

1 Biologically informed cortical models predict optogenetic 2 perturbations

3 Christos Sourmpis^{1,2}, Carl C.H. Petersen², Wulfram Gerstner¹, and
4 Guillaume Bellec^{1,3}

5 ¹Laboratory of Computational Neuroscience, Brain Mind Institute, School of
6 Computer and Communication Sciences and School of Life Sciences, École
7 Polytechnique Fédérale de Lausanne (EPFL)

8 ²Laboratory of Sensory Processing, Brain Mind Institute, School of Life
9 Sciences, École Polytechnique Fédérale de Lausanne (EPFL)

10 ³Machine Learning Research Unit, Technical University of Vienna (TU Wien)

11 Abstract

12 A recurrent neural network fitted to large electrophysiological datasets may help us un-
13 derstand the chain of cortical information transmission. In particular, successful network
14 reconstruction methods should enable a model to predict the response to optogenetic pertur-
15 bations. We test recurrent neural networks (RNNs) fitted to electrophysiological datasets on
16 unseen optogenetic interventions, and measure that generic RNNs used predominantly in the
17 field generalize poorly on these perturbations. Our alternative RNN model adds biologically
18 informed inductive biases like structured connectivity of excitatory and inhibitory neurons,
19 and spiking neuron dynamics. We measure that some biological inductive biases improve
20 the model prediction on perturbed trials in a simulated dataset, and a dataset recorded in
21 mice *in vivo*. Furthermore, we show in theory and simulations that gradients of the fitted
22 RNN can be used to target micro-perturbations in the recorded circuits, and discuss the
23 potential utility to bias an animal’s behavior and study cortical circuit mechanisms.

24 1 Introduction

25 A fundamental question in neuroscience is how cortical circuit mechanisms drive percep-
26 tion and behavior. To tackle this question, experimental neuroscientists have been collect-
27 ing large-scale electrophysiology datasets under reproducible experimental settings (Siegle
28 et al., 2021; Esmaeli et al., 2021; Urai et al., 2022; International Brain Laboratory et al.,
29 2023). However, neuroscience lacks data-grounded modeling approaches to generate and
30 test hypotheses on the causal role of neuronal and circuit-level mechanisms. To leverage the
31 high information density of contemporary recordings, we need both (1) modeling approaches
32 that scale well with data, and (2) metrics to quantify when the models provide a plausible
33 mechanism for the observed phenomena.

34 Biophysical simulations have been crucial for our understanding of single-cell mech-
35 anisms (Hodgkin, 1958), and have been used to describe interactions across cortical layers,
36 columns, and areas (Markram et al., 2015; Billeh et al., 2020; Isbister et al., 2023; Chen
37 et al., 2022; Rimehaug et al., 2023; Fraile et al., 2023; Spieler et al., 2023). A promising
38 approach to constrain models to electrophysiological data lies in the optimization of the sim-
39 ulation parameters by gradient descent. These methods were successful in quantitatively
40 classifying functional cell types (Pozzorini et al., 2015; Teeter et al., 2018), and modeling

41 micro-circuit interactions (Pillow et al., 2008; Deny et al., 2017; Mahuas et al., 2020). To
42 bridge the gap from single neurons or small retinal networks to cortical recordings in vivo,
43 recent studies made substantial progress towards data-constrained recurrent neural network
44 (RNN) models (Perich et al., 2020; Bellec et al., 2021; Arthur et al., 2023; Valente et al.,
45 2022; Kim et al., 2023; Dinc et al., 2023; Sourmpis et al., 2023; Pals et al., 2024). In this
46 line of work, neurons in the RNN are mapped one-to-one to recorded cells and optimized
47 by gradient descent to predict recorded activity at large scale.

48 An important question is whether these data-constrained RNNs can reveal a truthful
49 mechanism of neuronal activity and behavior. By construction, the RNNs can generate
50 brain-like network activity, but how can we measure whether the reconstructed network
51 faithfully represents the biophysical mechanism? To answer this question, we submit a range
52 of RNN reconstruction methods to a difficult *perturbation test*: we measure the similarity
53 of the network response to unseen perturbations in the RNN and the recorded biological
54 circuit.

55 Optogenetics is a powerful tool to induce precise causal perturbations in vivo (Esmaeili
56 et al., 2021; Guo et al., 2014). It involves the expression of light-sensitive ion channels
57 (Boyden et al., 2005), such as channelrhodopsins, in specific populations of neurons (e.g.,
58 excitatory/pyramidal or inhibitory/parvalbumin-expressing). In this paper, we use datasets
59 including both dense electrophysiological recordings and optogenetic perturbations to eval-
60 uate RNN reconstruction methods. Since the neurons in our RNNs are mapped one-to-one
61 to the recorded cells, we can model optogenetic perturbations targeting the same cell-types
62 and areas as done in vivo. Yet, we observe that the similarity between the simulated and
63 recorded perturbations varies greatly depending on the reconstruction methods.

64 Most prominently, we study two opposite types of RNN specifications. First, as a con-
65 trol model, we consider a traditional sigmoidal RNN (σ RNN) which is arguably the most
66 common choice for contemporary data-constrained RNNs (Perich et al., 2020; Arthur et al.,
67 2023; Pals et al., 2024); and second, we develop a model with biologically informed inductive
68 biases (bioRNN): (1) neuronal dynamics follow a simplified spiking neuron model, and (2)
69 neurons associated with fast-spiking inhibitory cells have short-distance inhibitory projec-
70 tions (other neurons are excitatory with both local and long-range interareal connectivity).
71 Following Neftci et al. (2019); Bellec et al. (2018b, 2021); Sourmpis et al. (2023), we adapt
72 gradient descent techniques to optimize the bioRNN parameters of neurons and synapses to
73 explain the recorded neural activity and behavior.

74 Strikingly, we find that the bioRNN is more robust to perturbations than the σ RNN.
75 This is nontrivial because it is in direct contradiction with other metrics often used in the
76 field: the σ RNN simulation achieves higher similarity with unseen recorded trials before per-
77 turbation, but lower than the bioRNN on perturbed trials. This contradiction is confirmed
78 both on synthetic and in vivo datasets. To analyze this result, we submit a spectrum of in-
79 termediate bioRNN models to the same *perturbation tests*, and identify two bioRNN model
80 features that are most important to improve robustness to perturbation: (1) Dale's law (the
81 cell type constrains the sign of the connections (Eccles, 1976)), and (2) local-only inhibi-
82 tion (inhibitory neurons do not project to other cortical areas). In contrast, other model
83 features are penalizing, or do not improve significantly the prediction of the optogenetically
84 perturbed response in this out-of-distribution fashion. It indicates that perturbation tests
85 can validate biophysical modeling strategies in data-constrained deep learning models of
86 neural mechanisms.

87 Beyond the optogenetic area inactivation available in the in vivo dataset, we investigate
88 how perturbation-robust RNNs could enable targeted optogenetic protocols for the discovery
89 of detailed neuronal circuit mechanisms in future experiments. Targeted causal interventions
90 will become decisive in studying smaller circuit mechanisms. Acute optogenetic inactivations
91 of genetically defined laminar sub-populations were used to characterize the causal role of
92 specific neurons in the sensory motor pathways (Tamura et al., 2025; Wyart et al., 2025),
93 and upcoming technology will make these experiments easier (Lakunina et al., 2025). To
94 illustrate how RNN reconstruction can help to target neuronal stimulation, we consider

95 micro-perturbations (μ -perturbation) targeting dozens of neurons in a small time window.
96 Inspired by recent read-write all-optical setups (Packer et al., 2015), we imagine a model-
97 informed μ -perturbation protocol, where neurons are targeted based on their functional
98 rather than genetic properties. While previous work has used linear models to produce
99 targeted stimulations (Wyart et al., 2025; Minai et al., 2024), we show that back-propagated
100 gradients of perturbation-robust RNNs provide a sensitivity map to predict the effect of
101 μ -perturbations. Concretely, in a closed-loop experimental setup in silicon, we can use
102 RNN gradients to target a μ -perturbation and change the movement in a simulated mouse.
103 The gradients are used to identify the few neurons having the strongest causal effect on
104 behavior. Conceptually, it means that our RNN reconstructions enable an estimation of
105 “circuit gradients”, bringing numerical and theoretical concepts from deep learning (LeCun
106 et al., 2015; Richards and Kording, 2023) to study biological network computation.

107 2 Results

108 2.1 Reconstructed Networks: biological inductive biases 109 strengthen robustness to perturbations

110 **Synthetic dataset for challenging causal inference** We build a toy synthetic dataset
111 to formalize how we intend to reverse engineer the mechanism of a recorded circuit using op-
112 togenetic perturbations and RNN reconstruction methods. It also serves as the first dataset
113 to evaluate our network reconstruction methods. This toy example represents a simplified
114 version of large-scale cortical recordings from multiple brain areas during a low-dimensional
115 instructed behavior (Steinmetz et al., 2019; Esmaeili et al., 2021; International Brain Lab-
116 oratory et al., 2023), similarly to the in vivo dataset of a GO/No-Go task Esmaeili et al.
117 (2021) analyzed in the next section. Let’s consider two areas A and B which are either
118 transiently active together (“hit trial” occurring with frequency p) or quiescent together
119 (“miss trial” occurring with probability $1 - p$). Since the two areas are active or in-active
120 together, it is hard to infer if they are connected in a feedforward or recurrent fashion. In
121 Methods 4.1, we describe a theoretical example where it is impossible to decide between
122 opposing mechanistic hypothesis (feedforward or recurrent) when recording only the macro-
123scopic activations of areas A and B . In this case, performing optogenetic inactivation of one
124 area is decisive to distinguish between the feedforward or recurrent hypothesis.

125 To generate artificial spike train recordings that capture this problem, we design two
126 reference circuits (RefCircs) from which we can record the spike trains. Each RefCirc consist
127 of two populations of 250 spiking neurons (80% are excitatory) representing areas A and
128 B . To highlight the importance of optogenetic perturbations as in the Methods 4.1, the
129 first circuit RefCirc1 is feedforward and the second RefCirc2 is recurrent: RefCirc1 (and
130 not RefCirc2) has strictly zero feedback connections from B to A . Yet, the two RefCircs
131 are almost identical without optogenetic perturbations: each neuron in RefCirc1 has been
132 constructed to have an almost identical trial-averaged activity as the corresponding neuron
133 in RefCirc2; and in response to a stimulus, the circuits display a similar bi-modal hit-or-
134 miss response with a hit trial frequency $p \approx 50\%$. We consider that a trial is a hit if area
135 A is active (averaged firing rate above 8Hz)¹. To simulate optogenetic inactivations of an
136 area in the RefCircs, we inject a transient current into the inhibitory neurons, modeling the
137 opening of light-sensitive ion channels². Figure S1 shows that optogenetic perturbations
138 in area B reflect the presence or absence of feedback connections which differs in RefCirc
139 1 and 2. Methods 4.4 provides more details on the construction of the artificial circuits.
140 Our *perturbation test* will consist of the comparison of optogenetic perturbations in the

¹Defining a “hit” trial based on area A is equivalent to saying that both areas need to be active during unperturbed trials with this dataset. But excluding B in this definition avoids that the hit rate is trivially impacted when manipulating the activity of B with optogenetic perturbations.

²Symmetrically, an optogenetic activation is simulated as a positive current injected into excitatory cells.

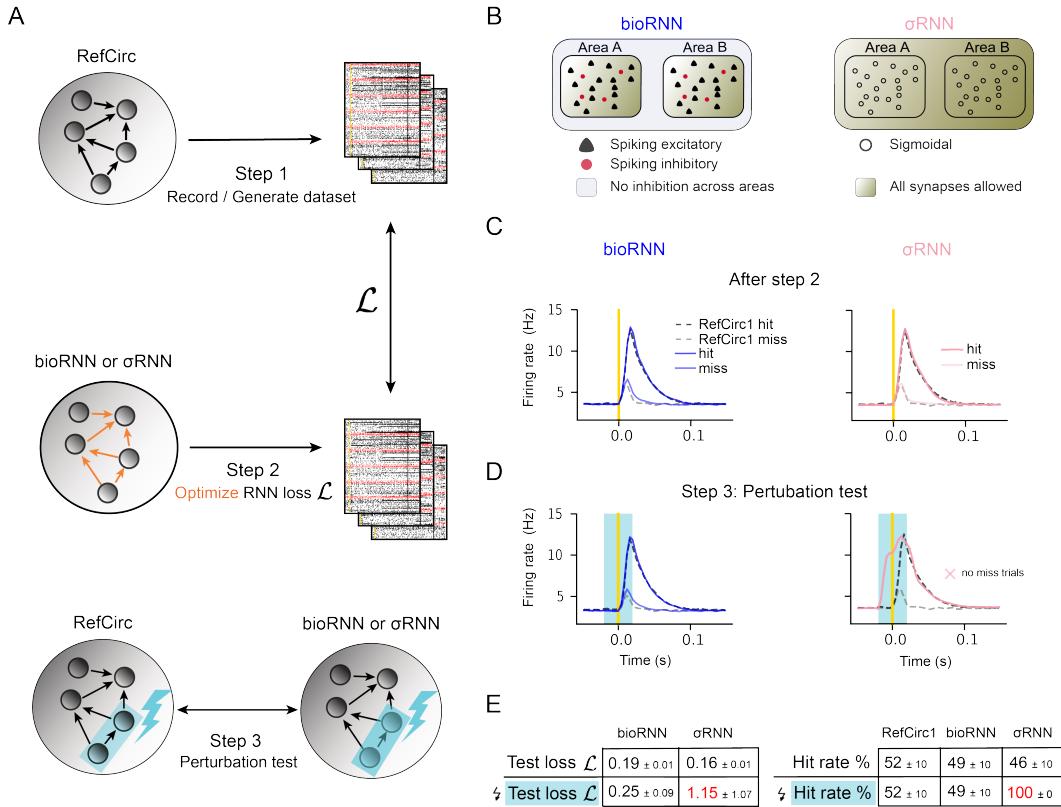


Figure 1: Network reconstruction and perturbation tests. **A.** The three steps to reconstruct the reference circuit (RefCirc) using a biologically informed RNN (bioRNN) or a sigmoidal RNN (σ RNN) and evaluate the reconstruction based on perturbation tests. **B.** Summary of the differences between a bioRNN and a σ RNN. **C.** Trial-averaged activity of area A of the two circuits during hit (black-dashed: RefCirc1; blue: bioRNN1; pink: σ RNN1) and miss (grey-dashed: RefCirc1; light blue: bioRNN1; light pink: σ RNN1) trials. All models display a hit rate of $p \approx 50\%$. **D.** Same as **C** during inactivation of area B. $\Delta p^D = 0$ is the recorded change of hit rate for the feedforward circuit RefCirc1, so a successful reconstruction achieves $\hat{p} \approx 0\%$. **E.** Quantitative results on perturbation tests showing that σ RNN achieves the lowest loss function on the unperturbed test trials, but only the bioRNN retains an accurate fit to the perturbed trials.

