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

Feng Zhu^{1,2}, Harrison A. Grier³, Raghav Tandon¹, Changjia Cai⁴, Andrea Giovannucci^{4,5,6,*}, Matthew T. Kaufman^{7,8,^,*}, Chethan Pandarinath^{1,9,10,^,*}

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

[^] 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 populations¹. These technologies are enabling new insights into how neural populations implement the computations necessary for motor, sensory, and cognitive processes². However, different recording technologies impose distinct tradeoffs in the types of questions that may be asked³. 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 millions⁴⁻⁶ - in 3-D, often with identified layers and cell types of interest^{7,8}. 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.

In recent years, a large amount of effort has been dedicated to improving the inference of spike trains from 2p data by detecting calcium influx events, *i.e.*, time points where single spikes, or multiple spikes in close succession, produce detectable fluorescence transients⁹. Ideally the spikes-to-fluorescence transformation would be invertible, such that analyzing calcium events would be equivalent to analyzing spiking activity³. However, recent benchmarks illustrate that a variety of algorithms to infer calcium events reach a similar ceiling of performance and make consistent predictions, and all achieve limited correspondence to ground truth spiking activity obtained with electrophysiology, particularly on fine timescales^{10,11}. Comparisons of calcium imaging and electrophysiology suggest that the two methods may therefore lead to divergent scientific findings^{3,12–14}, largely due to limitations when inferring spikes from calcium traces.

Rather than focusing on the responses of individual neurons, an alternative approach is to characterize patterns of covariation across a neuronal population to reveal the internal state of the underlying network. These “latent variable models”, or simply “latent models”, describe each neuron’s activity as a reflection of the network’s state. For example, when applied to electrophysiological data, latent models assume that an individual neuron’s spiking is a noisy observation of a latent “firing rate”, which fluctuates in a coordinated way with the firing rates of other neurons in the population. Despite their abstract nature, the network states inferred by latent models can reveal key insights into the computations being performed by the brain areas of interest². Inferred network states can also enhance our ability to relate neural activity to behavior. For example, one state-of-the-art deep learning method to estimate network states from electrophysiological spiking data is Latent Factor Analysis via Dynamical Systems (LFADS)^{15,16}. In applications to data from motor, sensory, and cognitive regions, LFADS reveals rules that govern how network states progress over time and that are consistent across behavioral conditions, while also revealing tight correspondences with single-trial behavior on a millisecond timescale^{16,17}.

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

Results

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

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

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

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

Novel features of RADICaL

RADICaL incorporates two major innovations over LFADS and AutoLFADS. First, we modified RADICaL’s observation model to better account for the statistics of deconvolved events. In LFADS, discrete spike count data are modeled as samples from an underlying time-varying Poisson process for each neuron. However, deconvolving 2p calcium signals results in continuous-valued, time-varying events, with limited correspondence to the actual spike times¹⁰. In RADICaL, deconvolved events are therefore modeled as samples from a time-varying zero-inflated gamma (ZIG) distribution, which has been shown to be more appropriate for calcium data²², and whose parameters are taken as the output of the generator RNN (**Fig. 1c**; details in *Methods*). We then define the network state at any given time point as a vector containing the inferred (i.e., de-noised) event rates of all neurons, where the de-noised event rate is taken as the mean of each neuron’s inferred ZIG distribution at each time point. The de-noised event rates are latent variables that are tied to the underlying network state at each time point. We note that in RADICaL, the generator RNN must produce multiple parameters that control the shape of the ZIG distributions, which may result in activity at the level of the generator and factors that does not directly correspond to the biological network’s activity. To avoid this complication, rather than using the factors as an estimate of the biological network’s state, we used the de-noised event rates. Doing so for both RADICaL and AutoLFADS allowed us to compare methods as directly as possible.

