task-related behavioral variables of interest<sup>36</sup>. Finally, this de-noising also reverses some of the distortions of neural activity introduced by calcium imaging, enabling greatly improved inference of neural dynamics (**Fig. 2**) when compared with known failures using imaging data<sup>3</sup>.

In recent years, a variety of computational methods have been developed to analyze 2p imaging data<sup>9</sup>. 2p preprocessing pipelines<sup>5,26</sup> normally include methods that correct for motion, localize and demix neurons' fluorescence signals, and infer event rates from fluorescence traces. Several studies have applied deep learning in attempts to improve signal quality<sup>37-39</sup>, while a few others have focused on uncovering population-level structure<sup>40-45</sup> or locally linear dynamics underlying population activity, in particular via switching linear dynamical systems-based methods<sup>46,47</sup>. Here we build RADICaL on the AutoLFADS architecture, which leverages deep learning and large-scale distributed training. This enables the integration of more accurate observation models (ZIG) and powerful optimization strategies (SBTT), while potentially inheriting the high performance and generalized applicability previously demonstrated for AutoLFADS<sup>17</sup>.

Many behaviors are performed on fast timescales (e.g., saccades, reaches, movement correction, etc), and thus previous work has made steps in overcoming the limits of modest 2p frame rates in attempts to infer the fast changes in neural firing rates that relate to these fast behaviors. Efforts to chip away at this barrier have relied on regularities imposed by repeated stimuli or highly stereotyped behavior<sup>48,49</sup>, or jittered inferred events on sub-frame timescales to minimize the reconstruction error of the associated fluorescence<sup>37</sup>. RADICaL takes a different approach. In particular, it links subframe timing to neural population dynamics, representing a more powerful and generalizable approach that does not require stereotypy in the behavior or neural response and which could therefore be applied to datasets with more naturalistic or flexible behaviors.

Though we made an effort to test with realistic simulations and on real 2p data from both M1 and S1, it remains untested how RADICaL would generalize to other experimental settings. Noise level can span a wide range in real experiments, depending on the optics, calcium indicators, expression levels, and other factors. Behaviors can vary in complexity and population dynamics can be high-dimensional. Though it is not guaranteed that RADICaL would work in all possible settings, it provides a solution to the spatiotemporal tradeoff that is inherent to any scanning technique, which enables retaining temporal resolution while increasing the spatial area of sampling.

While RADICaL operates on deconvolved calcium events, future work to eliminate this deconvolution step may allow further improvements in temporal resolution. Because deconvolution outputs an event rate for each frame - which is a summary of the cumulative effect of the spikes within the frame - it necessarily discards some high-frequency features in the data. Instead it may be possible to build an end-to-end model that integrates the generative rates-to-fluorescence process and operates on the fluorescence traces directly. Complementary work has begun exploring in this direction<sup>50</sup>, but our unique innovation of selective backprop through time presents an opportunity to greatly improve the quality of recovering high-frequency features when the sampling rate is limited. More broadly, carefully-designed benchmarking efforts for network state inference

367 from 2p data could provide an invaluable resource for systematically comparing methods and building on advances from  
368 various different developers.

369  
370 The ability to achieve high-quality network state inference despite limited neuronal population size opens the door to testing  
371 new choices about how to scan during the experiments themselves, e.g., by scanning smaller and more disparate FOVs  
372 (across layers<sup>7,8</sup> or brain areas<sup>51,52</sup>) to understand how spatially-segregated populations interact, while potentially preserving  
373 the ability to infer network states from each FOV. When the number of neurons within each FOV is limited, one advantage  
374 that RADICaL inherits from LFADS is that it allows for multi-session stitching<sup>16</sup>, which could provide an avenue to combine  
375 data from different sessions to improve inference of the underlying dynamics for each FOV.

376  
377 In sum, RADICaL provides a framework to push back the limits of the space-time tradeoff in 2p calcium imaging, enabling  
378 accurate inference of population dynamics in vast populations and with identified neurons. Future work will explore how best  
379 to exploit these capabilities for different experimental paradigms, and to link the power of dynamics with the anatomical detail  
380 revealed with calcium imaging.

## 381 382 **Acknowledgements**

383 We thank M. Rivers and R. Vecovi for help with the real-time camera setup, and D. Sabatini for contributions to the  
384 behavioral control software. This work was supported by the Emory Neuromodulation and Technology Innovation Center  
385 (ENTICe), NSF NCS 1835364, NIH Eunice Kennedy Shriver NICHD K12HD073945, the Simons Foundation as part of the  
386 Simons-Emory International Consortium on Motor Control (CP), the Alfred P. Sloan Foundation (CP, MTK), NSF NCS  
387 1835390, The University of Chicago, the Neuroscience Institute at The University of Chicago (MTK), and a Beckman Young  
388 Investigators Award (AG). The work was also supported by the following collaborative awards (PI: Prof. Ellen Hess, Emory):  
389 NIH NINDS R21 NS116311, Imagine, Innovate and Impact (I3) Funds from the Emory School of Medicine and through the  
390 Georgia CTSA NIH UL1-TR002378, and a pilot grant from the Emory Udall Center of Excellence for Parkinson's Research.

## 391 392 **Code availability**

393 Code will be made available upon publication.

## 394 395 **Data availability**

396 Data will be made available upon reasonable request at the time of publication.

## 397 398 **Author Contributions**

399 F.Z. and C.P. designed the study, with input from A.G. and M.K.. C.P. and M.T.K. conceptualized the SBTT approach.  
400 F.Z. and C.P. performed analyses and wrote the manuscript with input from all other authors. F.Z. and C.P. developed  
401 the algorithmic approach. F.Z., C.C., and A.G. developed the simulation pipeline. H.G. and M.K. designed and performed  
402 experiments with mice, and developed the real data preprocessing pipeline with input from F.Z. and C.P.. R.T. contributed  
403 to initial simulations and data analysis. F.Z., A.G., M.K., and C.P. edited and revised the manuscript with input from all  
404 other authors.