141 reconstructed RNN and in their references RefCirc1 and 2, without retraining the RNN on
142 these perturbations.

143 **Network reconstruction methodology (synthetic dataset)** To reconstruct the
144 recorded circuits with an RNN, we record activity from the spiking RefCircs, and opti-
145 mize the parameters of an RNN to generate highly similar network activity. The whole
146 reconstruction method is summarized graphically in panel A of Figure 1. In the simplest
147 cases, the RNN is specified as a sigmoidal network model (Rosenblatt, 1960; Elman, 1990):
148 σ RNN1 and σ RNN2 are optimized to reproduce the recording from RefCirc1 and RefCirc2
149 respectively. In this synthetic dataset, the reconstructed σ RNNs have the same size as the
150 RefCircs (500 neurons) and sigmoidal neurons are mapped one-to-one with RefCirc neu-
151 rons (20% are mapped to inhibitory RefCirc neurons). They are initialized with all-to-all
152 connectivity and are therefore blind to the structural difference of the RefCirc1 and 2 (feed-
153 forward or recurrent). From each of the two RefCircs, we store 2,000 trials of simultaneous
154 spike train recordings of all 500 neurons (step 1 in Fig. 1A). Half of the trials are used
155 as training set and will be the basis for our data-driven RNN optimization. The second
156 half of the recorded trials form the testing set and are used to evaluate the quality of the
157 reconstruction before perturbations.

158 We optimize the synaptic “weights” of the σ RNN to minimize the difference between
159 its activity and that of the RefCirc (step 2 in Fig 1A, see Methods). The optimization
160 combines three loss functions defined mathematically in Methods 4.5: (i) the neuron-specific
161 loss function $\mathcal{L}_{\text{neuron}}$ is the mean-square error of the *trial-averaged* neural activity (e.g.
162 the PSTH) between σ RNN and RefCirc neurons. (ii) To account for fluctuations of the
163 single-trial network activity, we use a trial-specific loss function $\mathcal{L}_{\text{trial}}$, which is the distance
164 between the distribution of single trial population-averaged activity of σ RNN and RefCirc
165 (see Sourmpis et al. (2023)). (iii) Finally, we add a regularization loss function \mathcal{L}_{reg} to
166 penalize unnecessarily large weights.

167 We also developed a biologically informed RNN model (bioRNN) for which we have
168 designed a successful optimization technique. The main differences between σ RNNs and
169 bioRNNs consist in the following biological inductive biases. Firstly, the bioRNN neuron
170 model follows a simplified leaky integrate and fire dynamics (see Methods 4.2) yielding
171 strictly binary spiking activity. Secondly, we constrain the recurrent weight matrix to de-
172 scribe cell-type specific connectivity constraints: following Dale’s law, neurons have either
173 non-negative or non-positive outgoing connections; moreover, since cortical inhibitory neu-
174 rons rarely project across areas, we assume that inhibitory neurons project only locally
175 within the same area. Thirdly, we add a term to the regularization loss \mathcal{L}_{reg} to imple-
176 ment the prior knowledge that cross-area connections are more sparse than within an area.
177 Adding these biological features into the model requires an adapted gradient descent algo-
178 rithm and matrix initialization strategies (Methods 4.5). The reconstruction method with
179 σ RNNs and bioRNNs is otherwise identical: the models have the same size, and are opti-
180 mized on the same data, for the same number of steps and using the same loss functions.
181 The two models bioRNN1 and bioRNN2 are optimized to explain recordings from RefCirc1
182 and RefCirc2, respectively. Importantly, the structural difference between RefCirc1 (feedfor-
183 ward) and RefCirc2 (feedback) is assumed to be unknown during parameter optimization:
184 at initialization, excitatory neurons in bioRNN1 or bioRNN2 project to any neuron in the
185 network with transmission efficacies (aka as synaptic weights) initialized randomly.

186 After parameter optimization, we have four models, σ RNN1, σ RNN2, bioRNN1 and
187 bioRNN2, that we call “reconstructed” models. To validate the reconstructed models, we
188 verify that the network trajectories closely match the data on the test set in terms of (i) the
189 “behavioral” hit-trial frequency, (ii) the peristimulus time histogram (PSTH) mean-square
190 error of single neurons as evaluated by $\mathcal{L}_{\text{neuron}}$, and (iii) the distance between single-trial
191 network dynamics as evaluated by $\mathcal{L}_{\text{trial}}$ (see Suppl. Fig. S2 and Table 1). At first sight, the
192 σ RNN displays a better data fitting when comparing with the non-perturbed trials of the
193 testing set: $\mathcal{L}_{\text{trial}}$ is for instance lower with σ RNN (see Table 1). This is expected considering

194 that the optimization of bioRNNs is less flexible and numerically efficient because of the
195 sign-constrained weight matrix and the imperfect surrogate gradient approximation through
196 spiking activity. However, the two bioRNNs are drastically more robust when evaluating
197 the models with *perturbation tests*.

198 **Perturbation test** To test which of the reconstructed RNNs capture the causal mechanisms
199 of the RefCircs, we simulate optogenetic activations and inactivations of area *B*
200 (step 3 in Fig. 1A). We first compare the change of hit probability after perturbations in
201 the reconstructed RNN ($\Delta\hat{p}$) and recorded in RefCirc (Δp^D) in Figure 2. For the σ RNN
202 the activation or inactivation of area *B* changes drastically the peak firing rate in area *A*:
203 all trials become a hit during inactivation of area *B*. This drastic increase of hit rate is
204 not consistent with the reference where the effect of the optogenetic inactivations is mild: the
205 distribution of network responses remains bi-modal (hit versus miss) with only a moderate
206 change of hit frequency for RefCirc2 $\Delta p^D = -3\%$. For RefCirc1 we even expect $\Delta p^D = 0\%$
207 by design because of the absence of feedback connections from *B* to *A*. In contrast, the
208 bioRNN models capture these changes more accurately (see Fig. 1 and 2). Quantitative
209 results are summarized in Fig. 2E, the error of hit probability changes $|\Delta p^D - \Delta\hat{p}|$ is 7%
210 with bioRNNs when averaged over all conditions (bioRNN1 and bioRNN2, with optogenetic
211 inactivations and activations). The corresponding error is 48.5% on average for σ RNNs.
212 In this sense, we argue that the bioRNN provides a better prediction of the perturbed hit
213 frequency than the σ RNN. We also performed spike train recordings in the area that is
214 not directly targeted by the light to compare the perturbed network dynamics in the fitted
215 RNNs and the RefCirc. The perturbed dynamics are displayed in Fig. 2D. The quantity
216 $\mathcal{L}_{\text{trial}}^{\text{light}}$ is a distance between the network dynamics (RNN versus reference) and is reported
217 in Fig. 2D-E and Table 1. Again, the perturbed dynamics of the bioRNN are more similar
218 to those of the reference circuits $\mathcal{L}_{\text{trial}}^{\text{light}} = 0.26$, than with the σ RNN $\mathcal{L}_{\text{trial}}^{\text{light}} = 1.19$ (t-test
219 p-value is 0.0003).

220 To analyze which features of bioRNN explain this robustness to perturbation, we then
221 derive a family of models where only one feature of the reconstruction is omitted. Namely,
222 the “No Dale’s law” model does not have excitatory and inhibitory weight constraints, the
223 “Non-local inhibition” model allows inhibitory neurons to project outside of their areas, the
224 “No Spike” model replaces the spiking dynamics with a sigmoidal neuron model, and the “No
225 Sparsity” model omits the cross-area sparsity penalty in \mathcal{L}_{reg} . Omitting all these features
226 in bioRNN would be equivalent to using a σ RNN. The accuracy metrics on the testing sets
227 before perturbation are reported for all RNN variants Fig. 2E and G. For reference, we also
228 include the model “No TM” (trial-matching), which omits the loss function $\mathcal{L}_{\text{trial}}$ during
229 training.

230 The strongest effect measured with this analysis is that the Dale’s law and local inhibition
231 explain most of the improved robustness of bioRNNs. This is visible in Fig. 2 as the
232 perturbed trajectories of “No Dale’s law” and “Non-local inhibition” are most distant from
233 the reference in Fig. 2D. This is confirmed numerically where both the hit-rate error and the
234 distance of network dynamics increase the most when lifting these constraints (Fig. 2E-G
235 and Table 1). We explain this result as follows: the mono-synaptic effect of a cell stimulated
236 by the light are always correct in bioRNN (according to Dale’s law, and inhibition locality),
237 but often wrong in the alternative models (see Fig. 2A). For instance, a simple explanation
238 may justify the failure of the “Non-local inhibition” model: the stimulation of inhibitory
239 neurons in *B* induces (via the erroneous mono-synaptic inhibition) a reduction in the base-
240 line activity in area *A* (see the green trace during inactivation in Fig. 2D). More generally
241 for *perturbation testing*, we speculate that these features are measured to be important here
242 because they are central to the biophysical nature of the perturbation considered: optoge-
243 netic perturbation targets specific cell types, and these features incorporate a biophysical
244 connectivity priors which is hard to infer entirely from the unperturbed data.

245 Not all the biological features that we implemented in bioRNN made comparable im-
246 provements in the prediction of optogenetic perturbations. We implemented simple spiking

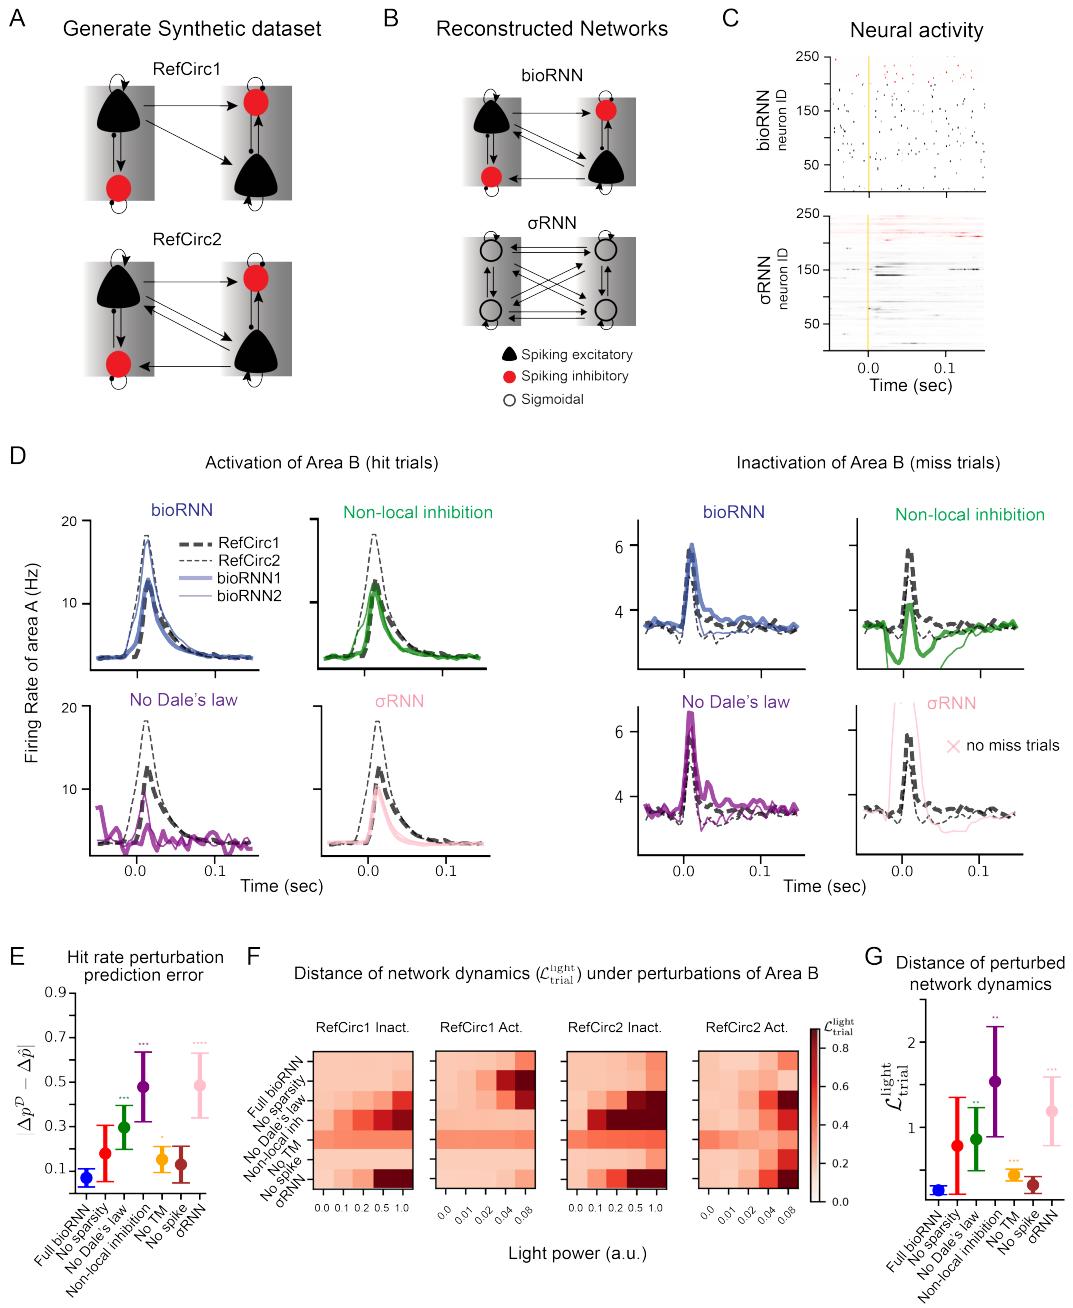


Figure 2: Reconstruction of network mechanisms. **A.** RefCirc1 is feedforward and RefCirc2 is recurrent. **B.** The fitted RNNs are blind to the structural difference of RefCirc1 and 2 and must infer this from the spiking data. **C.** Raster plot showing an example trial of the bioRNN and σ RNN models, neurons in red are mapped to inhibitory neurons. **D.** To study which model feature matters, bioRNN variants are defined by removing one of the features, for instance “No Dale’s law” refers to a bioRNN without weight sign constraints. Trial-averaged activity in area A under activation/inactivation of area B . All the RNNs are tested with the same reference circuit and training data (No spike and No Sparsity models are shown in Suppl. Fig. S4). **E.** Error between the change of hit probability after perturbations in the RNN $\Delta\hat{p}$ and in the RefCirc Δp^D . **F.** The distance of network dynamics $\mathcal{L}_{\text{trial}}^{\text{light}}$ between each RNN and RefCirc (horizontal axis: light power in arbitrary units). **G.** Same quantity as **D** but averaged for each RNN under the strongest light power condition (averaging activations and inactivations of area B). Statistical significance in comparison with bioRNN is computed using the mean over multiple network initializations and is indicated with 0 to 4 stars corresponding to p-values thresholds: 0.05 , 10^{-2} , 10^{-3} and 10^{-4} .