Second, we developed a novel neural network training strategy, selective backpropagation through time (SBTT), that leverages the precise sampling times of individual neurons to enable recovery of high-frequency network dynamics. Since standard multiphoton acquisition systems rely on point-by-point raster scanning of a laser beam to acquire frames, it is possible to increase temporal precision of individual neurons’ sample times beyond the timing of individual frames, by exploiting the relationship between scanning position and accumulated scanning time within the frame (**Fig. 1d**). To leverage this information to improve inference of high-frequency network dynamics on single trials, we reframe the underlying interpolation problem as a missing data problem: we treat low-sampling rate data from each neuron as high-sampling rate data at the level of the population. In this framing, each neuron is effectively sampled sparsely in time, *i.e.*, the majority of time points for each neuron do not contain valid data (**Fig. 1e**). Such sparsely sampled data creates a challenge when training the underlying neural network: briefly, neural networks are trained by adjusting their parameters (weights), and performing this adjustment requires evaluating the gradient of a cost function with respect to weights. SBTT allows us to compute this gradient using only the valid data, and ignore the many missing samples (**Fig. 1f**; see *Methods*). Because this feature only affects how we compute the gradient and update the weights, the network still infers event rates for every neuron at every time point, regardless of whether samples exist at that time point or not. This allows the trained network to accept 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. ²³). (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) SBT 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

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

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

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

Figure 2 | Application of RADiCaL to synthetic data. (a) Example firing rates and spiking activity from a Lorenz system simulated at 7 Hz, deconvolved calcium events (inputs to RADiCaL), and the corresponding rates and factors inferred by RADiCaL. Simulation parameters were tuned so that the performance in inferring spikes using OASIS matched previous benchmarks¹⁰ (see *Methods*). (b) True

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

These synthetic results provide an important proof-of-principle that RADICaL can infer high-frequency features of the network activity underlying 2p signals, which is readily validated when ground truth is known. However, it is important to acknowledge limitations of the simulation process that might constrain the generality of these results when applying to real data. In particular, the parameter space is very large, especially considering the variety of calcium indicators, protein expression patterns, imaging settings, cell types and firing rate patterns. An exhaustive search of this parameter space is infeasible, and we chose parameters so that the resulting signal-to-noise regime produced similar correlations between real and inferred spike trains as those typically observed¹⁰ (see *Methods*). However, it is difficult to know whether the results from any particular choice of simulation parameters (or a variety of choices) can be extrapolated to real experimental conditions. Thus, we next benchmarked performance on real data to demonstrate RADICaL's utility in the real world.

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

We next tested RADICaL on 2p recordings from mice performing a forelimb water grab task (**Fig. 3a, top**). We analyzed data from four experiments: two mice, and two sessions from each mouse in which different brain areas were imaged (M1, S1). Our task is a variant of the water-reaching task of Galiñanes & Huber²⁷. In each trial, the mouse was cued by the pitch of an auditory tone to reach to a left or right spout and retrieve a droplet of water with its right forepaw (**Fig. 3a, bottom**; see *Methods*). The forepaw position was tracked at 150 frames per second with DeepLabCut²⁸ for 420-560 trials per experiment. To test whether each method could reveal structure in the neural activity at finer resolution than left vs. right reaches, we divided trials from each condition into subgroups based on forepaw height during the reach (**Fig. 3a, top right**; see *Methods*). Two-photon calcium imaging from GCaMP6f transgenic mice was performed at 31 Hz, with 430-543 neurons within the FOV in each experiment (**Fig. 3b**).

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

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

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

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

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 0.91 across all experiments, versus 0.75 and 0.85 for s-deconv and AutoLFADS, respectively; velocity: average R^2 of 0.62 across the mice/areas, versus 0.37 and 0.51 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**).

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

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

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

In previous demonstrations on electrophysiological spiking data, LFADS maintained accurate performance in reconstructing single-trial neural activity and decoding even when reducing the number of sampled neurons¹⁶. Enabling the same capabilities for 2p imaging could help mitigate the effects of the tradeoffs in sampling frequency or signal-to-noise that occur when laser scanning over large FOVs. To evaluate whether this holds for RaDICAL, we performed a neuron-downsampling experiment where we gradually reduced the number of neurons used in training RADICaL or AutoLFADS (**Fig. 5a**). RADICaL retained relatively high decoding performance as the population size was reduced (**Fig. 5b**; data from Mouse2 M1). Decoding performance declined gradually, with a steeper slope for velocity. Notably, however, performance when only 121/439 neurons were used for training RADICaL was similar to that of AutoLFADS - and higher than for s-deconv - even when those methods were applied to the full population of 439 neurons. Note that this analysis represents a lower bound on performance: for this proof-of-concept, we simply artificially excluded from our analysis data collected as the laser scanned outside the restricted FOVs, which resulted in substantial time periods that lacked data entirely (e.g., 2/3 of the total sampling time for the smallest FOV considered). In a real application, those time periods that were artificially excluded could instead be used to monitor other brain areas or layers, or to monitor the same neurons with higher sampling rates, either of which might be expected to provide additional information. These results provide an avenue to retain information by scanning smaller areas when capturing multiple layers or regions, opening opportunities to study interesting questions such as communication between layers or interactions between regions (see *Discussion*).

