


מכון ויצמן למדע

WEIZMANN INSTITUTE OF SCIENCE



Dimensionality reduction of calcium-imaged neuronal population activity

Document Version:

Early version, also known as pre-print

Citation for published version:

Koh, TH, Bishop, WE, Kawashima, T, Jeon, BB, Srinivasan, R, Kuhlman, SJ, Ahrens, MB, Chase, SM & Yu, BM 2022, 'Dimensionality reduction of calcium-imaged neuronal population activity', *Nature Computational Science*. <https://doi.org/10.1038/s43588-022-00390-2>

Total number of authors:

9

Digital Object Identifier (DOI):

[10.1038/s43588-022-00390-2](https://doi.org/10.1038/s43588-022-00390-2)

Published In:

Nature Computational Science

License:

CC BY-NC

General rights

@ 2020 This manuscript version is made available under the above license via The Weizmann Institute of Science Open Access Collection is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognize and abide by the legal requirements associated with these rights.

How does open access to this work benefit you?

Let us know @ library@weizmann.ac.il

Take down policy

The Weizmann Institute of Science has made every reasonable effort to ensure that Weizmann Institute of Science content complies with copyright restrictions. If you believe that the public display of this file breaches copyright please contact library@weizmann.ac.il providing details, and we will remove access to the work immediately and investigate your claim.

1 Dimensionality reduction of calcium-imaged 2 neuronal population activity

3 **Tze Hui Koh**^{1,2}, **William E. Bishop**^{2,3,4}, **Takashi Kawashima**^{4,5}, **Brian B. Jeon**^{1,2}, **Ranjani**
4 **Srinivasan**^{1,6}, **Sandra J. Kuhlman**^{7,8}, **Misha B. Ahrens**⁴, **Steven M. Chase**^{1,7,*}, and **Byron**
5 **M. Yu**^{1,7,9,*}

6 ¹Department of Biomedical Engineering, Carnegie Mellon University, PA

7 ²Center for the Neural Basis of Cognition, PA

8 ³Department of Machine Learning, Carnegie Mellon University, PA

9 ⁴Janelia Research Campus, Howard Hughes Medical Institute, VA

10 ⁵Department of Brain Sciences, Weizmann Institute of Science, Israel

11 ⁶Department of Electrical and Computer Engineering, Johns Hopkins University, MD

12 ⁷Carnegie Mellon Neuroscience Institute, Carnegie Mellon University, PA

13 ⁸Department of Biological Sciences, Carnegie Mellon University, PA

14 ⁹Department of Electrical and Computer Engineering, Carnegie Mellon University, PA

15 *These authors contributed equally to this work

16 **ABSTRACT**

Calcium imaging has been widely adopted for its ability to record from large neuronal populations. To summarize the time course of neural activity, dimensionality reduction methods, which have been applied extensively to population spiking activity, may be particularly useful. However, it is unclear if the dimensionality reduction methods applied to spiking activity are appropriate for calcium imaging. We thus carried out a systematic study of design choices based on standard dimensionality reduction methods. We also developed a novel method to perform deconvolution and dimensionality reduction simultaneously (termed CILDS). CILDS most accurately recovered the single-trial, low-dimensional time courses from calcium imaging that would have been recovered from spiking activity. CILDS also outperformed the other methods on calcium imaging recordings from larval zebrafish and mice. More broadly, this study represents a foundation for summarizing calcium imaging recordings of large neuronal populations using dimensionality reduction in diverse experimental settings.

18 **Introduction**

19 Computations in the brain occur through the coordinated, time-varying activity of populations of neurons. Dimen-
20 sionality reduction is a class of statistical methods commonly used for summarizing neural population activity¹⁻³.
21 It transforms high-dimensional neural recordings, such as spiking activity from a population of recorded neurons,

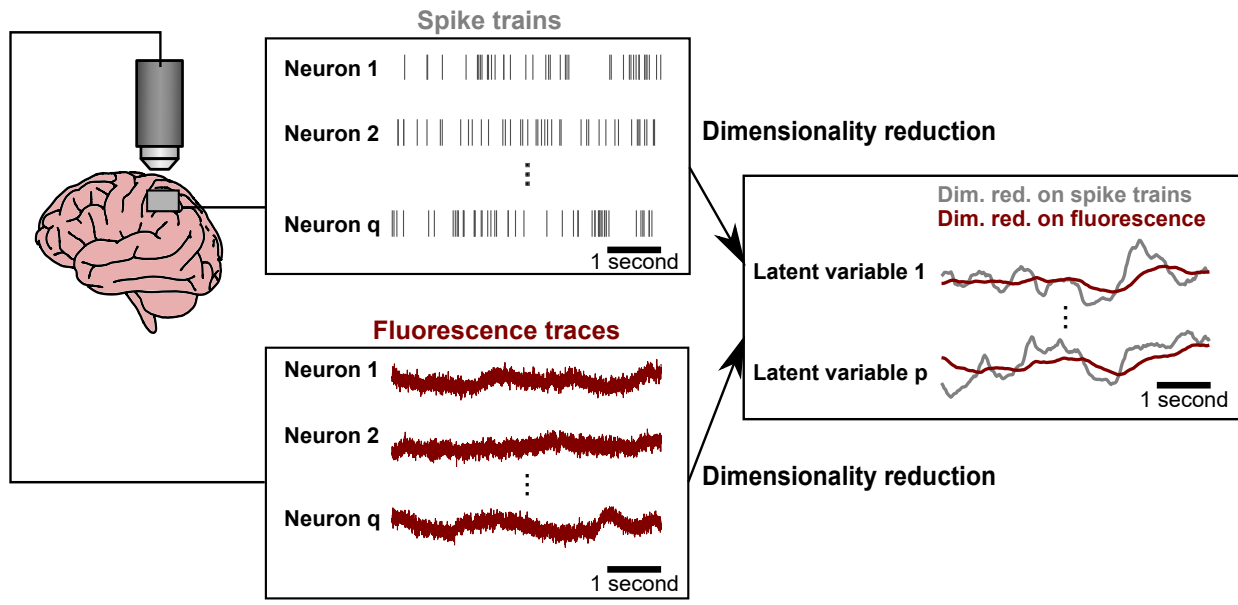


Figure 1. Dimensionality reduction of neuronal population activity based on electrophysiological recordings versus calcium imaging can yield different low-dimensional representations. Left, spike trains (grey) and fluorescence traces (maroon) recorded from the same hypothetical population of neurons. Right, different latent trajectories resulting from dimensionality reduction applied to spike trains (grey) versus fluorescence traces (maroon).

22 into compact low-dimensional representations termed *latent variables*. These low-dimensional representations
23 facilitate the investigation of how neural population activity varies over time, across experimental conditions,
24 and across repeated experimental trials of the same condition. In particular, dimensionality reduction has been
25 used to uncover neural mechanisms underlying decision making⁴, motor control⁵, learning⁶, working memory⁷,
26 sensorimotor timing⁸, attention⁹, olfaction¹⁰, speech¹¹, and more.

27 Dimensionality reduction has typically been applied to electrophysiological recordings. In the last decade,
28 optical imaging has been widely adopted to record from large populations of neurons given its ability to sample
29 neurons densely within the field of view, track neurons over long periods of time, and label neurons by cell type or
30 projection, among other advantages¹². A leading type of optical imaging is calcium imaging, which uses calcium
31 indicators to track the transient increase in intracellular calcium levels that accompanies electrical spiking activity¹³.
32 These changes in calcium levels are then optically recorded via changes in fluorescence. Calcium imaging has the
33 capability of imaging even the whole brain of some small animals (e.g., larval zebrafish) at single neuron resolution¹⁴,
34 albeit at a lower temporal resolution than electrical recordings.

35 With the increasing use of calcium imaging, many studies are now beginning to apply dimensionality reduction to

36 calcium imaging recordings^{15–26}. A critical question is whether the same dimensionality reduction methods applied
37 to study spiking activity are also appropriate for calcium imaging recordings²⁷. Here, we focus on single-trial time
38 courses of latent variables, termed *neural trajectories*. This enables the study of how trial-to-trial differences in the
39 time course of neural activity relate to trial-to-trial differences in perception, decision making, and behavior^{8, 16, 28–31}.
40 We seek to understand if the neural trajectories extracted from calcium imaging recordings match the neural
41 trajectories that would have been extracted from the spiking activity underlying these recordings. A key reason why
42 the neural trajectories would be different is the slow, indicator-dependent decay of measured calcium levels after each
43 spiking event¹³. This decay introduces temporal correlations in the calcium imaging recordings that would not be
44 present with spiking activity alone. Deconvolution is a technique that aims to recover spiking activity from calcium
45 imaging recordings³². However, deconvolution techniques do not, as yet, recover the underlying spikes exactly³³.
46 Thus, the neural trajectories that are extracted from calcium imaged activity may be quite different from the neural
47 trajectories extracted from spiking activity (Fig. 1), which may limit the use of calcium imaging for studying
48 population activity time courses. In this work, our central goals are to i) systematically study the appropriateness of
49 dimensionality reduction methods for summarizing the time course of calcium imaging recordings, and ii) propose
50 a new dimensionality reduction method that is tailored for extracting neural trajectories from calcium imaging
51 recordings.

52 We sought to address three questions. First, we asked if deconvolution should be used with dimensionality
53 reduction when extracting neural trajectories, and if so, how it should be applied. Second, we asked how different
54 experimental variables (e.g., the decay constant of the calcium indicator, the timescale of the latent time courses, and
55 the number of imaged neurons) impact the ability to recover neural trajectories from calcium imaging recordings.
56 Third, we asked if it is necessary for the dimensionality reduction method to employ a dynamics model for the latent
57 variables, as such a model might enable the time course of the neural trajectory to be more cleanly separated from
58 the time course of calcium decay.

59 We addressed these questions by comparing several approaches, including i) standard dimensionality reduction
60 applied directly to the recorded fluorescence; ii) a two stage-method in which deconvolution is applied separately
61 to each neuron's fluorescence trace to estimate spiking activity, then standard dimensionality reduction is applied
62 to the estimated spiking activity; and iii) a novel unified method that we propose here (Calcium Imaging Linear
63 Dynamical System, CILDS), which performs deconvolution and dimensionality reduction jointly. We first applied
64 these methods to simulated fluorescence traces, in which we systematically varied several experimental variables
65 over a wide range. We then applied these methods to calcium imaging recordings from the dorsal raphe nucleus of
66 larval zebrafish and the primary visual cortex of mice. Across these settings, we found that CILDS outperformed the
67 other methods. This was especially true when the neural activity fluctuated more quickly over time (timescale of tens
68 to hundreds of milliseconds). We also found that, to accurately peer through the calcium decay to obtain accurate

69 neural trajectories, it is necessary to include a latent dynamical model, as in CILDS. Overall, our work provides a
70 foundation for using dimensionality reduction to summarize the time course of calcium imaging recordings.

71 Results

72 Dimensionality reduction is typically applied to spike trains recorded from a population of neurons, yielding a
73 low-dimensional representation of that activity. Calcium imaging provides a transformed view of those spike trains.
74 Our central goal is to develop dimensionality reduction methods appropriate for calcium imaging recordings to
75 recover the same low-dimensional representation as that obtained from spike trains. To do so, we systematically
76 compare three approaches.

77 For the first approach, we applied a standard dimensionality reduction method directly to the recorded fluores-
78 cence traces from calcium imaging (Fig. 2a, top). Here we chose to use a latent Linear Dynamical System (LDS),
79 which is among the most basic methods for extracting neural trajectories. Conceptually, an LDS seeks to explain the
80 temporal structure in the data using latent variables that vary smoothly over time. We will examine the necessity of
81 using a latent dynamical model, like LDS, in a later section.

82 Each time a neuron spikes, intracellular free calcium increases, then decays slowly over time. The calcium
83 indicator kinetics influence the measured decay time, resulting in fluorescence traces whose intensity decays over
84 hundreds of milliseconds to seconds, depending on the particular calcium indicator used¹³. This decay transient
85 induces temporal correlations in the fluorescence measurements which are input to the LDS, which might attempt
86 to capture these correlations in its latent variable estimates. This motivates our second approach, which first
87 deconvolves each fluorescence trace separately, and then applies a LDS to the resulting estimated spiking activity.
88 The deconvolution serves to remove a substantial portion of the calcium decay transient, producing activity traces
89 similar to spike trains (or time varying firing rates). We term this two-stage method *deconv-LDS* (Fig. 2a, middle).
90 Here we deconvolved the fluorescence traces with OASIS^{32,34}, a deconvolution method that has been widely used in
91 calcium imaging studies^{20,23,24,35}.

92 In deconv-LDS, each neuron is deconvolved independently, and the stages of deconvolution and dimensionality
93 reduction are performed sequentially. We asked if performing these stages jointly would lead to more accurate
94 recovery of the latent variables (Fig. 2a, bottom). More specifically, we hypothesized that a sequential method
95 (e.g., deconv-LDS) may inadvertently discard some of the shared activity amongst neurons due to the independent
96 deconvolution of each neuron (Fig. 2b, left). Since the latent variables are intended to capture the shared spiking
97 activity among neurons (and not the calcium decay, which is independent across neurons), it might be possible to
98 better separate the calcium decay from the latent dynamics by considering all the neurons together, and performing
99 the two stages of deconvolution and dimensionality reduction jointly. This allows the dimensionality reduction
100 component to influence the deconvolution estimates, and vice versa (Fig. 2b, right). Thus, for our third method we

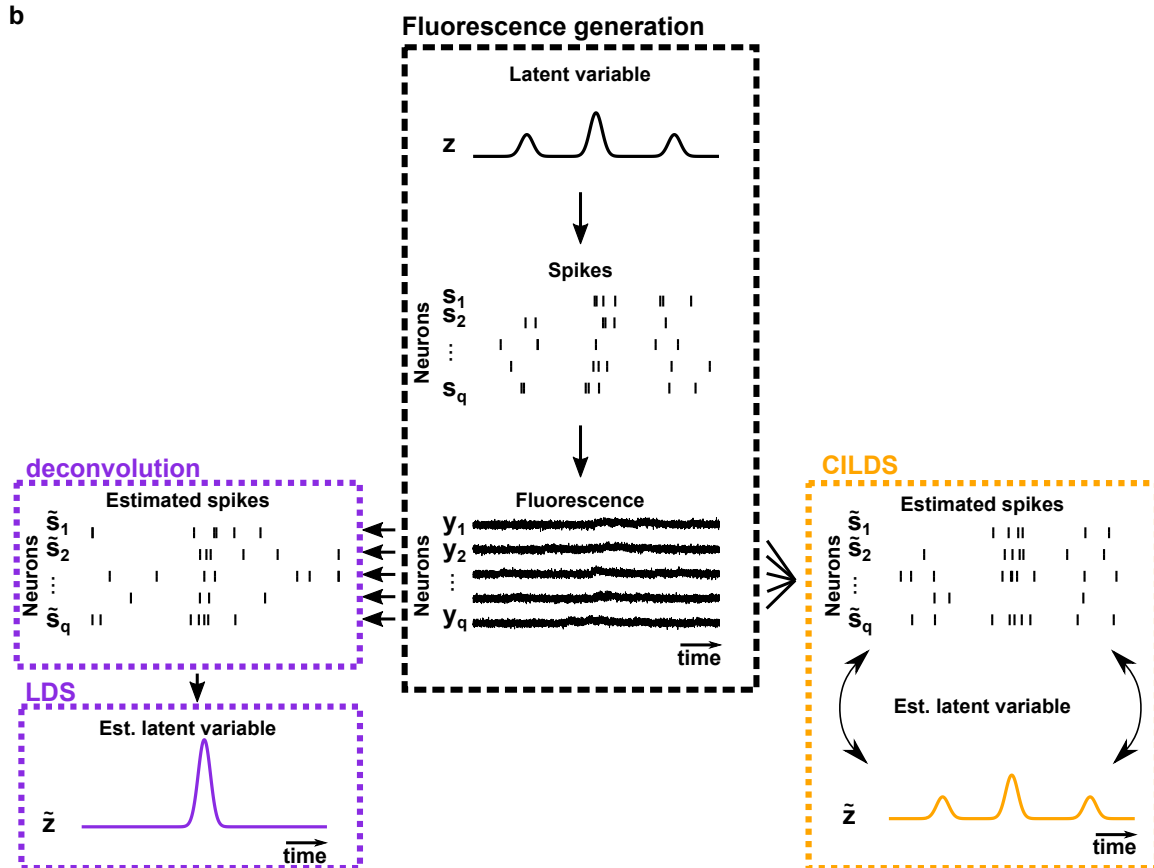
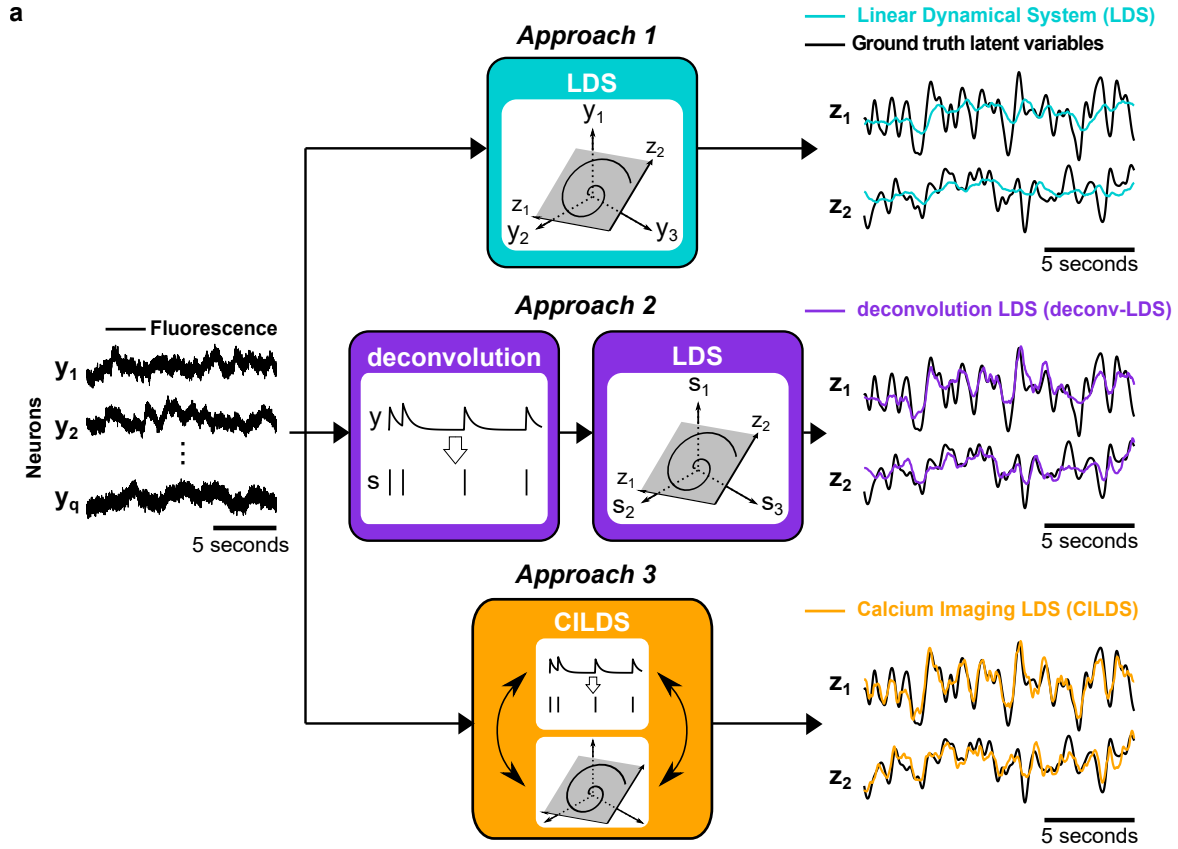


Figure 2. Comparison of three classes of dimensionality reduction methods. (a) Each of the three classes of methods was applied to the simultaneously-recorded fluorescence of a population of neurons (y_1, y_2, \dots, y_q) to extract latent variables. Top, Approach 1: a standard dimensionality reduction method (e.g., LDS) applied directly to calcium imaging recordings, extracting corresponding low-dimensional latent variables at each time point (illustrated here with two dimensions, z_1 and z_2). Middle, Approach 2: deconvolution is applied separately to each neuron's fluorescence trace to estimate its underlying spiking activity (s_1, s_2, \dots, s_q). A standard dimensionality reduction method (e.g., LDS) is then applied to the estimated spiking activity to extract latent variables (z_1 and z_2). Bottom, Approach 3: A unified method (e.g., CILDS) that takes calcium imaging recordings as input and performs deconvolution and dimensionality reduction simultaneously to extract the latent variables (z_1 and z_2). (b) Cartoon depicting the intuition behind the difference between Approaches 2 and 3. Center column: a latent variable z (representing, for example, common input) is used to generate spike trains which, in turn, are used to generate fluorescence traces. Left column: if deconvolution is performed neuron by neuron (Approach 2, deconv-LDS), it is unable to leverage the shared activity fluctuations among neurons to dissociate the calcium transients from the underlying shared spiking activity (i.e., the estimated latent variable). Right column: A unified method (Approach 3, CILDS) is applied to all neurons together and is therefore better able to dissociate the calcium transients from the underlying shared spiking activity among neurons (i.e., the estimated latent variable). As a result, CILDS more accurately estimates the ground truth latent variable than deconv-LDS. Note that the estimated spiking activity is depicted here as spike trains for visual clarity, even though they are in fact continuous-valued time courses.

