

1 Training deep neural density estimators to 2 identify mechanistic models of neural dynamics

3 **Pedro J. Gonçalves**^{1,2*}, **Jan-Matthis Lueckmann**^{1,2*}, **Michael Deistler**^{1*}, **Marcel Nonnenmacher**^{1,2},
4 **Kaan Öcal**^{2,3}, **Giacomo Bassetto**^{1,2}, **Chaitanya Chintaluri**⁴, **William F. Podlaski**⁴, **Sara A. Haddad**⁵,
5 **Tim P. Vogels**⁴, **David S. Greenberg**¹, **Jakob H. Macke**^{1,2}

6 ¹Computational Neuroengineering, Department of Electrical and Computer Engineering, Technical
7 University of Munich, Germany; ²Max Planck Research Group Neural Systems Analysis, Center of
8 Advanced European Studies and Research (caesar), Bonn, Germany; ³Mathematical Institute, University of
9 Bonn, Bonn, Germany; ⁴Centre for Neural Circuits and Behaviour, University of Oxford; ⁵Max Planck
10 Institute for Brain Research, Frankfurt, Germany

11
12 **Abstract** Mechanistic modeling in neuroscience aims to explain neural or behavioral phenomena in terms of
13 underlying causes. A central challenge in building mechanistic models is to identify which models and parameters can
14 achieve an agreement between the model and experimental data. The complexity of models and data characterizing
15 neural systems makes it infeasible to solve model equations analytically or tune parameters manually. To overcome
16 this limitation, we present a machine learning tool that uses density estimators based on deep neural networks—
17 trained using model simulations—to infer data-compatible parameters for a wide range of mechanistic models. Our
18 tool identifies all parameters consistent with data, is scalable both in the number of parameters and data features,
19 and does not require writing new code when the underlying model is changed. It can be used to analyze new data
20 rapidly after training, and can be applied to either raw data or selected data features. We demonstrate our approach
21 for parameter inference on ion channels, receptive fields, and Hodgkin–Huxley models. Finally, we use it to explore the
22 space of parameters which give rise to the same rhythmic activity in a network model of the crustacean stomatogastric
23 ganglion and to search for potential compensation mechanisms. The approach presented here will help close the gap
24 between data-driven and theory-driven models of neural dynamics.

26 Introduction

27 New experimental technologies allow us to observe neurons, networks, brain regions and entire systems at un-
28 precedented scale and resolution, but using these data to understand how behavior arises from neural processes
29 remains a challenge. To test our understanding of a phenomenon, we often take to rebuilding it in the form of a
30 computational model that incorporates the mechanisms we believe to be at play, based on scientific knowledge,
31 intuition, and hypotheses about the components of a system and the laws governing their relationships. The goal of
32 such mechanistic models is to investigate whether a proposed mechanism can explain experimental data, uncover
33 details that may have been missed, inspire new experiments, and eventually provide insights into the inner workings
34 of an observed neural or behavioral phenomenon [1–4]. Examples for such a symbiotic relationship between model
35 and experiments range from the now classical work of Hodgkin and Huxley [5], to population models investigating
36 rules of connectivity, plasticity and network dynamics [6–10], network models of inter-area interactions [11, 12], and
37 models of decision making [13, 14].

*These authors contributed equally to this work

For correspondence: pedro.goncalves@caesar.de; jan-matthis.lueckmann@tum.de; michael.deistler@tum.de; macke@tum.de

38 A crucial step in building a model is adjusting its free parameters to be consistent with experimental observations.
39 This is essential both for investigating whether the model agrees with reality and for gaining insight into processes
40 which cannot be measured experimentally. For some models in neuroscience, it is possible to identify the relevant
41 parameter regimes from careful mathematical analysis of the model equations. But as the complexity of both neural
42 data and neural models increases, it becomes very difficult to find well-fitted parameters by inspection, and *automated*
43 identification of data-consistent parameters is required.

44 Furthermore, to understand how a model quantitatively explains data, it is necessary to find not only the *best*,
45 but *all* parameter settings consistent with experimental observations. This is especially important when modeling
46 neural data, where highly variable observations can lead to broad ranges of data-consistent parameters. Elucidating
47 these ranges can reveal which combinations of parameters are well-constrained by data, and helps us design further
48 experiments to be maximally informative [15]. Moreover, many models in biology are inherently robust to some
49 perturbations of parameters, but highly sensitive to others [3, 16], e.g. because of processes such as homeostatic
50 regulation. For these systems, identifying the full range of data-consistent parameters can reveal how multiple distinct
51 parameter settings give rise to the same model behavior [7, 17, 18]. Yet despite the clear benefits of mechanistic
52 models in providing scientific insight, identifying their parameters given data remains a challenging open problem that
53 demands new algorithmic strategies.

54 The gold standard for automated parameter identification is *statistical inference*, which uses the likelihood $p(\mathbf{x}|\theta)$
55 to quantify the match between parameters θ and data \mathbf{x} . Likelihoods can be derived for purely statistical models
56 commonly used in neuroscience [19–25], but are unavailable for most mechanistic models. Mechanistic models
57 are designed to reflect knowledge about biological mechanisms, and not necessarily to be amenable to efficient
58 inference: Many mechanistic models are defined implicitly through stochastic computer simulations (e.g. a simulation
59 of a network of spiking neurons), and likelihood calculation would require the ability to integrate over all potential
60 paths through the simulator code. Similarly, a common goal of mechanistic modeling is to capture selected summary
61 features of the data (e.g. a certain firing rate, bursting behavior, etc...), *not* the full dataset in all its details. The same
62 feature (such as a particular average firing rate) can be produced by infinitely many realizations of the simulated
63 process (such as a time-series of membrane potential). This makes it impractical to compute likelihoods, as one would
64 have to average over all possible realizations which produce the same output.

65 Since the toolkit of statistical inference is inaccessible for mechanistic models, parameters are typically tuned
66 ad-hoc (often through laborious, and subjective, trial-and-error), or by computationally expensive parameter search:
67 A large set of models is generated, and grid search [26–28], genetic algorithms [29–32], or Approximate Bayesian
68 Computation (ABC) [33–35] are used to filter out models whose simulations do not match the data. Parameter search
69 methods require the user to define a heuristic rejection criterion to decide which simulations to keep, and typically
70 end up discarding most simulations. They struggle when models have many parameters or data features, cannot cope
71 with large datasets or high-throughput applications, and (except for ABC) yield only a single *best-fitting model*, rather
72 than the full range of data-compatible models. Thus, computational neuroscientists face a dilemma: Either create
73 carefully designed, highly interpretable mechanistic models (but rely on ad-hoc parameter tuning), or resort to purely
74 statistical models offering sophisticated parameter inference but limited mechanistic insight.

75 Here we propose a new approach using machine learning to combine the advantages of mechanistic and statistical
76 modeling. We present SNPE (Sequential Neural Posterior Estimation), a tool that rapidly identifies all mechanistic
77 model parameters consistent with observed experimental data (or summary features). SNPE builds on recent advances
78 in simulation-based Bayesian inference [36–39]: Given observed experimental data (or summary features) \mathbf{x}_o , and a
79 mechanistic model with parameters θ , it expresses both prior knowledge and the range of data-compatible parameters
80 through probability distributions. It returns a posterior distribution $p(\theta|\mathbf{x}_o)$ which is high for parameters θ consistent
81 with both the data \mathbf{x}_o and prior knowledge but approaches zero for θ inconsistent with either (Fig. 1).