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

Discussion

2p imaging is a widely-used method for interrogating neural circuits, with the potential to monitor vast volumes of neurons and provide new circuit insights that elude electrophysiology. To date, however, it has proven challenging to precisely infer network states from imaging data, due in large part to the inherent noise, indicator dynamics, and low temporal resolution 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^{17,29–32}. Third, this de-noising capability will enable greater insight into processes that fundamentally differ from trial to trial, such as learning from errors^{33,34}, variation in internal states such as arousal^{35,36}, or paradigms in which tuning to uninstructed movements contaminates measurement of the task-related behavioral variables of interest³⁷. 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 data³.

In recent years, a variety of computational methods have been developed to analyze 2p imaging data⁹. 2p preprocessing pipelines^{5,26} 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^{38–40}, while a few others have focused on uncovering population-level structure^{41–46} or locally linear dynamics underlying population activity, in particular via switching linear dynamical systems-based methods^{47,48}. 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 AutoLFADS¹⁷.

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^{49,50}, or jittered inferred events on sub-frame timescales to minimize the reconstruction error of the associated fluorescence³⁸. 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 direction⁵¹, 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

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

The ability to achieve high-quality network state inference despite limited neuronal population size opens the door to testing new choices about how to scan during the experiments themselves, e.g., by scanning smaller and more disparate FOVs (across layers^{7,8} or brain areas^{52,53}) to understand how spatially-segregated populations interact, while potentially preserving the ability to infer network states from each FOV. When the number of neurons within each FOV is limited, one advantage that RADICaL inherits from LFADS is that it allows for multi-session stitching¹⁶, which could provide an avenue to combine data from different sessions to improve inference of the underlying dynamics for each FOV.

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

Acknowledgements

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

Code availability

Code will be made available upon publication.