101 developed a unified approach, *CILDS*, in which dimensionality reduction and deconvolution are performed jointly
102 (Fig. 2a, bottom).

103 **Should deconvolution be used with dimensionality reduction, and if so, how?**

104 One might postulate that first deconvolving fluorescence to estimate spiking activity would be beneficial for
105 recovering the same underlying latent variables that one would have recovered from the spikes themselves²⁷.
106 Indeed, multiple studies have applied dimensionality reduction to deconvolved spiking estimates^{19–24,35}. However,
107 deconvolution is subject to particular statistical modeling assumptions (as is any statistical method) and usually does
108 not recover the underlying spikes exactly. Therefore it is unclear how, or even if, deconvolution should be used with
109 dimensionality reduction. We addressed this question by comparing the three approaches described above (Fig. 2a).

110 With calcium imaging recordings, we typically do not have simultaneous electrophysiological recordings from
111 the same neurons, and thus the "ground truth" latent variables are unknown. To directly compare each method's
112 ability to extract latent variables, we designed a simulation framework in which we created known ground truth
113 latent variables with smoothly-varying time courses. These latent variables were used to generate spike trains which,

114 in turn, were used to generate fluorescence traces (see Methods and Supplementary Fig. 1). We then applied each of
115 the three approaches to these simulated fluorescence traces to assess how accurately they recovered the ground truth
116 latent variables. Examples of two combinations of experimental variables are illustrated in Fig. 3a.

117 We found that in both settings, CILDS outperformed the other two methods, returning more accurate estimates
118 of the ground truth latent variables (Fig. 3b-c, points above the diagonal). By operating on the entire population of
119 neurons together, CILDS was better able to separate the calcium transients from the shared activity among neurons
120 (i.e., the latent variables) compared to deconv-LDS, which deconvolves the activity of each neuron individually, and
121 LDS, which makes no attempt at this separation. Taken together, when extracting single-trial neural trajectories, one
122 should use deconvolution (Fig. 3b) jointly with dimensionality reduction, as in CILDS (Fig. 3c).

123 **How do different experimental variables impact the ability to recover latent activity?**

124 We next asked how different experimental variables affect the accuracy of the recovered latent variables. We designed
125 the simulation framework to enable us to systematically vary experimentally relevant variables. These variables
126 comprise four axes along which we can explore different combinations of values that might mimic a particular
127 experimental paradigm. We varied the timescale of the latent time courses from 50ms to 5000ms, the number of
128 neurons from 20 to around 100, the calcium decay timescales to match GCaMP6f, GCaMP6m, and GCaMP6s (fast,
129 medium, slow)¹³, and the variance of the added imaging noise that was independent of the calcium and spiking
130 activity (Supplementary Table 1). We then assessed how the accuracy of the dimensionality reduction methods
131 changed as we systematically changed these variables.

132 As we increased the timescale of the latent variables, all dimensionality reduction methods improved their
133 accuracy in estimating the ground truth latent variables (Fig. 3d). This occurs because with slower latent fluctuations,
134 the latent variables become less independent across time. As a result, all methods can leverage future and past time
135 points to better estimate the latent variables at the current time point. When the latent timescale is slow (order of
136 seconds), the calcium indicator decay is shorter relative to the latent timescale, and influences the neural activity
137 less. Therefore a method that is not able to disambiguate between the time course of the latent variables and the
138 calcium decay (e.g., LDS) can still accurately recover the latent variables (see latent timescale of 5000 ms in Fig.
139 3d). However, at faster latent timescales (tens to hundreds of milliseconds) which reflect the timescales of many
140 sensory, cognitive, and motor functions^{35,36}, it is critical to use a method that accounts for the calcium decay, as
141 both CILDS and deconv-LDS do (see latent timescale of 50 ms in Fig. 3d).

142 Next, when we increased the number of "recorded" neurons, all three methods improved in their ability to
143 reconstruct the ground truth latent variables (Fig. 3e). This makes sense because each neuron provides a different,
144 noisy view of the underlying latent variables. With more neurons, all methods are better able to "triangulate" the
145 values of the latent variables. Across the entire range of neurons tested, CILDS outperformed deconv-LDS, which

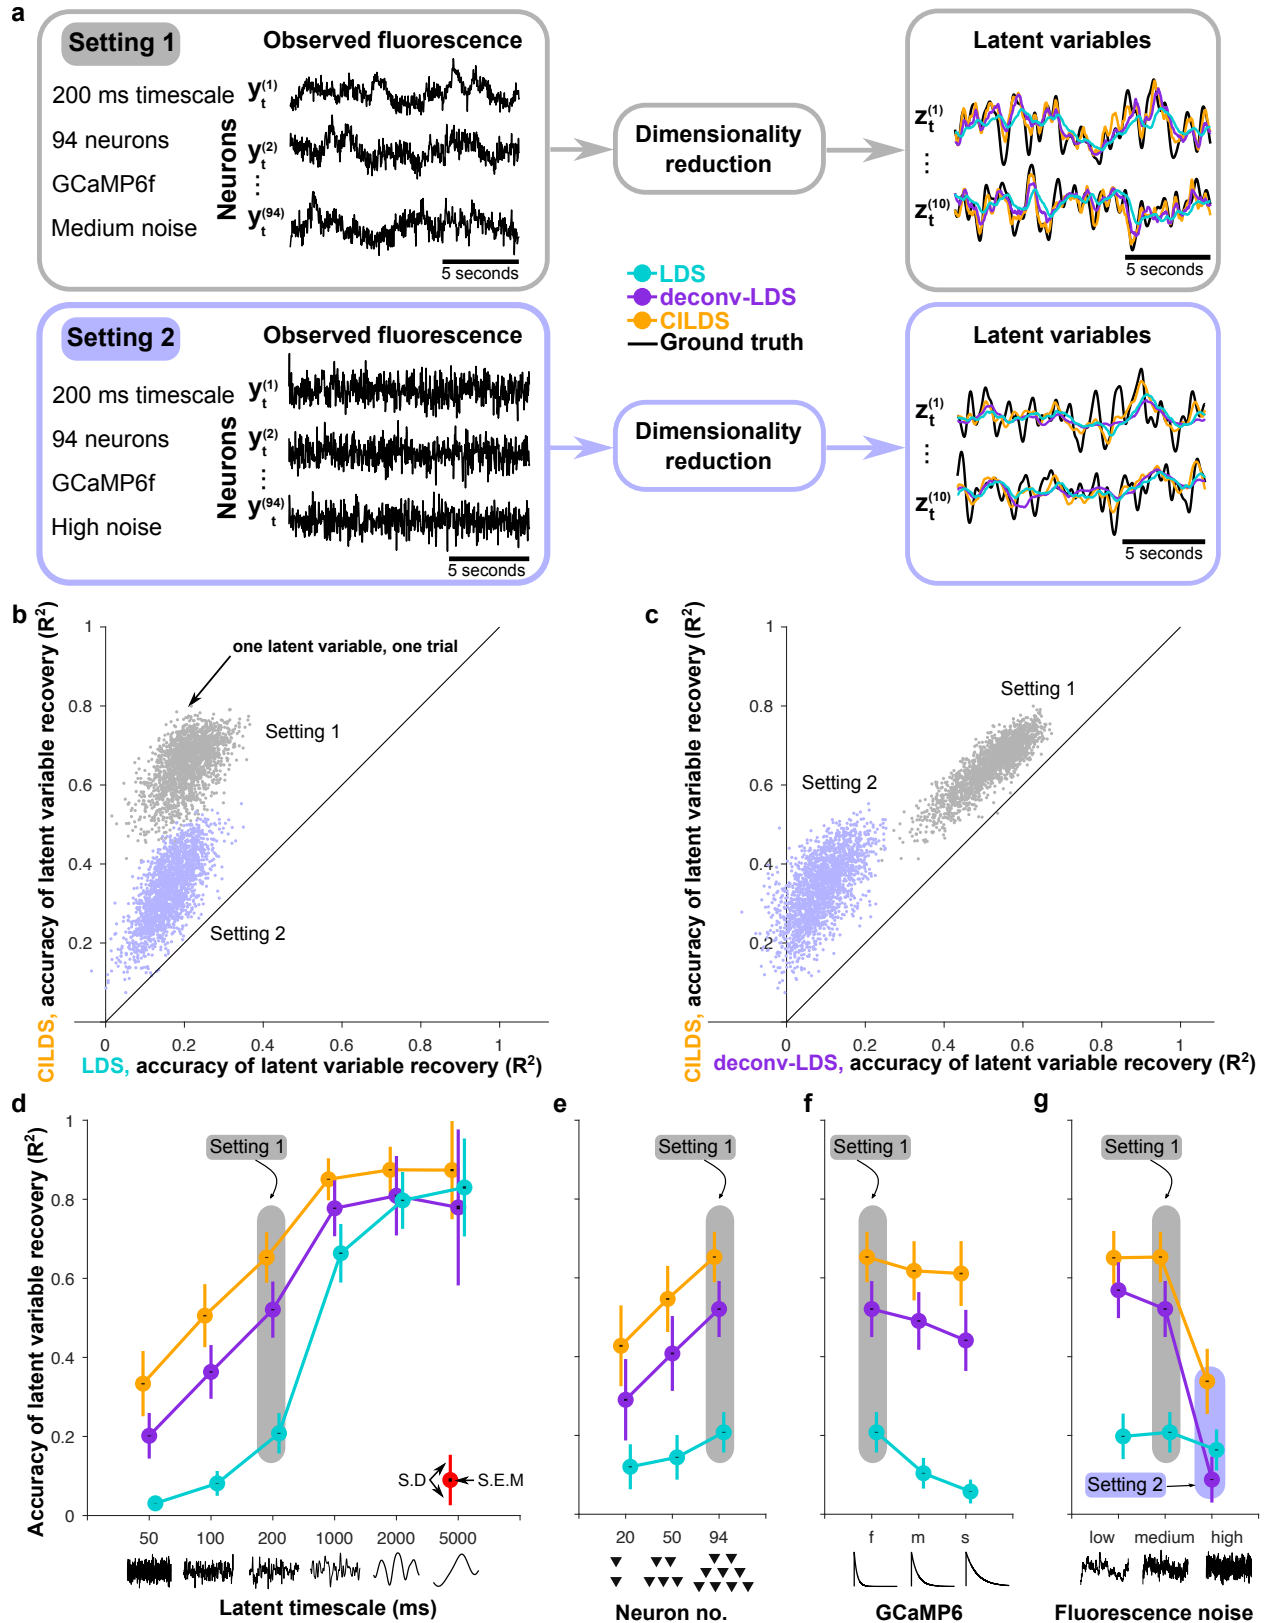


Figure 3. CILDS more accurately recovers the ground truth latent variables than the other two approaches

in simulation. (a) Example simulated fluorescence traces (left panels) and estimated latent variables (right panels) for two combinations of experimental variables. Setting 1 corresponds to a latent timescale of 200ms, 94 neurons, calcium decay corresponding to GCaMP6f, and medium fluorescence noise (see Methods). Setting 2 is the same as Setting 1, but with high fluorescence noise. Each of the three dimensionality reduction approaches introduced in Fig. 2 (LDS, cyan; deconv-LDS, purple; CILDS, orange) is applied to the simulated fluorescence traces. The latent variables extracted by each method can then be compared to the ground truth latent variables (black). (b-c) Accuracy of latent variables estimated by CILDS versus that of (b) LDS and (c) deconv-LDS. Accuracy is measured by the R^2 between the estimated and ground truth latent variables. Each point represents one latent variable on one trial. (d-g) Accuracy of latent variable recovery, as the (d) latent timescale, (e) number of neurons, (f) GCaMP6 indicator decay time constant, and (g) fluorescence noise level was varied. In each panel (d-g), one of the experimental variables was varied, while the other three variables were held constant at the Setting 1 values. The common point across the four panels is Setting 1 (shaded gray). Setting 2 (shaded purple) only appears in panel (g) because panels (d)-(f) correspond to medium rather than high fluorescence noise. The R^2 for other combinations of experimental variables are shown in Supplemental Fig. 2. Colored error bars indicate standard deviation, and black error bars indicate standard error across $n=2000$ latent variables (see Methods). The points are horizontally offset for visual clarity.

146 outperformed LDS (Fig. 3e).

147 We also varied the time constant of the GCaMP calcium indicator decay to match GCaMP6f, 6m, and 6s (from
148 fast to slow) (Fig. 3f). All methods performed worse as the decay time constant increased. This occurred because
149 the slower the calcium indicator is, the less the resulting fluorescence signal resembles the original spike train, which
150 increases the difficulty in disambiguating between the latent time course and the calcium decay (Fig. 3f).

151 Finally, we varied the amount of noise added to the fluorescence, which reflects imaging noise independent of
152 calcium and spiking activity (Fig. 3g). As the variance of the noise increased, all three methods performed worse as
153 expected, although the extent of the performance degradation differed across methods. As the fluorescence noise
154 variance increased, CILDS continued to outperform the other two methods (Fig. 3g). This indicates that leveraging
155 the population of neurons for simultaneous deconvolution and dimensionality reduction, as done by CILDS, provides
156 statistical power to mitigate a loss in accuracy due to increased fluorescence noise (Fig. 2b).

157 Overall, CILDS performed as well or better than the other two methods in every simulated setting we tested (Fig.
158 3d-g, orange higher than purple and cyan, also see Supplementary Fig. 2 for additional combinations of simulation
159 parameter settings). We additionally found that deconv-LDS usually outperformed LDS in accuracy of recovered
160 latent variables, consistent with Wei et al., which applied PCA to trial-averaged activity²⁷ (Fig. 3d-g, purple higher

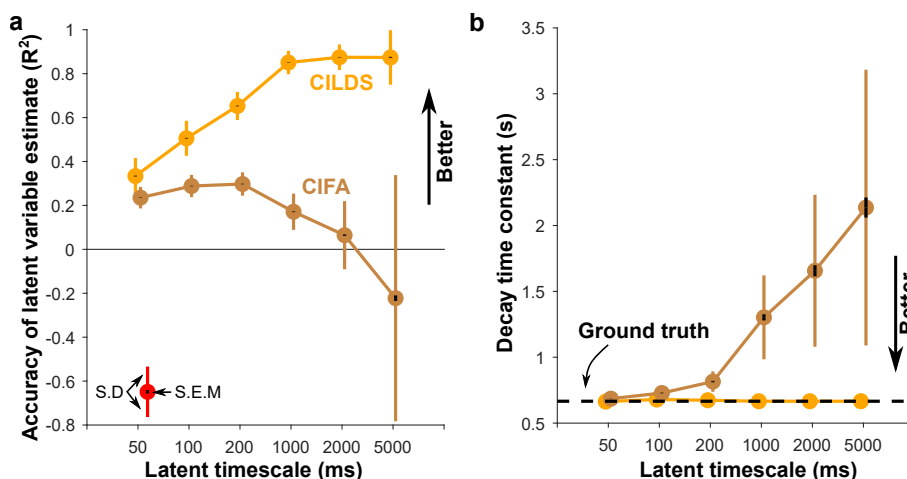


Figure 4. A latent dynamical model is necessary for accurately recovering neural trajectories. Comparison of a method that has no latent dynamics (Calcium Imaging Factor Analysis, CIFA) to CILDS using simulated fluorescence traces. The simulation parameters are GCaMP6f with 94 neurons and medium noise, as in Fig. 3d. **(a)** As the latent timescale increases, the ability of CILDS (orange) to accurately recover the neural trajectories increases, whereas that for CIFA (brown) decreases. Note that the R^2 can be < 0 because these results are cross-validated. The CILDS curve shown here is the same as in Fig. 3d. **(b)** Calcium decay time constant estimated using CIFA (brown) and CILDS (orange) for different simulated latent timescales. Dashed black line indicates ground truth decay time constant. CILDS accurately estimates the decay time constant across all latent timescales tested, whereas CIFA overestimates the decay time constant as the latent timescale increases. Coloured error bars indicate standard deviation, and black error bars indicate standard error across $n = 2000$ latent variables (see Methods).

161 than cyan). Using deconvolution is particularly important in regimes where the time scales of neural activity (i.e.,
162 the latent timescales) are faster than that of the calcium decay, which is the case for many commonly-studied brain
163 functions.

164 **Is it necessary to include a latent dynamical model in the dimensionality reduction method?**

165 All three dimensionality reduction methods considered so far explicitly attempt to extract latent variables that
166 evolve smoothly over time via a latent dynamical model. We asked if the latent dynamical model was necessary for
167 separating the calcium decay from the neural trajectories. To address this, we developed a method called Calcium
168 Imaging Factor Analysis (CIFA, see Methods). CIFA is identical to CILDS, except that CIFA's latent variables are
169 independent across time (i.e., there is no dynamical model) by analogy to conventional factor analysis.

170 We then compared the performance of CIFA to CILDS. We found that CILDS, which uses a latent dynamical
171 model, recovered the latent variables more accurately than CIFA, which does not use a latent dynamical model (Fig.

172 4a; additional cases are shown in Supplementary Fig. 3). As the latent timescale increased, CILDS more accurately
173 recovered the latent variables (Fig. 4a, also shown in Fig. 3d). Correspondingly, CILDS correctly ascertained that
174 the increased smoothness in the fluorescence was due to a longer latent timescale and not a longer calcium decay
175 time constant (Fig. 4b, Supplementary Fig. 3). In contrast, as the latent timescale increased, CIFA recovered the
176 latent variables less accurately (Fig. 4a). This occurs because CIFA erroneously attributed the increased smoothness
177 in the fluorescence to a slower calcium decay (Fig. 4b, Supplementary Fig. 3), rather than a longer latent timescale.
178 Unlike CILDS which can attribute temporal smoothness in the fluorescence to two possible sources (latent variables
179 that evolve smoothly over time and calcium decay), CIFA can attribute temporal smoothness to only one possible
180 source (calcium decay). This indicates that using a latent dynamical model is necessary for peering through the
181 calcium decay transients to accurately recover the neural trajectories.

182 **CILDS outperforms other methods on calcium imaging recordings**

183 To assess whether the advantages of CILDS also hold in real data, we applied each of the dimensionality reduction
184 methods described above (CILDS, deconv-LDS, LDS, and CIFA) to calcium imaging recordings in two experimental
185 contexts: larval zebrafish and mice. To emphasize the generality of our findings, these two experimental settings
186 involve not only different animal species, but also different brain areas, behavioral tasks, and properties of the
187 recorded fluorescence (see below). In these experiments, the spiking activity of the neurons was not recorded,
188 and thus the ground truth latent variables are unknown. To quantify the accuracy of each method, we adopted a
189 leave-neuron-out procedure used in previous studies^{37,38}, where we estimate the latent variables using all-but-one
190 neuron, and then assess how well these latent variables predict the recorded fluorescence of the held-out neuron (see
191 Methods). A more accurate prediction of the fluorescence of the held-out neuron indicates that the latent variables
192 are a better summary of the population activity.

193 The first experimental context involves larval zebrafish engaged in a "fictive swimming" motosensory gain
194 adaptation task (Fig. 5a)³⁹. Calcium imaging was performed on neurons expressing GCaMP6f in dorsal raphe
195 nucleus (DRN) using light-sheet microscopy in three fish (19 to 22 neurons; mean: 20) (Fig. 5a). We applied each
196 dimensionality reduction method to these recordings and assessed their performance using the leave-neuron-out
197 prediction procedure (Fig. 5b). We found that CILDS more accurately predicted the fluorescence of the held-out
198 neurons than the other methods, as quantified by the correlation between the predicted and recorded fluorescence.
199 CILDS outperformed LDS (Fig. 5c; 63% of neurons above the diagonal), deconv-LDS (Fig. 5d; 85% of neurons
200 above the diagonal), and CIFA (Fig. 5e; 73% points above the diagonal). These results were also true for each fish
201 individually (Supplementary Fig. 5).