82 Similar to parameter search methods, SNPE uses simulations instead of likelihood calculations, but instead of
83 filtering out simulations, it uses *all* simulations to train a multi-layer artificial neural network to identify admissible
84 parameters (Fig. 1). By incorporating modern deep neural networks for conditional density estimation [40, 41], it can
85 capture the full *distribution* of parameters consistent with the data, even when this distribution has multiple peaks or
86 lies on curved manifolds. Critically, SNPE decouples the design of the model and design of the inference approach,
87 giving the investigator maximal flexibility to design and modify mechanistic models. Our method makes minimal
88 assumptions about the model or its implementation, and can e.g. also be applied to non-differentiable models, such

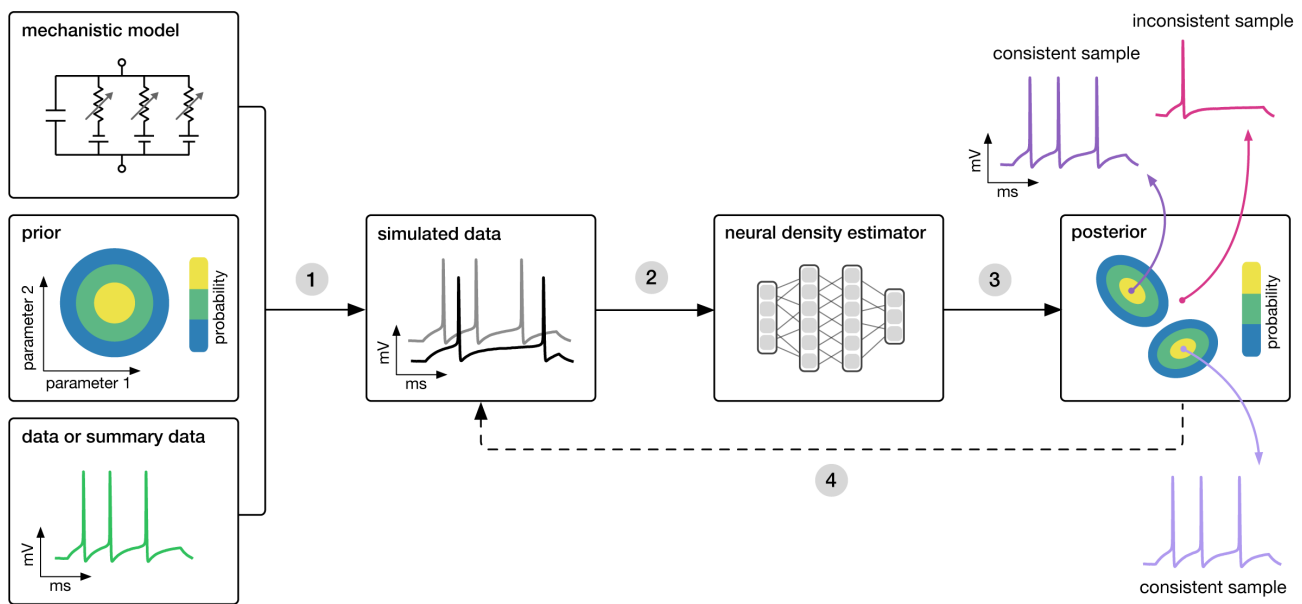


Figure 1. Goal: Algorithmically identify mechanistic models which are consistent with data. Our algorithm (SNPE) takes three inputs: A candidate mechanistic model, prior knowledge or constraints on model parameters, and data (or summary statistics). SNPE proceeds by 1) sampling parameters from the prior and simulating synthetic datasets from these parameters, and 2) using a deep density estimation neural network to learn the (probabilistic) association between data (or data features) and underlying parameters, i.e. to learn statistical inference from simulated data. 3) This density estimation network is then applied to empirical data to derive the full space of parameters consistent with the data and the prior, i.e. the posterior distribution. High posterior probability is assigned to parameters which are consistent with both the data and the prior, low probability to inconsistent parameters. 4) If needed, an initial estimate of the posterior can be used to adaptively generate additional informative simulations.

89 as networks of spiking neurons. Its only requirement is that one can run model simulations for different parameters,
90 and collect the resulting synthetic data or summary features of interest.

91 We test SNPE using mechanistic models expressing key neuroscientific concepts. Beginning with a simple neural
92 encoding problem with a known solution, we progress to more complex data types, large datasets and many-
93 parameter models inaccessible to previous methods. We estimate visual receptive fields using many data features,
94 demonstrate rapid inference of ion channel properties from high-throughput voltage-clamp protocols, and show how
95 Hodgkin–Huxley models are more tightly constrained by increasing numbers of data features. Finally, we explore how
96 multiple network models can explain the activity in the stomatogastric ganglion [7], and provide hypotheses for which
97 compensation mechanisms might be at play.

98 Concurrently with our work, Bittner and colleagues [42] developed an alternative approach to parameter identifica-
99 tion for mechanistic models, and showed how it can be used to characterize neural population models which exhibit
100 specific emergent computational properties. Both studies differ in their methodology and domain of applicability
101 (see descriptions of underlying algorithms in our [37, 38] and their [43] prior work), as well in the focus of their
102 neuroscientific contributions, but they share the overall goal of using deep probabilistic inference tools to build more
103 interpretable models of neural data. These complementary and concurrent advances will expedite the cycle of building,
104 adjusting and selecting mechanistic models in neuroscience.

105 Results

106 Estimating stimulus-selectivity in linear-nonlinear encoding models

107 We first illustrate SNPE on linear-nonlinear (LN) encoding models, a special case of generalized linear models (GLMs).
108 These are simple, commonly used phenomenological models for which likelihood-based parameter estimation is
109 feasible [44–49], and which can be used to validate the accuracy of our approach. We will show that SNPE returns the
110 correct posterior distribution over parameters, that it can cope with high-dimensional observation data and that it can
111 can recover multiple solutions to parameter inference problems.