Data availability

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

Author Contributions

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

References

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

- (2017).
6. Peron, S. P., Freeman, J., Iyer, V., Guo, C. & Svoboda, K. A cellular resolution map of barrel cortex activity during tactile behavior. *Neuron* **86**, 783–799 (2015).
7. Chen, J. L., Carta, S., Soldado-Magraner, J., Schneider, B. L. & Helmchen, F. Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. *Nature* **499**, 336–340 (2013).
8. Chen, S. X., Kim, A. N., Peters, A. J. & Komiyama, T. Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. *Nature neuroscience* **18**, 1109–1115 (2015).
9. Pnevmatikakis, E. A. Analysis pipelines for calcium imaging data. *Current Opinion in Neurobiology* **55**, 15–21 (2019).
10. Berens, P. *et al.* Community-based benchmarking improves spike rate inference from two-photon calcium imaging data. *PLoS computational biology* **14**, e1006157 (2018).
11. Pachitariu, M., Stringer, C. & Harris, K. D. Robustness of spike deconvolution for neuronal calcium imaging. *Journal of Neuroscience* **38**, 7976–7985 (2018).
12. Evans, M. H., Petersen, R. S. & Humphries, M. D. On the use of calcium deconvolution algorithms in practical contexts. *bioRxiv* 871137 (2019).
13. Huang, L. *et al.* Relationship between simultaneously recorded spiking activity and fluorescence signal in GCaMP6 transgenic mice. *eLife* **10**, e51675 (2021).
14. Siegle, J. H. *et al.* Reconciling functional differences in populations of neurons recorded with two-photon imaging and electrophysiology. *bioRxiv* (2020).
15. Sussillo, D., Jozefowicz, R., Abbott, L. & Pandarinath, C. LFADS-latent factor analysis via dynamical systems. *arXiv preprint arXiv:1608.06315* (2016).
16. Pandarinath, C. *et al.* Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat Methods* **15**, 805–815 (2018).
17. Keshtkaran, M. R. *et al.* A large-scale neural network training framework for generalized estimation of single-trial population dynamics. *bioRxiv* (2021).
18. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat Neurosci* **17**, 1500–1509 (2014).
19. Pandarinath, C. *et al.* Latent Factors and Dynamics in Motor Cortex and Their Application to Brain–Machine Interfaces. *J. Neurosci.* **38**, 9390–9401 (2018).
20. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
21. Keshtkaran, M. R. & Pandarinath, C. Enabling hyperparameter optimization in sequential autoencoders for spiking neural data. in *Advances in Neural Information Processing Systems* 15937–15947 (2019).
22. Wei, X.-X. *et al.* A zero-inflated gamma model for deconvolved calcium imaging traces. *arXiv preprint arXiv:2006.03737* (2020).
23. Chen, T.-W. *et al.* Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300 (2013).
24. Zhao, Y. & Park, I. M. Variational latent gaussian process for recovering single-trial dynamics from population spike trains. *Neural computation* **29**, 1293–1316 (2017).
25. Friedrich, J., Zhou, P. & Paninski, L. Fast online deconvolution of calcium imaging data. *PLoS Comput Biol* **13**, e1005423 (2017).
26. Giovannucci, A. *et al.* CalmAn an open source tool for scalable calcium imaging data analysis. *eLife* **8**, e38173 (2019).
27. Galiñanes, G. L., Bonardi, C. & Huber, D. Directional reaching for water as a cortex-dependent behavioral framework for mice. *Cell reports* **22**, 2767–2783 (2018).
28. Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience* **21**, 1281–1289 (2018).
29. Hatsopoulos, N. G., Xu, Q. & Amit, Y. Encoding of movement fragments in the motor cortex. *Journal of Neuroscience* **27**, 5105–5114 (2007).
30. Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A. & Poeppel, D. Neuroscience needs behavior: correcting a reductionist bias. *Neuron* **93**, 480–490 (2017).
31. Whishaw, I. Q. *et al.* Organization of the reach and grasp in head-fixed vs freely-moving mice provides support for multiple motor channel theory of neocortical organization. *Experimental brain research* **235**, 1919–1932 (2017).