	RefCirc1 vs. RNN1		RefCirc2 vs. RNN2	
	no light	light ↘	no light	light ↘
bioRNN	0.19 ± 0.01	0.25 ± 0.09	0.18 ± 0.01	0.28 ± 0.13
σ RNN	0.16 ± 0.01	1.15 ± 1.07	0.17 ± 0.01	1.22 ± 0.64
No sparsity	0.20 ± 0.01	1.37 ± 1.42	0.19 ± 0.01	0.19 ± 0.13
Non-local inhibition	0.20 ± 0.02	0.54 ± 0.42	0.18 ± 0.01	1.19 ± 0.91
No Dale’s law	0.18 ± 0.01	0.86 ± 0.23	0.18 ± 0.01	2.21 ± 1.60
No spike	0.17 ± 0.00	0.19 ± 0.04	0.18 ± 0.00	0.46 ± 0.19
No Trial Matching (TM)	0.33 ± 0.01	0.44 ± 0.19	0.35 ± 0.03	0.44 ± 0.09

Table 1: **BioRNN is more robust to optogenetic perturbations than σ RNN**. The table reports the trial matching (TM) loss $\mathcal{L}_{\text{trial}}$ on test trials, it measures the distance between the distributions of single trial network dynamics Sourmpis et al. (2023) in area A when stimulating area B. Column “no light” indicates values on the unperturbed test trials, and “light” the perturbation trials. ± indicates the 95% confidence interval, best values are shown in bold and major failure with distance above 0.5 is in red.

247 neuron dynamics and fitted the spiking network as any other RNN using surrogate gradients
 248 Neftci et al. (2019) as in Bellec et al. (2021); Sourmpis et al. (2023). On perturbed
 249 data, the spiking bioRNN achieves slightly better performance than its “No spike” variant,
 250 but without significant margins, t-test p-value is 0.31 for $\mathcal{L}_{\text{trial}}^{\text{light}}$ (see Table 1 and Fig 2E-G).
 251 We speculate that simulating spikes is not advantageous here, because optogenetic pertur-
 252 bations are relatively broad in space and time, and it might become more relevant for other
 253 perturbation experiments where precise timing matters or at a microcircuit level. The quan-
 254 titative of sparse connectivity regularization also did not make a significant improvement in
 255 all cases (see Table 1).

256 Besides predicting the response to optogenetic perturbations, we wondered if we could
 257 recover the connectivity structure of the recorded circuit. Our method would not be appro-
 258 priate to recover individual synaptic connections, but we tested whether the fitted RNNs
 259 reflected the optogenetic signature of the structural difference between the “feedforward”
 260 RefCirc1 and the “recurrent” RefCirc2. Our criteria is here qualitative: the early increase in
 261 the PSTH response in area A characteristic of mono-synaptic feedback from area B should
 262 not exist for RefCirc1 (see Figure 2 and Suppl. Fig. S4A). To reveal this difference in the
 263 fitted bioRNN1 and bioRNN2 models, not only Dale’s law and local-inhibition are necessary,
 264 but also spiking dynamics and sparsity appear to be helpful. For instance, the erroneous
 265 early onset on the perturbed trial in area A for the “No Sparsity” model is corrected with
 266 the sparsity prior (Suppl. Fig. S4, red versus blue curves). Yet, these are subtle qualita-
 267 tive results that are likely to be less impactful and reproducible than the clear qualitative
 268 improvement obtained with perturbation testing when modeling cell-type connectivity. In-
 269 deed, we will see in the next section that we obtain consistent qualitative perturbation
 270 testing results on the larger in vivo dataset.

2.2 Predicting perturbations on in vivo electrophysiology data

271 To test whether our reconstruction method with biological inductive biases can predict
 272 optogenetic perturbations in large-scale recordings, we used the dataset from Esmaeili et al.
 273 (2021). In this study, water-deprived mice were trained to perform a whisker tactile detection
 274 task. In 50% of the trials (Go trials), a whisker is deflected, followed by a 1-second delay,
 275 after which an auditory cue signals that the mice can lick a water spout to receive a water
 276 reward. In the other 50% of trials (No-Go trials), no whisker deflection occurs, and licking
 277 after the auditory cue results in a penalty with an extended time-out period. While the
 278 mice performed the task, experimenters recorded 6,182 units from 12 areas across 18 mice.
 279 Using this dataset, we focused on the 6 most relevant areas for executing this task (Esmaeili
 280

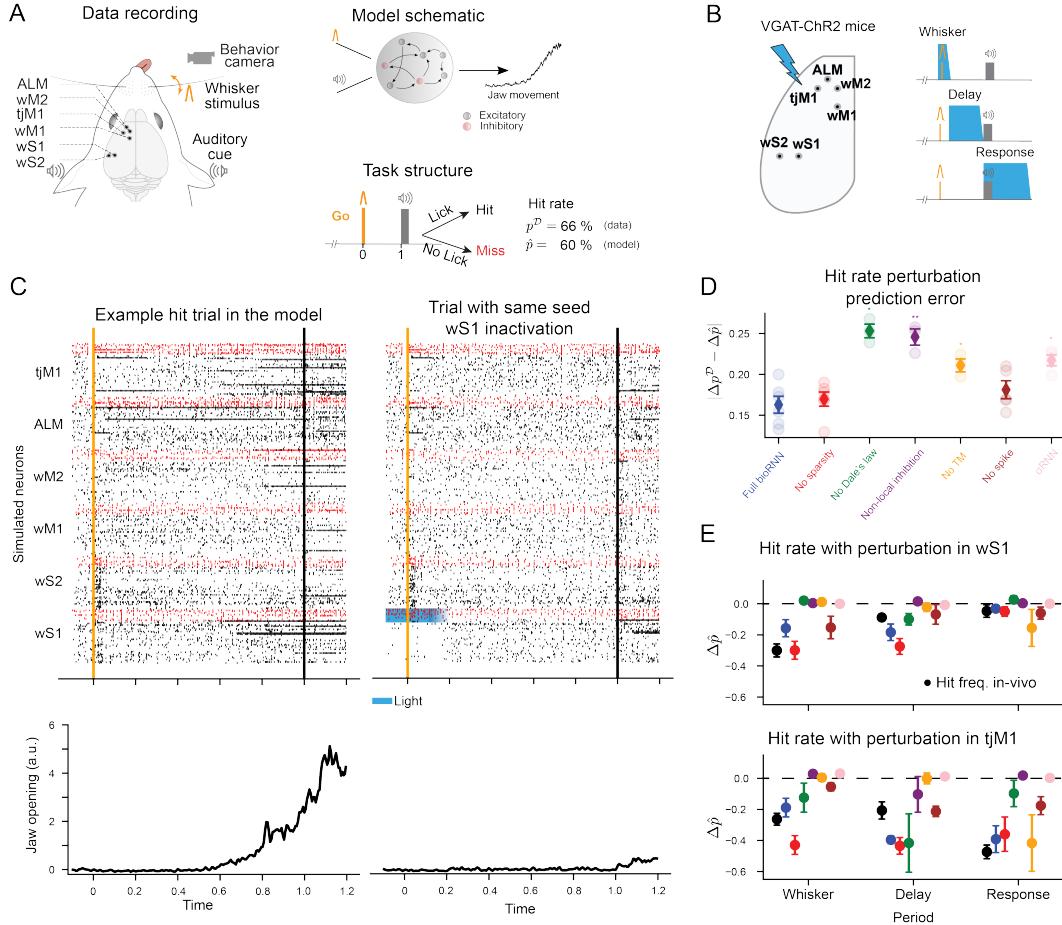


Figure 3: Predicting optogenetic perturbations for in vivo electrophysiology data

A. During a delayed whisker detection task, the mouse reports a whisker stimulation by licking to obtain a water reward. Jaw movements are recorded by a camera. Our model simulates the jaw movements and the neural activity from six areas. **B.** The experimentalists performed optogenetic inactivations of cortical areas (one area at a time) in three temporal windows. **C.** Example hit trial of a reconstructed network (left). Using the same random seed, the trial turns into a miss trial if we inactivate area wS1 (right, light stimulus indicated by blue shading) during the whisker period by stimulation of inhibitory neurons (red dots). **D.** Error of the change in lick frequency caused by the perturbation, $\Delta\hat{p}$ is predicted by the model, and Δp^D is recorded in mice. Light-shaded circles show individual reconstructed networks with different initializations. The whiskers are the standard error of means. **E.** Examples of $\Delta\hat{p}$ hit rate changes under perturbation for wS1 (Top) and tjM1 (Bottom). The black circles refer to the hit rate change from the recordings, Δp^D . See Suppl. Fig. S7 for the other areas.

281 et al., 2021). From each area, we randomly selected 250 neurons (200 putative excitatory
282 and 50 putative inhibitory), which correspond to 1500 neurons in total. These areas, all
283 shown to be causally involved in the task (Esmaeili et al., 2021), include the primary and
284 secondary whisker sensory cortex (wS1, wS2), the primary and secondary whisker motor
285 cortex (wM1, wM2), the anterior lateral motor cortex (ALM), and the primary tongue-jaw
286 motor cortex (tjM1). We fit the neuronal activity using the same reconstruction method as
287 used for the synthetic dataset. In the model, we simulate the jaw movement of the mouse
288 as a linear readout driven by the model’s neural activity. This readout is regressed with
289 the real jaw movement extracted from video footage. The parameter optimization of the
290 behavioral readout is performed jointly with fitting the synaptic weights to the neuronal
291 recordings, see Methods 4.5. After training, our reconstructed model can generate neural
292 activity with firing rate distribution, trial-averaged activity, single-trial network dynamics
293 and behavioral outcome which are all consistent with the recordings (see Suppl. Fig S6).
294 Before perturbations, we observe again that the σ RNN model fits the testing set data better
295 than the bioRNN model (see Table 2 and Fig. S6).

296 We then submit the reconstructed σ RNNs and bioRNNs models to *perturbation tests*.
297 For the sessions of the in vivo dataset with optogenetic perturbation that we considered,
298 only the behavior of an animal is recorded during inactivation of an area at a given time
299 window (stimulus, delay, or choice periods). For each of the six areas and time windows, we
300 extract the averaged hit frequency under optogenetic inactivation, and attempt to predict
301 this perturbed behavior by inducing the same inactivations to the fitted RNNs. These per-
302 turbations are acute spatiotemporal optogenetic inactivations of each area during different
303 time periods (see Figure 3B). As an example, we show the effect of an inactivation of wS1
304 during the whisker period in the model in Fig. 3. In panel C, we display the simulated
305 trial of a fitted bioRNN with and without perturbations side by side. The two trials are
306 simulated with the same random seed, and this example shows that an early perturbation
307 in wS1 can change a lick decision from hit to miss in the model (Fig. 3C).

308 Consistent with the synthetic dataset, we now find with this in vivo dataset that modeling
309 cell-type connectivity yields better prediction of the causal effect of optogenetic perturbation.
310 We denote by Δp^P the in vivo change in lick probability across Go trials in response to
311 optogenetic perturbations. The perturbations were performed in different periods for each
312 area in Esmaeili et al. (2021) (stimulation, delay, or choice periods). For all areas and time
313 windows, we measure the corresponding $\Delta \hat{p}$ in the model. On average, the error change
314 probability obtained with the σ RNN model is $|\Delta p^P - \Delta \hat{p}| = 21\%$ which is significantly
315 worse than the bioRNN model’s 16% (t-test p-value is 0.014, see Figure 3D). As in the
316 synthetic dataset, we find this to be consistent over multiple bioRNN model variants, and
317 we find that imposing Dale’s law and local inhibition best explain the improvement in
318 perturbation-robustness. We also measure that the spiking bioRNN predicts the change in
319 lick probability slightly better than the “No Spike” bioRNN model. Conversely, adding the
320 sparsity prior does not seem to improve the perturbed hit-rate prediction on the real data as
321 seen in the recurrent artificial dataset (RefCirc2) and not in the feedforward case (RefCirc1)
322 as shown in Suppl. Fig. S5. In this sense, in vivo perturbation testing emerges as a hard
323 test to evaluate modeling strategies combining deep learning and biophysical modeling.