112 An LN model describes how a neuron’s firing rate is modulated by a sensory stimulus through a linear filter θ ,

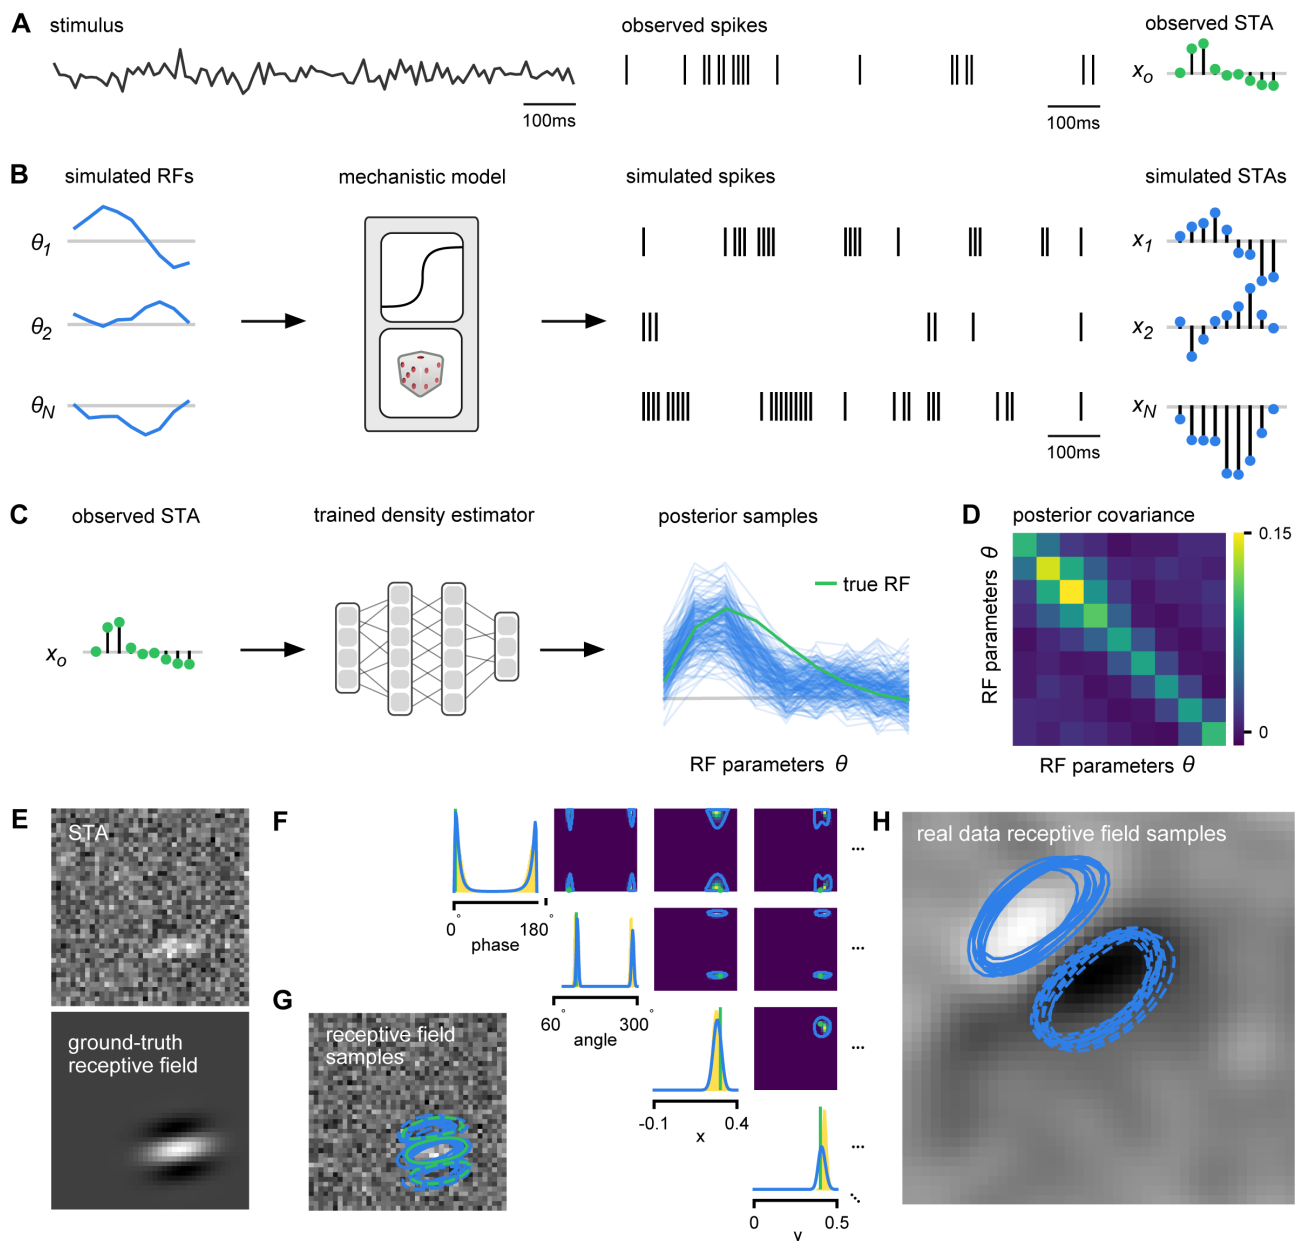


Figure 2. Estimating receptive fields in linear-nonlinear models of single neurons with statistical inference (A) Schematic of a time-varying stimulus, associated observed spike train and resulting spike-triggered average (STA) (B) SNPE proceeds by first randomly generating simulated receptive fields θ , and using the mechanistic model (here an LN model) to generate simulated spike trains and simulated STAs. (C) These simulated STAs and receptive fields are then used to train a deep neural density estimator to identify the distribution of receptive fields consistent with a given observed STA x_o . (D) Posterior covariance over receptive fields, indicating which parameters, or combinations of parameters, are constrained by the data. (E) Example of spatial receptive field. We simulated responses and an STA of a LN-model with oriented receptive field. (F) We used SNPE to recover the distribution of receptive-field parameters. Univariate and pairwise marginals for four parameters of the spatial filter (MCMC, yellow histograms; SNPE, blue lines; full posterior in Supplementary Fig. 4). Non-identifiabilities of the Gabor parameterization lead to multimodal posteriors. (G) Posterior samples from SNPE posterior (SNPE, blue) compared to ground-truth receptive field (green), overlaid on spike-triggered average. (H) Posterior samples for V1 data; full posterior in Supplementary Fig. 5.

113 often referred to as the *receptive field* [50, 51]. We first considered a model of a retinal ganglion cell (RGC) driven by
 114 full-field flicker (Fig. 2A). A statistic that is often used to characterize such a neuron is the *spike-triggered average* (STA)
 115 (Fig. 2A, right). We therefore used the STA, as well as the firing rate of the neuron, as input x_o to SNPE. (Note that, in
 116 the limit of infinite data, and for white noise stimuli, the STA will converge to the receptive field [45]—for finite, and

117 non-white data, the two will in general be different.) By randomly drawing receptive fields θ , we generated synthetic
118 spike trains and calculated STAs from them (Fig. 2B), and subsequently trained a neural conditional density estimator
119 to recover the underlying receptive-field model (Fig. 2C). This allowed us to estimate the posterior distribution over
120 receptive fields, i.e. to estimate which receptive fields are consistent with the data (and prior) (Fig. 2C-D). For LN
121 models, likelihood-based inference is possible, allowing us to validate the SNPE posterior by comparing it to a reference
122 posterior obtained via Markov Chain Monte Carlo (MCMC) sampling [48, 49] (Supplementary Fig. 1 and Supplementary
123 Fig. 2).

124 As a more challenging inference problem, we inferred the receptive field of a neuron in primary visual cortex
125 (V1) [52, 53]. Using a model composed of a bias (related to the spontaneous firing rate) and a Gabor function with 8
126 parameters [54] to describe location, shape and strength of the receptive field, we simulated responses to 5-minute
127 random noise movies of 41×41 pixels. In this case, the STA has 1681 dimensions (Fig. 2E), causing classical ABC
128 methods to fail (Supplementary Fig. 3). This problem admits multiple solutions (as e.g. rotating the receptive field by
129 180°). As a result, the posterior distribution has multiple peaks ('modes'). Starting from a simulation result x_o with
130 known parameters, we used SNPE to estimate the respective posterior distribution. To deal with the high-dimensional
131 data x_o in this problem, we used a convolutional neural network (CNN), as this architecture excels at learning relevant
132 features from image data [55, 56]. To deal with the multiple peaks in the posterior, we fed the CNN's output into
133 a mixture density network (MDN) [57], which can learn to assign probability distributions with multiple peaks as a
134 function of its inputs (full details in Methods). Using this strategy, SNPE was able to infer a posterior distribution that
135 tightly enclosed the ground truth simulation parameters which generated the original simulated data x_o , and closely
136 matched a reference MCMC posterior (Fig. 2F, full posterior in Supplementary Fig. 4). We also applied this approach
137 to electrophysiological data from a V1 cell [53], identifying a sine-shaped Gabor receptive field consistent with the
138 original spike-triggered average (Fig. 2H; posterior distribution in Supplementary Fig. 5).