32. Wilschko, A. B. *et al.* Revealing the structure of pharmacobehavioral space through motion sequencing. *Nature neuroscience* **23**, 1433–1443 (2020).
33. Herzfeld, D. J., Kojima, Y., Soetedjo, R. & Shadmehr, R. Encoding of error and learning to correct that error by the Purkinje cells of the cerebellum. *Nature neuroscience* **21**, 736–743 (2018).
34. Vyas, S., O'Shea, D. J., Ryu, S. I. & Shenoy, K. V. Causal role of motor preparation during error-driven learning. *Neuron* **106**, 329–339 (2020).
35. Steinmetz, N. A., Zarka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
36. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, (2019).
37. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nature neuroscience* **22**, 1677–1686 (2019).
38. Hoang, H. *et al.* Improved hyperacuity estimation of spike timing from calcium imaging. *Sci Rep* **10**, 17844 (2020).
39. Rupprecht, P. *et al.* Database and deep learning toolbox for noise-optimized, generalized spike inference from calcium imaging. <http://biorxiv.org/lookup/doi/10.1101/2020.08.31.272450> (2020) doi:10.1101/2020.08.31.272450.
40. Sebastian, J., Sur, M., Murthy, H. A. & Magimai-Doss, M. Signal-to-signal neural networks for improved spike estimation from calcium imaging data. *PLoS Comput Biol* **17**, e1007921 (2021).
41. Dechery, J. B. & MacLean, J. N. Functional triplet motifs underlie accurate predictions of single-trial responses in populations of tuned and untuned V1 neurons. *PLoS computational biology* **14**, e1006153 (2018).
42. Kirschbaum, E. *et al.* LeMoNADe: Learned Motif and Neuronal Assembly Detection in calcium imaging videos. *arXiv:1806.09963 [q-bio]* (2019).
43. Mackevicius, E. L. *et al.* Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience. *eLife* **8**, e38471 (2019).
44. Triplett, M. A., Pujic, Z., Sun, B., Avitan, L. & Goodhill, G. J. Model-based decoupling of evoked and spontaneous neural activity in calcium imaging data. *PLoS Comput Biol* **16**, e1008330 (2020).
45. Williams, A. H. *et al.* Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis. *Neuron* **98**, 1099–1115.e8 (2018).
46. Wu, A., Pashkovski, S., Datta, S. R. & Pillow, J. W. Learning a latent manifold of odor representations from neural responses in piriform cortex. in *Advances in Neural Information Processing Systems* (eds. Bengio, S. *et al.*) vol. 31 (Curran Associates, Inc., 2018).
47. Costa, A. C., Ahamed, T. & Stephens, G. J. Adaptive, locally linear models of complex dynamics. *Proc Natl Acad Sci USA* **116**, 1501–1510 (2019).
48. Glaser, J., Whiteway, M., Cunningham, J. P., Paninski, L. & Linderman, S. Recurrent Switching Dynamical Systems Models for Multiple Interacting Neural Populations. in *Advances in Neural Information Processing Systems* (eds. Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F. & Lin, H.) vol. 33 14867–14878 (Curran Associates, Inc., 2020).
49. Picardo, M. A. *et al.* Population-Level Representation of a Temporal Sequence Underlying Song Production in the Zebra Finch. *Neuron* **90**, 866–876 (2016).
50. Mano, O. *et al.* Using slow frame rate imaging to extract fast receptive fields. *Nature communications* **10**, 1–13 (2019).
51. Prince, L. Y., Bakhtiari, S., Gillon, C. J. & Richards, B. A. *Parallel inference of hierarchical latent dynamics in two-photon calcium imaging of neuronal populations.* <http://biorxiv.org/lookup/doi/10.1101/2021.03.05.434105> (2021) doi:10.1101/2021.03.05.434105.
52. Minderer, M., Brown, K. D. & Harvey, C. D. The spatial structure of neural encoding in mouse posterior cortex during navigation. *Neuron* **102**, 232–248 (2019).
53. Sofroniew, N. J., Flickinger, D., King, J. & Svoboda, K. A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *Elife* **5**, e14472 (2016).
54. Jaderberg, M. *et al.* Population Based Training of Neural Networks. *arXiv:1711.09846 [cs]* (2017).
55. Dana, H. *et al.* High-performance calcium sensors for imaging activity in neuronal populations and microcompartments. *Nat Methods* **16**, 649–657 (2019).
56. Art, J. Photon detectors for confocal microscopy. in *Handbook of biological confocal microscopy* 251–264 (Springer, 2006).

57. Starck, J.-L., Murtagh, F. D. & Bijaoui, A. *Image processing and data analysis: the multiscale approach*. (Cambridge University Press, 1998).
58. Vogelstein, J. T. *et al.* Fast Nonnegative Deconvolution for Spike Train Inference From Population Calcium Imaging. *Journal of Neurophysiology* **104**, 3691–3704 (2010).
59. Heikkilä, J. & Silvén, O. A four-step camera calibration procedure with implicit image correction. in *Proceedings of IEEE computer society conference on computer vision and pattern recognition* 1106–1112 (IEEE, 1997).
60. Zhang, Z. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence* **22**, 1330–1334 (2000).

Methods

AutoLFADS and RADICaL architecture and training

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

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

AutoLFADS is a recent implementation of the population based training (PBT) approach⁵⁴ on LFADS to perform automatic, large-scale hyperparameter (HP) search. A detailed overview of AutoLFADS is in refs. ^{17,21}. Briefly, PBT distributes training across dozens of models in parallel, and uses evolutionary algorithms to tune HPs over many generations. At the end of each generation, a selection process was performed to choose higher performing models and replace the poor models with the higher performing models. The HPs of the higher performing models were perturbed before the next generation to increase the HP search space. After many generations (~30-150), the PBT process converges upon a high performing model with optimized HPs.