324 To further analyze the consistency of the perturbations in the model, we can compare
325 the perturbation map showing changes in lick probability obtained from acute inactivation
326 in the data and the model. The Suppl. Fig. S7 summarizes visually which area has a critical
327 role at specific time points. The changes of lick probability in area wS1, ALM and tjM1
328 are accurately predicted by the bioRNN. In contrast, our model tends to underestimate the
329 causal effect induced by the inactivations of wS2, wM1 and wM2 (Suppl Fig S7). Overall,
330 our model is consistent with a causal chain of interaction from wS1 to ALM and continuing
331 to tjM1.

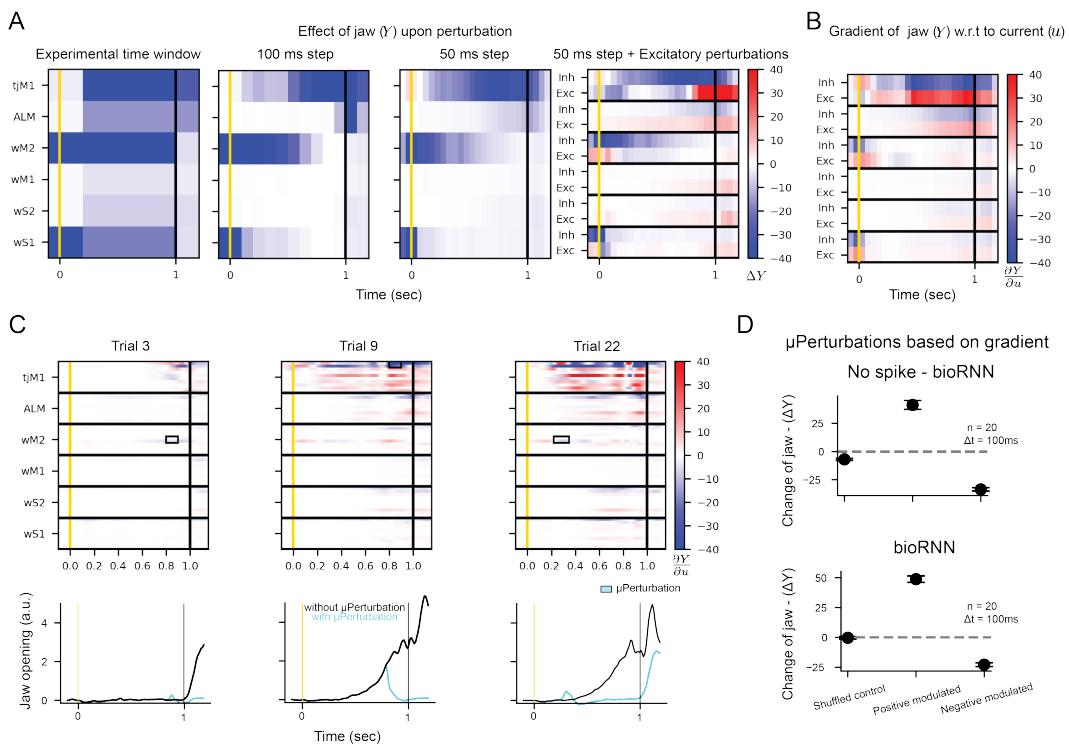


Figure 4: Measuring circuit gradients with μ -perturbations **A-B.** Numerical verification for equation (1). **A** shows the change of jaw movement ΔY following inactivations in a “No Spike” bioRNN. From left to right, we reduce the size of the spatiotemporal window for the optogenetic stimulation. **B**. Gradients values $\sum_{i,t} \frac{\partial Y}{\partial u}$ that approximate ΔY from **A** using equation (1). **C-D**. Verification that gradients predict the change of movement on single trials. In **C**, we display the gradients and jaw movement for three different trials, the neurons targeted by the μ -perturbation are boxed and the perturbed jaw movement is blue. Results averaged for every 100ms stimulation windows are shown in **D**: positive (resp. negative) modulated means that the 20 neurons with highest (resp. lowest) gradients are targeted, random neurons are selected for the shuffled case.

332 2.3 Applications for experimental electrophysiology

333 With future progress in recording technology and reconstruction methods, network recon-
 334 struction may soon predict the effect of optogenetic perturbation with even higher accu-
 335 racy. In this section, we explore possible consequences and applications for experimental
 336 electrophysiology. We demonstrate in the following that (1) perturbation-robust bioRNNs
 337 enable us to estimate gradients of the recorded circuits, (2) which in turn enable us to
 338 target μ -perturbations in the recorded circuit and optimally increase (or decrease) induced
 339 movements in our simulated mouse. The so-called “recorded circuit” is a bioRNN trained
 340 on the in vivo dataset that we use as a proxy experimental preparation. Its mathemati-
 341 cal underpinnings enable us to make rigorous theoretical considerations and the design of
 342 forward-looking in silico experiments.

343 **μ -perturbations measure brain gradients** We first prove a mathematical relationship
 344 between gradients in the recorded circuit and μ -perturbations. We define the integrated
 345 movement as $Y = \sum_t y_t$ where y_t is the movement of the jaw at time t generated by the
 346 model, and we denote ΔY^μ as the change of movement caused by the μ -perturbation. If
 347 the circuit has well-defined gradients (e.g. say a “No spike” bioRNN model trained on the

348 in vivo recordings in the previous section), using a Taylor expansion, we find that:

$$\Delta Y^* = \sum_{i,t \in \mathcal{I}} \frac{dY}{du_i^t} \Delta u_i^t + \epsilon, \quad (1)$$

349 where \mathcal{I} are the neuron and time indices selected for the optogenetic intervention. The
 350 error term ϵ is negligible when the current Δu_i^t induced by the light is small. We first
 351 confirm this approximation with numerical visualization in Fig. 4A: we display movement
 352 perturbations $\langle \Delta Y^* \rangle$ in the circuit with time windows of decreasing sizes ($\langle \cdot \rangle$ indicates a
 353 trial average). When the time window is small, and the perturbation is only applied to
 354 excitatory or inhibitory cells in Fig. 4A, one can appreciate visually the similarity with the
 355 binned gradient $\langle \sum_{i,t} \frac{dY}{du_i^t} \rangle$ in Fig. 4B. Proceeding to a quantitative verification of equation
 356 (1), we now compare the effect of small perturbations targeting only 20 neurons on a single-
 357 trial. We use the gradient $\sum_{i,t} \frac{dY}{du_i^t}$ (see Fig. 4C) to predict the outcome of μ -perturbation
 358 as follows: for each trial, and each 100ms time window, we identify 20 neurons in the
 359 model with highest (or lowest) gradients $\sum_{i,t} \frac{dY}{du_i^t}$. We then re-simulate the exact same trial
 360 with identical random seed, but induce a μ -perturbation on selected neurons (see rectangles
 361 in Figure 4). If we target neurons with strongly positive gradients, the perturbed jaw
 362 movements are strongly amplified $\Delta Y^* > 0$; conversely, if we target neurons with negative
 363 gradients the jaw movements are suppressed $\Delta Y^* < 0$. Although the equation (1) is only
 364 rigorously valid for models with well-defined gradients like the “No Spike” model, we also
 365 confirm in Fig. 4D that this numerical verification also holds in a spiking circuit model
 366 where the gradients are replaced with surrogate gradients (Neftci et al., 2019).

367 An implication of equation (1) is that the measurements $\langle \Delta Y^* \rangle$ that can be recorded in
 368 vivo are estimates of the gradients $\langle \sum_{i,t} \frac{dY}{du_i^t} \rangle$ in the recorded circuit. Yet, measuring detailed
 369 gradient maps (or perturbation maps) as displayed in Fig. 4 would be costly in vivo as it
 370 requires to average ΔY^* over dozens of trials for each spatio-temporal window. Instead,
 371 gradient calculation in a bioRNN model (that was fitted to the experimental preparation) is
 372 a rapid mathematical exercise. If the extracted model is valid, then the gradients $\sum_{i,t} \frac{dY}{du_i^t}$
 373 in the bioRNN approximate (1) the effect of μ -perturbations ΔY^* in the experimental
 374 preparation; (2) the gradient $\sum_{i,t} \frac{dY}{du_i^t}$ in the recorded circuit.

375 **Targeting in vivo μ -perturbations with bioRNN gradients** Building on this theo-
 376 retical finding, we build a speculative experimental setup where the bioRNN gradients are
 377 used to target a μ -perturbation and increase (or decrease) the movements Y in the
 378 experimental preparation in real time. We show a schematic of this speculative closed-loop
 379 experiment in Fig. 5C extending contemporary read-write elecrophysiology setups (Packer
 380 et al., 2015; Adesnik and Abdeladim, 2021; Grosenick et al., 2015; Papagiakoumou et al.,
 381 2020). We demonstrate in silico in Fig. 5A-B how this experiment could use bioRNN gradi-
 382 ents to bias the simulated mouse movement Y . As a preparation step, and before applying
 383 perturbations, we assume that the bioRNN is well fitted to the recorded circuit and we
 384 collect a large databank \mathcal{B} of simulated trials from the fitted bioRNN. Then in real-time,
 385 we record the activity from the experimental preparation until the time t^* at which the
 386 stimulation will be delivered (Step 1 in Fig. 5A, t^* is 100ms before the decision period).
 387 Rapidly, we find the trial with the closest spike trains in the databank of simulated trials
 388 (Step 2) and use the corresponding gradient maps to target neurons with the highest gradi-
 389 ent $\frac{dY}{du}$ in the model (Step 3). The targeted stimulation is then delivered immediately
 390 at t^* to the experimental preparation (Step 4). When testing this in silico on our artificial
 391 experimental preparation, we show in Fig. 5C that this approach can bias the quantity of
 392 jaw movement Y driven by the circuit in a predictable way. The amount of movement is
 393 consistently augmented if we target neurons with the highest $\frac{dY}{du}$ (or diminished if we target
 394 neurons with the lowest $\frac{dY}{du}$).

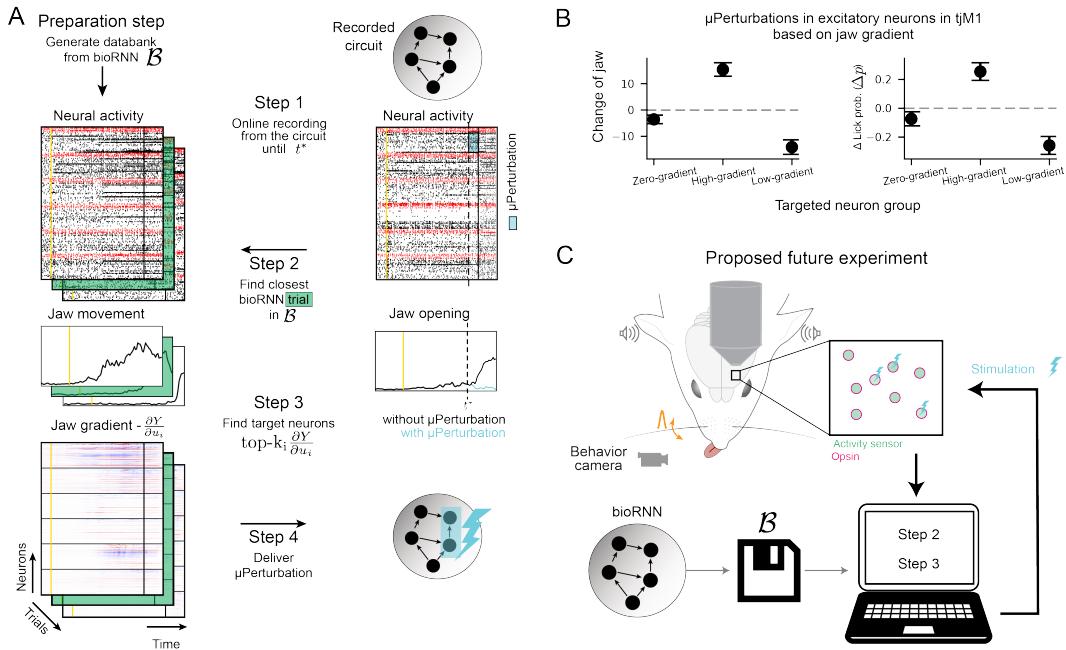


Figure 5: Gradient targeted μ -perturbations could precisely bias an animal behavior **A.** Protocol to deliver an optimal μ -perturbation on the experimental preparation based on jaw gradients. (Step 1) The circuit is recorded until stimulation time t^* . (Step 2) The closest bioRNN trial to the ongoing recorded trial is retrieved from the databank \mathcal{B} . (Step 3) We select the neurons with the highest (or lowest) gradient value for the μ -perturbation. (Step 4) The μ -perturbation is delivered at t^* . **B.** Effect of the μ -perturbation using the artificial setup **A** under different light protocols. Practically, for “High gradient”, we keep step 3 as it is, for “Low gradient”, we change the sign of the gradient, and for “Zero gradient”, we pick the 40 neurons with lowest gradient norm. **C.** Speculative schematic of a close-up setup implementing the protocol **A** inspired by the all optical “read-write” setup from [Aravanis et al. \(2007\)](#); [Packer et al. \(2015\)](#).