## 405 406 **References**

- 407 1. Stevenson, I. H. & Kording, K. P. How advances in neural recording affect data analysis. *Nat Neurosci* **14**,  
408 139–142 (2011).
- 409 2. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation Through Neural Population Dynamics.  
410 *Annual Review of Neuroscience* **43**, 249–275 (2020).
- 411 3. Wei, Z. *et al.* A comparison of neuronal population dynamics measured with calcium imaging and  
412 electrophysiology. *PLoS computational biology* **16**, e1008198 (2020).
- 413 4. Demas, J. *et al.* Volumetric Calcium Imaging of 1 Million Neurons Across Cortical Regions at Cellular  
414 Resolution using Light Beads Microscopy. <http://biorxiv.org/lookup/doi/10.1101/2021.02.21.432164> (2021)  
415 doi:10.1101/2021.02.21.432164.
- 416 5. Pachitariu, M. *et al.* Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *BioRxiv*

- 417 (2017).
- 418 6. Peron, S. P., Freeman, J., Iyer, V., Guo, C. & Svoboda, K. A cellular resolution map of barrel cortex activity  
419 during tactile behavior. *Neuron* **86**, 783–799 (2015).
- 420 7. Chen, J. L., Carta, S., Soldado-Magraner, J., Schneider, B. L. & Helmchen, F. Behaviour-dependent  
421 recruitment of long-range projection neurons in somatosensory cortex. *Nature* **499**, 336–340 (2013).
- 422 8. Chen, S. X., Kim, A. N., Peters, A. J. & Komiyama, T. Subtype-specific plasticity of inhibitory circuits in  
423 motor cortex during motor learning. *Nature neuroscience* **18**, 1109–1115 (2015).
- 424 9. Pnevmatikakis, E. A. Analysis pipelines for calcium imaging data. *Current Opinion in Neurobiology* **55**, 15–  
425 21 (2019).
- 426 10. Berens, P. *et al.* Community-based benchmarking improves spike rate inference from two-photon calcium  
427 imaging data. *PLoS computational biology* **14**, e1006157 (2018).
- 428 11. Pachitariu, M., Stringer, C. & Harris, K. D. Robustness of spike deconvolution for neuronal calcium  
429 imaging. *Journal of Neuroscience* **38**, 7976–7985 (2018).
- 430 12. Evans, M. H., Petersen, R. S. & Humphries, M. D. On the use of calcium deconvolution algorithms in  
431 practical contexts. *bioRxiv* 871137 (2019).
- 432 13. Huang, L. *et al.* Relationship between simultaneously recorded spiking activity and fluorescence signal in  
433 GCaMP6 transgenic mice. *eLife* **10**, e51675 (2021).
- 434 14. Siegle, J. H. *et al.* Reconciling functional differences in populations of neurons recorded with two-photon  
435 imaging and electrophysiology. *bioRxiv* (2020).
- 436 15. Sussillo, D., Jozefowicz, R., Abbott, L. & Pandarinath, C. LFADS-latent factor analysis via dynamical  
437 systems. *arXiv preprint arXiv:1608.06315* (2016).
- 438 16. Pandarinath, C. *et al.* Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat*  
439 *Methods* **15**, 805–815 (2018).
- 440 17. Keshtkaran, M. R. *et al.* A large-scale neural network training framework for generalized estimation of  
441 single-trial population dynamics. *bioRxiv* (2021).
- 442 18. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat Neurosci*  
443 **17**, 1500–1509 (2014).
- 444 19. Pandarinath, C. *et al.* Latent Factors and Dynamics in Motor Cortex and Their Application to Brain–  
445 Machine Interfaces. *J. Neurosci.* **38**, 9390–9401 (2018).
- 446 20. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical Control of Arm Movements: A Dynamical Systems  
447 Perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
- 448 21. Keshtkaran, M. R. & Pandarinath, C. Enabling hyperparameter optimization in sequential autoencoders for  
449 spiking neural data. in *Advances in Neural Information Processing Systems* 15937–15947 (2019).
- 450 22. Wei, X.-X. *et al.* A zero-inflated gamma model for deconvolved calcium imaging traces. *arXiv preprint*  
451 *arXiv:2006.03737* (2020).
- 452 23. Chen, T.-W. *et al.* Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300  
453 (2013).
- 454 24. Zhao, Y. & Park, I. M. Variational latent gaussian process for recovering single-trial dynamics from  
455 population spike trains. *Neural computation* **29**, 1293–1316 (2017).
- 456 25. Friedrich, J., Zhou, P. & Paninski, L. Fast online deconvolution of calcium imaging data. *PLoS Comput Biol*  
457 **13**, e1005423 (2017).
- 458 26. Giovannucci, A. *et al.* CalmAn an open source tool for scalable calcium imaging data analysis. *eLife* **8**,  
459 e38173 (2019).
- 460 27. Galiñanes, G. L., Bonardi, C. & Huber, D. Directional reaching for water as a cortex-dependent behavioral  
461 framework for mice. *Cell reports* **22**, 2767–2783 (2018).
- 462 28. Hatsopoulos, N. G., Xu, Q. & Amit, Y. Encoding of movement fragments in the motor cortex. *Journal of*  
463 *Neuroscience* **27**, 5105–5114 (2007).
- 464 29. Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., Maclver, M. A. & Poeppel, D. Neuroscience needs  
465 behavior: correcting a reductionist bias. *Neuron* **93**, 480–490 (2017).
- 466 30. Whishaw, I. Q. *et al.* Organization of the reach and grasp in head-fixed vs freely-moving mice provides  
467 support for multiple motor channel theory of neocortical organization. *Experimental brain research* **235**,  
468 1919–1932 (2017).
- 469 31. Wiltschko, A. B. *et al.* Revealing the structure of pharmacobehavioral space through motion sequencing.  
470 *Nature neuroscience* **23**, 1433–1443 (2020).