Twenty LFADS models were trained in parallel for 50 epochs per generation for both AutoLFADS and RADICaL runs and for both simulation and real 2P data. KL and L2 regularization penalties were linearly ramped for the first 80 epochs of training during the first generation. Training was stopped when there was no improvement in performance after 25 generations. The HPs optimized by PBT were the model's learning rate and six regularization HPs: scaling weights for the L2 penalties on the generator, controller, and initial condition encoder RNNs, scaling weights for the KL penalties on the initial conditions and controller outputs, and two dropout probabilities ("keep ratio" for coordinated dropout²¹; and RNN network dropout probability). The HP search ranges are detailed in **Table 1**. The magnitudes of the HP perturbation were controlled by weights and specified for different HPs (a weight of 0.3 results in perturbation factors between 0.7 and 1.3; **Table 1**). The learning rate and dropout probabilities were restricted to their specified search ranges and were sampled from uniform distributions. The KL and L2 HPs were sampled from log-uniform distributions and could be perturbed outside of the

initial search ranges. Identical hyperparameter settings were used for both RADICaL and AutoLFADS and for both synthetic datasets and real 2P datasets.

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) distribution²²:

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

where $x_n(t)$ is the distribution of observed deconvolved events, $\alpha_n(t)$, $k_n(t)$, and 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 . 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 (sig_n). sig_n is a positive parameter that scales the outputs of the sigmoid to be in a range between 0 and sig_n , and is optimized alongside network weights. An L2 penalty is applied between 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:

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

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

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

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.

Simulation experiments

Generating spike trains from an underlying Lorenz system

Synthetic data were generated using the Lorenz system as described in the original LFADS work^{15,16}. Lorenz parameters were set to standard values (σ : 10, ρ : 28, and β : 8/3), and Δt 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.

Generating fluorescence signals from synthetic spike trains

Realistic fluorescence signals were generated from the spike trains by convolving them with a kernel for an autoregressive process of order 2 and passing the results through a nonlinearity that matched values extracted from the literature for the calcium indicator GCaMP6f^{3,55} (**Supp Fig. 2a & b**). Three noise sources were added to reproduce variability present in real data^{56–58}: Gaussian noise to the size of the calcium spike, and Gaussian and Poisson noise to the final trace (**Supp Fig. 2a & b**). This fluorescence generation process was realized as follows: First, spike trains $s(t)$ were generated from the Lorenz system as mentioned above. Independent Gaussian noise ($sd = 0.1$) was added to each spike in the spike train to model the variability in spike amplitude. Next, we modeled the calcium concentration dynamics $c(t)$ as an autoregressive process of order 2:

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

with $s(t)$ representing the number of spikes at time t . The autoregressive coefficients γ_1 and γ_2 were computed based on the rise time, decay time ($\tau_{on} = 20$ ms, $\tau_{off} = 400$ ms for GCaMP6f) of the calcium indicators, and the sampling frequency. Note that while there is substantial variability in τ s across neurons in real data³, selecting and mimicking this variability was not relevant in our work, because we compared the methods (i.e., RADICAL, AutoLFADS, and s-deconv) after deconvolution. The calcium concentration dynamics were further normalized so that the peak height of the calcium dynamics generated from a single spike equalled one, regardless of the sampling frequency. Subsequently, we computed the noiseless fluorescence signals by passing the calcium dynamics through a nonlinear transformation estimated from the literature⁵⁵ for the calcium indicator GCaMP6f (**Supp Fig. 2c & d**). After the nonlinear transformation, the relationship between spike size and trace size was corrupted, and therefore we assumed the baseline of fluorescence signals to be zero and the signals were rescaled to the range in $[0,1]$ using min-max normalization. Finally, Gaussian noise ($\sim N(0,sn)$) and Poisson noise (simulated as gaussian with mean 0 and variance proportional to the signal amplitude at each time point via a constant d) were added to the normalized traces. The resulting fluorescence traces had the same sampling frequency as the synthetic spike trains (100 Hz).

A crucial parameter is the noise level associated with each fluorescence trace. High noise levels lead to very poor spike detection and very low noise levels enable a near-perfect reconstruction of the spike train. In order to select a realistic level of noise we matched the correlations between real and inferred spike trains of the simulated data to those typically observed¹⁰. We found that a truncated normal distribution of noise level for Gaussian and Poisson noise best matched the correlations. More specifically, for each neuron, $sn=d$ was sampled independently from a truncated normal distribution $N(0.12, 0.02)$ truncated below 0.06. With the above noise setting, the mean correlation coefficient r between the deconvolved events and ground truth spikes was 0.32, which is consistent with the standard results reported in the spikefinder paper¹⁰ for OASIS. It is worth stressing that real data feature a broad range of noise levels that depend on the imaging conditions, depth, expression level, laser power and other factors. Here we did not attempt to investigate all possible noise conditions, but instead, we aimed to create a simulation with known latent variables (i.e., low-dimensional factors and event rates) that reasonably approximated realistic signal-to-noise levels, in order to provide a tractable test case to compare RADICAL to other methods before attempting comparisons on real data.