139 **Functional diversity of ion channels: efficient high-throughput inference**

140 We next show how SNPE can be efficiently applied to estimation problems in which we want to identify a large number
141 of models for different observations in a database. We considered a flexible model of ion channels [59], which we
142 here refer to as the *Omnimodel*. This model uses 8 parameters to describe how the dynamics of currents through
143 non-inactivating potassium channels depend on membrane voltage (Fig. 3A). For various choices of its parameters θ ,
144 it can capture 350 specific models in publications describing this channel type, cataloged in the IonChannelGenealogy
145 (ICG) database [58]. We aimed to identify these ion channel parameters θ for each ICG model, based on 11 features
146 of the model's response to a sequence of 5 voltage clamp protocols, resulting in a total of 55 characteristic different
147 features per model (Fig. 3B, see Methods for details).

148 Because this model's output is a typical format for functional characterization of ion channels both in simulations
149 [58] and in high-throughput electrophysiological experiments [60–63], the ability to rapidly infer different parameters
150 for many separate experiments is advantageous. Existing approaches for model fitting based on numerical optimiza-
151 tion [59, 63] must repeat all computations anew for a new experiment or data point (Fig. 3C). However, for SNPE the
152 only heavy computational tasks are carrying out simulations to generate training data, and training the neural network.
153 We therefore reasoned that by training a network once using a large number of simulations, we could subsequently
154 carry out rapid "amortized" parameter inference on new data using a single pass through the network (Fig. 3D) [64, 65].
155 To test this idea, we used SNPE to train a neural network to infer the posterior from any data x . To generate training
156 data, we carried out 1 million *Omnimodel* simulations, with parameters randomly chosen across ranges large enough
157 to capture the models in the ICG database [58]. In this case, SNPE was run using a single round, i.e. it learned to
158 perform inference for all data from the prior (rather than a specific observed datum). Generating these simulations
159 took around 1000 CPU-hours and training the network 150 CPU-hours, but afterwards a full posterior distribution
160 could be inferred for new data in less than 10 ms.

161 As a first test, SNPE was run on simulation data, generated by a previous characterization of a non-inactivating
162 potassium channel (Fig. 3B). Simulations of the *Omnimodel* using parameter sets sampled from the obtained posterior
163 distribution (Fig. 3E) closely resembled the input data on which the SNPE-based inference had been carried out, while
164 simulations using "outlier" parameter sets with low probability under the posterior generated current responses
165 that were markedly different from the data x_o (Fig. 3F). Taking advantage of SNPE's capability for rapid amortized
166 inference, we further evaluated its performance on all 350 non-inactivating potassium channel models in ICG. In each