- 471 32. Herzfeld, D. J., Kojima, Y., Soetedjo, R. & Shadmehr, R. Encoding of error and learning to correct that  
472 error by the Purkinje cells of the cerebellum. *Nature neuroscience* **21**, 736–743 (2018).
- 473 33. Vyas, S., O'Shea, D. J., Ryu, S. I. & Shenoy, K. V. Causal role of motor preparation during error-driven  
474 learning. *Neuron* **106**, 329–339 (2020).
- 475 34. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and  
476 engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
- 477 35. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, (2019).
- 478 36. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are  
479 dominated by richly varied movements. *Nature neuroscience* **22**, 1677–1686 (2019).
- 480 37. Hoang, H. *et al.* Improved hyperacuity estimation of spike timing from calcium imaging. *Sci Rep* **10**, 17844  
481 (2020).
- 482 38. Rupprecht, P. *et al.* Database and deep learning toolbox for noise-optimized, generalized spike inference  
483 from calcium imaging. <http://biorxiv.org/lookup/doi/10.1101/2020.08.31.272450> (2020)  
484 doi:10.1101/2020.08.31.272450.
- 485 39. Sebastian, J., Sur, M., Murthy, H. A. & Magimai-Doss, M. Signal-to-signal neural networks for improved  
486 spike estimation from calcium imaging data. *PLoS Comput Biol* **17**, e1007921 (2021).
- 487 40. Dechery, J. B. & MacLean, J. N. Functional triplet motifs underlie accurate predictions of single-trial  
488 responses in populations of tuned and untuned V1 neurons. *PLoS computational biology* **14**, e1006153  
489 (2018).
- 490 41. Kirschbaum, E. *et al.* LeMoNADe: Learned Motif and Neuronal Assembly Detection in calcium imaging  
491 videos. *arXiv:1806.09963 [q-bio]* (2019).
- 492 42. Mackevicius, E. L. *et al.* Unsupervised discovery of temporal sequences in high-dimensional datasets, with  
493 applications to neuroscience. *eLife* **8**, e38471 (2019).
- 494 43. Triplett, M. A., Pujic, Z., Sun, B., Avitan, L. & Goodhill, G. J. Model-based decoupling of evoked and  
495 spontaneous neural activity in calcium imaging data. *PLoS Comput Biol* **16**, e1008330 (2020).
- 496 44. Williams, A. H. *et al.* Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across  
497 Multiple Timescales through Tensor Component Analysis. *Neuron* **98**, 1099–1115.e8 (2018).
- 498 45. Wu, A., Pashkovski, S., Datta, S. R. & Pillow, J. W. Learning a latent manifold of odor representations from  
499 neural responses in piriform cortex. in *Advances in Neural Information Processing Systems* (eds. Bengio,  
500 S. *et al.*) vol. 31 (Curran Associates, Inc., 2018).
- 501 46. Costa, A. C., Ahamed, T. & Stephens, G. J. Adaptive, locally linear models of complex dynamics. *Proc Natl*  
502 *Acad Sci USA* **116**, 1501–1510 (2019).
- 503 47. Glaser, J., Whiteway, M., Cunningham, J. P., Paninski, L. & Linderman, S. Recurrent Switching Dynamical  
504 Systems Models for Multiple Interacting Neural Populations. in *Advances in Neural Information Processing*  
505 *Systems* (eds. Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F. & Lin, H.) vol. 33 14867–14878  
506 (Curran Associates, Inc., 2020).
- 507 48. Picardo, M. A. *et al.* Population-Level Representation of a Temporal Sequence Underlying Song  
508 Production in the Zebra Finch. *Neuron* **90**, 866–876 (2016).
- 509 49. Mano, O. *et al.* Using slow frame rate imaging to extract fast receptive fields. *Nature communications* **10**,  
510 1–13 (2019).
- 511 50. Prince, L. Y., Bakhtiari, S., Gillon, C. J. & Richards, B. A. *Parallel inference of hierarchical latent dynamics*  
512 *in two-photon calcium imaging of neuronal populations*.  
513 <http://biorxiv.org/lookup/doi/10.1101/2021.03.05.434105> (2021) doi:10.1101/2021.03.05.434105.
- 514 51. Minderer, M., Brown, K. D. & Harvey, C. D. The spatial structure of neural encoding in mouse posterior  
515 cortex during navigation. *Neuron* **102**, 232–248 (2019).
- 516 52. Sofroniew, N. J., Flickinger, D., King, J. & Svoboda, K. A large field of view two-photon mesoscope with  
517 subcellular resolution for in vivo imaging. *Elife* **5**, e14472 (2016).
- 518 53. Jaderberg, M. *et al.* Population Based Training of Neural Networks. *arXiv:1711.09846 [cs]* (2017).
- 519 54. Dana, H. *et al.* High-performance calcium sensors for imaging activity in neuronal populations and  
520 microcompartments. *Nat Methods* **16**, 649–657 (2019).
- 521 55. Art, J. Photon detectors for confocal microscopy. in *Handbook of biological confocal microscopy* 251–264  
522 (Springer, 2006).
- 523 56. Starck, J.-L., Murtagh, F. D. & Bijaoui, A. *Image processing and data analysis: the multiscale approach*.  
524 (Cambridge University Press, 1998).

- 525 57. Vogelstein, J. T. *et al.* Fast Nonnegative Deconvolution for Spike Train Inference From Population Calcium  
526 Imaging. *Journal of Neurophysiology* **104**, 3691–3704 (2010).  
527 58. Heikkilä, J. & Silvén, O. A four-step camera calibration procedure with implicit image correction. in  
528 *Proceedings of IEEE computer society conference on computer vision and pattern recognition* 1106–1112  
529 (IEEE, 1997).  
530 59. Zhang, Z. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and*  
531 *machine intelligence* **22**, 1330–1334 (2000).  
532

## 533 **Methods**

### 534 **AutoLFADS and RADICaL architecture and training**

535 The core model that AutoLFADS and RADICaL build on is LFADS. A detailed overview of the LFADS model is given in refs.  
536 <sup>15,16</sup>. Briefly, LFADS is a sequential application of a variational auto-encoder (VAE). A pair of bidirectional RNNs (the initial  
537 condition and controller input encoders) operate on the spike sequence and produce initial conditions for the generator RNN  
538 and time-varying inputs for the controller RNN. All RNNs were implemented using gated recurrent unit (GRU) cells. At each  
539 time step, the generator state evolves with input from the controller and the controller receives delayed feedback from the  
540 generator. The generator states are linearly mapped to factors, which are mapped to the firing rate of the neurons using a  
541 linear mapping followed by an exponential nonlinearity. The optimization objective is to maximize a lower bound on the  
542 likelihood of the observed spiking activity given the rates produced by the generator network, and includes KL and L2  
543 regularization penalties. During training, network weights are optimized using stochastic gradient descent and  
544 backpropagation through time.  
545

546 Identical network sizes were used for both AutoLFADS and RADICaL runs and for both simulation and real 2P data. The  
547 dimension of initial condition encoder, controller input encoder, and controller RNNs was 64. The dimension of the generator  
548 RNN was 100. The generator was provided with 64-dimensional initial conditions and 2-dimensional controller outputs (i.e.,  
549 inferred inputs  $u(t)$ ) and linearly mapped to 100-dimensional factors. The initial condition prior distribution was Gaussian with  
550 a trainable mean that was initialized to 0 and a variance that was fixed to 0.1. The minimum allowable variance of the initial  
551 condition posterior distribution was set to  $1e-4$ . The controller output prior was autoregressive with a trainable autocorrelation  
552  $\tau$  and noise variance, initialized to 10 and 0.1, respectively. The Adam optimizer (epsilon:  $1e-8$ ; beta1: 0.9; beta2: 0.99;  
553 initial learning rate:  $1e-3$ , **Table 1**) was used to control weight updates. The loss was scaled by a factor of  $1e4$  prior to  
554 computing the gradients for numerical stability. To prevent potential pathological training, the GRU cell hidden states were  
555 clipped at 5 and the global gradient norm was clipped at 300.  
556

557 AutoLFADS is a recent implementation of the population based training (PBT) approach<sup>53</sup> on LFADS to perform automatic,  
558 large-scale hyperparameter (HP) search. A detailed overview of AutoLFADS is in refs. <sup>17,21</sup>. Briefly, PBT distributes training  
559 across dozens of models in parallel, and uses evolutionary algorithms to tune HPs over many generations. At the end of  
560 each generation, a selection process was performed to choose higher performing models and replace the poor models with  
561 the higher performing models. The HPs of the higher performing models were perturbed before the next generation to  
562 increase the HP search space. After many generations (~30-150), the PBT process converges upon a high performing model  
563 with optimized HPs.  
564