Recreating variability in sampling times due to 2p laser scanning

The fluorescence traces were simulated at 100 Hz as mentioned above. A subsampling step was then performed with sampling times for each neuron staggered in time to simulate the variability in sampling times due to 2p laser scanning (as in **Fig. 1e**). This produced fluorescence traces where individual neurons were sampled at 33.3 Hz, with phases of 0, 11, 22 ms based on each neuron's location (top, middle and bottom of the FOV, respectively). To break this down, each neuron was sparsely sampled every three time points and the relative sampled times between neurons were fixed. For example, in 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 1 was sampled at time points 2, 5, 8, ... and neuron 2 was sampled at time points 3, 6, 9, Thus, the sampling frequency for each individual neuron was 33.3 Hz, while the sampling frequency for the population was retained at 100 Hz by filling the non-sampled time points with NaNs. The resulting 33.3 Hz simulated fluorescence signals for each individual neuron (i.e., with NaNs excluded) were deconvolved using OASIS²⁵ (as implemented in CalmAn²⁶) using an auto-regressive model of order 1 with s_{min} of 0.1.

Data preparation for each method

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

Mapping to ground truth Lorenz states

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

Real 2p experiments

Surgical procedures

All procedures were approved by the University of Chicago Animal Care and Use Committee. Two male Ai148D transgenic mice (TIT2L-GC6f-ICL-tTA2, stock 030328; Jackson Laboratory) were used. Each mouse underwent a single surgery. Mice were injected subcutaneously with dexamethasone (8 mg/kg) 24 hours and 1 hour before surgery. Mice were anesthetized with 2-2.5% inhaled isoflurane gas, then injected intraperitoneally with a ketamine-medetomidine solution (60 mg/kg ketamine, 0.25 mg/kg medetomidine), and maintained on a low level of supplemental isoflurane (0-1%) if they showed any signs that the depth of anesthesia was insufficient. Meloxicam was also administered subcutaneously (2 mg/kg) at the beginning of the surgery and for 1-3 subsequent days. The scalp was shaved, cleaned, and resected, the skull was cleaned and the wound margins glued to the skull with tissue glue (VetBond, 3M), and a 3 mm circular craniotomy was made with a 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 anterior and 1.6 mm lateral of bregma. The craniotomy was cleaned with SurgiFoam (Ethicon) soaked in phosphate-buffered solution (PBS), then virus (AAV9-CaMKII-Cre, stock 2.1×10^{13} particles/nL, 1:1 dilution in PBS, Addgene) was pressure injected (Nanoject III, Drummond Scientific) at two or four sites near the target site, with 140 nL injected at each of two depths per site (250 and 500 μ m below the pia) over 5 minutes each. The craniotomy was then sealed with a custom cylindrical glass plug (3 mm diameter, 660 μ m depth; Tower Optical) bonded (Norland Optical Adhesive 61, Norland) to a 4 mm #1 round coverslip (Harvard Apparatus), glued in place first with tissue glue (VetBond) and then with cyanoacrylate glue (Krazy Glue) mixed with dental acrylic powder (Ortho Jet; Lang Dental). A small craniotomy was also made using a dental drill over right CFA at 0.4 mm anterior and 1.6 mm lateral of bregma, where 140 nL of AAVretro-tdTomato (stock 1.02×10^{13}

particles/nL, Addgene) was injected at 300 μm below the pia. This injection labeled cells in left CFA projecting to the contralateral CFA. Here, this labeling was used solely for stabilizing the imaging plane (see below). The small craniotomy was sealed with a drop of Kwik-Cast (World Precision Instruments). Two layers of MetaBond (C & B) were applied, then a custom laser-cut titanium head bar was affixed to the skull with black dental acrylic. Animals were awoken by administering atipamezole via intraperitoneal injection and allowed to recover at least 3 days before water restriction.