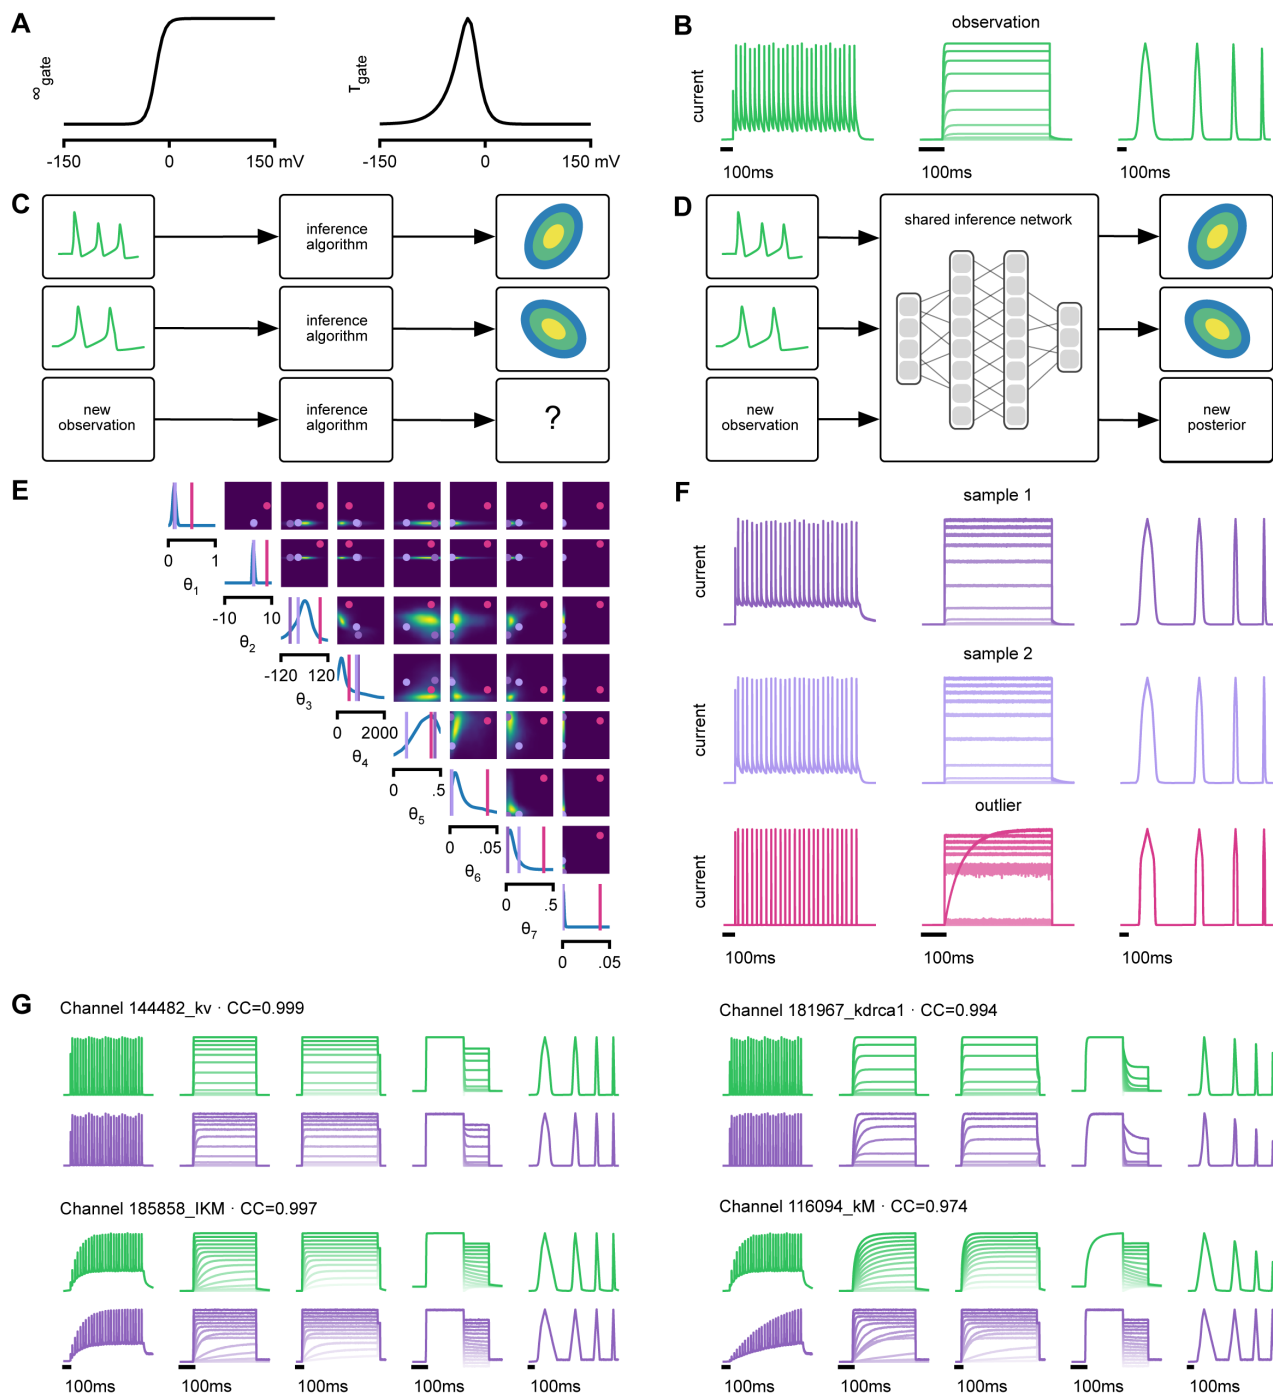


Figure 3. Inference on a database of ion-channel models. (A) We perform inference over the parameters of non-inactivating potassium channel models. Channel kinetics are described by steady-state activation curves, ∞_{gate} , and time-constant curves, τ_{gate} . (B) Observation generated from a channel model from ICG database: Normalized current responses to three (out of five) voltage-clamp protocols (action potentials, activation, and ramping). Details in [58]. (C) Classical approach to parameter identification: Inference is optimized on each datum separately, requiring new computations for each new datum. (D) Amortized inference: An inference network is learned which can be applied to multiple data, enabling rapid inference on new data. (E) Posterior distribution over eight model parameters, θ_1 to θ_8 . (F) Traces obtained by sampling from the posterior in E. Purple: traces sampled from posterior, i.e. with high posterior probability. Magenta: trace from parameters with low probability. (G) Observations (green) and traces generated by posterior samples (purple) for four models from the database.

167 case, we carried out a simulation to generate initial data from the original ICG model, used SNPE to calculate the
168 posterior given the Omnimodel, and then generated a new simulation x using parameters sampled from the posterior
169 (Fig. 3F). This resulted in high correlation between the original ICG model response and the Omnimodel response,
170 in every case (>0.98 for more than 90% of models, see Supplementary Fig. 6). However, this approach was not able
171 to capture all traces perfectly, as e.g. it failed to capture the shape of the onset of the bottom right model in Fig.
172 3G. Additional analysis of this example revealed that this example is not a failure of SNPE, but rather a limitation of
173 the Omnimodel—thus, SNPE can be used to reveal limitations of candidate models and aid the development of more
174 verisimilar mechanistic models.

175 Calculating the posterior for all 350 ICG models only took a few seconds, and was fully automated, i.e. did not
176 require user interactions. These results show how SNPE allows fast and accurate identification of biophysical model
177 parameters on new data, and shows how these approaches could scale to high-throughput or online applications
178 which require rapid and automated inference.

179 **Hodgkin–Huxley model: stronger constraints from additional data features**

180 The Hodgkin–Huxley (HH) model [5] of action potential generation through ion channel dynamics is a highly influential
181 mechanistic model in neuroscience. A number of algorithms have been proposed for fitting HH models to electro-
182 physiological data [26, 31, 32, 66–68], but (with the exception of [69]) these approaches do not attempt to estimate
183 the full posterior. Given the central importance of the HH model in neuroscience, we sought to test how SNPE would
184 cope with this challenging non-linear model. As previous approaches for HH models concentrated on reproducing
185 specified features (e.g. the number of spikes) [66], we also sought to determine how various features provide different
186 constraints. We considered the problem of inferring 8 biophysical parameters in a HH single-compartment model,
187 describing voltage-dependent sodium and potassium conductances and other intrinsic membrane properties (Fig. 4A,
188 left). We simulated the neuron’s voltage response to the injection of a square wave of depolarizing current, and defined
189 the model output x used for inference as the number of evoked action potentials along with 6 additional features
190 of the voltage response (Fig. 4A, right, details in Methods). We first applied SNPE to observed data x_o created by
191 simulation from the model, calculating the posterior distribution using all 7 features in the observed data (Fig. 4B). The
192 posterior contained the ground truth parameters in a high probability-region, as in previous applications, indicating
193 the consistency of parameter identification. The variance of the posterior was narrower for some parameters than for
194 others, indicating that the 7 data features strongly constrain some parameters (such as the potassium conductance),
195 but only weakly others (such as the adaptation time constant). Additional simulations with parameters sampled from
196 the posterior closely resembled the observed data x_o , in terms of both the raw membrane voltage over time and the 7
197 data features (Fig. 4C, purple and green). Parameters with low posterior probability (outliers) generated simulations
198 that markedly differed from x_o (Fig. 4C, magenta).

199 To investigate how individual data features constrain parameters, we compared SNPE-estimated posteriors based
200 1) solely on the spike count, 2) on the spike count and 3 voltage-features, or 3) on all 7 features of x_o . This analysis
201 revealed that as more features are taken into account, the posterior became narrower and centered more closely on
202 the ground truth parameters (Fig. 4D, Supplementary Fig. 7). Posterior simulations matched the observed data only
203 in those features that had been used for inference, e.g. applying SNPE to spike counts alone identified parameters
204 that generated the correct number of spikes, but for which spike timing and subthreshold voltage time course were
205 off, unless these additional data features were also provided to SNPE (Fig. 4E). For some parameters, such as the
206 potassium conductance, providing more data features brought the peak of the posterior (the *posterior mode*) closer to
207 the ground truth and also decreased uncertainty. For other parameters, such as V_T , a parameter adjusting the spike
208 threshold [66], the peak of the posterior was already close to the correct value with spike counts alone, but adding
209 additional features reduced uncertainty. While SNPE can be used to study the effect of additional data features in
210 reducing parameter uncertainty, this would not be the case for methods that only return a single best-guess estimate
211 of parameters. These results show that SNPE can reveal how information from multiple data features imposes
212 collective constraints on channel and membrane properties in the HH model.

213 We also inferred HH parameters for 8 *in vitro* recordings from the Allen Cell Types database using the same current-
214 clamp stimulation protocol as in our model [60, 70] (Fig. 4F, Supplementary Fig. 8). In each case, simulations based
215 on the SNPE-inferred posterior closely resembled the original data (Fig. 4F). We note that while inferred parameters
216 differed across recordings, some parameters (the spike threshold, the density of sodium channels, the membrane