395 3 Discussion

396 Finding the right level of detail to model recorded phenomena has sparked intensive de-
397bates in computational neuroscience. When the goal is to achieve the strongest predictive
398 power, generalist deep learning models have proven successful across many scientific disci-
399 plines, questioning how biophysical modeling plays a role in this context. Our results show
400 that *perturbation testing* is a new approach to evaluate the implementation of biophysical
401 features in a deep learning system. Our key finding about *perturbation testing* relies on
402 the difficulty for deep learning models to predict the effect of optogenetic perturbations
403 out-of-distribution (meaning, the perturbed trials are not available in the training set of
404 the data-constrained model). We see that standard deep learning RNNs generalize poorly
405 to perturbed trials, even when they achieved the best fit on the unperturbed test set. In
406 contrast, this is alleviated with our bioRNN, which implements biophysical constraints that
407 are relevant to the nature of the perturbation. In our case, modeling cell type connectivity
408 is crucial because the optogenetic perturbations are targeted to these genetically encoded
409 cell types. In this sense, we believe that these features were successful on the perturbation
410 tests because they are central to modeling the perturbation of the deep learning system.
411 Perturbation testing emerges as a quantitative tool to search for data-constrained models
412 beyond two standard types of incomplete brain models in computational neuroscience: (1)
413 physiologically detailed models intended to explain brain mechanisms but do not enable
414 powerful quantitative predictions; (2) deep learning models with high predictive power but
415 capturing a wrong biophysical mechanism, causing erroneous generalizations. We view our
416 work as a simple and reasonable way to combine deep learning and biophysical modeling,
417 while rigorously evaluating the combined models.

418 Our reconstruction method and modeling choices when building the data-constrained
419 bioRNN are innovative and are validated on *perturbation tests*. We achieve a reconstruction
420 of the sensory-motor pathway in the mouse cortex during a sensory detection task from
421 electrophysiology data. The model is optimized to explain electrophysiological recordings
422 and generalizes better than standard models to *in vivo* optogenetic interventions. We found
423 unambiguously that anatomically informed sign and connectivity constraints for dominant
424 excitatory and inhibitory cell types improve the model robustness to optogenetic perturba-
425 tions. We also find that assuming that inhibitory connections are short and do not project
426 to other areas is crucial to pass our optogenetic *Perturbation test*. Modeling spiking neuron
427 dynamics and adding a sparsity prior yielded more nuanced results and was not decisive,
428 showing that making a difference on *Perturbation testing* is challenging. In hindsight, we
429 conclude that adding biological constraints becomes beneficial when (1) they model the
430 interaction between the circuit and the perturbation mechanism; (2) their implementation
431 should not impair the efficiency of the optimization process.

432 Broadly speaking, this hindsight is also supported by other results elsewhere in neuro-
433 science. For instance, biologically inspired topological networks having higher correlation
434 for neighboring neurons are more consistent with comparable causal interventions in the
435 Monkey's visual system [Schrimpf et al. \(2024\)](#), and detailed cell-type distribution and con-
436 nectome improve models of vision in the fly brain [Lappalainen et al. \(2023\)](#); [Cowley et al.](#)
437 ([2024](#)). For future work, there is a dense knowledge of unexploited physiological data at the
438 connectivity, laminar or cell-type level that could be added to improve a cortical model like
439 ours ([Harris et al., 2019](#); [Liu et al., 2022](#); [Udvary et al., 2022](#); [Staiger and Petersen, 2021](#);
440 [Rimehaug et al., 2023](#)). By submitting the extended models to the relevant *perturbation*
441 *tests*, it becomes possible to measure quantitatively the goodness of their biological mecha-
442 nism implementations. We do not rule out, that significant improvements on *perturbation*
443 *tests* can also be achieved with other means (e.g. by training deep learning architectures [Az-](#)
444 [abou et al. \(2024\)](#); [Pandarinath et al. \(2018\)](#); [Ye et al. \(2023\)](#) on larger datasets to enable
445 generalization, or with generic regularization techniques like low-rank connectivity [Dubreuil](#)
446 [et al. \(2022\)](#); [Valente et al. \(2022\)](#)). However, in a similar way as the σ RNN was apriori
447 an excellent predictor on our initial test-set, any powerful brain model will likely have fail-

448 ure modes that can be well characterized and measured with an appropriate perturbation
449 test. So *perturbation tests* could become a central component of an iterative loop to identify
450 needed data collection or model improvements towards robust brain models.

451 To highlight the importance of perturbation-robust circuit models, we have discussed
452 possible implications for experimental neuroscience in section 2.3. We build the RNN twin
453 of the biological circuit from unperturbed electrode recordings. By implementing the correct
454 biophysical constraints, the RNN becomes perturbation robust (i.e. it predicts the effect of
455 causal perturbation) even without including perturbation data in the RNN training. We
456 then demonstrated in silico that gradients of this RNN produce sensitivity maps to target
457 micro-stimulation of the biological circuit. As a result, we could design a hypothetical closed-
458 loop setup combining read-write electrophysiology with a brain model to influence the brain
459 activity or behavior, having potentially important practical and ethical consequences. More
460 conceptually, we have shown theoretically that the gradients of a perturbation robust RNN
461 are also consistent with the gradients of the recorded biological circuits. In perspective
462 with the foundational role of gradients in machine learning theory LeCun et al. (2015);
463 Richards and Kording (2023), it enables the measurement of “brain gradients” and lays a
464 computational link between in vivo experimental research and decades of theoretical results
465 on artificial learning and cognition.

466 4 Methods

467 4.1 Mathematical toy model of the difficult causal inference be- 468 tween H1 and H2

469 Let's consider two simplistic mathematical models that both depend on two binary random
470 variables A and B which represent that putative area A is active as $A = 1$ and area B as
471 $B = 1$. With this notation, we can construct two hypothetical causal mechanisms $H1$ ("a
472 feedforward hypothesis") and $H2$ ("a recurrent hypothesis"), which are radically different.
473 The empirical frequency $p(A, B)$ of the outcome does not allow us to differentiate whether
474 the system was generated by a feedforward mechanism $H1$ or a recurrent mechanism $H2$.
475 Schematically, we can represent the two mechanism hypotheses as follows:

$$(H1) \quad A \longrightarrow B, \quad (2)$$

$$(H2) \quad A \longleftrightarrow B. \quad (3)$$

476 For hypothesis $H1$: we assume that external inputs are driving the activity of area A such
477 that $A = 1$ is active with probability p_0 , and there are strong feed-forward connections from
478 A to B causing systemically $B = 1$ as soon as $A = 1$. Alternatively, in $H2$, we assume
479 that areas A and B receive independent external inputs with probability $p_1 = 1 - \sqrt{1 - p_0}$.
480 Each of these two inputs is sufficient to cause $A = 1$ or $B = 1$, and the two areas are
481 also strongly connected, so $A = 1$ always causes $B = 1$ and vice versa. Under these
482 hypothetical mechanisms $H1$ and $H2$, one finds that the empirical probability table $p(A, B)$
483 is identical ³: $p_{H2}(A = 1, B = 1) = 2p_1 - p_1^2 = p_0$ ("Hit trial", both areas are active),
484 $p(A = 0, B = 0) = 1 - p_0$ ("Miss trial", the areas are quiescent). In both cases, the
485 possibility that only one area is active is excluded by construction. So for any A and B
486 $p_{H1}(A, B) = p_{H2}(A, B)$ and in other words, even if we observe an infinite number of trials
487 and compute any statistics of the binary activations A and B , discriminating the two possible
488 causal interactions ($H1$ versus $H2$) is impossible.

489 A solution to discriminate between hypotheses $H1$ and $H2$ is to induce a causal pertur-
490 bation. We can discriminate between our two hypotheses if we can impose a perturbation
491 that forces the inactivation of area B in both mathematical models. In mathematical terms
492 we refer to the *do* operator from causality theory. Under the feedforward mechanism $H1$
493 and inactivation of B , A is not affected $p_{H1}(A = 1 \mid \text{do}(B = 0)) = p_0$. Under the recurrent
494 hypothesis, $H2$, and inactivation of B , A is activated only by its external input such that
495 $p_{H2}(A = 1 \mid \text{do}(B = 0)) = p_1 \neq p_0$. So the measurement of the frequency of activation
496 of area A under inactivation of B can discriminate between $H1$ and $H2$ which illustrates
497 mathematically how a causal perturbation can be decisive to discriminate between those
498 two hypothetical mechanisms.

499 4.2 Neuron and jaw movement model

500 We model neurons as leaky-integrate and fire (LIF) neurons. The output of every neuron
501 j at time t is a binary outcome z_j^t (spike if $z_j^t = 1$, no spike if $z_j^t = 0$) generated from its
502 membrane voltage v_j^t . The following equations give the dynamics of the membrane voltage
503 v_j^t :

³To prove this: a and b to denote the binary external inputs into A and B , so we have: $p_{H2}(A = 1, B = 1) = \sum_{a,b} p(A = 1, B = 1 | a, b) p(a, b) = p(a = 1, b = 1) + p(b = 0, a = 1) + p(b = 1, a = 0)$ where we used that $p(A = 1, B = 1 | a, b)$ is 0 or 1, then using $p(a = 1) = p(b = 1) = p_1$ and the independence between a and b we find: $p(A = 1, B = 1) = 2p_1 - p_1^2 = p_0$

$$v_j^t = \alpha_j v_j^{t-1} + (1 - \alpha_j) u_j^t - v_{\text{thr},j} z_j^{t-1} + \xi_j^t \quad (4)$$

$$u_j^t = \sum_{d,i} W_{ij}^{rec,d} \frac{z_i^{t-d}}{\delta t} + \sum_i W_{ij}^{in} \frac{x_i^t}{\delta t} \quad (5)$$

504 where W_{ij}^d , and $W_{ij}^{in,d}$ are the recurrent and input weight matrices. The timestep of the
 505 simulation δt is 2 ms when we simulate the real dataset and 1 ms otherwise. The superscript
 506 d denotes the synaptic delay; every synapse has one synaptic delay of either 2 or 3 ms.
 507 With $\alpha_j = \exp\left(-\frac{\delta t}{\tau_{m,j}}\right)$, we define the integration speed of the membrane voltage, where
 508 $\tau_m = 30$ ms for excitatory and $\tau_m = 10$ ms for inhibitory neurons. The noise source
 509 ξ_j^t is a Gaussian random variable with zero mean and standard deviation $\beta_j v_{\text{thr},j} \sqrt{\delta t}$ (β_j
 510 is typically initialized at 0.14). The input x_i^t is a binary pulse signal with a duration of
 511 10 ms. For the real dataset, we have two binary pulse input signals, one for the whisker
 512 deflection and one for the auditory cue. The spikes are sampled with a Bernoulli distribution
 513 $z_j^t \sim \mathcal{B}(\exp(\frac{v_j^t - v_{\text{thr},j}}{v_0}))$, where v_0 is the temperature of the exponential function and $v_{\text{thr},j}$
 514 is the effective membrane threshold. After each spike, the neuron receives a reset current
 515 with an amplitude of $v_{trh,j}$ and enters an absolute refractory period of 4 ms, during which
 516 it cannot fire.

517 For networks fitted to the real dataset, we also simulate the jaw movement. The jaw
 518 movement trace y is controlled by a linear readout from the spiking activity of all excitatory
 519 neurons. Specifically, y is computed as $y = \exp(\tilde{y}) + b$, where b is a scaling parameter and
 520 \tilde{y}^t is given by $\tilde{y}^t = \alpha_{jaw} \tilde{y}^{t-1} + (1 - \alpha_{jaw}) \sum_{d,j} W_j^{jaw} z_j^{t-d}$. Here, W_j^{jaw} is the output weight
 521 matrix (linear readout) for the jaw, and $\tau_{jaw} = 5$ ms defines $\alpha_{jaw} = \exp(-\frac{\delta t}{\tau_{jaw}})$, which
 522 controls the integration velocity of the jaw trace.

523 4.3 Session-stitching and network structure

524 As in (Sourmpis et al., 2023), we simulate multi-area cortical neuronal activity fitted to
 525 electrophysiology neural recordings. Before we start the optimization, we define and fix
 526 each neuron's area and cell type in the model by uniquely assigning them to a neuron from
 527 the recordings. For the real dataset from Esmaili et al. (2021), the cell type is inferred from
 528 the cell's action potential waveform (with fast-spiking neurons classified as inhibitory and
 529 regular-spiking neurons as excitatory). Most electrophysiology datasets include recordings
 530 from multiple sessions, and our method would typically require simultaneous recordings of all
 531 neurons. To address this challenge, similarly to (Sourmpis et al., 2023) we use the technique
 532 called “session-stitching” which allows neighboring modeled neurons to be mapped with
 533 neurons recorded across multiple sessions. This effectively creates a “collage” that integrates
 534 data from multiple sessions within our model. This approach has practical implications for
 535 our optimization process. Specifically, the trial-matching loss includes a term for each
 536 session, with the overall loss calculated as the average across all sessions (see 4.5).

537 For both the real and the synthetic datasets, we simulate each area with 250 LIF neurons
 538 and impose that each area has 200 excitatory neurons and 50 inhibitory. Respecting the ob-
 539 servation that inhibitory neurons mostly project in the area that they belong to (Tamamaki
 540 and Tomioka, 2010; Markram et al., 2004), we don't allow for across-area inhibitory con-
 541 nections. The “thalamic” input is available to every neuron of the circuit, and the “motor”
 542 output for the real dataset, i.e., jaw movement, is extracted with a trained linear readout
 543 from all the excitatory neurons of the network, see 4.2.

544 4.4 Reference circuits for hypotheses 1 and 2

545 To build a synthetic dataset that illustrates the difficulty of separating the feedforward (H1)
 546 and recurrent hypotheses (H2), we construct two reference spiking circuit models RefCirc1

547 and RefCirc2. The two networks consists of two areas A and B, and their activity follows the
 548 hard causal inference problem from method 4.1, making it hard to distinguish A1 and A2
 549 when recording the co-activation of A and B. Moreover, to make the problem even harder,
 550 the two networks are constructed to make it almost impossible to distinguish between H1
 551 and H2 with dense recordings: the two circuits are designed to have the same PSTH and
 552 single-trial network dynamics despite their structural difference, one is feedforward and the
 553 other is recurrent.