202 The second experimental context involves awake, head-fixed mice passively viewing static visual gratings
203 (Fig. 6a)⁴⁰. Two-photon calcium imaging was performed on neurons expressing GCaMP6f in the primary visual

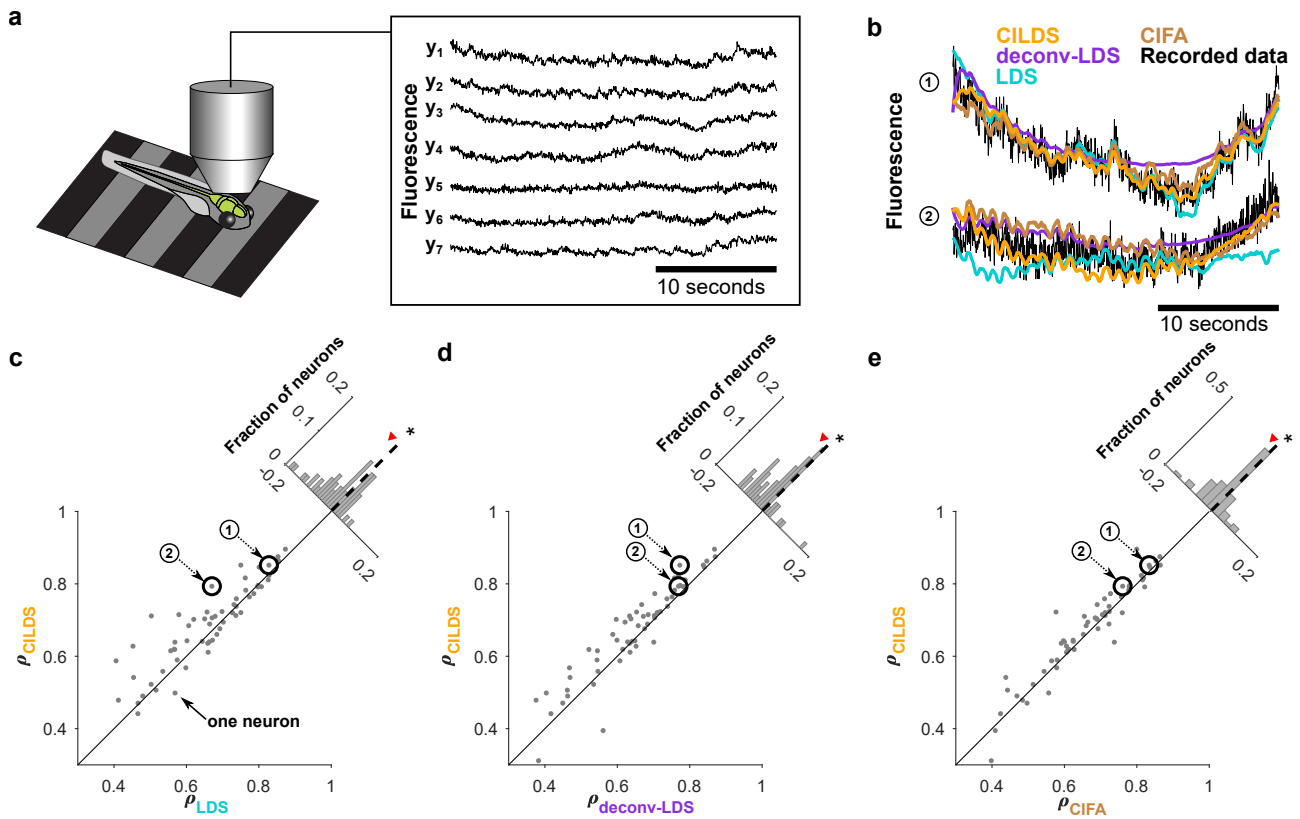


Figure 5. In larval zebrafish DRN recordings, CILDS captures the latent time courses better than the other three methods. (a) Two-photon calcium imaging using GCaMP6f at 30Hz was performed on three larval zebrafish in a virtual reality environment. Shown are representative fluorescence traces from seven of the imaged neurons. (b) Example recorded fluorescence traces (black) and leave-neuron-out predicted fluorescence using CILDS (orange), deconv-LDS (purple), LDS (cyan), and CIFA (brown). (c-e) Correlation between the recorded fluorescence and the leave-neuron-out predicted fluorescence for CILDS versus each of the other methods. The correlation is higher for CILDS than (c) LDS ($p < 1 \times 10^{-7}$, $n = 60$ neurons, paired two-tailed t-test), (d) deconv-LDS ($p < 1 \times 10^{-4}$, $n = 60$ neurons), and (e) CIFA ($p < 0.05$, $n = 60$ neurons). Each point represents one neuron, where the correlation is computed for each trial (27 seconds long) then averaged across all 15 trials. The numbered points (black circles) correspond to the examples shown in panel (b). Diagonal histogram shows the paired difference in correlation between CILDS and one of the other methods, as indicated. A paired two-tailed t-test is applied to assess statistical significance, with statistical significance indicated by an asterisk. Note that the histogram is zoomed-in for visual clarity, and therefore the ends of the histograms are not shown.

204 cortex (V1) of three mice (133 to 319 neurons; mean: 234.7). Comparing the raw recorded fluorescence of the two
 205 experimental contexts (Fig. 5a versus Fig. 6a), the neurons in the DRN of the larval zebrafish tend to exhibit slower
 206 fluctuations in fluorescence that are correlated across neurons (Fig. 5a). In contrast, the neurons in mouse V1 exhibit

207 faster changes in fluorescence, which appear less correlated across neurons (Fig. 6a). We applied the same analyses
 208 in Fig. 5 to these mouse recordings (Fig. 6b). Despite the stark differences between the fish DRN and mouse V1
 209 fluorescence traces, we again found that CILDS more accurately predicted the fluorescence of held-out neurons than
 210 LDS (Fig. 6c; 63% of neurons above the diagonal), deconv-LDS (Fig. 6d; 68% of neurons above the diagonal),
 211 and CIFA (Fig. 6e; 82% of neurons above the diagonal). These results were also true for each mouse individually
 212 (Supplementary Fig. 6).

213 Thus far, we have shown that CILDS extracts latent variables that more accurately predict the fluorescence of
 214 held-out neurons than the other methods. The leave-neuron-out evaluation assesses each method's ability to capture

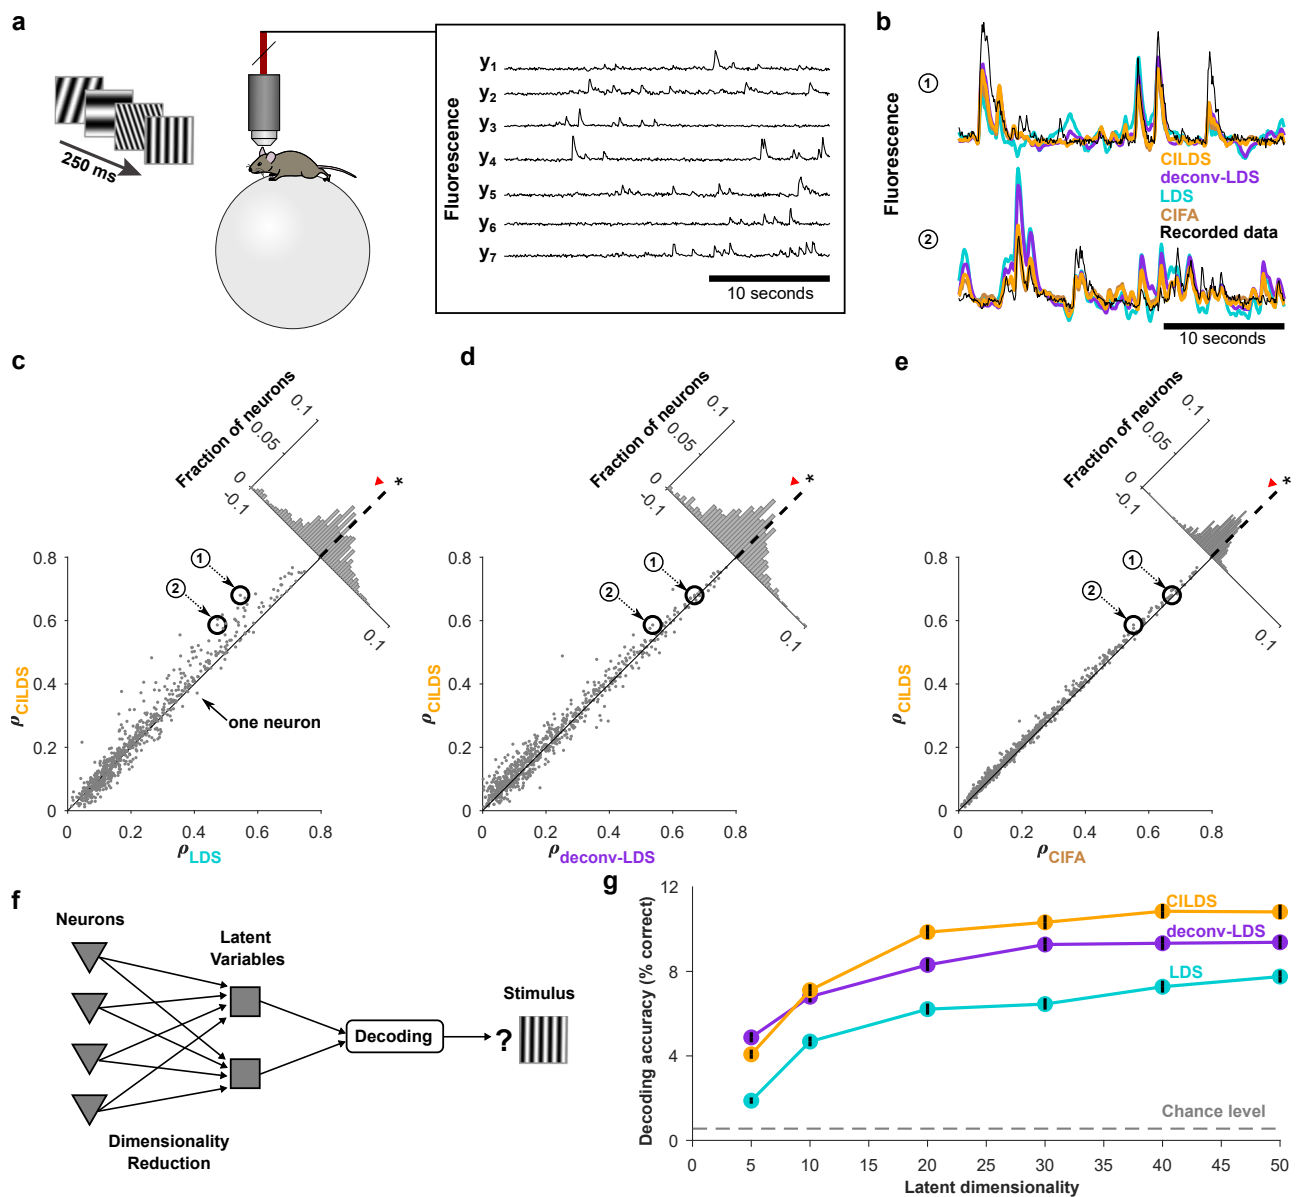


Figure 6. In mice V1 recordings, CILDS captures the latent time courses better than the other three

methods. (a) Two-photon calcium imaging performed on awake mice viewing static gratings with different spatial frequencies and orientations (180 total stimuli) using GCaMP6f at 15.5Hz. Shown are representative fluorescence traces from seven of the imaged neurons. (b) Example segment of recorded fluorescence traces (black) and leave-neuron-out predicted fluorescence traces using CILDS (orange), deconv-LDS (purple), LDS (cyan), and CIFA (brown). (c-e) Correlation between recorded fluorescence and leave-neuron-out predicted fluorescence for CILDS versus each of the other methods. The correlation is higher for CILDS than (c) LDS ($p < 1 \times 10^{-16}$, $n = 704$ neurons, paired two-tailed t-test), (d) deconv-LDS ($p < 1 \times 10^{-26}$, $n = 704$ neurons), and (e) CIFA ($p < 1 \times 10^{-80}$, $n = 704$ neurons). Each point represents one neuron, where the correlation is computed for each trial (196.7 seconds long) then averaged across all 15 trials. The numbered points (black circles) correspond to the examples shown in panel (b). Diagonal histogram shows the paired difference in correlation between CILDS and one of the other methods, as indicated. Asterisk denotes $p < 0.05$ for the paired two-tailed t-test above. Note that the histogram is zoomed-in for visual clarity, and therefore the ends of the histograms are not shown. (f) Flow diagram depicting decoding of visual stimuli using low-dimensional latent variables, which are obtained by applying a dimensionality reduction method to the recorded fluorescence traces. (g) Classification accuracy of the visual stimulus based on latent variables extracted using CILDS (orange), deconv-LDS (purple), and LDS (cyan). Classification was performed using a Gaussian Naive Bayes decoder, where the number of latent variables extracted by each dimensionality reduction method was systematically varied (horizontal axis). There were 180 total gratings (with different orientations and spatial frequencies) shown during the experiment, so the chance classification accuracy is 1/180 (gray dashed). The decoding window was 250 ms, which is the duration of stimulus presentation. Black error bars indicate 95% confidence intervals (Bernoulli process).

215 the shared activity changes among neurons. Another way to assess how meaningful the latent variables extracted
216 by each method are is to measure how strongly they reflect external variables, such as the sensory stimulus⁴¹. We
217 thus performed a decoding analysis, whereby we classified the orientation and spatial frequency of the presented
218 grating using the latent variables extracted by each of three methods (Fig. 6f). The stimuli could be more accurately
219 decoded from the latent variables extracted by CILDS than from the latent variables extracted by either deconv-LDS
220 or LDS (Fig. 6g). At the highest decoding accuracy of each method, CILDS outperformed deconv-LDS by 1.16
221 times and LDS by 1.43 times. This demonstrates that CILDS is better at capturing the shared activity changes
222 among neurons that are relevant to the visual stimulus than the other two methods. Furthermore, the latent variables
223 extracted by deconv-LDS provided more accurate decoding of the stimuli than the latent variables extracted by LDS,
224 which indicates that accounting for the calcium transients is important when the variable of interest (here, the visual

225 stimulus, which changes every 250 ms) changes on the timescale of tens to hundreds of milliseconds.

226 Taken together, the results based on calcium imaging recordings from two different recording regimes are
227 consistent with what we identified in simulation. Namely, deconvolution should be used jointly with a dimensionality
228 reduction method that incorporates a latent dynamical model (as in CILDS), particularly if the neural process of
229 interest changes on a timescale of tens to hundreds of milliseconds.

230 Discussion

231 In this work, we investigated the use of dimensionality reduction to summarize the time course of calcium imaging
232 recordings. We considered dimensionality reduction approaches with and without deconvolution, and developed a
233 novel method (CILDS) which jointly performs deconvolution and dimensionality reduction. We compared these
234 methods using simulated fluorescence time courses, where we systematically varied experimental variables, as well
235 as using calcium imaging recordings from larval zebrafish and mice. We found that: 1) CILDS, which leverages
236 the population of neurons to jointly perform deconvolution and dimensionality reduction, outperformed the other
237 methods, 2) using deconvolution was increasingly important with faster latent timescales, and 3) the use of a latent
238 dynamical model, as in CILDS, was important for peering through the calcium decay transients to accurately
239 recover the neural trajectories. Overall, this work provides a foundation for using dimensionality reduction to extract
240 single-trial neural trajectories from calcium imaging recordings.

241 This work focused on the question of which dimensionality reduction method is most appropriate for extracting
242 single-trial neural trajectories from calcium imaging recordings. There are two other settings in which dimensionality
243 reduction is commonly used. First, dimensionality reduction can be applied to analyze how trial-averaged activity
244 differs across experimental conditions (e.g., ref.^{4,5,7,17}). In a recent study, Wei et al., applied principal components
245 analysis (PCA) to trial-averaged electrophysiological recordings and calcium imaging with GCaMP6s²⁷. They found
246 important differences in the low-dimensional PCA trajectories obtained from electrophysiological recordings versus
247 calcium imaging. This difference was mitigated by first deconvolving the calcium imaging recordings before applying
248 dimensionality reduction, consistent with our findings. It may be possible to further improve the correspondence
249 by applying CILDS to single-trial fluorescence recordings, then averaging the extracted low-dimensional neural
250 trajectories across trials.

251 Second, dimensionality reduction is often used to analyze the trial-to-trial variability of neural population activity
252 without time courses, i.e., using one time point or time window per trial (e.g., ref.^{9,42-45}). In this case, there would be
253 no information about how the calcium decays and so one would not be able to make use of a method that incorporates
254 deconvolution. One might consider using CIFA, by analogy to the use of factor analysis to study the trial-to-trial
255 variability of spike counts without time courses. However, it is important to note that, like CILDS, CIFA also
256 requires multiple time points to be able to deconvolve, even though the latent variables in CIFA are independent

257 from one time point to the next. If the original time series of calcium imaging is available, one can apply CILDS to
258 the time series first, then average across the time points of the extracted latent variables. If the original time series of
259 calcium imaging is not available, then a standard dimensionality reduction method such as factor analysis might be
260 more suitable.

261 Previous studies have proposed methods for analyzing calcium imaging recordings that include latent variables
262 and deconvolution in the same statistical model. Triplett et al.⁴⁶ developed a method to study the interaction between
263 evoked and spontaneous activity using calcium imaging recordings in sensory systems. In their model, the latent
264 variables represent activity fluctuations shared amongst neurons that are not explained by the sensory stimulus,
265 where these activity fluctuations are defined to be spontaneous activity. Aitchison et al.⁴⁷ developed a method to infer
266 spiking activity and neural connectivity from calcium imaging experiments that involve optogenetic stimulation. In
267 their model, the latent variables represent shared activity amongst neurons that are not explained by the optogenetic
268 stimulation or the activity of other neurons recorded simultaneously, and are intended to represent input from other
269 brain areas. We developed CILDS for extracting latent variable time courses to summarize the population activity
270 on individual experimental trials. In contrast to the two methods above which have more specific analysis goals
271 (to separate evoked from spontaneous activity in the case of Triplett et al., or to infer spiking activity and neural
272 connectivity in the case of Aitchison et al.), CILDS is general-purpose and well-suited for exploratory data analysis.
273 This is akin to the use of methods such as LDS⁴⁸, GPFA³⁷, TCA⁴⁹, LFADS⁴¹, and dPCA⁵⁰ for exploratory analysis
274 of population spiking activity to extract latent variable time courses, whose results can then lead to the use of
275 population analysis methods with more specific goals (see ref.⁵¹⁻⁵³ for examples).

276 CILDS can be extended in the following ways. First, the order of the autoregressive process for the calcium
277 dynamics determines how quickly the calcium concentration (and consequently the fluorescence) rises after each
278 spiking event³⁴. Here, we used an autoregressive order of one, which corresponds to an instantaneous rise in calcium
279 after each spiking event, as was done in previous work^{32,34,54}. Although this is a reasonable first approximation,
280 when imaging rates are fast or the calcium indicator is slow, it may be desirable to use an autoregressive order greater
281 than one to better capture the non-instantaneous rise in calcium¹³. Second, different latent time series models can be
282 used in the place of the linear dynamical system to achieve different analysis goals. For example, if one seeks only
283 temporal smoothing in the latent variables without explicit dynamics, one can replace the linear dynamical system
284 with Gaussian processes³⁷.

285 Dimensionality reduction on population spiking activity has led to many insights about brain function. With
286 the development of dimensionality methods that are tailored for calcium imaging, such as CILDS, we can leverage
287 the advantages of calcium imaging such as being able to obtain information about neuron type or knowledge about
288 which neurons project to other brain areas. For example, one could use dimensionality reduction to understand
289 how populations of different neuron types interact⁵⁵, or one could utilize information about where neurons project,

290 coupled with dimensionality reduction, to understand how the projections contribute to the coordination of activity
291 between brain areas^{56,57}. This can enable novel insights about neural population activity recorded using calcium
292 imaging that go beyond what is currently possible with electrophysiology.

293 **Methods**

294 **Dimensionality reduction methods**

295 Here we describe mathematically the dimensionality reduction methods used in this work: LDS, deconv-LDS,
296 CILDS, and CIFA. For the purposes of this work, we assume that the spatial footprint of each neuron has already
297 been identified from the raw calcium imaging data (a procedure known as image segmentation^{54,58}), resulting in a
298 fluorescence time course for each neuron. The dimensionality reduction methods presented here are applied to these
299 fluorescence time courses.