565 Twenty LFADS models were trained in parallel for 50 epochs per generation for both AutoLFADS and RADICaL runs and  
566 for both simulation and real 2P data. KL and L2 regularization penalties were linearly ramped for the first 80 epochs of  
567 training during the first generation. Training was stopped when there was no improvement in performance after 25  
568 generations. The HPs optimized by PBT were the model's learning rate and six regularization HPs: scaling weights for the  
569 L2 penalties on the generator, controller, and initial condition encoder RNNs, scaling weights for the KL penalties on the  
570 initial conditions and controller outputs, and two dropout probabilities ("keep ratio" for coordinated dropout<sup>21</sup>; and RNN  
571 network dropout probability). The HP search ranges are detailed in **Table 1**. The magnitudes of the HP perturbation were  
572 controlled by weights and specified for different HPs (a weight of 0.3 results in perturbation factors between 0.7 and 1.3;  
573 **Table 1**). The learning rate and dropout probabilities were restricted to their specified search ranges and were sampled from  
574 uniform distributions. The KL and L2 HPs were sampled from log-uniform distributions and could be perturbed outside of the  
575 initial search ranges. Identical hyperparameter settings were used for both RADICaL and AutoLFADS and for both synthetic  
576 datasets and real 2P datasets.

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

RADICaL is an adaptation of AutoLFADS for 2P calcium imaging. RADICaL operates on sequences of deconvolved calcium events  $x(t)$ .  $x(t)$  are modeled as a noisy observation of an underlying time-varying Zero-Inflated Gamma (ZIG) distribution<sup>22</sup>:

$$x_n(t) \sim (1 - q_n(t)) \cdot \delta(0) + q_n(t) \cdot \text{gamma}(\alpha_n(t), k_n(t), \text{loc}_n),$$

where  $x_n(t)$  is the distribution of observed deconvolved events,  $\alpha_n(t)$ ,  $k_n(t)$ , and  $\text{loc}_n$  are the scale, shape, and location parameters, respectively, of the gamma distribution, and  $q_n(t)$  denotes the probability of non-zeros, for neuron  $n$  at time  $t$ .  $\text{loc}_n$  was fixed as the minimum nonzero deconvolved event ( $s_{\min}$ ). In the original AutoLFADS model, factors were mapped to a single time-varying parameter for each neuron (the Poisson firing rate) via a linear transformation followed by an exponential nonlinearity. RADICaL instead infers the three time-varying parameters for each neuron,  $\alpha_n(t)$ ,  $k_n(t)$ , and  $q_n(t)$ , by linearly transforming the factors followed by a trainable scaled sigmoid nonlinearity ( $\text{sig}_n$ ).  $\text{sig}_n$  is a positive parameter that scales the outputs of the sigmoid to be in a range between 0 and  $\text{sig}_n$ , and is optimized alongside network weights. An L2 penalty is applied between  $\text{sig}_n$  and a PBT-searchable prior (**Table 1**) to prevent extreme values. The training objective is to minimize the negative log-likelihood of the deconvolved events given the inferred parameters:

593

594

$$\prod p(x_n(t) | \text{ZIG}(\hat{\alpha}_n(t), \hat{k}_n(t), \hat{q}_n(t)))$$

595

596

The event rate for neuron  $n$  at time  $t$  was estimated by taking the mean of the inferred ZIG distribution:

597

598

$$\hat{q}_n(t) \cdot (\hat{k}_n(t) \cdot \hat{\alpha}_n(t) + s_{\min})$$

599

600

601

602

603

604

605

606

607

RADICaL uses an SBTT training strategy to achieve subframe modeling resolution. RADICaL operates on binned deconvolved calcium events, with bin size smaller than the frame timebase of imaging. Bins where the neurons were sampled were filled with the corresponding event rates, while bins where the neurons were not sampled were filled with NaNs. The networks still output the time-varying ZIG distribution at each timestep; however, a mask was applied to the timesteps where the NaN samples were to prevent the cost computed from these timesteps being backpropagated during gradient calculation. As a result, the model weights were only updated based on the cost at the sampled timesteps. The reconstruction cost also excluded the cost calculated at the non-sampled timesteps so the PBT model selection was not affected by the cost computed from the non-sampled timesteps.

608

## Simulation experiments.

609

### *Generating spike trains from an underlying Lorenz system*

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

Synthetic data were generated using the Lorenz system as described in the original LFADS work<sup>15,16</sup>. Lorenz parameters were set to standard values ( $\sigma$ : 10,  $\rho$ : 28, and  $\beta$ : 8/3), and  $\Delta t$  was set to 0.01. Datasets with different speeds of dynamics were generated by downsampling the original generated Lorenz states by different factors. The speed of the Lorenz dynamics was quantified based on the peak location of the power spectra of the Lorenz Z dimension, with a sampling frequency of 100 Hz. The downsampling factors were 3, 5, 7, 9, 11 and 14 for speeds 4, 7, 10, 13, 15 and 20 Hz, respectively. Each dataset/speed consisted of 8 conditions, with 60 trials per condition. Each condition was obtained by starting the Lorenz system with a random initial state vector and running it for 900 ms. The trial length for the 4 Hz dataset was longer (1200 ms) than that of other datasets (900 ms) to ensure that all conditions had significant features to be modeled - with shorter windows, the extremely low frequency oscillations caused the Lorenz states for some conditions to have little variance across the entire window, making it trivial to approximate the essentially flat firing rates. We simulated a population of 278 neurons with firing rates given by linear readouts of the Lorenz state variables using random weights, followed by an exponential nonlinearity. Scaling factors were applied so the baseline firing rate for all neurons was 3 spikes/sec. Each bin represents 10 ms and an arbitrary frame time was set to be 30 ms (i.e., one "imaging frame" takes 3 bins). Spikes from the firing rates were then generated by a Poisson process.