Behavioral task

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

Behavior was also recorded using a pair of cameras (BFS-U3-16S2M-CS, FLIR; varifocal lenses COZ2813CSIR2, Computar) mounted 150 mm from the right paw rest at 10° apart to enable 3D triangulation. Infrared illuminators enabled behavioral imaging while performing 2p imaging in a darkened microscope enclosure. Cameras were synchronized and recorded at 150 frames per second with real-time image cropping and JPEG compression, and streamed to one HDF5 file per camera (areaDetector module of EPICS, CARS). The knuckles and wrist of the reaching paw were tracked in each camera using DeepLabCut²⁸ and triangulated into 3D using camera calibration parameters obtained from the MATLAB Stereo Camera Calibration toolbox^{59,60}. To screen the tracked markers for quality we created distributions of all inter-marker distances in 3D across every labeled frame and identified as problematic frames with any inter-marker distance exceeding the 99.9th percentile of its respective distribution. Trials with more than one problematic frame in the period of -200 ms to 800 ms after the raw reach onset were discarded (where reach onset was taken as the first 60 Hz tick after the paw rest touch sensor fell below contact threshold). The kinematics of all trials that passed this screening procedure were visualized to confirm quality. Centroid marker kinematics were obtained by averaging the kinematics of all paw markers, locking them to behavioral events and then smoothing using a Gaussian filter (15 ms s.d.). To obtain velocity and acceleration, centroid data was numerically differentiated with MATLAB’s *diff* function and then smoothed again using a Gaussian filter (15 ms s.d.).

Two-photon imaging

Calcium imaging was performed with a Neurolabware two-photon microscope and pulsed Ti:sapphire laser (Vision II, Coherent). Depth stability of the imaging plane was maintained using a custom plugin that acquired an image stack at the beginning of the session (1.4 μm spacing), then compared a registered rolling average of the red-channel data to each plane 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 acquired, the objective was automatically moved to compensate. This typically resulted in a slow and steady upward (outward) movement of the objective over the course of the session.

Offline, images were run through Suite2p to perform motion correction, region-of-interest (ROI) detection, and fluorescence extraction from both ROIs and neuropil. ROIs were manually curated using the Suite2p GUI to retain only those corresponding to somas. We then subtracted the neuropil signal scaled by 0.7²³. Neuropil-subtracted ROI fluorescence was then detrended by performing a running 10th percentile operation, smoothing with a Gaussian filter (20 s s.d.), then

subtracting the result from the trace. This result was fed into OASIS²⁵ using the ‘thresholded’ method, AR1 event model, and limiting the tau parameter to be between 300 and 800 ms. Neurons were discarded if they did not meet a minimum signal-to-noise (SNR) criterion. To compute SNR, we took the fluorescence at each time point when OASIS identified an “event” (non-zero), computed (fluorescence - neuropil) / neuropil, and computed the median of the resulting distribution. ROIs were 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 event sizes, smoothed the distribution (ksdensity in MATLAB, with an Epanechnikov kernel and log transform), found the peak of the smoothed distribution, and divided all event sizes by this value. This rescales the peak of the distribution to have a value of unity. Data from two mice and two brain areas (4 sessions in total) were used (Mouse1 M1: 510 neurons, 560 trials; Mouse1 S1: 543 neurons, 506 trials; Mouse2 M1: 439 neurons, 475 trials; Mouse2 S1: 509 neurons, 421 trials).

Data preparation for modeling with RADICaL and AutoLFADS

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

Trial grouping

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

PSTH analysis and comparing RADICaL and AutoLFADS single-trial rates

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

Low-D analysis

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

Decoding analysis

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

Coherence analysis

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

t-SNE analysis on the weights mapping from factors to ZIG parameters

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

Neuron downsampling

The neuron downsampling experiment was performed on the Mouse2 M1 dataset. The number of neurons included when training RADICaL or AutoLFADS was gradually reduced by limiting the area of FOV that the neurons were sampled from. 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 of included neurons to be 439, 321, 262, 121 and 59. An individual RADICaL model and AutoLFADS model were trained for each number of neurons. Decoding was performed using ridge regression (see above).

			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	-

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