300 **Linear Dynamical System (LDS)**

301 We first considered a standard dimensionality reduction method for summarizing the time course of spiking activity,
302 a latent Linear Dynamical System (LDS), here applied to calcium imaging recordings. Let $\mathbf{y}_t \in \mathbb{R}^{q \times 1}$ be a high-
303 dimensional vector of fluorescence values recorded at time point t , where q is the number of neurons imaged
304 simultaneously. The goal is to extract a corresponding low-dimensional latent variable $\mathbf{z}_t \in \mathbb{R}^{p \times 1}$ at each time point,
305 where p is the number of latent dimensions ($p < q$). The observation model is

$$\mathbf{y}_t = A\mathbf{z}_t + \mathbf{b} + \boldsymbol{\varepsilon}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, R) \quad (1)$$

306 where $A \in \mathbb{R}^{q \times p}$ is the loading matrix that specifies how each neuron's activity is related to the latent variables,
307 $\mathbf{b} \in \mathbb{R}^{q \times 1}$ is an offset vector that accounts for constant background fluorescence, $R \in \mathbb{R}^{q \times q}$ is the observation noise
308 covariance, and $t = 1, \dots, T$. We constrained R to be diagonal, thereby capturing activity variability and imaging
309 noise independent to each neuron. The time-evolution of the latent variables is described as a linear dynamical
310 system

$$\mathbf{z}_t = D\mathbf{z}_{t-1} + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, P) \quad (2)$$

$$\mathbf{z}_1 \sim \mathcal{N}(\mathbf{h}_1, G_1) \quad (3)$$

311 where $D \in \mathbb{R}^{p \times p}$ is the dynamics matrix that determines the timescale of the latent variables, $P \in \mathbb{R}^{p \times p}$ is the
312 dynamics noise covariance, $\mathbf{h}_1 \in \mathbb{R}^{p \times 1}$ and $G_1 \in \mathbb{R}^{p \times p}$ describe the mean and covariance of the latent variable at
313 the first time point, and $t = 2, \dots, T$. We constrained D , P , and G_1 here to be diagonal as a form of regularization,
314 although a general LDS with these parameters unconstrained could also be used.

315 Equations (1), (2), and (3) together define the latent LDS⁵⁹. We fit the model parameters ($A, \mathbf{b}, R, D, P, \mathbf{h}_1, G_1$)
316 using the expectation-maximization (EM) algorithm⁶⁰. To initialize the model parameters, we first performed factor
317 analysis (FA) on \mathbf{y}_t to obtain A , \mathbf{b} , and R . We initialized D to be $0.999I$ (a stable system), which we found to work

318 well in practice in the simulations. We ran the EM algorithm until convergence (defined as a log data likelihood
319 increase of $< 10^{-6}$ or 1500 iterations, whichever came first). This maximum number of iterations was chosen
320 heuristically, by noting empirically that the latent variables do not change significantly beyond this point.

321 ***Deconvolution - Linear Dynamical System (deconv-LDS)***

322 Since fluorescence traces are an indirect measure of spiking activity, we also considered a two-stage approach,
323 whereby we first deconvolve the fluorescence traces one neuron at a time to estimate the underlying spiking activity.
324 Then we applied a standard dimensionality reduction method, in this case LDS, to those deconvolved estimates. We
325 refer to this two-stage method as deconv-LDS.

326 For the deconvolution stage of deconv-LDS, we used the Online Active Set method to Infer Spikes (OASIS),
327 developed by Friedrich et al.³⁴, using their L_1 regularization. We also tested their L_0 regularization and found the L_1
328 regularization to work better for our datasets. As per Friedrich et al., this first order autoregressive model for OASIS
329 is described for each neuron as

$$y_t = ac_t + b + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2) \quad (4)$$

$$c_t = \gamma c_{t-1} + s_t \quad (5)$$

330 where y_t is the recorded fluorescence at time t , c_t represents the calcium concentration at time t , ε_t captures
331 imaging noise independent of the calcium and spiking activity, and s_t is the spiking activity. The parameter a relates
332 the calcium concentration to fluorescence, b accounts for the baseline fluorescence, σ^2 captures the variance of
333 the imaging noise, and γ specifies how quickly the calcium trace decays, which depends on the calcium indicator.
334 Additionally, there is a hyperparameter in the OASIS model, minimum spike size, that sets the minimum value
335 of s_t that would be identified. In this model, all variables are scalars and a is constrained to be non-negative. We
336 initialized OASIS with γ values that are typical for the calcium indicators used¹³ (see Supplementary Table 1), and
337 allowed OASIS to optimize a , b , γ , σ^2 and the minimum spike size.

338 Applying deconvolution to the recorded fluorescence traces returned estimates of the time course of spiking
339 activity for each neuron, s_t . We then used the estimated spiking activity of all the neurons as the observations
340 $\mathbf{y}_t \in \mathbb{R}^{q \times 1}$ for time points $t = 1, \dots, T$ in the LDS model defined in equations (1)-(3).

341 ***Calcium Imaging Linear Dynamical System (CILDS)***

342 CILDS unifies the approaches described above by allowing estimates of shared activity among neurons (i.e., the
343 latent variables) to influence the estimates of deconvolved spiking activity, and vice versa. In other words CILDS
344 performs deconvolution for all neurons and dimensionality reduction jointly, in a unified framework. This is in
345 contrast to deconv-LDS, which deconvolves the activity of each neuron independently. With low-dimensional latent

346 variables that are jointly estimated with the model of calcium decay, CILDS is better able to peer through the calcium
 347 decay to more clearly identify the shared activity among neurons, as compared to deconv-LDS and LDS applied
 348 directly on fluorescence.

349 Let $\mathbf{y}_t \in \mathbb{R}^{q \times 1}$ be the high-dimensional vector of fluorescence traces recorded at each time point t , where q is the
 350 number of neurons imaged simultaneously. The goal is to extract a corresponding low-dimensional latent variable
 351 $\mathbf{z}_t \in \mathbb{R}^{p \times 1}$ at each time point, where p is the number of latent dimensions ($p < q$). The observation model follows
 352 the multivariate form of equation (4) which was used for deconvolution

$$\mathbf{y}_t = B\mathbf{c}_t + \boldsymbol{\varepsilon}_t, \quad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, R) \quad (6)$$

353 where $B \in \mathbb{R}^{q \times q}$ maps the calcium concentration to the recorded fluorescence, $R \in \mathbb{R}^{q \times q}$ is the fluorescence noise
 354 covariance, and $t = 1, \dots, T$. We constrained B and R to be diagonal to allow each dimension of \mathbf{c}_t to represent the
 355 calcium concentration of one neuron. B accounts for all experimental variables influencing the scale of the signal
 356 from each neuron, such as the amplification of the imaging system³², and R accounts for fluorescence fluctuations
 357 independent of calcium concentration. Here we omit the additive offset found in equation (4) without loss of
 358 generality due to the offset included in equation (7).

359 Similar to equation (5), the calcium decay for each neuron is described using a first order autoregressive model

$$\mathbf{c}_t = \Gamma\mathbf{c}_{t-1} + A\mathbf{z}_t + \mathbf{b} + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, Q) \quad (7)$$

$$\mathbf{c}_1 \sim \mathcal{N}(\boldsymbol{\mu}_1, V_1) \quad (8)$$

360 where $\Gamma \in \mathbb{R}^{q \times q}$ captures the calcium decay, $A \in \mathbb{R}^{q \times p}$ is the loading matrix that describes how latent variable maps
 361 to calcium concentrations, $\mathbf{b} \in \mathbb{R}^{q \times 1}$ is a constant vector, $Q \in \mathbb{R}^{q \times q}$ captures the spiking variability independent to
 362 each neuron, $\boldsymbol{\mu}_1 \in \mathbb{R}^{q \times 1}$ and $V_1 \in \mathbb{R}^{q \times q}$ describe the mean and variance of the calcium concentration at the first time
 363 point, respectively, and $t = 2, \dots, T$. Our key innovation is to replace the spiking activity s_t from equation (5) with
 364 $A\mathbf{z}_t + \mathbf{b} + \mathbf{w}_t$. By analogy to factor analysis, $A\mathbf{z}_t$ describes the shared activity among neurons, and \mathbf{w}_t describes the
 365 activity independent to each neuron. We constrain Γ , Q , and V_1 to be diagonal to prevent intermixing among neurons
 366 outside of the loading matrix. The form of Γ allows each neuron to have a different calcium decay as determined by
 367 the fitting procedure (see below), which can depend on the extent of calcium buffering within a cell and the calcium
 368 indicator used^{13,32,34}.

369 Similar to LDS, the low-dimensional latent variables \mathbf{z}_t evolve over time according to a linear dynamical system

$$\mathbf{z}_t = D\mathbf{z}_{t-1} + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, P) \quad (9)$$

$$\mathbf{z}_2 \sim \mathcal{N}(\mathbf{h}_2, G_2) \quad (10)$$

370 where $D \in \mathbb{R}^{p \times p}$ is the dynamics matrix, $P \in \mathbb{R}^{p \times p}$ is the dynamics noise covariance, $\mathbf{h}_2 \in \mathbb{R}^{p \times 1}$ and $G_2 \in \mathbb{R}^{p \times p}$ are
371 the mean and covariance of the latent variable at the first time point respectively, and $t = 3, \dots, T$. We constrained D ,
372 P , and G_2 here to be diagonal as a form of regularization, although a model with these parameters unconstrained
373 could also be used. Note that according to equation (7), \mathbf{z}_2 is the first latent variable in the time series (i.e., there is
374 no \mathbf{z}_1).

375 Equations (6) - (10) define CILDS. The joint estimation of the parameters $B, R, \Gamma, A, \mathbf{b}, Q, \boldsymbol{\mu}_1, V_1, D, P, \mathbf{h}_2$, and
376 G_2 allows CILDS to leverage the entire recorded neural population to perform deconvolution and estimate latent
377 variables in a unified fashion. CILDS can be viewed as a special case of the standard LDS, where the parameters are
378 constrained in a specific way. We fit CILDS using the EM algorithm, initialized using deconv-LDS run for 100 EM
379 iterations. The EM algorithm was run until convergence, defined as a log data likelihood increase of $< 10^{-6}$ or 1500
380 iterations, whichever came first. The maximum number of iterations was chosen by noting empirically that the latent
381 variables do not change substantially beyond this point. The EM equations for CILDS are provided in Supplemental
382 Information.

383 **Calcium Imaging Factor Analysis (CIFA)**

384 To evaluate the importance of incorporating a latent dynamical system in dimensionality reduction methods for
385 calcium imaging, we developed a method (CIFA) identical to CILDS, with the only difference being that CIFA does
386 not enforce latent dynamics. Like CILDS, CIFA also uses equations (6), (7), and (8). The only difference between
387 CIFA and CILDS is that we replaced the latent dynamical system equations (9) and (10) with

$$\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, I) \quad (11)$$

388 for $t = 2, \dots, T$. In other words, CIFA defines latent variables that are independent across time, whereas CILDS
389 defines latent variables that evolve smoothly over time. Like CILDS, we fit CIFA using the EM algorithm, initialized
390 using parameters from deconv-FA run for 100 EM iterations. The EM algorithm was run until convergence, defined
391 as a log data likelihood increase of $< 10^{-6}$ or 1500 iterations, whichever came first. The EM equations for CIFA
392 are the same as for CILDS (see Supplemental Information), without the equations estimating D, P, \mathbf{h}_2, G_2 . Note
393 that there is no loss of generality by setting the prior distribution of \mathbf{z}_t to $\mathcal{N}(\mathbf{0}, I)$, compared to a general Gaussian
394 distribution. Additionally, although CIFA has FA in its name, there is one key difference from FA. Whereas FA can
395 be fit on data with no concept of time, CIFA requires a time series for deconvolution (equations (7) and (8)).

396 **Simulation framework**

397 We created a framework to simulate fluorescence recordings from calcium imaging for two reasons. First, in calcium
 398 imaging recordings, the ground truth latent variables are unknown. A simulation of fluorescence traces from known
 399 latent variables allows us to directly evaluate our dimensionality reduction methods by comparing the estimated
 400 latent variables with the ground truth latent variables. Second, we wanted a simulation framework in which we could
 401 systematically vary various experimentally relevant parameters to see their effects on the estimated latent variables.
 402 Specifically, we evaluated our ability to recover the ground truth latent variables as a function of the timescale of
 403 the latent variables, the number of neurons, the calcium decay rate, and the size of the fluorescence noise. The
 404 simulation procedure consists of first generating fluorescence traces from known latent variables while varying the
 405 experimentally relevant parameters listed above. Then, we applied each dimensionality reduction method to estimate
 406 latent variables from the simulated fluorescence traces. The estimated latent variables were then compared to the
 407 ground truth latent variables. The steps of this procedure are detailed below.

408 **Generating fluorescence traces**

409 To simulate fluorescence traces, we first drew p Gaussian processes (GP)⁶¹, where each GP has T time points at 1 ms
 410 time resolution. We denote the i th GP as $\mathbf{z}_{i,:} \in \mathbb{R}^{1 \times T}$, where $i = 1, \dots, p$. The GP allows us to specify the covariance
 411 $K_i \in \mathbb{R}^{T \times T}$ for the i th GP across the T time points as

$$\mathbf{z}_{i,:} \sim \mathcal{N}(\mathbf{0}, K_i), \text{ where the } (t_1, t_2) \text{ element of } K_i \text{ is} \quad (12)$$

$$K_i(t_1, t_2) = \sigma_{f,i}^2 \exp\left(-\frac{(t_1 - t_2)^2}{2\tau_i^2}\right) + \sigma_{n,i}^2 \delta_{t_1, t_2} \quad (13)$$

412 and $t_1, t_2 = 1, \dots, T$. Here we chose the commonly used squared exponential covariance function. The squared
 413 exponential covariance is defined by its signal variance $\sigma_{f,i}^2 \in \mathbb{R}_+$, characteristic timescale $\tau_i \in \mathbb{R}_+$, and noise
 414 variance $\sigma_{n,i}^2 \in \mathbb{R}_+$. The Kronecker delta δ_{t_1, t_2} equals 1 if $t_1 = t_2$ and 0 otherwise. We set $\sigma_{f,i}^2 = 1 - \sigma_{n,i}^2$ so that the
 415 latent variable $\mathbf{z}_t \in \mathbb{R}^{p \times 1}$ at every time point has mean $\mathbf{0}$ and a variance I . The vector \mathbf{z}_t comprises the t th time point
 416 from each of the p Gaussian processes. The noise variance $\sigma_{n,i}^2$ must be nonzero to ensure that K_i is invertible, hence
 417 for practical purposes we set $\sigma_{n,i}^2 = 10^{-9}$. We note that an LDS can also be used to introduce latent dynamics, but
 418 we chose to use a GP due to the ease by which we can generate stationary time series with specified time scales.
 419 Furthermore, using a GP introduced model mismatch for all of the dimensionality reduction methods, which enables
 420 a more meaningful comparison across methods.

421 Next, we projected the low-dimensional latent variables \mathbf{z}_t into the high-dimensional neural space to obtain
 422 neural firing rates for each of q neurons at each time point $t = 1, \dots, T$. We then imposed a rectifying nonlinearity
 423 applied element-by-element using $\log(\mathbf{1} + \exp(W\mathbf{z}_t + \boldsymbol{\mu}))$, where $W \in \mathbb{R}^{q \times p}$ is the loading matrix and $\boldsymbol{\mu} \in \mathbb{R}^{q \times 1}$ is a

424 constant offset, to ensure that firing rates are non-negative. We generated binary spikes $\mathbf{s}_t \in \mathbb{R}^{q \times 1}$ at each time point
425 using an inhomogeneous Poisson process with time-varying rates defined by the output of this rectifying nonlinearity.
426 Finally, we obtained the calcium concentration $\mathbf{c}_t \in \mathbb{R}^{q \times 1}$ using a first order autoregressive model and fluorescence
427 $\mathbf{y}_t \in \mathbb{R}^{q \times 1}$, as in Friedrich et al. 2017³⁴

$$\mathbf{c}_t = \Gamma \mathbf{c}_{t-1} + \mathbf{s}_t \quad (14)$$

$$\mathbf{y}_t = \mathbf{B} \mathbf{c}_t + \mathbf{b} + \boldsymbol{\varepsilon}_t, \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, R) \quad (15)$$

428 where $\Gamma \in \mathbb{R}^{q \times q}$ determines how quickly the calcium decays, $\mathbf{B} \in \mathbb{R}^{q \times q}$ relates the calcium concentration to the
429 fluorescence, $\mathbf{b} \in \mathbb{R}^{q \times 1}$ is the baseline fluorescence, $R \in \mathbb{R}^{q \times q}$ is the imaging noise covariance, and $t = 1, \dots, T$, at
430 the same (1ms) resolution as the GP. We set Γ , \mathbf{B} , and R to be diagonal. We specified Γ such that the decay constants
431 approximately matched the decay constants of GCaMP6f, GCaMP6m, and GCaMP6s, found in Chen et al.¹³ (See
432 Supplementary Table 1). \mathbf{B} represents experimental variables influencing the scale of the calcium signal of each
433 neuron, such as the amplification of the imaging system. \mathbf{B} is set as the identity matrix and \mathbf{b} is set to be $\mathbf{0}$ in our
434 simulations. We varied the signal-to-noise ratio by varying R to simulate low, medium, and high noise regimes
435 (Supplementary Table 1).

436 For each simulation run, we simulated $p = 10$ latent variables, where each latent dimension had the same latent
437 timescale for ease of interpretation. Across different settings of experimentally relevant parameters, we explored a
438 range of timescales $\tau \in \{50, 100, 200, 1000, 2000, 5000\}$ ms. To define loading matrices W that were realistic for
439 spiking activity, we used electrophysiological recordings with 94 neurons⁶. Across runs, we tested $q \in \{20, 50, 94\}$
440 neurons. For the $q = 20$ and $q = 50$ cases, we randomly subsampled the electrophysiological recordings. To avoid
441 a start-up transient at the start of every trial from the calcium model, we generated two long fluorescence traces
442 for each neuron, each 6,000,000 time points long (1 ms resolution). One fluorescence trace was used for training,
443 and the other for testing. We divided each trace into 100 trials. Each trial was 60 s long, comprising 60,000 time
444 points. Finally, having generated these fluorescence traces at 1 kHz, we down-sampled the fluorescence rate to a
445 more typical imaging rate of 40 Hz (2,400 time points per trial). The simulation parameters are summarized in
446 Supplementary Table 1. See Practical Considerations below for how we generated GPs with 6,000,000 time points
447 within the computer's memory constraints.

448 **Evaluation of dimensionality reduction methods**

449 We applied the four dimensionality reduction methods to the simulated fluorescence traces to evaluate how accurately
450 each method reconstructed the ground truth latent variables. We used two-fold cross validation in these estimates,
451 so that the model parameters were fit using half of the data (100 training trials, each with 2,400 time points) and
452 those parameters were then used to estimate the latent variables in the other half of the data (100 held-out trials).

453 For all methods, the latent variables are only unique up to an arbitrary linear transformation. In order to compare
454 the estimated and ground truth latent variables, we aligned them using the following procedure. First we split the
455 estimated latent variables for each test fold into two further inner halves. We concatenated the trials over time, such
456 that the estimated latent variables are defined as $\tilde{Z} \in \mathbb{R}^{p \times \tilde{T}}$, and the ground truth latent variables are defined as
457 $Z \in \mathbb{R}^{p \times \tilde{T}}$, where $p = 10$ is the dimensionality of the latent variables and $\tilde{T} = 120,000$ for one inner half. Then we
458 applied linear regression to relate \tilde{Z} and the Z corresponding to that inner half. This yielded a transformation matrix
459 W , where $W = (Z\tilde{Z}')(Z\tilde{Z}')^{-1}$. Finally we used W to align the \tilde{Z} corresponding to the other inner half to obtain the
460 transformed estimated latent variables $\hat{Z} = W\tilde{Z}$, where $\hat{Z} \in \mathbb{R}^{p \times \tilde{T}}$. We performed the same procedure to align each
461 inner half of each cross-validation fold.