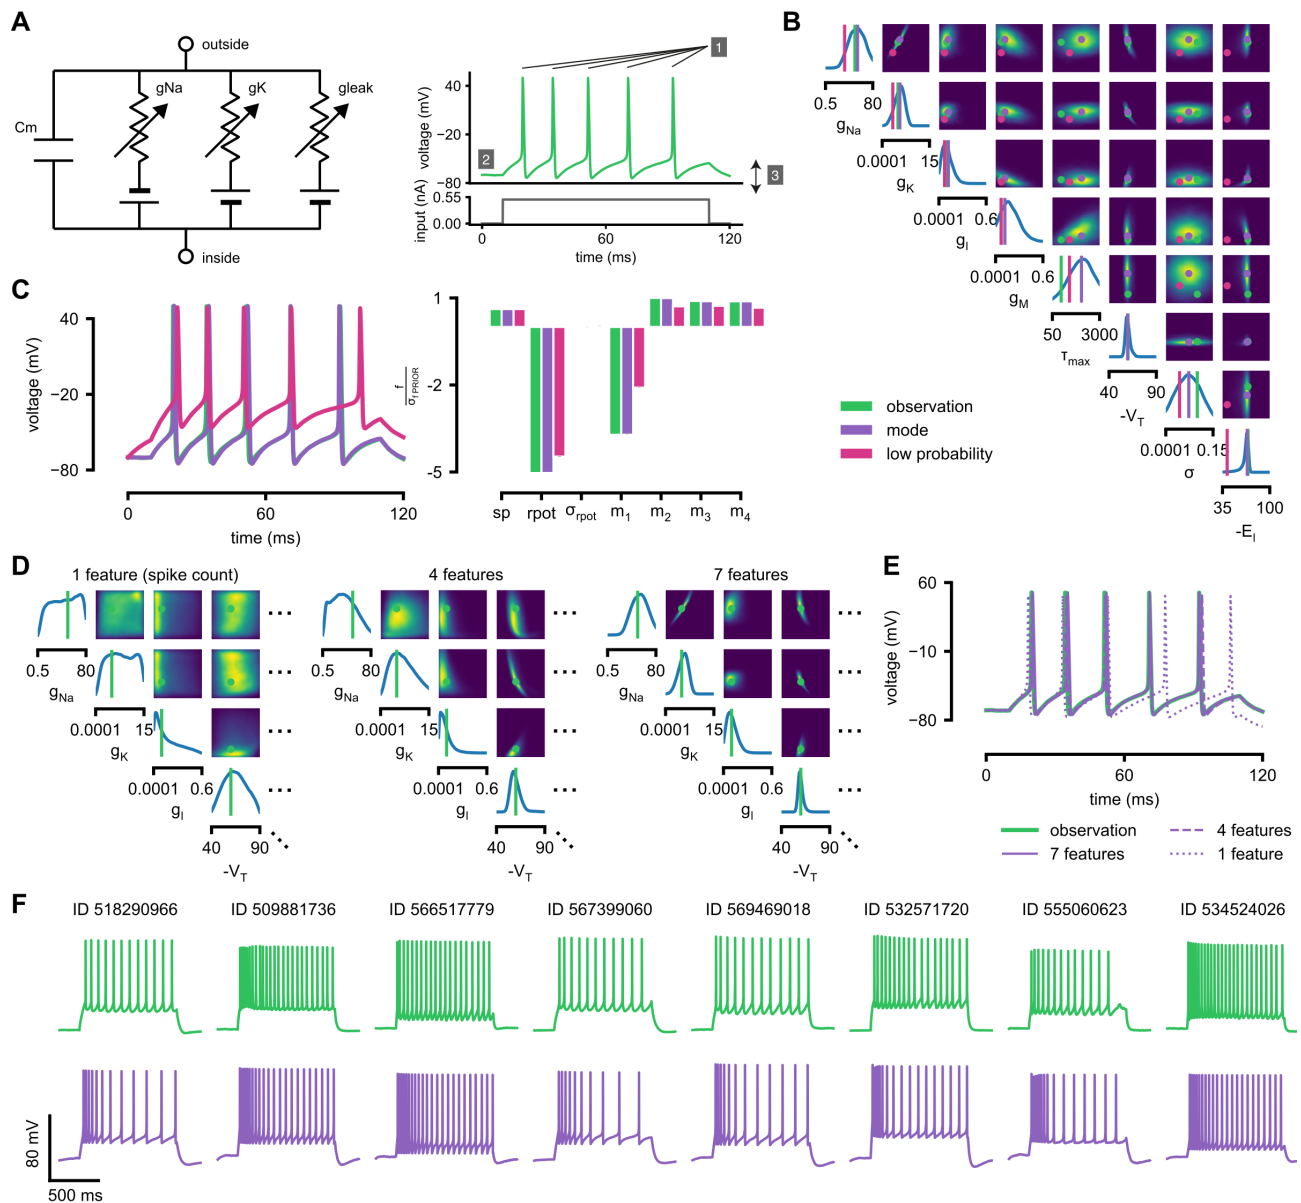


Figure 4. Inference for single compartment Hodgkin-Huxley model. (A) Circuit diagram describing the Hodgkin-Huxley model (left), and simulated voltage-trace given a current input (right). 3 out of 7 voltage features are depicted: (1) number of spikes, (2) mean resting potential and (3) standard deviation of the pre-stimulus resting potential. (B) Inferred posterior for 8 parameters given 7 voltage features. (C) Traces (left) and associated features f (right) for the desired output (observation), the mode of the inferred posterior, and a sample with low posterior probability. The voltage features are: number of spikes sp , mean resting potential $rpot$, standard deviation of the resting potential σ_{rpot} , and the first 4 voltage moments, mean m_1 , standard deviation m_2 , skewness m_3 and kurtosis m_4 . Each feature is normalized by $\sigma_{F \text{ PRIOR}}$, the standard deviation of the respective feature of simulations sampled from the prior. (D) Partial view of the inferred posteriors (4 out of 8 parameters) given 1, 4 and 7 features (full posteriors over 8 parameters in Supplementary Fig. 7). (E) Traces for posterior modes given 1, 4 and 7 features. Increasing the number of features leads to posterior traces that are closer to the observed data. (F) Observations from Allen Cell Types Database (green) and corresponding mode samples (purple). Posteriors in Supplementary Fig. 8.

217 reversal potential and the density of potassium channels) were consistently more strongly constrained than others
 218 (the intrinsic neural noise, the adaptation time constant, the density of slow voltage-dependent channels and the leak
 219 conductance) (Supplementary Fig. 8). Overall, these results suggest that the electrophysiological responses measured
 220 by this current-clamp protocol can be approximated by a single-compartment HH model, and that SNPE can identify
 221 the admissible parameters.

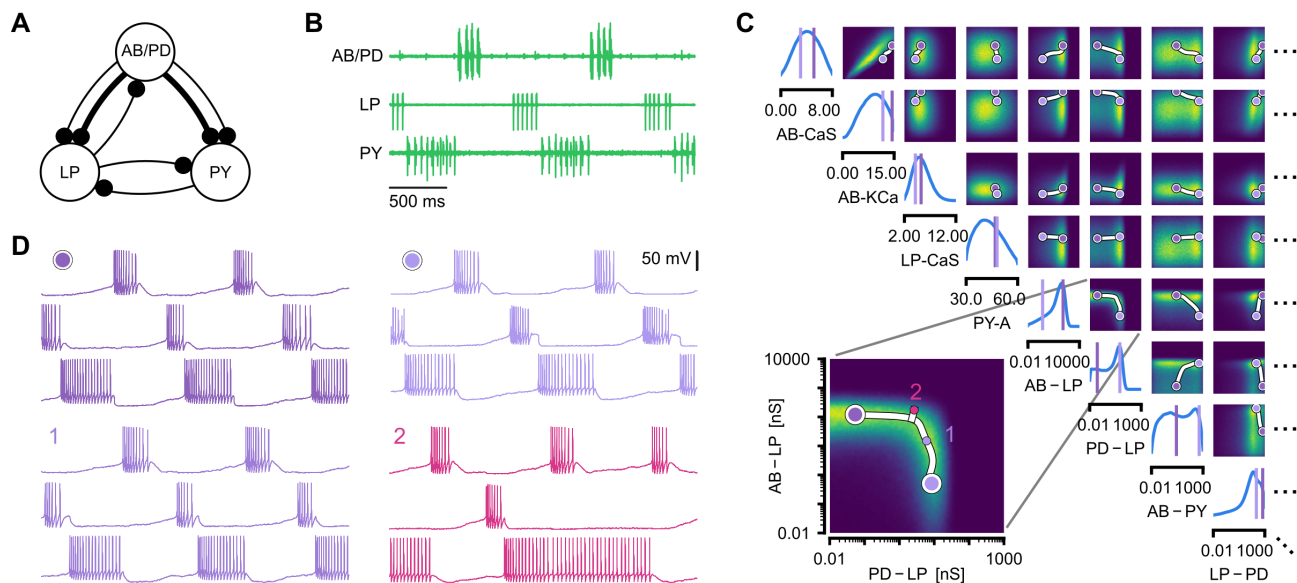


Figure 5. Identifying network models underlying pyloric rhythms in the crustacean stomatogastric ganglion. (A) Circuit diagram of the stomatogastric ganglion. Thin connections are fast glutamatergic, thick connections are slow cholinergic. (B) Extracellular recordings from the crab *Cancer borealis* [71]. (C) Posterior over 24 membrane and 7 synaptic conductances given the experimental observation shown in panel B (8 parameters shown, full posterior in Supplementary Fig. 9). Inset: magnified marginal posterior for the synaptic strengths AB to LP neuron vs. PD to LP neuron. (D) Identifying directions of sloppiness and stiffness. Two samples from the posterior both show similar network activity as the experimental observation (top left and top right), but have very different parameters (purple dots in panel C). Along the high-probability path between these samples, network activity is preserved (trace 1). When perturbing the parameters orthogonally off the path, network activity changes abruptly and becomes non-pyloric (trace 2).