625

### *Generating fluorescence signals from synthetic spike trains*

626 Realistic fluorescence signals were generated from the spike trains by convolving them with a kernel for an autoregressive  
627 process of order 2 and passing the results through a nonlinearity that matched values extracted from the literature for the  
628 calcium indicator GCaMP6f<sup>3,54</sup> (**Supp Fig. 2a & b**). Three noise sources were added to reproduce variability present in real  
629 data<sup>55-57</sup>: Gaussian noise to the size of the calcium spike, and Gaussian and Poisson noise to the final trace (**Supp Fig. 2a**  
630 **& b**). This fluorescence generation process was realized as follows: First, spike trains  $s(t)$  were generated from the Lorenz  
631 system as mentioned above. Independent Gaussian noise (sd = 0.1) was added to each spike in the spike train to model  
632 the variability in spike amplitude. Next, we modeled the calcium concentration dynamics  $c(t)$  as an autoregressive process  
633 of order 2:  
634

$$c(t) = \gamma_1 c(t-1) + \gamma_2 c(t-2) + s(t)$$

635  
636  
637 with  $s(t)$  representing the number of spikes at time  $t$ . The autoregressive coefficients  $\gamma_1$  and  $\gamma_2$  were computed based on the  
638 rise time, decay time ( $\tau_{on} = 20$  ms,  $\tau_{off} = 400$  ms for GCaMP6f) of the calcium indicators, and the sampling frequency. Note  
639 that while there is substantial variability in taus across neurons in real data<sup>3</sup>, selecting and mimicking this variability was not  
640 relevant in our work, because we compared the methods (i.e., RADICaL, AutoLFADS, and s-deconv) after deconvolution.  
641 The calcium concentration dynamics were further normalized so that the peak height of the calcium dynamics generated  
642 from a single spike equalled one, regardless of the sampling frequency. Subsequently, we computed the noiseless  
643 fluorescence signals by passing the calcium dynamics through a nonlinear transformation estimated from the literature<sup>54</sup> for  
644 the calcium indicator GCaMP6f (**Supp Fig. 2c & d**). After the nonlinear transformation, the relationship between spike size  
645 and trace size was corrupted, and therefore we assumed the baseline of fluorescence signals to be zero and the signals  
646 were rescaled to the range in  $[0,1]$  using min-max normalization. Finally, Gaussian noise ( $\sim N(0,sn)$ ) and Poisson noise  
647 (simulated as gaussian with mean 0 and variance proportional to the signal amplitude at each time point via a constant  $d$ )  
648 were added to the normalized traces. The resulting fluorescence traces had the same sampling frequency as the synthetic  
649 spike trains (100 Hz).  
650

651 A crucial parameter is the noise level associated with each fluorescence trace. High noise levels lead to very poor spike  
652 detection and very low noise levels enable a near-perfect reconstruction of the spike train. In order to select a realistic level  
653 of noise we matched the correlations between real and inferred spike trains of the simulated data to those typically  
654 observed<sup>10</sup>. We found that a truncated normal distribution of noise level for Gaussian and Poisson noise best matched the  
655 correlations. More specifically, for each neuron,  $sn=d$  was sampled independently from a truncated normal distribution  
656  $N(0.12, 0.02)$  truncated below 0.06. With the above noise setting, the mean correlation coefficient  $r$  between the deconvolved  
657 events and ground truth spikes was 0.32, which is consistent with the standard results reported in the spikefinder paper<sup>10</sup>  
658 for OASIS. It is worth stressing that real data feature a broad range of noise levels that depend on the imaging conditions,  
659 depth, expression level, laser power and other factors. Here we did not attempt to investigate all possible noise conditions,  
660 but instead, we aimed to create a simulation with known latent variables (i.e., low-dimensional factors and event rates) that  
661 reasonably approximated realistic signal-to-noise levels, in order to provide a tractable test case to compare RADICaL to  
662 other methods before attempting comparisons on real data.  
663

#### 664 *Recreating variability in sampling times due to 2p laser scanning*

665 The fluorescence traces were simulated at 100 Hz as mentioned above. A subsampling step was then performed with  
666 sampling times for each neuron staggered in time to simulate the variability in sampling times due to 2p laser scanning (as  
667 in **Fig. 1e**). This produced fluorescence traces where individual neurons were sampled at 33.3 Hz, with phases of 0, 11, 22  
668 ms based on each neuron's location (top, middle and bottom of the FOV, respectively). To break this down, each neuron  
669 was sparsely sampled every three time points and the relative sampled times between neurons were fixed. For example, in  
670 trial 1, neuron 1 was sampled at time points 1, 4, 7, ... and neuron 2 was sampled at time points 2, 5, 8, ...; in trial 2, neuron  
671 1 was sampled at time points 2, 5, 8, ... and neuron 2 was sampled at time points 3, 6, 9, ... . Thus, the sampling frequency  
672 for each individual neuron was 33.3 Hz, while the sampling frequency for the population was retained at 100 Hz by filling the  
673 non-sampled time points with NaNs. The resulting 33.3 Hz simulated fluorescence signals for each individual neuron (i.e.,  
674 with NaNs excluded) were deconvolved using OASIS<sup>25</sup> (as implemented in CalmAn<sup>26</sup>) using an auto-regressive model of  
675 order 1 with  $s_{min}$  of 0.1.  
676



## 677 *Data preparation for each method*

678 Four methods (RADICaL, AutoLFADS, s-deconv and s-sim-fluor) were compared by their performance on recovering the  
679 ground truth latent states across different datasets/speeds. Trials (480 total for each simulated dataset) were split into 80/20  
680 training and validation sets for modeling AutoLFADS and RADICaL. To prepare data for non-RADICaL methods, non-  
681 sampled bins were removed so all the sampled bins were treated as if they were sampled at the same time and each bin  
682 then represented 30 ms (i.e., sampling frequency = 33.3 Hz). Preparing the data for AutoLFADS required discretizing the  
683 deconvolved events into spike count estimates, because AutoLFADS was primarily designed to model discrete spiking data.  
684 In the discretizing step, if the event rate was 0, it was left as 0; if the event rate was between 0 and 2, it was cast to 1 (to  
685 bias toward the generally higher probability of fewer spikes). If the event rate was greater than 2, it was rounded down to  
686 the nearest integer. We note that this is one of many possible patches to convert continuously-valued event intensities to  
687 natural numbers for compatibility with the Poisson distribution and AutoLFADS; a more principled solution would be to modify  
688 the network to use the ZIG distribution, as we have done in RADICaL. With s-deconv, the deconvolved events were  
689 smoothed by convolution with a Gaussian filter (6 ms s.d.) to produce event rates. With s-sim-fluor, the generated  
690 fluorescence signals were smoothed by convolution with a Gaussian filter (6 ms s.d.) to produce event rates. The choice of  
691 filter width was optimized by sweeping values ranging from 3 to 40 ms. Smoothing with a 6 ms s.d. filter gave the highest  
692 performance in recovering the ground truth Lorenz states for experiments with higher Lorenz frequencies (i.e.,  $\geq 10$  Hz).  
693 The event rates produced from RADICaL had a sampling frequency of 100 Hz, while the event rates produced from the non-  
694 RADICaL methods had a sampling frequency of 33.3 Hz. The non-RADICaL rates were then resampled at 100 Hz using  
695 linear interpolation.