462 We computed the accuracy of the i th transformed estimated latent variables

$$R^2 = 1 - \frac{\sum_{t=1}^{\tilde{T}} (z_t^{(i)} - \hat{z}_t^{(i)})^2}{\sum_{t=1}^{\tilde{T}} (z_t^{(i)} - \bar{z}^{(i)})^2} \quad (16)$$

463 where $z_t^{(i)} \in \mathbb{R}$ is the (i, t) entry of Z , and $\hat{z}_t^{(i)} \in \mathbb{R}$ is the (i, t) entry of \hat{Z} . The mean of the i th ground truth latent
464 variable across time is defined as $\bar{z}^{(i)} \in \mathbb{R}$, where $i = 1, \dots, p$. A larger R^2 means a better match between the estimated
465 latent variables and the ground truth latent variables, based on the proportion of total variance (of the ground truth
466 latent variables) explained. R^2 has an upper limit of 1, and any value below 0 indicates that the estimate is poorer
467 than the using the ground truth mean. We repeated this process for each cross-validation fold and averaged the
468 results across all p latent dimensions and folds. It is important to note that since the results are cross-validated,
469 a dimensionality reduction method with more parameters will not necessarily outperform a method with fewer
470 parameters.

471 **Experimental data**

472 ***Larval Zebrafish***

473 Neurons were imaged from the dorsal raphe nucleus (DRN) of larval zebrafish while they engaged in a "fictive
474 swimming" motosensory gain adaptation task³⁹. Calcium imaging was performed using light-sheet microscopy
475 at 30 Hz on a single plane of narrow area around the DRN. This was performed with *Tg(elavl3:GCaMP6f)^{if1}* fish
476 expressing GCaMP6f in the cytosol⁶². We analyzed recordings from three fish. In the task, the fish underwent an
477 initialization period of 20 seconds to increase locomotor drive, a training period in which the fish attenuated their
478 locomotor drive, and a delay period of 10 seconds which stopped the fish from swimming. Finally, there was a test
479 period of 5 seconds to probe the extent to which the attenuated locomotor drive persisted throughout the delay period.
480 There were three different training period lengths of 7, 15, or 30 seconds. Here we combined the data from the
481 different training periods. We included only the 20-second initialization period and the first 7 seconds of the training

482 period of each trial. Thus, the analyzed portion of each trial is nominally identical. For each fish, this yielded 15
483 trials, each 27 seconds long. We analysed 22, 19, and 19 neurons imaged from the DRN of each fish, respectively.

484 **Mouse**

485 Two-photon calcium imaging was performed in the binocular zone of V1 in awake head-fixed mice resting atop a
486 floating spherical treadmill⁴⁰. GCaMP6f was expressed in excitatory neurons and imaged at 15.5 Hz. We analyzed
487 recordings from three mice. Mice were positioned to passively view static sinusoidal gratings, without reward. There
488 were 180 gratings presented, comprising 12 different orientations equally spaced with range $\{0 - 165\}^\circ$ and 15
489 different spatial frequencies equally spaced with range $\{0.02 - 0.30\}$ cycles/ $^\circ$. Each "trial" was a 196.7 seconds
490 long recording (3049 time points), comprising 4 presentations of each of the 180 possible gratings in random order.
491 Each presentation lasted 250 ms without an intervening grey screen. The onset time of the first stimulus relative to
492 the beginning of the recording was varied. This means that there is a short period of time recorded before the first
493 stimulus is shown, and a short period of time recorded after the last stimulus is shown. The experiment comprised 15
494 trials for each mouse. We analysed 133, 252, and 319 neurons from V1 of mouse (labeled mouse 1-3, respectively).
495 These mice correspond to mouse2317, mouse2320, and mouse2209 in the experiments.

496 **Data analysis**

497 ***Leave-neuron-out fluorescence prediction***

498 We sought to compare the four dimensionality reduction methods using experimental data. In experimental data,
499 ground truth latent variables are unknown, and so we could not use the same evaluation procedure of comparing
500 estimated and ground truth latent variables as in the simulations. Furthermore, the cross-validated data likelihoods
501 are not comparable across all methods. Hence to compare the four methods, we performed a leave-neuron-out
502 fluorescence prediction test to determine which method best summarizes the neuronal activity with low-dimensional
503 latent variables^{37,38}. The intuition is that a method that provides a better summary of the population activity using
504 the latent variables would be better able to predict the activity of held-out neuronal fluorescence traces.

505 For the leave-neuron-out fluorescence prediction, we performed 5-fold cross-validation. We first split the trials
506 into five equal-sized folds. For a given dimensionality reduction method, we fit the model parameters using four
507 of the folds. With the remaining validation fold, we estimated the latent variables using all but one neuron, and
508 then predicted the activity of that held-out neuron using those estimated latent variables. We did this for every
509 neuron in the validation fold. In the same manner, we performed leave-neuron-out fluorescence predictions for the
510 remaining folds and thus obtained predictions of the fluorescence activity for all the neurons at all time points. We
511 then computed the Pearson's correlation coefficient between the predicted and recorded fluorescence traces.

512 Performing the leave-neuron-out prediction procedure requires selecting the latent dimensionality p for each
513 method. For the larval zebrafish DRN recordings, we used nested cross-validation to select the optimal dimensionality

514 for each method. The leave-neuron-out prediction then used this optimal latent dimensionality, which can be different
515 for each method. Nested cross-validation uses the 5-fold cross-validation described earlier as the outer folds. Within
516 each outer fold, the training portions are then used to perform an inner 4-fold cross-validation where we determined
517 the optimal latent dimensionality using cross-validated data likelihood. We used this optimal latent dimensionality
518 to estimate the model parameters using the same four training folds. These model parameters were then used for the
519 leave-neuron-out prediction procedure of the remaining (validation) outer fold. Note that although the cross-validated
520 data likelihood is not comparable across methods, it is comparable across different latent dimensionalities for the
521 same method.

522 For the mouse V1 recordings, we found that the performance of all dimensionality reduction methods increased
523 as latent dimensionality increased within the range of dimensionalities tested (5-50). Thus we used a latent
524 dimensionality of 50 for all methods in our 5-fold cross validation procedure for leave-neuron-out prediction.

525 **Decoding analysis**

526 Another way to assess how meaningful the extracted latent variables are is by decoding external variables from the
527 latent variables⁴¹. For the larval zebrafish recordings, there is no moment by moment behavior that can be decoded
528 from the neural activity. In the mouse recordings, we can decode the orientation and spatial frequency of the grating
529 stimuli. To begin, we applied the dimensionality reduction methods to the mouse recordings. We then applied a
530 linear Gaussian Naïve Bayes classifier⁶³ to the extracted latent variables

$$P(\mathbf{z}|C_k) = \frac{1}{(2\pi)^{p/2} |\Sigma|^{1/2}} \exp[-(\mathbf{z} - \boldsymbol{\mu}_k)' \Sigma^{-1} (\mathbf{z} - \boldsymbol{\mu}_k)] \quad (17)$$

531 where $\mathbf{z} \in \mathbb{R}^{p \times 1}$ is the latent variable averaged across the time points in a 250 ms window corresponding to a given
532 stimulus C_k , $k = 1, \dots, 180$. The parameters of the classifier are $\boldsymbol{\mu}_k$, the mean of the latent variables corresponding
533 to class k , and Σ , the covariance of the latent variables across trials. We constrain Σ to be diagonal and the same
534 across all classes. The parameters $\boldsymbol{\mu}_k$ and Σ are fit by maximizing the likelihood of the training data using 5-fold
535 cross-validation.

536 We then used the parameters found from the training data to classify the held-out data according to

$$\hat{y} = \arg \max_y P(C_k = y | \mathbf{z}) \quad (18)$$

537 where \hat{y} is the predicted class label (1,...,180). We then computed the accuracy of the predicted class labels
538 against the true class labels (chance level is $\frac{1}{180}$). A higher accuracy indicates that the latent variables are better at
539 capturing the shared modulations among neurons that are relevant to the visual stimulus. For latent dimensionalities

540 $\{5, 10, 20, 30, 40, 50\}$, we fit each dimensionality reduction method using all data. Then, we performed 5-fold
541 cross-validation in the decoding stage.

542 To decode visual stimuli, there are two considerations. First, visual information takes time to arrive in the visual
543 cortex, hence there is a need to shift the window of neural activity relative to the stimulus presentation. Second,
544 there is an additional time delay introduced by the calcium indicators. Unlike deconv-LDS and CILDS, LDS does
545 not attempt to remove the calcium decay. Hence, we expected a longer latency for LDS than for deconv-LDS and
546 CILDS. To determine the appropriate window, we considered a range of time lags and evaluated their cross-validated
547 classification accuracy. We found that the best cross-validated accuracy was obtained for a 4 time point (260 ms)
548 shift for LDS, and a 3 time point (190 ms) shift for deconv-LDS and CILDS. Thus we used these time lags to report
549 the classification accuracy of our decoding analysis (Fig. 6g).

550 **Practical considerations**

551 ***Approximating a long Gaussian process***

552 To carry out computations using a GP, one would need to represent a covariance matrix of size $T \times T$ in memory,
553 where $T = 6,000,000$ in our simulations. For such values of T , the memory requirement can exceed the memory
554 capacity of the computer. To overcome this, we employed two strategies in tandem. First, we generated the GPs one
555 segment at a time, rather than all T time points at once. For example, to generate a GP with 6,000,000 time points,
556 we can first generate a GP with 5,000 time points according to equations (12) and (13). Then, we can generate a GP
557 for the next 5,000 time points conditioned on the first 5,000 time points using Gaussian conditioning⁶¹

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} A & C \\ C^T & B \end{bmatrix} \right) \quad (19)$$

$$\mathbf{x} | \mathbf{y} \sim \mathcal{N}(CB^{-1}\mathbf{y}, A - CB^{-1}C^T) \quad (20)$$

558 where $A \in \mathbb{R}^{5000 \times 5000}$ and $B \in \mathbb{R}^{5000 \times 5000}$ are the covariances of the second and first half of the GP, respectively.
559 $C \in \mathbb{R}^{5000 \times 5000}$ is the covariance between the first and second half of the GP. We can continue this procedure, where
560 for each new segment we condition on all of the segments that have been generated thus far. Statistically this
561 procedure is equivalent to if we had generated the entire GP time series at once.

562 As we continue this procedure, the number of time points being conditioned on will grow and the matrices
563 B and C in equations (19) and (20) can exceed the memory capacity of the computer. We thus employ a second
564 strategy that leverages the fact that, according to the squared exponential covariance (13), two time points covary
565 highly when they are close in time and almost independent when they are far apart in time. Thus, we make the
566 approximation that in equations (19) and (20), we condition only on the most recent 5,000 time points.

567 **Computational running time**

568 The following are representative running times to fit the different dimensionality reduction models. These running
569 times are based on single threads run on Matlab (2019a) using Intel(R) Xeon(R) CPU processors (Gold 6230, 2.1
570 GHz) with 250 GB of RAM. First, consider one cross-validation fold of the zebrafish recordings with 4 trials,
571 22 neurons, 10 latent variables, and 1950 time points per trial. For each EM iteration, CILDS takes on average
572 approximately 0.9 s, LDS (as well as the second-stage of deconv-LDS) takes 0.4 s. Second, consider one cross-
573 validation fold of the mouse recordings with 12 trials, 319 neurons, 30 latent variables, and 3049 time points per
574 trial. For each EM iteration, CILDS takes on average approximately 110 s, and LDS (as well as the second-stage
575 of deconv-LDS) takes 10 s. For all methods, the most expensive computations are the matrix inversions in the
576 expectation step of the EM algorithm. There is a $p \times p$ matrix inversion at each time point for LDS (and deconv-LDS),
577 and a $(p + q) \times (p + q)$ matrix inversion at each time point for CILDS, where p is the latent dimensionality and q is
578 the number of neurons.

579 **Data Availability**

580 The data that support the findings of this study are available from the authors upon reasonable request.

581 **Code Availability**

582 Matlab code for the simulations and dimensionality reduction methods in this work will be made publicly available
583 upon publication.

584 **References**

- 585 **1.** Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17**,
586 1500–1509, [10.1038/nn.3776](https://doi.org/10.1038/nn.3776) (2014).
- 587 **2.** Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural manifolds for the control of movement. *Neuron*
588 **94**, 978–984, <https://doi.org/10.1016/j.neuron.2017.05.025> (2017).
- 589 **3.** Gao, P. & Ganguli, S. On simplicity and complexity in the brave new world of large-scale neuroscience. *Curr.*
590 *Opin. Neurobiol.* **32**, 148–155, <https://doi.org/10.1016/j.conb.2015.04.003> (2015). Large-Scale Recording
591 Technology (32).
- 592 **4.** Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics
593 in prefrontal cortex. *Nature* **503**, 78–84, [10.1038/nature12742](https://doi.org/10.1038/nature12742) (2013).
- 594 **5.** Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56, [10.1038/nature11129](https://doi.org/10.1038/nature11129)
595 (2012).
- 596 **6.** Sadtler, P. T. *et al.* Neural constraints on learning. *Nature* **512**, 423–426, [10.1038/nature13665](https://doi.org/10.1038/nature13665) (2014).

- 597 **7.** Murray, J. D. *et al.* Stable population coding for working memory coexists with heterogeneous neural dynamics
598 in prefrontal cortex. *Proc. Natl. Acad. Sci.* **114**, 394–399, [10.1073/pnas.1619449114](https://doi.org/10.1073/pnas.1619449114) (2017). <https://www.pnas.org/content/114/2/394.full.pdf>.
599
- 600 **8.** Wang, J., Narain, D., Hosseini, E. A. & Jazayeri, M. Flexible timing by temporal scaling of cortical responses.
601 *Nat. Neurosci.* **21**, 102–110, [10.1038/s41593-017-0028-6](https://doi.org/10.1038/s41593-017-0028-6) (2018).
- 602 **9.** Cohen, M. R. & Maunsell, J. H. R. A neuronal population measure of attention predicts behavioral performance
603 on individual trials. *J. Neurosci.* **30**, 15241–15253, [10.1523/JNEUROSCI.2171-10.2010](https://doi.org/10.1523/JNEUROSCI.2171-10.2010) (2010). <https://www.jneurosci.org/content/30/45/15241.full.pdf>.
604
- 605 **10.** Mazor, O. & Laurent, G. Transient dynamics versus fixed points in odor representations by locust antennal lobe
606 projection neurons. *Neuron* **48**, 661–673, <https://doi.org/10.1016/j.neuron.2005.09.032> (2005).
- 607 **11.** Bouchard, K. E., Mesgarani, N., Johnson, K. & Chang, E. F. Functional organization of human sensorimotor
608 cortex for speech articulation. *Nature* **495**, 327–332, [10.1038/nature11911](https://doi.org/10.1038/nature11911) (2013).
- 609 **12.** Peron, S., Chen, T.-W. & Svoboda, K. Comprehensive imaging of cortical networks. *Curr. Opin. Neurobiol.* **32**,
610 115–123, <https://doi.org/10.1016/j.conb.2015.03.016> (2015). Large-Scale Recording Technology (32).
- 611 **13.** Chen, T.-W. *et al.* Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300,
612 [10.1038/nature12354](https://doi.org/10.1038/nature12354) (2013).
- 613 **14.** Ahrens, M. B., Orger, M. B., Robson, D. N., Li, J. M. & Keller, P. J. Whole-brain functional imaging at cellular
614 resolution using light-sheet microscopy. *Nat. Methods* **10**, 413–420, [10.1038/nmeth.2434](https://doi.org/10.1038/nmeth.2434) (2013).
- 615 **15.** Briggman, K., Abarbanel, H. & Kristan, W. Optical imaging of neuronal populations during decision-making.
616 *Sci. (New York, N.Y.)* **307**, 896–901, [10.1126/science.1103736](https://doi.org/10.1126/science.1103736) (2005).
- 617 **16.** Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation
618 decision task. *Nature* **484**, 62–68, [10.1038/nature10918](https://doi.org/10.1038/nature10918) (2012).
- 619 **17.** Ahrens, M. B. *et al.* Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature* **485**, 471–477,
620 [10.1038/nature11057](https://doi.org/10.1038/nature11057) (2012).
- 621 **18.** Kato, H., Chu, M., Isaacson, J. & Komiyama, T. Dynamic sensory representations in the olfactory bulb:
622 modulation by wakefulness and experience. *Neuron* **76**, 962–975, <https://doi.org/10.1016/j.neuron.2012.09.037>
623 (2012).
- 624 **19.** Daie, K., Goldman, M. & Aksay, E. Spatial patterns of persistent neural activity vary with the behavioral context
625 of short-term memory. *Neuron* **85**, 847–860, <https://doi.org/10.1016/j.neuron.2015.01.006> (2015).

- 626 **20.** Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumu-
627 lation in cortex. *Nat. Neurosci.* **19**, 1672–1681, [10.1038/nn.4403](https://doi.org/10.1038/nn.4403) (2016).
- 628 **21.** Makino, H. *et al.* Transformation of cortex-wide emergent properties during motor learning. *Neuron* **94**,
629 880–890.e8, <https://doi.org/10.1016/j.neuron.2017.04.015> (2017).
- 630 **22.** Driscoll, L. N., Pettit, N. L., Minderer, M., Chettih, S. N. & Harvey, C. D. Dynamic reorganization of neuronal
631 activity patterns in parietal cortex. *Cell* **170**, 986–999.e16, <https://doi.org/10.1016/j.cell.2017.07.021> (2017).
- 632 **23.** Stringer, C., Pachitariu, M., Steinmetz, N., Carandini, M. & Harris, K. D. High-dimensional geometry of
633 population responses in visual cortex. *Nature* **571**, 361–365, [10.1038/s41586-019-1346-5](https://doi.org/10.1038/s41586-019-1346-5) (2019).
- 634 **24.** Rummyantsev, O. I. *et al.* Fundamental bounds on the fidelity of sensory cortical coding. *Nature* **580**, 100–105,
635 [10.1038/s41586-020-2130-2](https://doi.org/10.1038/s41586-020-2130-2) (2020).
- 636 **25.** Pashkovski, S. L. *et al.* Structure and flexibility in cortical representations of odour space. *Nature* **583**, 253–258,
637 [10.1038/s41586-020-2451-1](https://doi.org/10.1038/s41586-020-2451-1) (2020).
- 638 **26.** Nieh, E. *et al.* Geometry of abstract learned knowledge in the hippocampus. *Nature* **595**, 80–84, [10.1038/
639 s41586-021-03652-7](https://doi.org/10.1038/s41586-021-03652-7) (2021).
- 640 **27.** Wei, Z. *et al.* A comparison of neuronal population dynamics measured with calcium imaging and electrophysi-
641 ology. *PLOS Comput. Biol.* **16**, e1008198, [10.1371/journal.pcbi.1008198](https://doi.org/10.1371/journal.pcbi.1008198) (2020).
- 642 **28.** Afshar, A. *et al.* Single-trial neural correlates of arm movement preparation. *Neuron* **71**, 555–564, [10.1016/j.
643 neuron.2011.05.047](https://doi.org/10.1016/j.neuron.2011.05.047) (2011).
- 644 **29.** Kiani, R., Cueva, C. J., Reppas, J. B. & Newsome, W. T. Dynamics of neural population responses in prefrontal
645 cortex indicate changes of mind on single trials. *Curr. Biol.* **24**, 1542–1547, [https://doi.org/10.1016/j.cub.2014.
646 05.049](https://doi.org/10.1016/j.cub.2014.05.049) (2014).
- 647 **30.** Gouvêa, T. S. *et al.* Striatal dynamics explain duration judgments. *eLife* **4**, e11386, [10.7554/eLife.11386](https://doi.org/10.7554/eLife.11386) (2015).
- 648 **31.** Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Vacillation, indecision and hesitation in
649 moment-by-moment decoding of monkey motor cortex. *eLife* **4**, e04677, [10.7554/eLife.04677](https://doi.org/10.7554/eLife.04677) (2015).
- 650 **32.** Vogelstein, J. T. *et al.* Fast nonnegative deconvolution for spike train inference from population calcium imaging.
651 *J. Neurophysiol.* **104**, 3691–3704, [10.1152/jn.01073.2009](https://doi.org/10.1152/jn.01073.2009) (2010). PMID: 20554834, [https://doi.org/10.1152/jn.
652 01073.2009](https://doi.org/10.1152/jn.01073.2009).
- 653 **33.** Theis, L. *et al.* Benchmarking spike rate inference in population calcium imaging. *Neuron* **90**, 471–482,
654 [10.1016/j.neuron.2016.04.014](https://doi.org/10.1016/j.neuron.2016.04.014) (2016).