554 To do so, RefCirc1 and 2 are circuit models that start from random network initializations
 555 following the specifications described in Methods 4.2 and 4.3. The only difference is that
 556 we do not allow feedback connections from A to B in RefCirc1, the construction below
 557 is otherwise identical. The synaptic weights of the two circuits are optimized with the
 558 losses described in Methods 4.5 to fit the identical target statistics in all areas: the same
 559 PSTH activity for each neuron and the same distribution of single-trial network dynamics.
 560 The target statistics are chosen so the activity in RefCirc1 and 2 resemble kinematics and
 561 statistics from a primary and a secondary sensory area. The baseline firing rates of the
 562 neurons is dictated by the target PSTH distribution and it follows a log-normal distribution,
 563 with excitatory neurons having a mean of 2.9 Hz and a standard deviation of 1.25 Hz and
 564 inhibitory neurons having a mean of 4.47 Hz and a standard deviation of 1.31 Hz. The
 565 distribution of single-trial activity is given by the targeted single-trial dynamics: in RefCirc1
 566 and 2, the areas A and B respond to input 50% of the time with a transient population
 567 average response following a double exponential kernel characterized by $\tau_{rise} = 5$ ms and
 568 $\tau_{fall} = 20$ ms. Mimicking a short signal propagation between areas, these transients have
 569 a 4 ms delay in area A and 12 ms delay in B (relative to the onset time of the stimulus).
 570 To impose a "behavioral" hit versus miss distribution that could emerge from a feedforward
 571 and recurrent hypothesis (see method 4.1), the targeted population-averaged response of
 572 each trial is either a double-exponential transient in both area A and B ("Hit trials"), or
 573 remains at a baseline level in both areas ("Miss trials") in the remaining trials. At the
 574 end of the training, we verified that RefCirc1 and RefCirc2 generate very similar network
 575 activity in the absence of perturbation (see Figure S1). The circuits are then frozen and
 576 used to generate the synthetic dataset. We generate 2000 trials from these RefCircs, 1000
 577 of which are used for the training set and 1000 for the testing set.

578 4.5 Optimization and loss function

579 The optimization method we use to fit our models is back-propagation through time
 580 (BPTT). To overcome the non-differentiability of the spiking function, we use surrogate
 581 gradients (Neftci et al., 2019). In particular, we use the piece-wise linear surrogate deriva-
 582 tive from Bellec et al. (2018b). For the derivative calculations, we use $\frac{v_j^t - v_{thr,j}}{v_0}$ and not
 583 $\exp(\frac{v_j^t - v_{thr,j}}{v_0})$. We use sample-and-measure loss functions that rely on summary statistics,
 584 as in (Bellec et al., 2021; Sourmpis et al., 2023), to fit the networks to the data. Our loss
 585 function has two main terms: one to fit the trial-averaged activity of every neuron (\mathcal{L}_{neuron}),
 586 and one to fit the single trial population average activity (\mathcal{L}_{trial}), $\mathcal{L} = \mathcal{L}_{neuron} + \mathcal{L}_{trial}$. The
 587 two terms of the loss function are reweighted with a parameter-free multi-task method
 588 (Défossez et al., 2023) that enables the gradients to have comparable scales.

589 As in Sourmpis et al. (2023): (1) To calculate the **trial-averaged loss**, we first filter the
 590 trial-averaged spiking activity $\mathcal{T}_{neuron,j}^t(\mathbf{z}) = \frac{1}{K} \sum_k \mathbf{z}_{j,k}^t * f$ using a rolling average window
 591 (f) of 8 ms. We then normalize it by the trial-averaged filtered data activity, (\mathbf{z}^D are
 592 recorded spike trains)

$$593 \mathcal{T}_{neuron,j}^t(\mathbf{z}) = (\mathcal{T}_{neuron,j}^t(\mathbf{z}) - \langle \mathcal{T}_{neuron,j}^t(\mathbf{z}^D) \rangle_t) / (\sigma_t(\mathcal{T}_{neuron,j}^t(\mathbf{z}^D))), \quad (6)$$

593 where $\langle \cdot \rangle_t$ is the time average, and σ_t the standard deviation over time. The trial-averaged
 594 loss function is defined as:

$$\mathcal{L}_{\text{neuron}} = \sum_j^N \sum_t^T \|\mathcal{T}_{\text{neuron},j}^t(\mathbf{z}) - \mathcal{T}_{\text{neuron},j}^t(\mathbf{z}^D)\|^2, \quad (7)$$

595 where T is the number of time points in a trial and N is the number of neurons. For the
 596 real dataset, where we want to fit also the jaw movement, we have an additional term for
 597 the trial-averaged filtered and normalized jaw, $\|\sum_t^T \mathcal{T}_{\text{neuron}}^t(\mathbf{y}) - \mathcal{T}_{\text{neuron}}^t(\mathbf{y}^D)\|^2$, where \mathbf{y} is
 598 the simulated jaw movement and \mathbf{y}^D the recorded jaw movement.

599 (2) To calculate the **trial-matching loss**, we first filter the population-average activity
 600 of each area A , $\mathcal{T}_{A,k}^t(\mathbf{z}) = \frac{1}{|A|} \sum_{j \in A} \mathbf{z}_{j,k}^t * f$, using a rolling average window of 32 ms. We
 601 then normalize it by the population-averaged filtered activity of the same area from the
 602 recordings, $\mathcal{T}_{A,k}^t(\mathbf{z}) = (\mathcal{T}_{A,k}^t(\mathbf{z}) - \langle \mathcal{T}_{A,k}^t(\mathbf{z}^D) \rangle_k) / \sigma_k(\mathcal{T}_{A,k}^t(\mathbf{z}^D))$, and concatenate all the areas
 603 that were simultaneously recorded, $\mathcal{T}_{\text{trial},k}^t(\mathbf{z}) = (\mathcal{T}_{A1,k}^t, \mathcal{T}_{A2,k}^t)$, where $\langle \cdot \rangle_k$ is the trial average,
 604 and σ_k the standard deviation over trials. The trial-matching loss is defined as:

$$\mathcal{L}_{\text{trial}} = \min_{\pi} \sum_{k=1}^K \sum_{t=1}^T \|\mathcal{T}_{\text{trial},k}^t(\mathbf{z}) - \mathcal{T}_{\text{trial},\pi(k)}^t(\mathbf{z}^D)\|^2, \quad (8)$$

605 where π is an assignment between pairs of K recorded and generated trials $\pi : \{1, \dots, K\} \rightarrow \{1, \dots, K\}$. Note that the minimum over π is a combinatorial optimization that needs to be
 606 calculated for every evaluation of the loss function. For the real dataset, we consider the jaw
 607 movement as an additional area, and we concatenate it to the $\mathcal{T}_{\text{trial},k}^t = (\mathcal{T}_{A1,k}^t, \mathcal{T}_{A2,k}^t, \mathcal{T}_{\text{jaw},k}^t)$.

608 Based on this loss function, we optimize the following parameters: $W_{ij}^{rec,d}$, $W_{ij}^{in,d}$, $v_{thr,j}$,
 609 and β for the RefCircs. For the RNNs, we optimize only the recurrent connectivity $W_{ij}^{rec,d}$,
 610 and the rest are fixed from the RefCircs. For the real dataset, additionally to the parameters
 611 optimized in the RefCircs, we also optimize the jaw's linear readout W_j^{jaw} and its scaling
 612 parameter b .

614 **Implementing Dale's law and local inhibition** In our network, the recurrent weights
 615 W^{rec} are computed as the elementwise product of two matrices: \tilde{W}^{rec} , which encodes the
 616 strength of synaptic efficacies and is always positive, and W_{sign}^{rec} , which has a fixed sign
 617 determined by the neurotransmitter type of the presynaptic neuron and $|W_{sign}^{rec}| = 1$:

$$W^{rec} = \tilde{W}^{rec} \circ W_{sign}^{rec} \quad (9)$$

618 To enforce Dale's law during optimization, we set any negative values of \tilde{W}^{rec} to zero
 619 at each iteration as in [Bellec et al. \(2018a\)](#). Similarly, to constrain only local inhibitory
 620 connections during optimization, we zero out any changes in the synaptic efficacies of across-
 621 areas inhibitory connections at each iteration. In simplified models, Dale's law or the local
 622 inhibition constraint can be disrupted by omitting this correction step.

623 The success of the network optimization highly depends on the initialization of the re-
 624 current weight matrices. To initialize signed matrices we follow the theoretical [Rajan and](#)
 625 [Abbott \(2006\)](#) and practical insights [Bellec et al. \(2018b\)](#); [Cornford et al. \(2020\)](#) developed
 626 previously. After defining the constraints on the weight signs W_{sign}^{rec} , the initialization am-
 627 plitude \tilde{W}^{rec} for each target neuron is adjusted to a zero-summed input weights (the sum of
 628 incoming excitatory inputs is equal to the sum of inhibitory inputs). Then the weight ampli-
 629 tude is re-normalized by the modulus of its largest eigenvalue of W^{rec} , so all the eigenvalues
 630 of this matrix W^{rec} have modulus 1 or smaller.

631 **Stopping criterion for the optimization** For the synthetic dataset, we train the models
 632 for 4000 gradient descent steps. For the real dataset, due to limited data and a noisy test

633 set, we select the final model based on the optimization step that yields the best trial-
634 type accuracy (closest to the trial-type accuracy from the data), derived from the jaw trace
635 and whisker stimulus, along with the highest trial-matched Pearson correlation between the
636 model and the recordings.

637 **Sparsity regularization** There is a plethora of ways to enforce sparsity. In this work,
638 we use weight regularization. In particular, we use the l_1 norm of the recurrent and input
639 weights that promote a high level of sparsity (Xu et al., 2012). To avoid numerical instabil-
640 ities, we apply this regularization only for synaptic weights above α and prune all synapses
641 below α . (we set $\alpha = 1e^{-7}$). The regularized loss function becomes:

$$\mathcal{L}_{all} = \mathcal{L} + \lambda_1 ||W^{in}||_2^{\frac{1}{2}} + \lambda_2 ||W^{rec,d}||_2^{\frac{1}{2}} + \lambda_3 ||W^{across,d}||_2^{\frac{1}{2}}, \quad (10)$$

642 where $W^{across,d}$ are the connections from one area to the other.

643 For the synthetic dataset, we choose the level of across-area sparsity by performing a
644 small grid search for λ_3 . In particular, the sparsity level λ_3 is the maximum value λ_3 where
645 the performance remains as good as without sparsity, see Suppl. Fig S3. For the real
646 dataset, we use the same value λ_3 as the one we found for the full reconstruction method
647 of bioRNN1.

648 4.6 Perturbation test of in silico optogenetics

649 In systems neuroscience, a method to test causal interactions between brain regions uses
650 spatially and temporally precise optogenetic activations or inactivations (Esmaeili et al.,
651 2021; Guo et al., 2014). Usually, inactivations refer to the strong activation of inhibitory
652 neurons for cortical areas. These inhibitory neurons have strong intra-area connections that
653 effectively “silence” their local-neighborhood (Helmstaedter et al., 2009).

654 Our model can simulate these perturbations and allow us to compare the causal mech-
655 nisms of two networks based on their responses to optogenetic perturbations. We implement
656 activations and inactivations as a strong input current to all the neurons in one area’s excita-
657 tory or inhibitory population. For the RefCircs and reconstructed RNNs, we use a transient
658 current that lasts 40 ms, from 20 ms before to 20 ms after the input stimulus. The strength
659 of the current (light power) varies until there is an effect in the full reconstruction method
660 bioRNN1. For the synthetic dataset in Figure 2 (except for panel D), we inject a current of
661 $\Delta u_i^t = 0.08$ into excitatory neurons for activations and $\Delta u_i^t = 1$ into inhibitory neurons for
662 inactivations. For the real dataset, we perform optogenetics inactivations in three different
663 periods. As in Esmaeili et al. (2021), we silence the cortical circuit during the whisker
664 presentation, the time between the whisker and auditory stimulus, or when the animal was
665 licking for the reward. In particular, we use transient currents to the inhibitory neurons
666 during (i.) 100 ms before and after the whisker presentation, (ii.) 100 ms after the whisker
667 presentation till 100ms before the onset of the auditory cue, and (iii.) after the auditory
668 cue till the end of our simulation. For cases (i.) and (ii.), we linearly decreased the strength
669 of the current to avoid rebound excitation. The light power is chosen so that our model
670 has the best results in reproducing the lick probability of the recordings. It is important
671 to mention that the perturbation data are not used to optimize the network but to test
672 whether the resulting network has the same causal interactions with the recordings.

673 For the RefCircs and bioRNNs, we evaluate the effect of the perturbations directly from
674 the neural activity. We use the distance of network dynamics \mathcal{L}_{trial} to compare the two
675 perturbed networks. For the real dataset, we compare the effect of the inactivations on
676 the behavior; as behavior here, we mean whether the mouse/model licked. We classify the
677 licking action using a multilayer perceptron with two hidden layers with 128 neurons each.
678 The classifier is trained with the jaw movement of the real dataset, which was extracted from
679 video filming using custom software, to predict the lick action, which was extracted from a
680 piezo sensor placed in the spout. This classifier predicted lick correctly 94% of the time. We

681 then used the same classifier on the jaw movement from the model to determine whether
682 there was a “lick” or not. For the comparisons in both the artificial and real datasets, we
683 trained multiple models with different random seeds for each variant and aggregated the
684 results. The different random seeds affect both the weight initialization and the noise of our
685 model. In particular, we used from 3 to 6 different random seeds for each different model
686 variant.

687 5 Data availability statement

688 The code for this project is open sourced and published at
689 <https://github.com/Sourmpis/BiologicallyInformed>. The dataset for the artificial dataset
690 can be downloaded/generated on our code repository. The in vivo dataset was published
691 openly for the previous publication Esmaeili et al. (2021). The dataset is accessible at:
692 <https://zenodo.org/records/4720013>.

693 6 Acknowledgment

694 We thank Alireza Modirshanechi, Shuqi Wang, and Tâm Nguyen for their valuable feedback
695 on the manuscript. We are grateful to Vahid Esmaeili for collecting the dataset and ongoing
696 support throughout this project. This research is supported by the Sinergia project CR-
697 SII5_198612, the Swiss National Science Foundation (SNSF) project 200020_207426 awarded
698 to WG, SNSF projects TMAG-3_209271 and 31003A_182010 awarded to CP, and the Vienna
699 Science and Technology Fund (WWTF) project VRG24-018 awarded to GB.