Crustacean stomatogastric ganglion: sensitivity to perturbations

For some biological systems, multiple parameter sets give rise to the same system behavior [7, 18, 72–75]. In particular, neural systems can be robust to specific perturbations of parameters [75–77], yet highly sensitive to others, properties referred to as *sloppiness* and *stiffness* [3, 16, 78]. To demonstrate how SNPE can identify which parameter perturbations affect model outputs, we applied it to a model [7] and data [71] of the pyloric rhythm in the crustacean stomatogastric ganglion (STG). This model describes a triphasic motor pattern generated by a fully characterized circuit (Fig. 5A). The circuit consists of two electrically coupled pacemaker neurons (anterior burster and pyloric dilator, AB/PD), modeled as a single neuron, as well as two types of follower neurons (lateral pyloric (LP) and pyloric (PY)), all connected through inhibitory synapses (details in Methods). Eight membrane conductances are included for each modeled neuron, along with 7 synaptic conductances, for a total of 31 parameters. This model has been used to demonstrate that virtually indistinguishable activity can arise from vastly different membrane and synaptic conductances in the STG [7, 18].

We applied SNPE to an extracellular recording from the STG of the crab *Cancer borealis* [71] which exhibited pyloric activity (Fig. 5B), and inferred the posterior distribution over all 31 parameters based on 18 salient features of the voltage traces, including cycle period, burst durations, burst delays, and phase gaps (Fig. 5C, full posterior in Supplementary Fig. 9, details in Methods). Consistent with previous reports, the posterior distribution has high probability over extended value ranges for many membrane and synaptic conductances. To verify that parameter settings across these extended ranges are indeed capable of generating the experimentally observed network activity, we sampled two sets of membrane and synaptic conductances from the posterior distribution. These two samples have widely disparate parameters from each other (Fig. 5C, purple dots, details in Methods), but both exhibit activity highly similar to the experimental observation (Fig. 5D, top left and top right).

We then investigated the geometry of the parameter space producing these rhythms [17, 18]. First, we wanted to identify directions of sloppiness, and we were interested in whether parameter settings producing pyloric rhythms form a single connected region, as has been shown for single neurons [79], or whether they lie on separate ‘islands.’ Starting from the two above parameter settings showing similar activity, we examined whether they were connected

246 by a path through parameter space along which pyloric activity was maintained. To do this, we algorithmically
247 identified a path lying only in regions of high posterior probability (Fig. 5C, white, details in Methods). Along the path,
248 network output was tightly preserved, despite a substantial variation of the parameters (voltage trace 1 in Fig. 5D,
249 Supplementary Fig. 10A,C). Second, we inspected directions of stiffness by perturbing parameters off the path. We
250 applied perturbations that yield maximal drops in posterior probability (see Methods for details), and found that the
251 network quickly produced non-pyloric activity (voltage trace 2, Fig. 5D). In identifying these paths and perturbations,
252 we exploited the fact that SNPE provides a differentiable estimate of the posterior, as opposed to parameter search
253 methods which provide only discrete samples.

254 Overall, these results show that the pyloric network can be robust to specific perturbations in parameter space, but
255 sensitive to others, and that one can interpolate between disparate solutions while preserving network activity. This
256 analysis demonstrates the flexibility of SNPE in capturing complex posterior distributions, and how the differentiable
257 posterior can be used to study directions of sloppiness and stiffness.

258 **Predicting compensation mechanisms from posterior distributions**

259 Experimental and computational studies have shown that stable neural activity can be maintained despite variable
260 circuit parameters [7, 82, 83]. This behavior can emerge from two sources [82]: either the variation of a certain
261 parameter barely influences network activity at all, or alternatively, the variation of several parameters influence
262 network activity, but their effects compensate for one another. Here, we investigated these possibilities by using the
263 posterior distribution over membrane and synaptic conductances of the STG.

264 We begin by drawing samples from the posterior and inspecting their pairwise histograms (Fig. 6A, full posterior
265 over all parameters in Supplementary Fig. 9). Consistent with previously reported results [84], we found that most
266 pairs of parameters are only weakly correlated (Fig. 6B). However, in these histograms over two parameters, all other
267 parameters are fully unconstrained and can take on diverse values, which could blur out compensation mechanisms.
268 Therefore, we held all but 2 parameters constant at a given consistent circuit configuration (sampled from the
269 posterior), and observed the network activity across different values of the remaining pair of parameters. We can do
270 so by calculating the conditional posterior distribution (details in Methods), and do not have to generate additional
271 simulations (as would be required by parameter search methods). Doing so has a simple interpretation: when all but
272 2 parameters are fixed, what values of the remaining 2 parameters can then lead to pyloric activity? We found that
273 pyloric activity can emerge only from narrowly tuned and often highly correlated combinations of the remaining 2
274 parameters, showing how these parameters can compensate for one another (Fig. 6C). When repeating this analysis
275 across multiple network configurations, i.e. when holding the parameters fixed at different values, we found that
276 these 'conditional correlations' are often preserved (Fig. 6C, left and right). We calculated conditional correlations
277 for each parameter pair using 500 different circuit configurations sampled from the posterior (Fig. 6D). Compared to
278 correlations based on the pairwise histograms (Fig. 6B), these conditional correlations were substantially stronger.
279 They were particularly strong across membrane conductances of the same neuron, but primarily weak across different
280 neurons (black boxes in Fig. 6D).

281 Finally, we tested whether the conditional correlations were in line with experimental observations. For the PD
282 and the LP neuron, it has been reported that overexpression of the fast transient potassium current leads to a
283 compensating increase of the hyperpolarization current, suggesting a positive correlation between these two currents
284 [80, 85]. Also, using current injections into the LP neuron, a positive correlation has been reported between the
285 strength of the synaptic input and the maximal conductance of the hyperpolarization current [81]. These results are
286 qualitatively consistent with the conditional correlations (Fig. 6E), which were positive both between the fast transient
287 potassium and hyperpolarization currents for all three model neurons, as well as for 6 out of 7 correlations between
288 synaptic input strength and hyperpolarization current.

289 Overall, we showed how SNPE can be used to study parameter dependencies in circuits with parameter degeneracy,
290 and how the posterior distribution can be used to efficiently explore potential compensation mechanisms. We found
291 that our method can predict compensation mechanisms which are qualitatively consistent with experimental studies.