- 655 **34.** Friedrich, J., Zhou, P. & Paninski, L. Fast online deconvolution of calcium imaging data. *PLOS Comput. Biol.*
656 **13**, e1005423, [10.1371/journal.pcbi.1005423](https://doi.org/10.1371/journal.pcbi.1005423) (2017).
- 657 **35.** Runyan, C. A., Piasini, E., Panzeri, S. & Harvey, C. D. Distinct timescales of population coding across cortex.
658 *Nature* **548**, 92–96, [10.1038/nature23020](https://doi.org/10.1038/nature23020) (2017).
- 659 **36.** Murray, J. D. *et al.* A hierarchy of intrinsic timescales across primate cortex. *Nat. Neurosci.* **17**, 1661–1663,
660 [10.1038/nn.3862](https://doi.org/10.1038/nn.3862) (2014).
- 661 **37.** Yu, B. M. *et al.* Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population
662 activity. *J. neurophysiology* **102**, 614–635, [10.1152/jn.90941.2008](https://doi.org/10.1152/jn.90941.2008) (2009).
- 663 **38.** Pei, F. C. *et al.* Neural latents benchmark ‘21: Evaluating latent variable models of neural population activity.
664 In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round*
665 *2)* (2021).
- 666 **39.** Kawashima, T., Zwart, M. F., Yang, C.-T., Mensh, B. D. & Ahrens, M. B. The serotonergic system tracks the
667 outcomes of actions to mediate short-term motor learning. *Cell* **167**, 933–946.e20, [https://doi.org/10.1016/j.cell.](https://doi.org/10.1016/j.cell.2016.09.055)
668 [2016.09.055](https://doi.org/10.1016/j.cell.2016.09.055) (2016).
- 669 **40.** Jeon, B. B., Swain, A. D., Good, J. T., Chase, S. M. & Kuhlman, S. J. Feature selectivity is stable in primary
670 visual cortex across a range of spatial frequencies. *Sci. reports* **8**, 15288–15288, [10.1038/s41598-018-33633-2](https://doi.org/10.1038/s41598-018-33633-2)
671 (2018).
- 672 **41.** Pandarinath, C. *et al.* Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat.*
673 *Methods* **15**, 805–815, [10.1038/s41592-018-0109-9](https://doi.org/10.1038/s41592-018-0109-9) (2018).
- 674 **42.** Churchland, M. *et al.* Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nat.*
675 *neuroscience* **13**, 369–78, [10.1038/nn.2501](https://doi.org/10.1038/nn.2501) (2010).
- 676 **43.** Williamson, R. C. *et al.* Scaling properties of dimensionality reduction for neural populations and network
677 models. *PLOS Comput. Biol.* **12**, e1005141, [10.1371/journal.pcbi.1005141](https://doi.org/10.1371/journal.pcbi.1005141) (2016).
- 678 **44.** Umakantha, A. *et al.* Bridging neuronal correlations and dimensionality reduction. *Neuron* **109**, 2740–2754.e12,
679 [10.1016/j.neuron.2021.06.028](https://doi.org/10.1016/j.neuron.2021.06.028) (2021).
- 680 **45.** Huang, C. *et al.* Circuit models of low-dimensional shared variability in cortical networks. *Neuron* **101**,
681 337–348.e4, <https://doi.org/10.1016/j.neuron.2018.11.034> (2019).
- 682 **46.** Triplett, M. A., Pujic, Z., Sun, B., Avitan, L. & Goodhill, G. J. Model-based decoupling of evoked and
683 spontaneous neural activity in calcium imaging data. *PLOS Comput. Biol.* **16**, e1008330, [10.1371/journal.pcbi.](https://doi.org/10.1371/journal.pcbi.1008330)
684 [1008330](https://doi.org/10.1371/journal.pcbi.1008330) (2020).

- 685 **47.** Aitchison, L. *et al.* Model-based Bayesian inference of neural activity and connectivity from all-optical
686 interrogation of a neural circuit. In Guyon, I. *et al.* (eds.) *Advances in Neural Information Processing Systems*,
687 vol. 30 (Curran Associates, Inc., 2017).
- 688 **48.** Kao, J. C. *et al.* Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nat.*
689 *Commun.* **6**, 7759, [10.1038/ncomms8759](https://doi.org/10.1038/ncomms8759) (2015).
- 690 **49.** Williams, A. H. *et al.* Discovering precise temporal patterns in large-scale neural recordings through robust and
691 interpretable time warping. *Neuron* **105**, 246–259.e8, <https://doi.org/10.1016/j.neuron.2019.10.020> (2020).
- 692 **50.** Kobak, D. *et al.* Demixed principal component analysis of neural population data. *eLife* **5**, e10989, [10.7554/](https://doi.org/10.7554/eLife.10989)
693 [eLife.10989](https://doi.org/10.7554/eLife.10989) (2016).
- 694 **51.** Archer, E. W., Koster, U., Pillow, J. W. & Macke, J. H. Low-dimensional models of neural population activity in
695 sensory cortical circuits. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. & Weinberger, K. Q. (eds.)
696 *Advances in Neural Information Processing Systems*, vol. 27 (Curran Associates, Inc., 2014).
- 697 **52.** Lakshmanan, K. C., Sadtler, P. T., Tyler-Kabara, E. C., Batista, A. P. & Yu, B. M. Extracting low-dimensional
698 latent structure from time series in the presence of delays. *Neural Comput.* **27**, 1825–1856, [10.1162/NECO_a_](https://doi.org/10.1162/NECO_a_00759)
699 [00759](https://doi.org/10.1162/NECO_a_00759) (2015). https://direct.mit.edu/neco/article-pdf/27/9/1825/946601/neco_a_00759.pdf.
- 700 **53.** Elsayed, G. & Cunningham, J. Structure in neural population recordings: An expected byproduct of simpler
701 phenomena? *Nat. Neurosci.* **20**, [10.1038/nn.4617](https://doi.org/10.1038/nn.4617) (2017).
- 702 **54.** Pnevmatikakis, E. *et al.* Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron*
703 **89**, 285–299, [10.1016/j.neuron.2015.11.037](https://doi.org/10.1016/j.neuron.2015.11.037) (2016).
- 704 **55.** Bittner, S. R. *et al.* Population activity structure of excitatory and inhibitory neurons. *PLOS ONE* **12**, 1–27,
705 [10.1371/journal.pone.0181773](https://doi.org/10.1371/journal.pone.0181773) (2017).
- 706 **56.** Chen, J. L., Voigt, F. F., Javadzadeh, M., Krueppel, R. & Helmchen, F. Long-range population dynamics of
707 anatomically defined neocortical networks. *eLife* **5**, e14679, [10.7554/eLife.14679](https://doi.org/10.7554/eLife.14679) (2016).
- 708 **57.** Smedo, J. D., Gokcen, E., Machens, C. K., Kohn, A. & Yu, B. M. Statistical methods for dissecting interactions
709 between brain areas. *Curr. Opin. Neurobiol.* **65**, 59–69, <https://doi.org/10.1016/j.conb.2020.09.009> (2020).
710 Whole-brain interactions between neural circuits.
- 711 **58.** Pachitariu, M. *et al.* Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv* [10.1101/](https://doi.org/10.1101/061507)
712 [061507](https://doi.org/10.1101/061507) (2017). <https://www.biorxiv.org/content/early/2017/07/20/061507.full.pdf>.
- 713 **59.** Ghahramani, Z. & Hinton, G. E. Parameter estimation for linear dynamical systems. Tech. Rep. (1996).

- 714 **60.** Dempster, A. P., Laird, N. M. & Rubin, D. B. Maximum likelihood from incomplete data via the em algorithm. *J.*
715 *Royal Stat. Soc. Ser. B (Methodological)* **39**, 1–22, <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x> (1977).
716 <https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.2517-6161.1977.tb01600.x>.
- 717 **61.** Rasmussen, C. E. & Williams, C. K. I. *Gaussian processes for machine learning*. Adaptive computation and
718 machine learning (MIT Press, 2006).
- 719 **62.** Freeman, J. *et al.* Mapping brain activity at scale with cluster computing. *Nat. methods* **11**, [10.1038/nmeth.3041](https://doi.org/10.1038/nmeth.3041)
720 (2014).
- 721 **63.** Santhanam, G. *et al.* Factor-analysis methods for higher-performance neural prostheses. *J. Neurophysiol.* **102**,
722 1315–1330, [10.1152/jn.00097.2009](https://doi.org/10.1152/jn.00097.2009) (2009). PMID: 19297518, <https://doi.org/10.1152/jn.00097.2009>.
- 723 **64.** Yu, B., Shenoy, K. & Sahani, M. Derivation of Kalman filtering and smoothing equations. (2004).

724 **Acknowledgements**

725 The authors would like to thank Katrina P. Nguyen for the animal illustrations. This work was supported by the
726 Agency for Science, Technology and Research (A*STAR) Singapore (T. Koh), Howard Hughes Medical Institute
727 (W.E.B., T. Kawashima, and M.B.A.), Simons Foundation Simons Collaboration on the Global Brain Award 542943
728 (M.B.A.) and 543065 (B.M.Y.), the Shurl and Kay Curci Foundation (S.M.C. and S.J.K.), NIH R01 HD071686
729 (S.M.C. and B.M.Y.), NIH R01EY024678 (S.J.K.), NSF NCS BCS1533672 (S.M.C. and B.M.Y.), NSF CAREER
730 award IOS1553252 (S.M.C.), NSF NCS BCS1734916 (B.M.Y.), NIH CRCNS R01 NS105318 (B.M.Y.), NIH
731 CRCNS R01 MH118929 (B.M.Y.), NIH R01 EB026953 (B.M.Y.).

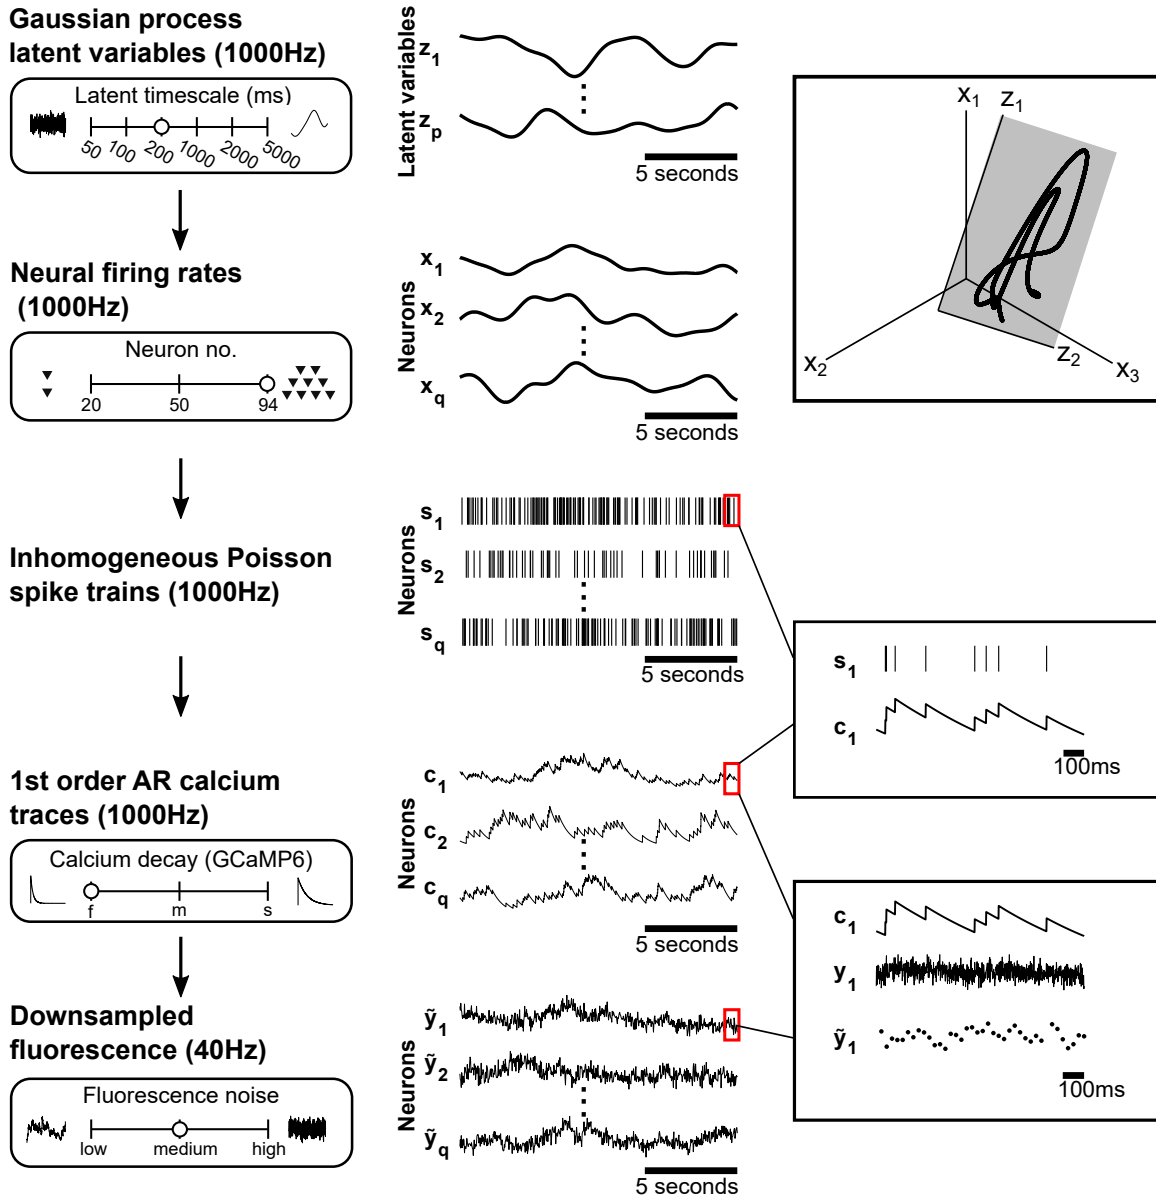
732 **Author contributions statement**

733 T. Koh, W.E.B., S.M.C., and B.M.Y. designed the dimensionality reduction and analysis methods. T. Koh derived
734 and implemented the dimensionality reduction methods. T. Koh performed the analyses, based on earlier analyses
735 by R.S. T. Kawashima and M.B.A. contributed to evaluation of the methods. T. Kawashima, B.B.J., S.J.K., M.B.A.
736 designed the animal experiments. T. Kawashima performed the larval zebrafish experiments, and B.B.J. performed
737 the mouse experiments. T. Koh, W.E.B., M.B.A., S.M.C., and B.M.Y. wrote the manuscript. All authors discussed
738 the results and commented on the manuscript. S.M.C. and B.M.Y. contributed equally to this work.