700 References

701 Adesnik, H. and Abdeladim, L. (2021). Probing neural codes with two-photon holographic
702 optogenetics. *Nature Neuroscience*, 24(10):1356–1366.

703 Aravanis, A. M., Wang, L.-P., Zhang, F., Meltzer, L. A., Mogri, M. Z., Schneider, M. B., and
704 Deisseroth, K. (2007). An optical neural interface: in vivo control of rodent motor cortex
705 with integrated fiberoptic and optogenetic technology. *Journal of Neural Engineering*,
706 4(3):S143.

707 Arthur, B. J., Kim, C. M., Chen, S., Preibisch, S., and Darshan, R. (2023). A scalable im-
708 plementation of the recursive least-squares algorithm for training spiking neural networks.
709 *Frontiers in Neuroinformatics*, 17:1099510.

710 Azabou, M., Arora, V., Ganesh, V., Mao, X., Nachimuthu, S., Mendelson, M., Richards,
711 B., Perich, M., Lajoie, G., and Dyer, E. (2024). A unified, scalable framework for neural
712 population decoding. *Advances in Neural Information Processing Systems*, 36.

713 Bellec, G., Kappel, D., Maass, W., and Legenstein, R. (2018a). Deep rewiring: Training
714 very sparse deep networks. *ICLR*.

715 Bellec, G., Salaj, D., Subramoney, A., Legenstein, R., and Maass, W. (2018b). Long short-
716 term memory and learning-to-learn in networks of spiking neurons. *Advances in Neural
717 Information Processing Systems*, 31.

718 Bellec, G., Wang, S., Modirshanechi, A., Brea, J., and Gerstner, W. (2021). Fitting summary
719 statistics of neural data with a differentiable spiking network simulator. *Advances in
720 Neural Information Processing Systems*, 34.

721 Billeh, Y. N., Cai, B., Gratiy, S. L., Dai, K., Iyer, R., Gouwens, N. W., Abbasi-Asl, R.,
722 Jia, X., Siegle, J. H., Olsen, S. R., et al. (2020). Systematic integration of structural
723 and functional data into multi-scale models of mouse primary visual cortex. *Neuron*,
724 106(3):388–403.

725 Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K. (2005). Millisecond-
726 timescale, genetically targeted optical control of neural activity. *Nature neuroscience*,
727 8(9):1263–1268.

728 Chen, G., Scherr, F., and Maass, W. (2022). A data-based large-scale model for primary
729 visual cortex enables brain-like robust and versatile visual processing. *Science Advances*,
730 8(44):eabq7592.

731 Cornford, J., Kalajdzievski, D., Leite, M., Lamarquette, A., Kullmann, D. M., and Richards,
732 B. (2020). Learning to live with dale's principle: Anns with separate excitatory and
733 inhibitory units. *bioRxiv*, pages 2020–11.

734 Cowley, B. R., Calhoun, A. J., Rangarajan, N., Ireland, E., Turner, M. H., Pillow, J. W.,
735 and Murthy, M. (2024). Mapping model units to visual neurons reveals population code
736 for social behaviour. *Nature*, 629(8014):1100–1108.

737 Défossez, A., Copet, J., Synnaeve, G., and Adi, Y. (2023). High fidelity neural audio
738 compression. *Transactions on Machine Learning Research*.

739 Deny, S., Ferrari, U., Mace, E., Yger, P., Caplette, R., Picaud, S., Tkačik, G., and Marre,
740 O. (2017). Multiplexed computations in retinal ganglion cells of a single type. *Nature
741 Communications*, 8(1):1964.

742 Dinc, F., Shai, A., Schnitzer, M., and Tanaka, H. (2023). Cormn: Convex optimization of
743 recurrent neural networks for rapid inference of neural dynamics. *Advances in Neural
744 Information Processing Systems*, 36.

745 Dubreuil, A., Valente, A., Beiran, M., Mastrogiovanni, F., and Ostoicic, S. (2022). The role
746 of population structure in computations through neural dynamics. *Nature Neuroscience*,
747 25(6):783–794.

748 Eccles, J. C. (1976). *From Electrical to Chemical Transmission in the Central Nervous
749 System: The Closing Address of the Sir Henry Dale Centennial Symposium*. Springer.

750 Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14(2):179–211.

751 Esmaeili, V., Tamura, K., Muscinelli, S. P., Modirshanechi, A., Boscaglia, M., Lee, A. B.,
752 Oryshchuk, A., Foustoukos, G., Liu, Y., Crochet, S., et al. (2021). Rapid suppression
753 and sustained activation of distinct cortical regions for a delayed sensory-triggered motor
754 response. *Neuron*, 109(13):2183–2201.

755 Fraile, J. G., Scherr, F., Ramasco, J. J., Arkhipov, A., Maass, W., and Mirasso, C. R.
756 (2023). Competition between bottom-up visual input and internal inhibition generates
757 error neurons in a model of the mouse primary visual cortex. *bioRxiv*, pages 2023–01.

758 Grosenick, L., Marshel, J. H., and Deisseroth, K. (2015). Closed-loop and activity-guided
759 optogenetic control. *Neuron*, 86(1):106–139.

760 Guo, Z. V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J. T., Feng, G., and Svo-
761 boda, K. (2014). Flow of cortical activity underlying a tactile decision in mice. *Neuron*,
762 81(1):179–194.

763 Harris, J. A., Mihalas, S., Hirokawa, K. E., Whitesell, J. D., Choi, H., Bernard, A., Bohn,
764 P., Caldejon, S., Casal, L., Cho, A., et al. (2019). Hierarchical organization of cortical
765 and thalamic connectivity. *Nature*, 575(7781):195–202.

766 Helmstaedter, M., Sakmann, B., and Feldmeyer, D. (2009). L2/3 interneuron groups de-
767 fined by multiparameter analysis of axonal projection, dendritic geometry, and electrical
768 excitability. *Cerebral Cortex*, 19(4):951–962.

769 Hodgkin, A. L. (1958). The croonian lecture-ionic movements and electrical activity in giant
770 nerve fibres. *Proceedings of the Royal Society of London. Series B-Biological Sciences*,
771 148(930):1–37.

772 International Brain Laboratory, Benson, B., Benson, J., Birman, D., Bonacchi, N., Caran-
773 dini, M., Catarino, J. A., Chapuis, G. A., Churchland, A. K., Dan, Y., et al. (2023). A
774 brain-wide map of neural activity during complex behaviour. *bioRxiv*, pages 2023–07.

775 Isbister, J. B., Ecker, A., Pokorny, C., Bolaños-Puchet, S., Santander, D. E., Arnaudon, A.,
776 Awile, O., Barros-Zulaica, N., Alonso, J. B., Boci, E., et al. (2023). Modeling and sim-
777 ulation of neocortical micro-and mesocircuitry. part ii: Physiology and experimentation.
778 *bioRxiv*, pages 2023–05.

779 Kim, C. M., Finkelstein, A., Chow, C. C., Svoboda, K., and Darshan, R. (2023). Dis-
780 tributing task-related neural activity across a cortical network through task-independent
781 connections. *Nature Communications*, 14(1):2851.

782 Lakunina, A., Socha, K. Z., Ladd, A., Bowen, A. J., Chen, S., Colonell, J., Doshi, A., Karsh,
783 B., Krumin, M., Kulik, P., Li, A., Neutens, P., O’Callaghan, J., Olsen, M., Putzeys, J.,
784 Reddy, C. B., Tilmans, H. A., Vargas, S., Welkenhuysen, M., Ye, Z., Häusser, M., Koch,
785 C., Ting, J. T., Consortium, N. O., Dutta, B., Harris, T. D., Steinmetz, N. A., Svoboda,
786 K., Siegle, J. H., and Carandini, M. (2025). Neuropixels opto: Combining high-resolution
787 electrophysiology and optogenetics. *bioRxiv*.

788 Lappalainen, J. K., Tschopp, F. D., Prakhy, S., McGill, M., Nern, A., Shinomiya, K.,
789 Takemura, S.-y., Gruntman, E., Macke, J. H., and Turaga, S. C. (2023). Connectome-
790 constrained deep mechanistic networks predict neural responses across the fly visual sys-
791 tem at single-neuron resolution. *bioRxiv*, pages 2023–03.

792 LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.

793 Liu, Y., Foustoukos, G., Crochet, S., and Petersen, C. C. (2022). Axonal and dendritic
794 morphology of excitatory neurons in layer 2/3 mouse barrel cortex imaged through whole-
795 brain two-photon tomography and registered to a digital brain atlas. *Frontiers in Neu-
796 roanatomy*, 15:791015.

797 Mahuas, G., Isacchini, G., Marre, O., Ferrari, U., and Mora, T. (2020). A new inference
798 approach for training shallow and deep generalized linear models of noisy interacting
799 neurons. *Advances in Neural Information Processing Systems*, 33.

800 Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A.,
801 Ailamaki, A., Alonso-Nanclares, L., Antille, N., Arsever, S., et al. (2015). Reconstruction
802 and simulation of neocortical microcircuitry. *Cell*, 163(2):456–492.

803 Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C.
804 (2004). Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*,
805 5(10):793–807.

806 Minai, Y., Soldado-Magraner, J., Smith, M., and Yu, B. M. (2024). Miso: Optimizing brain
807 stimulation to create neural activity states. *Advances in Neural Information Processing
808 Systems*, 37:24126–24149.

809 Neftci, E. O., Mostafa, H., and Zenke, F. (2019). Surrogate gradient learning in spiking
810 neural networks: Bringing the power of gradient-based optimization to spiking neural
811 networks. *IEEE Signal Processing Magazine*, 36(6):51–63.

812 Packer, A. M., Russell, L. E., Dagleish, H. W., and Häusser, M. (2015). Simultaneous
813 all-optical manipulation and recording of neural circuit activity with cellular resolution
814 in vivo. *Nature Methods*, 12(2):140–146.

815 Pals, M., Sağtekin, A. E., Pei, F., Gloeckler, M., and Macke, J. H. (2024). Inferring stochastic
816 low-rank recurrent neural networks from neural data. *arXiv:2406.16749*.

817 Pandarinath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C.,
818 Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., et al. (2018). Inferring
819 single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*,
820 15(10):805–815.

821 Papagiakoumou, E., Ronzitti, E., and Emiliani, V. (2020). Scanless two-photon excitation
822 with temporal focusing. *Nature Methods*, 17(6):571–581.

823 Perich, M. G., Arlt, C., Soares, S., Young, M. E., Mosher, C. P., Minxha, J., Carter,
824 E., Rutishauser, U., Rudebeck, P. H., Harvey, C. D., et al. (2020). Inferring brain-
825 wide interactions using data-constrained recurrent neural network models. *bioRxiv*, pages
826 2020–12.

827 Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E., and Si-
828 moncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete
829 neuronal population. *Nature*, 454(7207):995–999.

830 Pozzorini, C., Mensi, S., Hagens, O., Naud, R., Koch, C., and Gerstner, W. (2015). Auto-
831 mated high-throughput characterization of single neurons by means of simplified spiking
832 models. *PLoS Computational Biology*, 11(6):e1004275.

833 Rajan, K. and Abbott, L. F. (2006). Eigenvalue spectra of random matrices for neural
834 networks. *Physical Review Letters*, 97(18):188104.

835 Richards, B. A. and Kording, K. P. (2023). The study of plasticity has always been about
836 gradients. *The Journal of Physiology*, 601(15):3141–3149.

837 Rimehaug, A. E., Stasik, A. J., Hagen, E., Billeh, Y. N., Siegle, J. H., Dai, K., Olsen, S. R.,
838 Koch, C., Einevoll, G. T., and Arkhipov, A. (2023). Uncovering circuit mechanisms of
839 current sinks and sources with biophysical simulations of primary visual cortex. *eLife*,
840 12:e87169.

841 Rosenblatt, F. (1960). Perceptual generalization over transformation groups. *Self Organizing
842 Systems*, pages 63–96.

843 Schrimpf, M., McGrath, P., Margalit, E., and DiCarlo, J. J. (2024). Do topographic deep
844 ann models of the primate ventral stream predict the perceptual effects of direct cortical
845 interventions? *bioRxiv*, pages 2024–01.

846 Siegle, J. H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez,
847 T. K., Choi, H., Luviano, J. A., et al. (2021). Survey of spiking in the mouse visual system
848 reveals functional hierarchy. *Nature*, 592(7852):86–92.

849 Sourmpis, C., Petersen, C., Gerstner, W., and Bellec, G. (2023). Trial matching: capturing
850 variability with data-constrained spiking neural networks. *Advances in Neural Information
851 Processing Systems*, 36.

852 Spieler, A., Rahaman, N., Martius, G., Schölkopf, B., and Levina, A. (2023). The expres-
853 sive leaky memory neuron: an efficient and expressive phenomenological neuron model
854 can solve long-horizon tasks. In *The Twelfth International Conference on Learning Rep-
855 resentations*.

856 Staiger, J. F. and Petersen, C. C. (2021). Neuronal circuits in barrel cortex for whisker
857 sensory perception. *Physiological Reviews*, 101(1):353–415.

858 Steinmetz, N. A., Zatka-Haas, P., Carandini, M., and Harris, K. D. (2019). Distributed
859 coding of choice, action and engagement across the mouse brain. *Nature*, 576(7786):266–
860 273.

861 Tamamaki, N. and Tomioka, R. (2010). Long-range gabaergic connections distributed
862 throughout the neocortex and their possible function. *Frontiers in Neuroscience*, 4:202.

863 Tamura, K., Bech, P., Mizuno, H., Veautre, L., Crochet, S., and Petersen, C. C. (2025).
864 Cell-class-specific orofacial motor maps in mouse neocortex. *Current Biology*, 35(6):1382–
865 1390.