## 697 *Mapping to ground truth Lorenz states*

698 Since our goal was to quantify modeling performance by estimating the underlying Lorenz states, we trained a mapping from  
699 the output of each model (i.e., the event rates) to the ground truth Lorenz states using ridge regression. First, we split the  
700 trials into training (80%) and test (20%) sets. We used the training set to optimize the regularization coefficient using 5-fold  
701 cross-validation, and used the optimal regularization coefficient to train the mapping on the full training set. We then  
702 quantified state estimation performance by applying this trained mapping to the test set and calculating the coefficient of  
703 determination ( $R^2$ ) between the true and predicted Lorenz states. We repeated the above procedure five times with train/test  
704 splits drawn from the data in an interleaved fashion. We reported the mean  $R^2$  across the repeats, such that all reported  
705 numbers reflect held-out performance. We tested whether the difference of  $R^2$  between each pair of methods was significant  
706 by performing a paired, one-sided Student's t-test on the distribution of  $R^2$  across the five folds of predictions.

## 708 **Real 2p experiments**

### 709 *Surgical procedures*

710 All procedures were approved by the University of Chicago Animal Care and Use Committee. Two male Ai148D transgenic  
711 mice (TIT2L-GC6f-ICL-tTA2, stock 030328; Jackson Laboratory) were used. Each mouse underwent a single surgery. Mice  
712 were injected subcutaneously with dexamethasone (8 mg/kg) 24 hours and 1 hour before surgery. Mice were anesthetized  
713 with 2-2.5% inhaled isoflurane gas, then injected intraperitoneally with a ketamine-medetomidine solution (60 mg/kg  
714 ketamine, 0.25 mg/kg medetomidine), and maintained on a low level of supplemental isoflurane (0-1%) if they showed any  
715 signs that the depth of anesthesia was insufficient. Meloxicam was also administered subcutaneously (2 mg/kg) at the  
716 beginning of the surgery and for 1-3 subsequent days. The scalp was shaved, cleaned, and resected, the skull was cleaned  
717 and the wound margins glued to the skull with tissue glue (VetBond, 3M), and a 3 mm circular craniotomy was made with a  
718 3 mm biopsy punch centered over the left CFA/S1 border. The coordinates for the center of CFA were taken to be 0.4 mm  
719 anterior and 1.6 mm lateral of bregma. The craniotomy was cleaned with SurgiFoam (Ethicon) soaked in phosphate-buffered  
720 solution (PBS), then virus (AAV9-CaMKII-Cre, stock  $2.1 \times 10^{13}$  particles/nL, 1:1 dilution in PBS, Addgene) was pressure  
721 injected (NanoJect III, Drummond Scientific) at two or four sites near the target site, with 140 nL injected at each of two  
722 depths per site (250 and 500  $\mu$ m below the pia) over 5 minutes each. The craniotomy was then sealed with a custom  
723 cylindrical glass plug (3 mm diameter, 660  $\mu$ m depth; Tower Optical) bonded (Norland Optical Adhesive 61, Norland) to a 4  
724 mm #1 round coverslip (Harvard Apparatus), glued in place first with tissue glue (VetBond) and then with cyanoacrylate glue  
725 (Krazy Glue) mixed with dental acrylic powder (Ortho Jet; Lang Dental). A small craniotomy was also made using a dental  
726 drill over right CFA at 0.4 mm anterior and 1.6 mm lateral of bregma, where 140 nL of AAVretro-tdTomato (stock  $1.02 \times 10^{13}$   
727 particles/nL, Addgene) was injected at 300  $\mu$ m below the pia. This injection labeled cells in left CFA projecting to the  
728 contralateral CFA. Here, this labeling was used solely for stabilizing the imaging plane (see below). The small craniotomy

729 was sealed with a drop of Kwik-Cast (World Precision Instruments). Two layers of MetaBond (C & B) were applied, then a  
730 custom laser-cut titanium head bar was affixed to the skull with black dental acrylic. Animals were awoken by administering  
731 atipamezole via intraperitoneal injection and allowed to recover at least 3 days before water restriction.  
732

### 733 *Behavioral task*

734 The behavioral task (Fig. 3a) was a variant of the water reaching task of ref. <sup>27</sup> which we term the “Water Grab” task. This  
735 task was performed by water-restricted, head-fixed mice, with the forepaws beginning on paw rests (eyelet screws) and the  
736 hindpaws and body supported by a custom 3D printed clear acrylic tube enclosure. After holding the paw rests for 700-900  
737 ms, a tone was played by stereo speakers and a 2-3  $\mu$ L droplet of water appeared at one of two water spouts (22 gauge,  
738 90-degree bent, 1” blunt dispensing needles, McMaster) positioned on either side of the snout. The pitch of the tone indicated  
739 the location of the water, with a 4000 Hz tone indicating left and a 7000 Hz tone indicating right, and it lasted 500 ms or until  
740 the mouse made contact with the correct water spout. The mouse could grab the water droplet and bring it to its mouth to  
741 drink any time after the tone began. Both the paw rests and spouts were wired with capacitive touch sensors (Teensy 3.2,  
742 PJRC). Good contact with the correct spout produced an inter-trial interval of 3-6 s, while failure to make contact (or  
743 insufficiently strong contact) with the spout produced an inter-trial interval of 20 s. Because the touch sensors required good  
744 contact from the paw, this setup encouraged complex contacts with the spouts. The mice were trained to make all reaches  
745 with the right paw and to keep the left paw on the paw rest during reaching. Training took approximately two weeks, though  
746 the behavior continued to solidify for at least two more weeks. Data presented here were collected after 6-8 weeks’  
747 experience with the task. Control software was custom written in MATLAB R2018a using PsychToolbox 3.0.14, and for the  
748 Teensy. Touch event monitoring and task control were performed at 60 Hz.  
749

750 Behavior was also recorded using a pair of cameras (BFS-U3-16S2M-CS, FLIR; varifocal lenses COZ2813CSIR2,  
751 Computar) mounted 150 mm from the right paw rest at 10° apart to enable 3D triangulation. Infrared illuminators enabled  
752 behavioral imaging while performing 2p imaging in a darkened microscope enclosure. Cameras were synchronized and  
753 recorded at 150 frames per second with real-time image cropping and JPEG compression, and streamed to one HDF5 file  
754 per camera (areaDetector module of EPICS, CARS). The knuckles and wrist of the reaching paw were tracked in each  
755 camera using DeepLabCut and triangulated into 3D using camera calibration parameters obtained from the MATLAB Stereo  
756 Camera Calibration toolbox<sup>58,59</sup>. To screen the tracked markers for quality we created distributions of all inter-marker  
757 distances in 3D across every labeled frame and identified as problematic frames with any inter-marker distance exceeding  
758 the 99.9th percentile of its respective distribution. Trials with more than one problematic frame in the period of -200 ms to  
759 800 ms after the raw reach onset were discarded (where reach onset was taken as the first 60 Hz tick after the paw rest  
760 touch sensor fell below contact threshold). The kinematics of all trials that passed this screening procedure were visualized  
761 to confirm quality. Centroid marker kinematics were obtained by averaging the kinematics of all paw markers, locking them  
762 to behavioral events and then smoothing using a Gaussian filter (15 ms s.d.). To obtain velocity and acceleration, centroid  
763 data was numerically differentiated with MATLAB’s *diff* function and then smoothed again using a Gaussian filter (15 ms  
764 s.d.).  
765