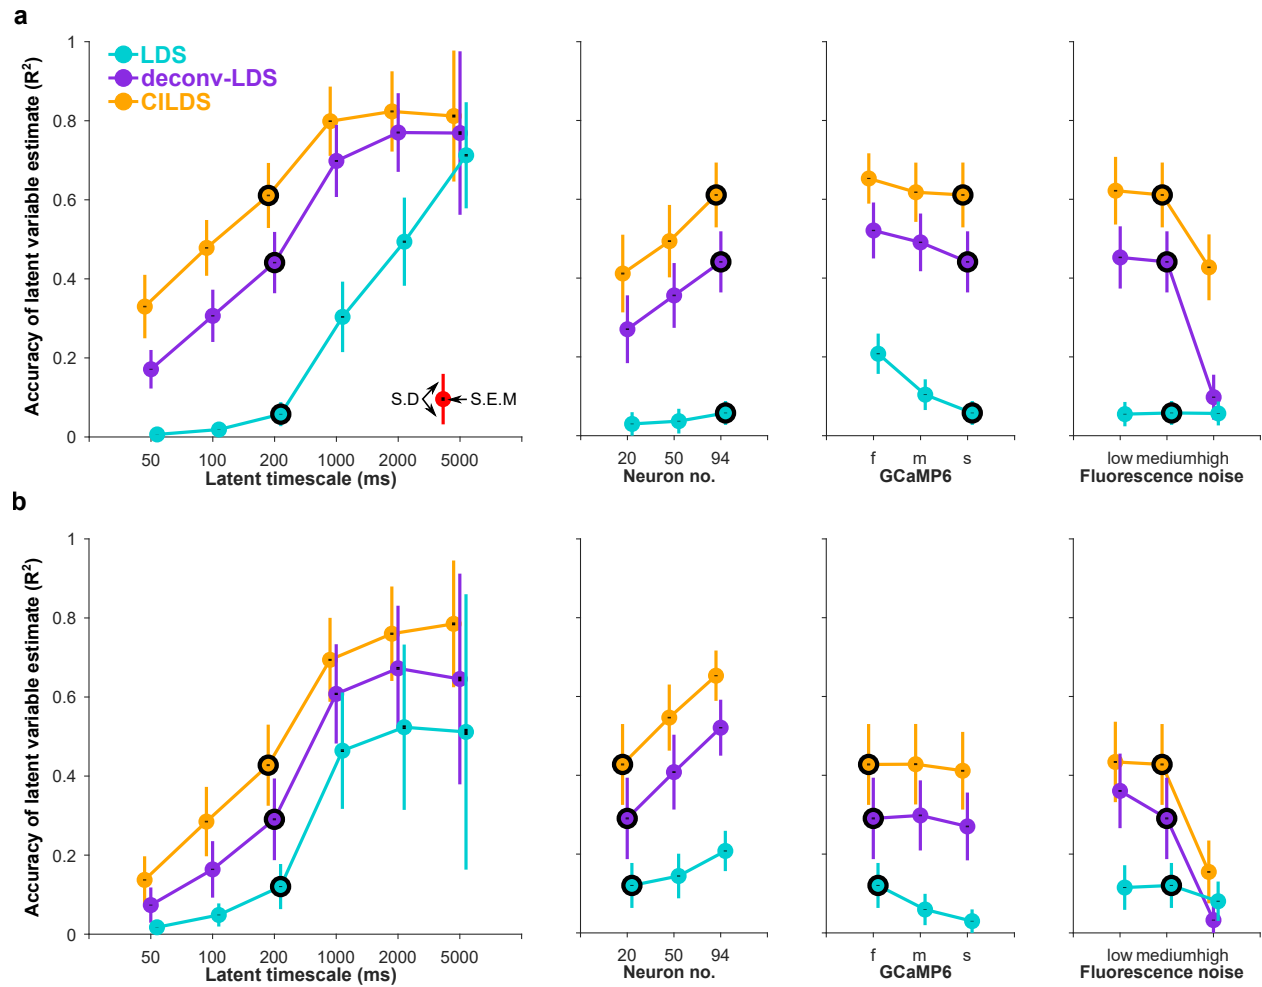
739 **Competing interests**

740 The authors declare no competing interests.

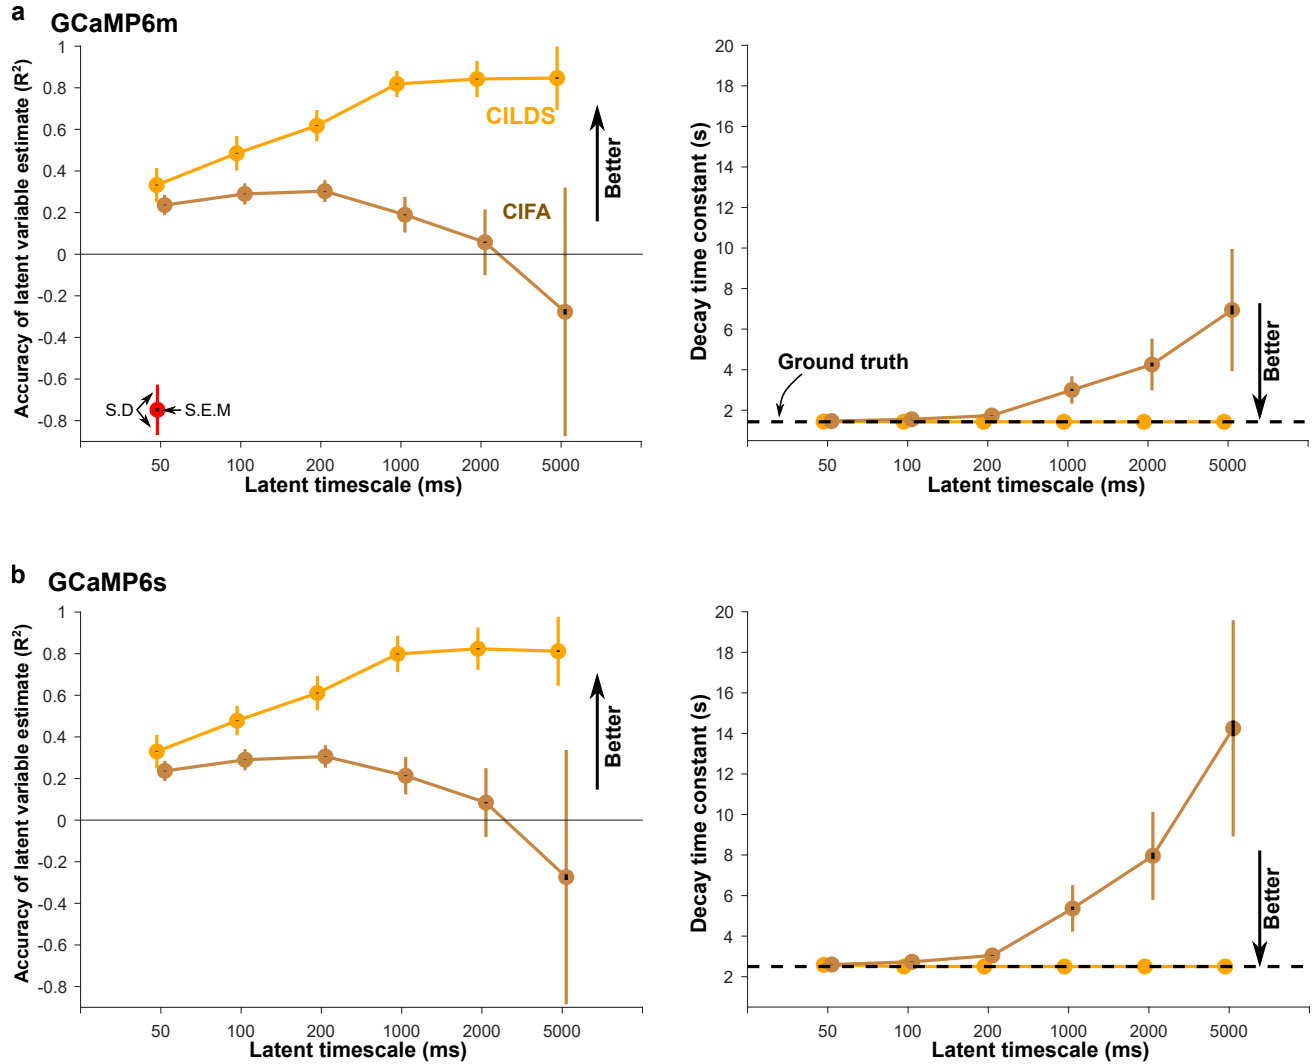
741 **Supplementary Figures**



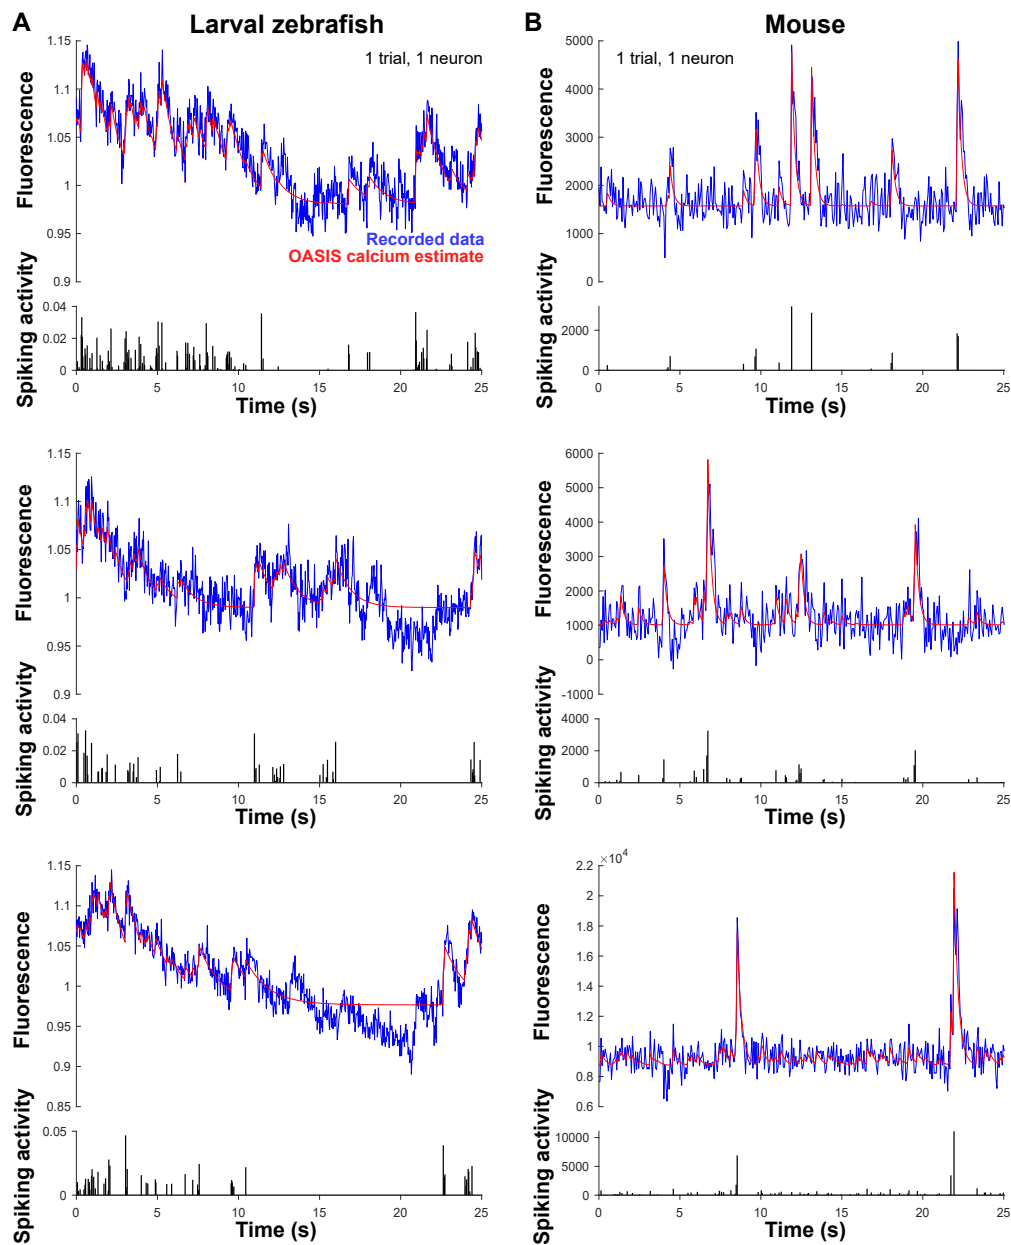
Supplementary Figure 1. Simulation framework and range of parameters. Detailed simulation framework depicting the generation of fluorescence traces from latent variables. The sliders show the range of parameters explored, and the open circles represent one example of a set parameter values. The middle column shows example signals generated at each stage. Gaussian process latent variables $\mathbf{z} \in \mathbb{R}^{p \times T}$ (sampled at 1000 Hz time resolution) are projected to form firing rates $\mathbf{x} \in \mathbb{R}^{q \times T}$, where p is the number of latent variables, q is the number of neurons, and T is the number of time points in each trace. The upper right panel shows an example of the latent space (\mathbf{z}) placed within a higher-dimensional neural space (\mathbf{x}). Firing rates are used to generate spike trains, $\mathbf{s} \in \mathbb{R}^{q \times T}$, according to an inhomogeneous Poisson process. Calcium traces, $\mathbf{c} \in \mathbb{R}^{q \times T}$, are generated from spike trains using a 1st order autoregressive process³², with the center right panel showing a close-up of the spike-to-calcium transformation. White noise is added to calcium traces to form fluorescence traces $\mathbf{y} \in \mathbb{R}^{q \times T}$. These traces are down-sampled to 40Hz, forming $\tilde{\mathbf{y}} \in \mathbb{R}^{q \times (T/25)}$, to more closely match typical sampling rates in calcium imaging recordings. The lower right panel shows a zoomed in view of this process.



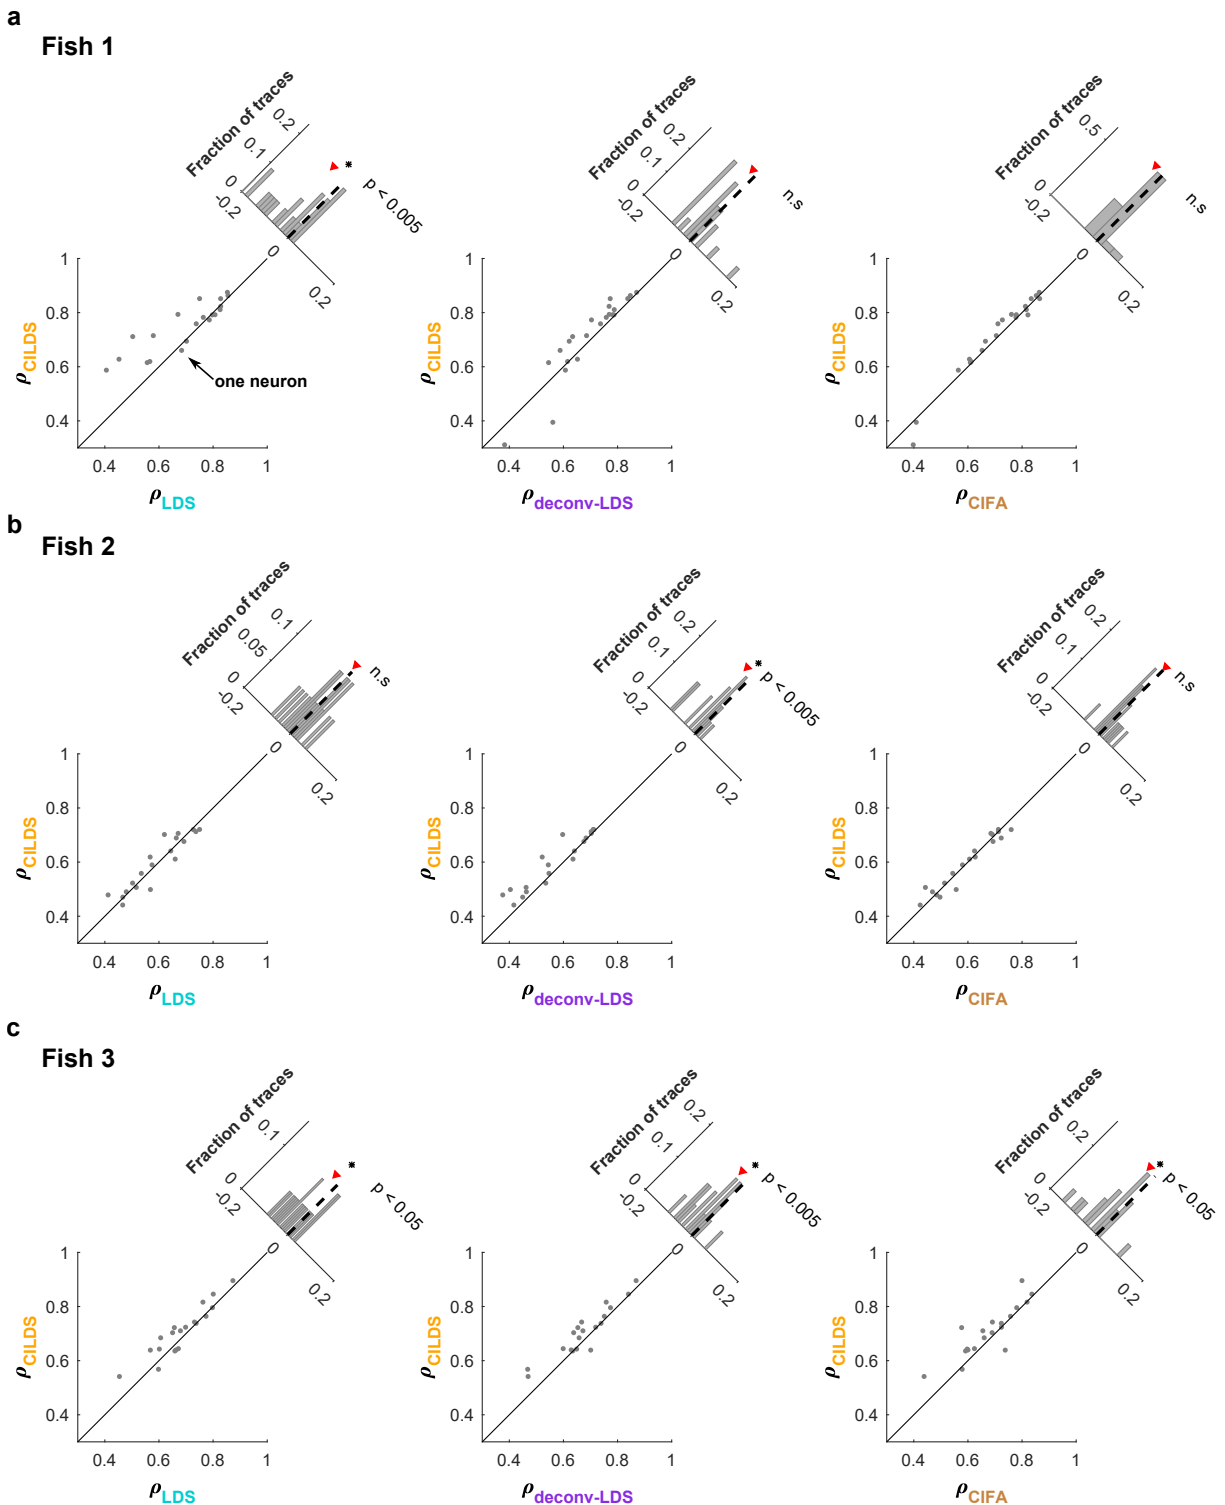
Supplementary Figure 2. CILDS continues to outperform other methods under different settings of the experimental variables. Same conventions as Fig. 3d-g. These experimental variables are different from Fig. 3 in the following ways: **(a)** a calcium decay constant matching the calcium indicator GCaMP6s instead of GCaMP6f and **(b)** a smaller number of neurons (20 instead of 94). Overall CILDS extracts latent variables that more closely match the ground truth simulated latent variables than the other two methods, consistent with Fig. 3. Black circles represent the parameter settings that are fixed across all panels in each row. Comparing to Fig. 3, there are two notable features. First, when the calcium indicator decay is slow, it becomes even more important to deconvolve. Recall that CILDS and deconv-LDS both include deconvolution, whereas LDS does not. CILDS performs similarly whether the calcium indicator is fast (Fig. 3d) or slow (here in panel a). The same is true for deconv-LDS. By contrast, LDS performs worse for slow compared to fast calcium indicator decay because it does not include deconvolution. Second with fewer neurons, the performance of all methods goes down. And as a result, there is a smaller difference in performance between methods (here in panel b). With less statistical power to leverage for separating calcium decay and latent timescale, the three methods show less distinction in performance.



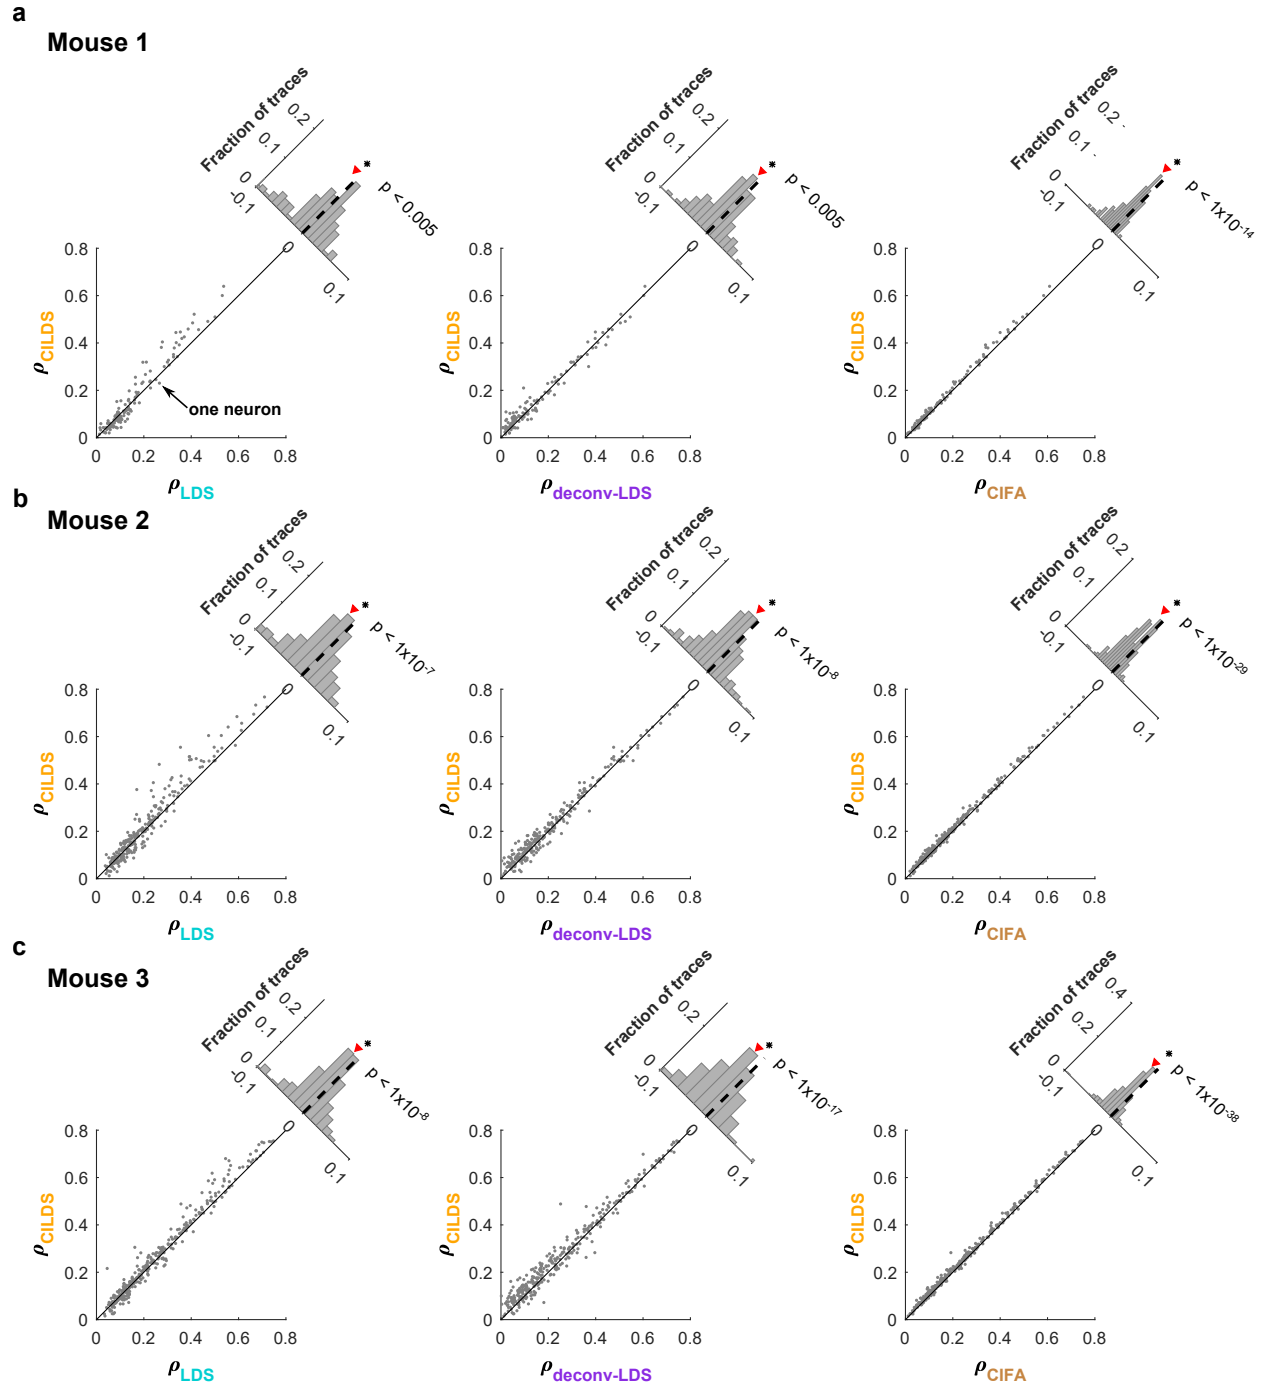
Supplementary Figure 3. CILDS continues to outperform CIFA using different calcium indicators. Same conventions as Fig. 4. Here we compare between CILDS and CIFA for (a) GCaMP6m and (b) GCaMP6s, instead of GCaMP6f (Fig. 4). Consistent with Fig. 4, the overestimation of the calcium indicator decay time constant for CIFA increases as the latent timescales increase. The overestimation also increases when progressing from GCaMP6f (Fig. 4, note different vertical scale) to GCaMP6m (panel a) to GCaMP6s (panel b).