866 Teeter, C., Iyer, R., Menon, V., Gouwens, N., Feng, D., Berg, J., Szafer, A., Cain, N., Zeng,
867 H., Hawrylycz, M., et al. (2018). Generalized leaky integrate-and-fire models classify
868 multiple neuron types. *Nature Communications*, 9(1):709.

869 Udvary, D., Harth, P., Macke, J. H., Hege, H.-C., de Kock, C. P., Sakmann, B., and Ober-
870 laender, M. (2022). The impact of neuron morphology on cortical network architecture.
871 *Cell Reports*, 39(2).

872 Urai, A. E., Doiron, B., Leifer, A. M., and Churchland, A. K. (2022). Large-scale neural
873 recordings call for new insights to link brain and behavior. *Nature Neuroscience*, 25(1):11–
874 19.

875 Valente, A., Pillow, J. W., and Ostoic, S. (2022). Extracting computational mechanisms
876 from neural data using low-rank rnns. *Advances in Neural Information Processing Sys-
877 tems*, 35.

878 Wyart, C., Jia, X., Kumar, A., Peysson, N., Carbo-Tano, M., Fidelin, K., Didelot, M.,
879 Lejeune, C., and Grosse-Wentrup, M. (2025). All-optical investigation reveals a hier-
880 archical organization of vsx2+ reticulospinal neurons coordinating steering and forward
881 locomotion.

882 Xu, Z., Chang, X., Xu, F., and Zhang, H. (2012). $l_{\cdot\{1/2\}}$ regularization: A threshold-
883 ing representation theory and a fast solver. *IEEE Transactions on neural networks and
884 learning systems*, 23(7):1013–1027.

885 Ye, J., Collinger, J., Wehbe, L., and Gaunt, R. (2023). Neural data transformer 2: multi-
886 context pretraining for neural spiking activity. *Advances in Neural Information Processing
887 Systems*, 36:80352–80374.

Method name	Real Dataset vs Reconstructed network
biRNN	0.76 ± 0.14
σ RNN	0.62 ± 0.12
No sparsity	0.77 ± 0.15
Non-local inhibition	0.79 ± 0.15
No Dale's law	0.68 ± 0.13
No TM	1.63 ± 0.55
No spike	0.64 ± 0.13

Table 2: Trial-matching loss test loss $\mathcal{L}_{\text{trial}}$ of the different reconstruction methods with the real recordings from (Esmaeili et al., 2021) \pm indicates the 95% confidence interval.

888

Supplemental information

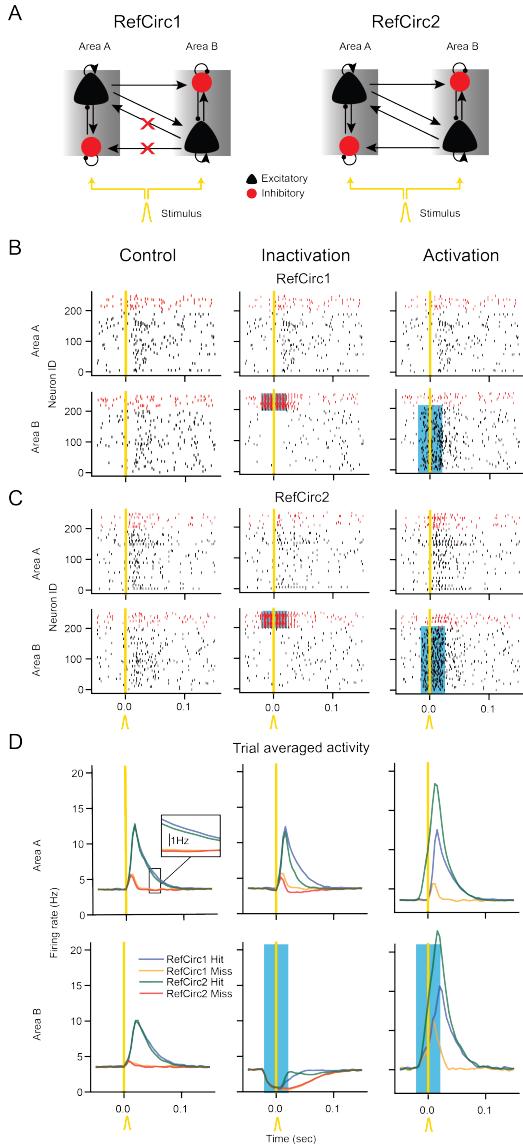


Figure S1: Modeling “optogenetic” perturbations. **A.** Two different network hypotheses for implementing a detection task. In RefCirc1, area A projects to area B but not vice versa. In RefCirc2, the areas are recurrently connected. **B.** Raster plots of all neurons in RefCirc1 during a single hit trial under normal conditions (control, left) and under optogenetic perturbation of excitatory (middle) and inhibitory (right) neurons. The duration of the light stimulus is shown with a blue shading. **C.** Same for RefCirc2 **D.** Trial-averaged activity of the two circuits during Hit (blue: RefCirc1; green: RefCirc2) and Miss (yellow: RefCirc 1; red: RefCirc2) trials. A trial is classified as “Hit” if area A reaches a transient firing rate above 8Hz; and otherwise as “Miss”. For the control case, the maximal difference between the trial average activity of the two networks is below 0.51 Hz (zoom inset).

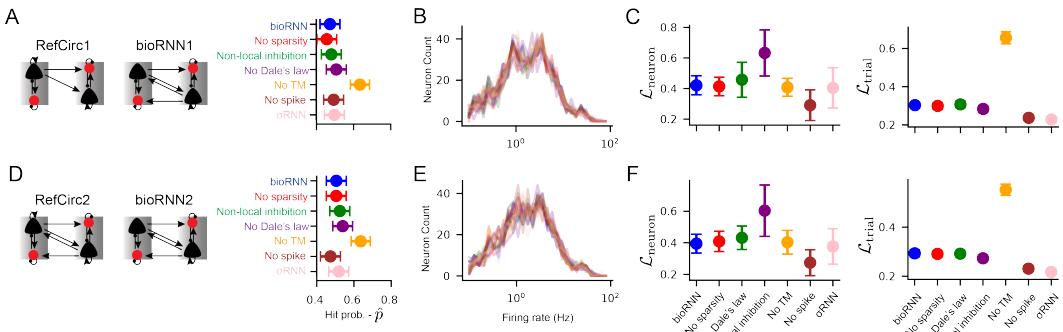


Figure S2: Fitting Reconstructed networks to the synthetic dataset. **A.** Schematic representation of the RefCirc1 and bioRNN1. and probability of hit trials. **B.** Histogram of the firing rate distribution of the RefCirc1 and all the RNN1 variants. We observe that all RNN1 versions fit well with the RefCirc1. **C.** Left: Neuron loss of the different RNN1 variants. Right: Trial-matching loss of the different RNN1 variants. We observe that the model without the trial-matching loss function behaves considerably worse. The whiskers show the 95% confidence interval of the mean across trials. **D-F.** Same as **A-B** for RefCirc2 and RNNs2.

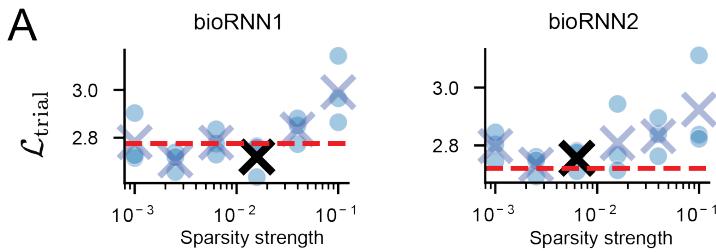


Figure S3: Picking the sparsity level. **A.** Grid search for the optimal maximum regularization strength (λ_3) without a drop in performance. As a performance measure, we used the trial-matching loss, $\mathcal{L}_{\text{trial}}$.

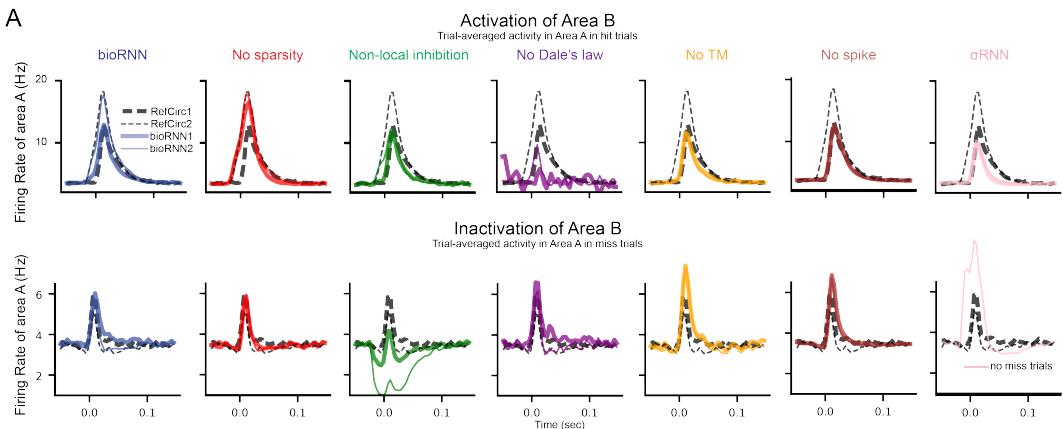


Figure S4: Trial averaged traces across RNN variants. Trial-averaged activity in area A under activation/inactivation of area B . Dashed black lines indicate the activity of RefCirc1 (thick dashed) and RefCirc2 (thin dashed). All the RNNs are tested with the same reference circuit and training data, each bioRNN model variant is shown with a different color.

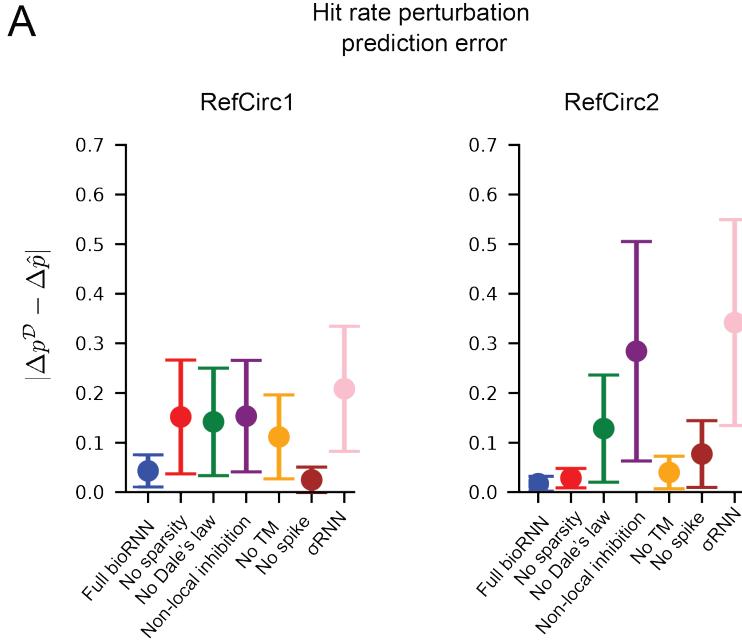


Figure S5: **A** Hit frequency prediction error $|\Delta p^D - \Delta \hat{p}|$ as in Figure 2E. In contrast to Figure 2E, here we show separately the change of hit probability for RefCirc1 (left) and RefCirc2 (right).

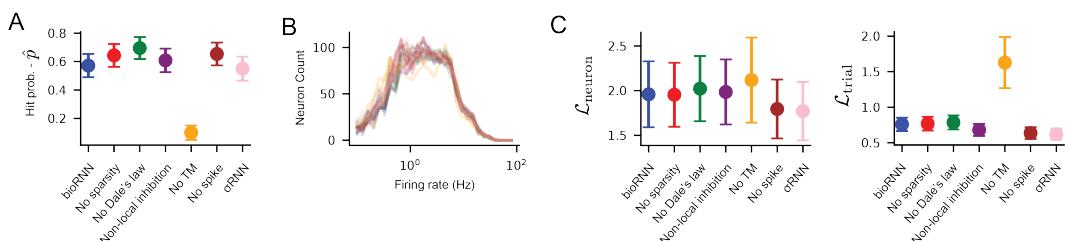


Figure S6: **Reconstruction of the real recordings.** **A.** Probability of hit trials of the different variant models. **B.** Histogram of the firing rate distribution from the real recordings and all the variants. **C.** Top: Neuron loss of the different RNN1 variants. All RNN versions have a similar loss value. Bottom: Trial-matching loss of the different model variants. We observe that the model without the trial-matching loss function behaves considerably worse. The whiskers show the 95% confidence interval.

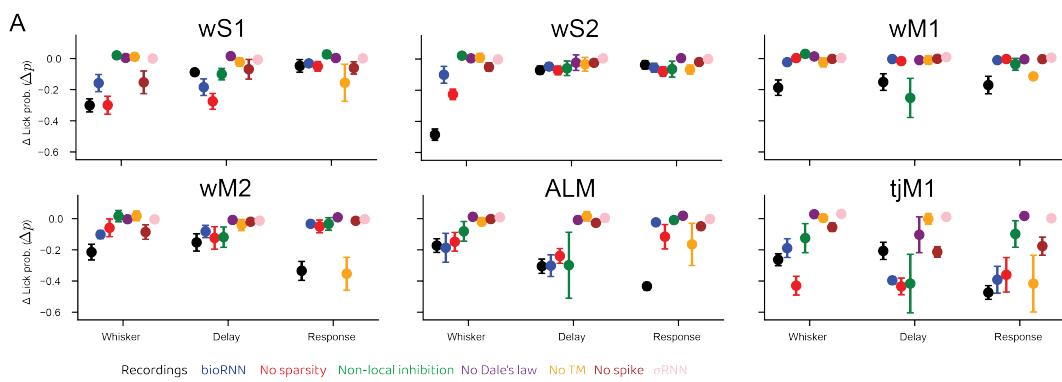


Figure S7: **A** Change of lick probability under inactivation of all areas in all the different temporal windows. We show the Δp from the data and reconstruction model variants.