### 766 *Two-photon imaging*

767 Calcium imaging was performed with a NeuroLabware two-photon microscope and pulsed Ti:sapphire laser (Vision II,  
768 Coherent). Depth stability of the imaging plane was maintained using a custom plugin that acquired an image stack at the  
769 beginning of the session (1.4  $\mu$ m spacing), then compared a registered rolling average of the red-channel data to each plane  
770 of the stack. If sufficient evidence indicated that a plane not at the center of the stack was a better match to the image being  
771 acquired, the objective was automatically moved to compensate. This typically resulted in a slow and steady upward  
772 (outward) movement of the objective over the course of the session.  
773

774 Offline, images were run through Suite2p to perform motion correction, region-of-interest (ROI) detection, and fluorescence  
775 extraction from both ROIs and neuropil. ROIs were manually curated using the Suite2p GUI to retain only those  
776 corresponding to somas. We then subtracted the neuropil signal scaled by 0.7<sup>23</sup>. Neuropil-subtracted ROI fluorescence was  
777 then detrended by performing a running 10th percentile operation, smoothing with a Gaussian filter (20 s s.d.), then  
778 subtracting the result from the trace. This result was fed into OASIS<sup>25</sup> using the ‘thresholded’ method, AR1 event model, and  
779 limiting the tau parameter to be between 300 and 800 ms. Neurons were discarded if they did not meet a minimum signal-  
780 to-noise (SNR) criterion. To compute SNR, we took the fluorescence at each time point when OASIS identified an “event”

781 (non-zero), computed (fluorescence - neuropil) / neuropil, and computed the median of the resulting distribution. ROIs were  
782 excluded if this value was less than 0.05. To put events on a more useful scaling, for each ROI we found the distribution of  
783 event sizes, smoothed the distribution (ksdensity in MATLAB, with an Epanechnikov kernel and log transform), found the  
784 peak of the smoothed distribution, and divided all event sizes by this value. This rescales the peak of the distribution to have  
785 a value of unity. Data from two mice and two brain areas (4 sessions in total) were used (Mouse1 M1: 510 neurons, 560  
786 trials; Mouse1 S1: 433 neurons, 502 trials; Mouse2 M1: 439 neurons, 475 trials; Mouse2 S1: 509 neurons, 421 trials).

#### 787 *Data preparation for modeling with RADICaL and AutoLFADS*

788 To prepare data for RADICaL, the deconvolved events were normalized by the  $s_{\min}$  value output by OASIS so that the  
789 minimal event size was 0.1 across all neurons. The deconvolved events for individual neurons had a sampling rate equal to  
790 the frame rate (31.08 Hz). For modeling with RADICaL, the deconvolved events were assigned into 10ms bins using the  
791 timing of individual measurements for each neuron to achieve sub-frame resolution (i.e., 100 Hz). The non-sampled bins  
792 were filled with NaNs. To prepare data for AutoLFADS, the deconvolved events were rescaled using the distribution-scaling  
793 method described above, and casted using the casting step described in the simulation section. For both AutoLFADS and  
794 s-deconv, the deconvolved events were assigned into a single time bin per frame (i.e., 32.17 ms bins) to mimic standard  
795 processing of 2p imaging data, where the sub-frame timing of individual measurements is discarded. Trials were created by  
796 aligning the data to 200 ms before and 800 ms after reach onset (100 time points per trial for RADICaL, and 31 time points  
797 per trial for AutoLFADS and s-deconv). An individual RADICaL model and AutoLFADS model were trained for each dataset  
798 (4 total). Failed trials (latency to contact with correct spout > 15 s for Mouse1, 20 s for Mouse2), or trials where the grab to  
799 the incorrect spout occurred before the grab to the correct spout, were discarded. For each dataset, trials (Mouse1 M1: 552  
800 total; Mouse1 S1: 500 total; Mouse2 M1: 467 total; Mouse2 S1: 413 total) were split into 80/20 training and validation.

#### 801 *Trial grouping*

802 PSTH analysis and low dimensional neural trajectory visualization were performed based on subgroups of trials. Trials were  
803 sorted into two subgroups per spout based on the Z dimension (height) of hand position. The hand position was obtained by  
804 smoothing the centroid marker position with a Gaussian filter (40 ms s.d.). Time windows where the height of hand was used  
805 to split trials were hand-selected to present a good separation between subgroups of hand trajectories. For mouse1 M1, a  
806 window of 30 ms to 50 ms after reach onset was used to split left condition trials and a window of 180 ms to 200 ms after  
807 reach onset was used to split right condition trials; for mouse1 S1, a window of 180 ms to 200 ms after reach onset was  
808 used to split left condition trials and a window of 140 ms to 160 ms after reach onset was used to split right condition trials;  
809 for both mouse2 M1 and mouse2 S1, a window of 30 ms to 50 ms after reach onset was used to split both left and right  
810 condition trials. For both left or right conditions and for all mice/areas (with the exception of mouse1 M1), 55 trials with the  
811 lowest and highest heights were selected as group 1 and group 2, respectively; trials with middle-range heights were  
812 discarded. For mouse1 M1, the first 25 trials with the lowest heights for right condition were discarded because these reaches  
813 were highly non-stereotyped and loopy; instead, the 26th to 80th trials with the lowest heights were selected as group 1 for  
814 the right condition.

#### 815 *PSTH analysis and comparing RADICaL and AutoLFADS single-trial rates*

816 RADICaL was first validated by comparing the PSTHs computed using RADICaL inferred event rates and the empirical  
817 PSTHs. Empirical PSTHs were computed by trial-averaging s-deconv rates (40 ms kernel s.d., 32.17 ms bins) within each  
818 of the 4 subgroups of trials. RADICaL inferred rates were first downsampled from 100 Hz to 31.08 Hz with an antialiasing  
819 filter applied, to match the sampling frequency (i.e., the frame rate) of the original deconvolved signals. RADICaL PSTHs  
820 were computed by similarly averaging RADICaL rates. Single-trial inferred rates were then compared to the empirical PSTHs  
821 to assess how well each method recapitulated the empirical PSTHs on single trials. The correlation coefficient ( $r$ ) was  
822 computed between inferred single-trial event rates and the corresponding empirical PSTHs in a cross-validated fashion, i.e.,  
823 each trial's inferred event rate was compared against an empirical PSTH computed using all other trials within the subgroup.  
824  $r$  was assessed for the time window spanning 200 ms before to 800 ms after reach onset, and computed by concatenating  
825 all trials across the four subgroups, yielding one  $r$  for each neuron. Neurons that had fewer than 40 nonzero events within  
826 this time window (across all trials) were excluded from the analysis.