Supplementary Figure 4. Examples of OASIS deconvolution, the first stage of deconv-LDS. (a) OASIS³⁴ applied to recorded fluorescence traces (blue) to estimate the calcium time course (red) and spiking activity (black) for larval zebrafish. Each row shows 25 seconds of a different example neuron and trial from the same fish. **(b)** Same conventions as (a), but for recordings in mice. Note that no dimensionality reduction has been applied at this stage.



Supplementary Figure 5. CILDS outperforms the other methods for individual larval zebrafish. Same conventions as Fig. 5c-e. **(a)** From left to right, correlations between recorded fluorescence and leave-neuron-out predicted fluorescence, comparing CILDS to LDS, CILDS to deconv-LDS, and CILDS to CIFA for Fish 1. **(b)** Same as (a), but for Fish 2. **(c)** Same as (a), but for Fish 3. These results are consistent with Fig. 5c-e, even at the level of individual fish. With an average of 20 points (i.e., neurons) per scatter plot, statistical significance is not attained in several plots, but in each case CILDS has a higher mean correlation than the other methods.



Supplementary Figure 6. CILDS outperforms the other methods for individual mice. Same conventions as Fig. 6c-e. (a) From left to right, correlations between recorded fluorescence and leave-neuron-out predicted fluorescence, comparing CILDS to LDS, CILDS to deconv-LDS, and CILDS to CIFA for Mouse 1. (b) Same as (a), but for Mouse 2. (c) Same as (a), but for Mouse 3. These results are consistent with Fig. 6c-e, for each mouse individually.

Parameters	Variables
q : No. of neurons	20, 50, 94
p : No. of latent variables	10
N : No. of trials	200
T : Length of trial	60 seconds
W : Loading matrix	Estimated from electrophysiological recordings using FA ⁶
μ : Mean firing rate	Estimated from electrophysiological recordings data using FA ⁶
B : Non-negative diagonal matrix	I
\mathbf{b} : Baseline fluorescence	$\mathbf{0}$
τ : Latent timescale	50, 100, 200, 1000, 5000 ms
Γ : GCaMP6 decay	f: $(0.9985)I$, m: $(0.9993)I$, s: $(0.9996)I$
R : Fluorescence noise	low: $(0.15)I$, medium: $(1.5)I$, high: $(15)I$

Supplementary Table 1. Parameter values used for simulations. Parameters within the simulation framework that were varied, and the range of parameter values tested. We chose parameter values to mimic those in real data.

742 Supplemental Information

743 EM algorithm for CILDS

744 The CILDS model is defined by equations (6) - (10). Only the fluorescence values \mathbf{y}_t are observed, whereas the
745 calcium concentrations \mathbf{c}_t and latent variables \mathbf{z}_t are not observed.

746 The goal of the EM algorithm is to maximise the probability of the observed fluorescence traces $P(\{\mathbf{y}\})$ with
747 respect to the model parameters $\theta := \{D, P, \mathbf{h}_2, G_2, \Gamma, A, \mathbf{b}, Q, B, R, \boldsymbol{\mu}_1, V_1\}$, where $\{\mathbf{y}\}$ is shorthand for $\mathbf{y}_1, \dots, \mathbf{y}_T$. To
748 perform this maximization, we iteratively perform an expectation step (E-step), then a maximization step (M-step),
749 as detailed below.

750 1 Expectation Step

751 The goal of the E-step is to compute the posterior distribution $s := P(\{\mathbf{c}\}, \{\mathbf{z}\} | \{\mathbf{y}\})$. Using this posterior distribution,
752 we can compute the following expectations

$$E_s \left[\mathbf{c}_{t-1}^{(i)} \right], E_s \left[\mathbf{c}_t^{(i)} \right], E_s \left[\mathbf{c}_t^{(i)} \mathbf{c}_t^{(i)'} \right], E_s \left[\mathbf{c}_t^{(i)} \mathbf{c}_{t-1}^{(i)'} \right], E_s \left[\mathbf{c}_t^{(i)} \mathbf{z}_t^{(i)'} \right]$$

$$E_s \left[\mathbf{c}_{t-1}^{(i)} \mathbf{z}_t^{(i)'} \right], E_s \left[\mathbf{z}_t^{(i)} \mathbf{z}_t^{(i)'} \right], E_s \left[\mathbf{z}_t^{(i)} \mathbf{z}_{t-1}^{(i)'} \right], E_s \left[\mathbf{z}_t^{(i)} \right], E_s \left[\mathbf{z}_{t-1}^{(i)} \right]$$

753 which are needed in the M-step. We start by rewriting equations (7) and (9) in block matrix notation

$$\begin{bmatrix} \mathbf{c}_t \\ \mathbf{z}_{t+1} \end{bmatrix} = \begin{bmatrix} \Gamma & A \\ 0 & D \end{bmatrix} \begin{bmatrix} \mathbf{c}_{t-1} \\ \mathbf{z}_t \end{bmatrix} + \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{w}_t \\ \mathbf{v}_{t+1} \end{bmatrix}.$$

754 The observation model (6) can be written as

$$\mathbf{y}_t = \Phi \begin{bmatrix} \mathbf{c}_t \\ \mathbf{z}_{t+1} \end{bmatrix} + \boldsymbol{\varepsilon}_t, \quad \text{where } \Phi := \begin{bmatrix} B & 0 \end{bmatrix}.$$

755 We define $\mathbf{l}_t := \begin{bmatrix} \mathbf{c}_t \\ \mathbf{z}_{t+1} \end{bmatrix}$, $\Lambda := \begin{bmatrix} \Gamma & A \\ 0 & D \end{bmatrix}$, $\mathbf{m} := \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix}$, $\mathbf{n}_t := \begin{bmatrix} \mathbf{w}_t \\ \mathbf{v}_{t+1} \end{bmatrix}$, $S := \begin{bmatrix} Q & 0 \\ 0 & P \end{bmatrix}$, $\boldsymbol{\eta}_1 := \begin{bmatrix} \boldsymbol{\mu}_1 \\ \mathbf{h}_2 \end{bmatrix}$, $\Sigma_1 := \begin{bmatrix} V_1 & 0 \\ 0 & G_2 \end{bmatrix}$.

756
757 The CILDS model can thus be written in block matrix notation as

$$\mathbf{y}_t = \Phi \mathbf{l}_t + \boldsymbol{\varepsilon}_t \tag{21}$$

$$\mathbf{l}_t = \Lambda \mathbf{l}_{t-1} + \mathbf{m} + \mathbf{n}_t \tag{22}$$

758 where $\mathbf{l}_1 \sim \mathbb{N}(\boldsymbol{\eta}_1, \Sigma_1)$, $\mathbf{n}_t \sim \mathbb{N}(\mathbf{0}, S)$, and $t = 1, \dots, T$. In other words, CILDS can be written as an LDS whose
 759 parameters are constrained in a specific way. We seek to compute $P(\mathbf{l}_t, \mathbf{l}_{t-1} | \{\mathbf{y}\}_1^T)$ for $t = 2, \dots, T$. This distribution
 760 is Gaussian, and thus it is sufficient to find its mean and covariance. We denote $E[\mathbf{l}_t | \{\mathbf{y}\}_1^t]$ by \mathbf{l}_t^τ and $\text{Var}[\mathbf{l}_t | \{\mathbf{y}\}_1^t]$
 761 by V_t^τ , as in ref.⁵⁹. To obtain the forward and backward recursion equations, we followed the steps outlined in ref.⁶⁴.
 762 For brevity, we only show the results of the derivations below.

763 1.1 Forward Recursions

764 To obtain \mathbf{l}_t^t and V_t^t , we recursively compute the following equations from $t = 1$ to $t = T$

$$\begin{aligned} \mathbf{l}_t^{t-1} &= \Lambda \mathbf{l}_{t-1}^{t-1} + \mathbf{m} \\ V_t^{t-1} &= \Lambda V_{t-1}^{t-1} \Lambda' + S \\ K_t &= V_t^{t-1} \Phi' (R + \Phi V_t^{t-1} \Phi')^{-1} \\ \mathbf{l}_t^t &= \mathbf{l}_t^{t-1} + K_t (\mathbf{y}_t - \Phi \mathbf{l}_t^{t-1}) \\ V_t^t &= V_t^{t-1} - K_t \Phi V_t^{t-1}. \end{aligned}$$

765 The recursions are initialized with $\mathbf{l}_1^0 = \boldsymbol{\eta}_1$, $V_1^0 = \Sigma_1$.

766 1.2 Backward Recursions

767 To obtain \mathbf{l}_t^T and V_t^T , we recursively compute the following equations from $t = T$ to $t = 2$. We also compute the
 768 covariance of the joint posterior distribution $P(\mathbf{l}_t, \mathbf{l}_{t-1} | \{\mathbf{y}\}_1^T)$, denoted as $V_{t,t-1}^T$

$$\begin{aligned} J_{t-1} &= V_{t-1}^{t-1} \Lambda' (V_t^{t-1})^{-1} \\ \mathbf{l}_{t-1}^T &= \mathbf{l}_{t-1}^{t-1} + J_{t-1} (\mathbf{l}_t^T - \Lambda \mathbf{l}_{t-1}^{t-1} - \mathbf{m}) \\ V_{t-1}^T &= V_{t-1}^{t-1} + J_{t-1} (V_t^T - V_t^{t-1}) J_{t-1}' \\ V_{t,t-1}^T &= V_t^T J_{t-1}'. \end{aligned}$$

769 2 Maximization Step

770 In the M-step, we seek to maximize the expected log joint distribution

$$\mathcal{Q} := E_s \left[\log P \left(\{\mathbf{y}\}^N, \{\mathbf{c}\}^N, \{\mathbf{z}\}^N \right) \right]$$

771 with respect to the model parameters, where $s := P(\{\mathbf{c}\}^N, \{\mathbf{z}\}^N | \{\mathbf{y}\}^N; \theta)$ and $\{\cdot\}^N$ represents all T time points
 772 across all N trials. The joint distribution for one trial can be factorized as

$$\begin{aligned}
 P(\{\mathbf{y}\}, \{\mathbf{c}\}, \{\mathbf{z}\}) &= P(\{\mathbf{y}\} | \{\mathbf{c}\}, \{\mathbf{z}\})P(\{\mathbf{c}\}, \{\mathbf{z}\}) \\
 &= P(\{\mathbf{y}\} | \{\mathbf{c}\})P(\{\mathbf{c}\} | \{\mathbf{z}\})P(\{\mathbf{z}\}) \\
 &= \prod_{t=1}^T P(\mathbf{y}_t | \mathbf{c}_t) \prod_{t=2}^T P(\mathbf{c}_t | \mathbf{c}_{t-1}, \mathbf{z}_t) \prod_{t=3}^T P(\mathbf{z}_t | \mathbf{z}_{t-1}) P(\mathbf{c}_1) P(\mathbf{z}_2)
 \end{aligned} \tag{23}$$

773 where these distributions are defined in equations (6) - (10).

$$\begin{aligned}
 \mathcal{Q} &= E_s \left[\log \prod_{i=1}^N P(\{\mathbf{y}\}^{(i)}, \{\mathbf{c}\}^{(i)}, \{\mathbf{z}\}^{(i)}) \right] \\
 &= E_s \left[\sum_{i=1}^N \left(-\frac{1}{2} \sum_{t=1}^T [\mathbf{y}_t^{(i)} - B\mathbf{c}_t^{(i)}]' R^{-1} [\mathbf{y}_t^{(i)} - B\mathbf{c}_t^{(i)}] - \frac{T}{2} \log |R| \right. \right. \\
 &\quad \left. \left. - \frac{1}{2} \sum_{t=2}^T [\mathbf{c}_t^{(i)} - \Gamma\mathbf{c}_{t-1}^{(i)} - A\mathbf{z}_t^{(i)} - \mathbf{b}]' Q^{-1} [\mathbf{c}_t^{(i)} - \Gamma\mathbf{c}_{t-1}^{(i)} - A\mathbf{z}_t^{(i)} - \mathbf{b}] - \frac{T-1}{2} \log |Q| \right. \right. \\
 &\quad \left. \left. - \frac{1}{2} \sum_{t=3}^T [\mathbf{z}_t^{(i)} - D\mathbf{z}_{t-1}^{(i)}]' P^{-1} [\mathbf{z}_t^{(i)} - D\mathbf{z}_{t-1}^{(i)}] - \frac{T-2}{2} \log |P| \right. \right. \\
 &\quad \left. \left. - \frac{1}{2} [(\mathbf{z}_2^{(i)} - \mathbf{h}_2)' G^{-1} (\mathbf{z}_2^{(i)} - \mathbf{h}_2)] - \frac{1}{2} \log |G| - \frac{1}{2} [\mathbf{c}_1^{(i)} - \boldsymbol{\mu}_1]' V_1^{-1} [\mathbf{c}_1^{(i)} - \boldsymbol{\mu}_1] - \frac{1}{2} \log |V_1| - \frac{T(2q+p)}{2} \log 2\pi \right) \right].
 \end{aligned}$$

774 To maximize \mathcal{Q} with respect to the model parameters θ , we compute the following partial derivatives

$$\frac{\partial \mathcal{Q}}{\partial B}, \frac{\partial \mathcal{Q}}{\partial R^{-1}}, \frac{\partial \mathcal{Q}}{\partial \Gamma}, \frac{\partial \mathcal{Q}}{\partial A}, \frac{\partial \mathcal{Q}}{\partial \mathbf{b}}, \frac{\partial \mathcal{Q}}{\partial Q^{-1}}, \frac{\partial \mathcal{Q}}{\partial D}, \frac{\partial \mathcal{Q}}{\partial P^{-1}}, \frac{\partial \mathcal{Q}}{\partial \boldsymbol{\mu}_1}, \frac{\partial \mathcal{Q}}{\partial V_1^{-1}}, \frac{\partial \mathcal{Q}}{\partial \mathbf{h}_2}, \frac{\partial \mathcal{Q}}{\partial G_2^{-1}}$$

775 and set them to zero to solve for the parameters. Doing so results in the following M-step parameter updates, all of
 776 which can be expressed in closed form.

$$B^{new} = \left(\sum_{i=1}^N \sum_{t=1}^T \text{diag}\{\mathbf{y}_t^{(i)} E_s [\mathbf{c}_t^{(i)'}]\} \right) \left(\sum_{i=1}^N \sum_{t=1}^T \text{diag}\{E_s [\mathbf{c}_t^{(i)} \mathbf{c}_t^{(i)'}]\} \right)^{-1}$$

$$R^{new} = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \left(\text{diag}\{\mathbf{y}_t^{(i)} \mathbf{y}_t^{(i)'}\} - 2 \text{diag}\{B E_s [\mathbf{c}_t^{(i)}] \mathbf{y}_t^{(i)'}\} + \text{diag}\{B E_s [\mathbf{c}_t^{(i)} \mathbf{c}_t^{(i)'}] B'\} \right)$$

$$\Gamma^{new} = \left(\sum_{i=1}^N \sum_{t=2}^T \text{diag}\{E_s [\mathbf{c}_t^{(i)} \mathbf{c}_{t-1}^{(i)'}]\} - \text{diag}\{A E_s [\mathbf{z}_t^{(i)} \mathbf{c}_{t-1}^{(i)'}]\} - \text{diag}\{\mathbf{b} E_s [\mathbf{c}_{t-1}^{(i)'}]\} \right) \left(\sum_{i=1}^N \sum_{t=2}^T \text{diag}\{E_s [\mathbf{c}_{t-1}^{(i)} \mathbf{c}_{t-1}^{(i)'}]\} \right)^{-1}$$

$$A^{new} = \left(\sum_{i=1}^N \sum_{t=2}^T \left(E_s \left[\mathbf{c}_t^{(i)} \mathbf{z}_t^{(i)'} \right] - \Gamma E_s \left[\mathbf{c}_{t-1}^{(i)} \mathbf{z}_t^{(i)'} \right] - \mathbf{b} E_s \left[\mathbf{z}_t^{(i)'} \right] \right) \right) \left(\sum_{i=1}^N \sum_{t=2}^T E_s \left[\mathbf{z}_t^{(i)} \mathbf{z}_t^{(i)'} \right] \right)^{-1}$$

$$\mathbf{b}^{new} = \frac{1}{N(T-1)} \sum_{i=1}^N \sum_{t=2}^T \left(E_s \left[\mathbf{c}_t^{(i)} \right] - \Gamma E_s \left[\mathbf{c}_{t-1}^{(i)} \right] - A E_s \left[\mathbf{z}_t^{(i)} \right] \right)$$

$$Q^{new} = \frac{1}{N(T-1)} \sum_{i=1}^N \sum_{t=2}^T \left(\text{diag}\{E_s \left[\mathbf{c}_t^{(i)} \mathbf{c}_t^{(i)'} \right]\} - 2\text{diag}\{\Gamma E_s \left[\mathbf{c}_{t-1}^{(i)} \mathbf{c}_t^{(i)'} \right]\} - 2\text{diag}\{A E_s \left[\mathbf{z}_t^{(i)} \mathbf{c}_t^{(i)'} \right]\} \right. \\ \left. - 2\text{diag}\{\mathbf{b} E_s \left[\mathbf{c}_t^{(i)'} \right]\} + 2\text{diag}\{A E_s \left[\mathbf{z}_t^{(i)} \mathbf{c}_{t-1}^{(i)'} \right] \Gamma'\} + \text{diag}\{\Gamma E_s \left[\mathbf{c}_{t-1}^{(i)} \mathbf{c}_{t-1}^{(i)'} \right] \Gamma'\} + 2\text{diag}\{\mathbf{b} E_s \left[\mathbf{c}_{t-1}^{(i)'} \right] \Gamma'\} \right. \\ \left. + \text{diag}\{A E_s \left[\mathbf{z}_t^{(i)} \mathbf{z}_t^{(i)'} \right] A'\} + 2\text{diag}\{\mathbf{b} E_s \left[\mathbf{z}_t^{(i)'} \right] A'\} + \text{diag}\{\mathbf{b} \mathbf{b}'\} \right)$$

$$D^{new} = \left(\sum_{i=1}^N \sum_{t=3}^T \left(\text{diag}\{E_s \left[\mathbf{z}_t^{(i)} \mathbf{z}_{t-1}^{(i)'} \right]\} \right) \right) \left(\sum_{i=1}^N \sum_{t=3}^T \text{diag}\{E_s \left[\mathbf{z}_{t-1}^{(i)} \mathbf{z}_{t-1}^{(i)'} \right]\} \right)^{-1}$$

$$P^{new} = \frac{1}{N(T-2)} \sum_{i=1}^N \sum_{t=3}^T \left(\text{diag}\{E_s \left[\mathbf{z}_t^{(i)} \mathbf{z}_t^{(i)'} \right]\} - 2\text{diag}\{D E_s \left[\mathbf{z}_{t-1}^{(i)} \mathbf{z}_t^{(i)'} \right]\} + \text{diag}\{D E_s \left[\mathbf{z}_{t-1}^{(i)} \mathbf{z}_{t-1}^{(i)'} \right] D'\} \right)$$

$$\boldsymbol{\mu}_1^{new} = \frac{1}{N} \sum_{i=1}^N E_s \left[\mathbf{c}_1^{(i)} \right]$$

$$V_1^{new} = \frac{1}{N} \sum_{i=1}^N \text{Var} \left[\mathbf{c}_1^{(i)} \right] + \frac{1}{N} \sum_{i=1}^N \left(E_s \left[\mathbf{c}_1^{(i)} \right] - \boldsymbol{\mu}_1 \right) \left(E_s \left[\mathbf{c}_1^{(i)} \right] - \boldsymbol{\mu}_1 \right)'$$

$$\mathbf{h}_2^{new} = \frac{1}{N} \sum_{i=1}^N E_s \left[\mathbf{z}_2^{(i)} \right]$$

$$G_2^{new} = \frac{1}{N} \sum_{i=1}^N \text{Var} \left[\mathbf{z}_2^{(i)} \right] + \frac{1}{N} \sum_{i=1}^N \left(E_s \left[\mathbf{z}_2^{(i)} \right] - \mathbf{h}_2 \right) \left(E_s \left[\mathbf{z}_2^{(i)} \right] - \mathbf{h}_2 \right)'$$