#### 827 *Low-D analysis*

832 To visualize the low-dimensional neural trajectories that RADICaL produced, principal component analysis (PCA) was  
833 performed on RADICaL inferred rates. RADICaL rates (aligned to 200 ms before and 800 ms after reach onset) were log-  
834 transformed (with  $1e-4$  added to prevent numerical precision issues) and normalized to have zero mean and unit standard  
835 deviation for each neuron. PCA was applied to the trial-averaged rates and the projection matrix was then used to project  
836 the log-transformed and normalized single-trial rates (aligned to 200 ms before and 400 ms after reach onset) onto the top  
837 3 PCs.

#### 838 839 *Decoding analysis*

840 RADICaL-inferred rates, AutoLFADS-inferred rates, and s-deconv (Gaussian kernel 40 ms s.d.) rates were used to decode  
841 hand position and velocity using ridge regression. The hand position and velocity were obtained as described above and  
842 binned at 10 ms (i.e., 100 Hz). The non-RADICaL rates were retained to a sampling frequency of 100 Hz using linear  
843 interpolation. For simplicity, we did not include a lag between the neural data and kinematics. However, additional analyses  
844 confirmed that adding a lag did not alter the results (data not shown). Trials with an interval between water presentation and  
845 reach onset that was longer than a threshold were discarded due to potential variations in behavior (e.g., inattention). The  
846 threshold was selected arbitrarily for different sessions based on the actual distribution of the intervals in the session (Mouse1  
847 M1: 500 ms; Mouse1 S1: 400 ms; Mouse2 M1: 400 ms; Mouse2 S1: 600 ms). The data were aligned to 50 ms before and  
848 350 ms after reach onset. The decoder was trained and tested using cross-validated ridge regression. First, we split the  
849 trials into training (80%) and test (20%) sets. We used the training set to optimize the regularization coefficient using 5-fold  
850 cross-validation, and used the optimal regularization coefficient to train the decoder on the full training set. This trained  
851 decoder was applied to the test set, and the coefficient of determination ( $R^2$ ) was computed and averaged across x-, y- and  
852 z- kinematics. We repeated the above procedure five times with train/test splits drawn from the data in an interleaved fashion.  
853 We reported the mean  $R^2$  across the repeats, such that all reported numbers reflect held-out performance. We tested  
854 whether the difference of  $R^2$  between each pair of methods was significant by performing paired, one-sided Student's t-Test  
855 on the distribution of  $R^2$  across the five folds of predictions.

#### 856 857 *Coherence analysis*

858 Coherence was computed between the true and predicted kinematics (window: 200 ms before and 500 ms after reach onset)  
859 across all trials and across all x-, y- and z- dimensions using magnitude-squared coherence (MATLAB: mscohere). The  
860 power spectral density estimation parameters within mscohere were specified to ensure a robust calculation on the single  
861 trial activity: Hanning windows with 35 timesteps (i.e., 350 ms) for the FFT and window size, and 25 timesteps (i.e., 250 ms)  
862 of overlap between windows.

#### 863 864 *t-SNE analysis on the weights mapping from factors to ZIG parameters*

865 RADICaL relies on subframe bins in which neurons are grouped based on their spatial locations within the FOV. Because  
866 this strategy results in consistent neuron grouping, it could potentially result in different groups of neurons corresponding to  
867 different latent factors. To test whether such an artifact existed, we visualized the transformation from latents to neurons by  
868 using t-SNE to reduce the 300-dimensional weights vector (100 factors \* 3 ZIG parameters) into a 2-D t-SNE space for each  
869 individual neuron (510 neurons total) (**Supp. Fig. 9**). We did not observe a relationship between neurons' position within the  
870 field of view (i.e., top, middle, and bottom) and the underlying factors. This suggested that the model did not use distinct  
871 factors for sets of neurons that were sampled with different phases, despite neurons in distant portions of the FOV never  
872 being grouped in the same bin.

#### 873 874 *Neuron downsampling*

875 The neuron downsampling experiment was performed on the Mouse2 M1 dataset. The number of neurons included when  
876 training RADICaL or AutoLFADS was gradually reduced by limiting the area of FOV that the neurons were sampled from.  
877 The area was shrunk from the entire FOV with an area-to-FOV ratio of 1, 25/36, 9/16, 1/4, and 1/9, resulting in the number  
878 of included neurons to be 439, 321, 262, 121 and 59. An individual RADICaL model and AutoLFADS model were trained for  
879 each number of neurons. Decoding was performed using ridge regression (see above).

832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863  
864  
865  
866  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883

			LR	CD	DO	KL CO	KL IC	L2 Con	L2 Gen	Sig Prior
Defaults	RADICaL	Ranges	(1e-5, 5e-3)	(0.01, 0.99)	(0.3, 1.0)	(1e-6, 1e-4)	(1e-6, 1e-4)	(1e-5, 0.1)	(1e-5, 0.1)	(1.0, 100.0)
		Initial values	1e-3	0.5	uniform	loguniform	loguniform	loguniform	loguniform	20.0
		Explore weight	0.3	0.3	0.3	0.8	0.8	0.8	0.8	0.2
		Limit explore	TRUE	TRUE	TRUE	FALSE	FALSE	FALSE	FALSE	FALSE
	AutoLFADS	Ranges	(1e-5, 5e-3)	(0.01, 0.99)	(0.3, 1.0)	(1e-6, 1e-4)	(1e-6, 1e-4)	(1e-5, 0.1)	(1e-5, 0.1)	-
		Initial values	1e-3	0.5	uniform	loguniform	loguniform	loguniform	loguniform	-
		Explore weight	0.3	0.3	0.3	0.8	0.8	0.8	0.8	-
		Limit explore	TRUE	TRUE	TRUE	FALSE	FALSE	FALSE	FALSE	-

884 **Table 1.** Hyperparameter ranges for RADICaL and AutoLFADS runs. Cells with a dash indicate “not applicable” for the  
885 method. LR is the learning rate. CD is the coordinated dropout rate (i.e., proportion of samples dropped at input). DO is the  
886 dropout probability for the RNN network. KL indicates the weight applied to the KL divergence of a posterior from its prior.  
887 CO indicates the controller output distributions and IC indicates the initial condition distributions. L2 indicates the weight  
888 applied to the Frobenius norm of the recurrent kernel of the GRU cell. Con indicates the controller GRU cell, Gen indicates  
889 the generator GRU cell, IC Enc indicates the initial condition encoder GRU cells. Sig Prior indicates the prior of the scaling  
890 factors applied to the sigmoid nonlinearity when mapping from factors to ZIG parameters.