292 **Discussion**

293 How can we build models which give insights into the causal mechanisms underlying neural or behavioral dynamics?
294 The cycle of building mechanistic models, generating predictions, comparing them to empirical data, and rejecting,

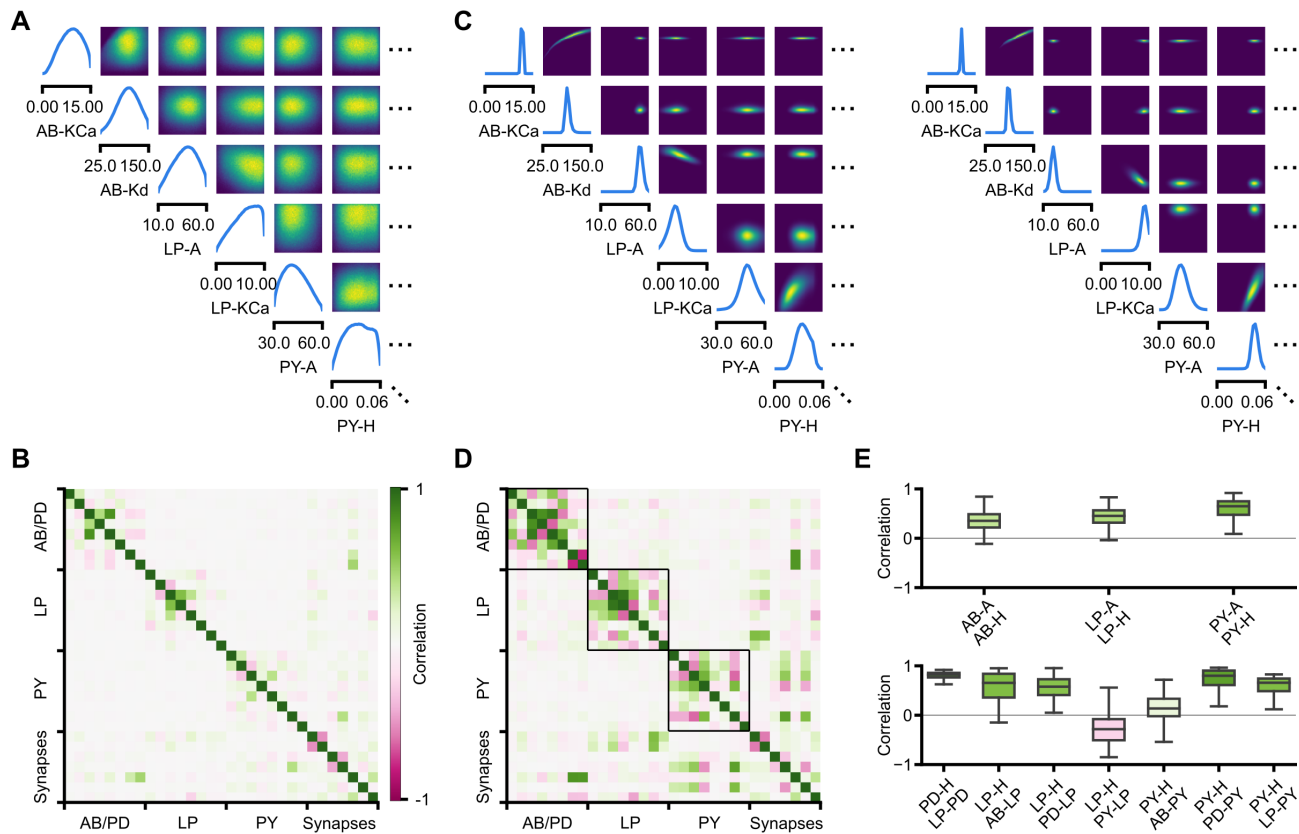


Figure 6. Investigating compensation mechanisms in the stomatogastric ganglion. (A) Inferred posterior. We show a subset of parameters which are weakly constrained (full posterior in Supplementary Fig. 9). Pyloric activity can emerge from a wide range of maximal membrane conductances, as the 1D and 2D posterior marginals cover almost the entire extent of the prior. (B) Correlation matrix, based on the samples shown in panel A. Almost all correlations are weak. (C) Conditional distributions given a particular circuit configuration: For the plots on the diagonal, we keep all but one parameter fixed. For plots above the diagonal, we keep all but two parameters fixed. The remaining parameter(s) are narrowly tuned, tuning across parameters is often highly correlated. When conditioning on a different parameter setting (right plot), the conditional posteriors change, but correlations are often maintained. (D) Conditional correlation matrix, averaged over 500 conditional distributions like the ones shown in panel C. Black squares highlight parameter-pairs within the same model neuron. (E) Consistency with experimental observations. Top: Maximal conductance of the fast transient potassium current and the maximal conductance of the hyperpolarization current are positively correlated for all three neurons. This has also been experimentally observed in the PD and the LP neuron [80]. Bottom: Maximal conductance of the hyperpolarization current is positively correlated with the strength of the synaptic input for six out of seven synapses in the circuit, as experimentally observed using current injection into the LP neuron [81]. The boxplots indicate the maximum, 75% quantile, median, 25% quantile, and minimum across 500 conditional correlations for different parameter pairs. Face color indicates mean correlation using the colorbar shown in panel B.

295 or refining models has been of crucial importance in the empirical sciences. However, a key challenge has been the
 296 difficulty of identifying mechanistic models which can quantitatively capture observed phenomena. We predict that
 297 a generally applicable tool to constrain mechanistic models by data is going to expedite progress in neuroscience.
 298 While many considerations should go into designing a model that is appropriate for a given question and level
 299 of description [2, 3, 86, 87], the question of whether and how one can perform statistical inference on the model
 300 should not compromise model-design. In our tool, SNPE, the process of model building and parameter inference
 301 are entirely decoupled. We illustrated the power of our approach on a diverse set of applications, highlighting the
 302 potential of SNPE to rapidly identify data-compatible mechanistic models, to investigate which data-features constrain
 303 which parameters, to reveal shortcomings of candidate-models, and to explore the parameter-landscape of a neural
 304 oscillator and provide hypotheses for compensation mechanisms.

305 Related work

306 SNPE builds on recent advances in machine learning, and in particular in density-estimation approaches to likelihood-
307 free inference [36–39, 88, 89]. The idea of learning inference networks on simulated data can be traced back to
308 *regression-adjustment* methods in ABC [33, 90]. Papamakarios and Murray (2016) [36] first proposed to use expressive
309 conditional density estimators in the form of deep neural networks [41, 57], and to optimize them sequentially
310 over multiple rounds with cost-functions derived from Bayesian inference principles. Compared to commonly
311 used rejection-based ABC methods [91, 92], such as MCMC-ABC [34], SMC-ABC [35, 93], Bayesian-Optimization ABC
312 [94], or ensemble methods [95, 96], SNPE approaches do not require one to define a distance function in data
313 space. In addition, by leveraging the ability of neural networks to learn informative features, they enable scaling
314 to high-dimensional estimation problems, as are common in neuroscience and other fields in biology. Alternative
315 likelihood-free approaches include *synthetic likelihood* methods [97–102], moment-based approximations of the
316 posterior [103, 104], inference compilation [105, 106], and density-ratio estimation [107]. For some mechanistic
317 models in neuroscience (e.g. for integrate-and-fire neurons), likelihoods can be computed via stochastic numerical
318 approximations [67, 108, 109] or model-specific analytical approaches [68, 110–113].

319 Finally, a complementary approach to mechanistic modeling is to pursue purely phenomenological models, which
320 are designed to have favorable statistical and computational properties: these data-driven models can be efficiently
321 fit to neural data [19–25, 44, 46] or to implement desired computations [114]. Although tremendously useful for a
322 quantitative characterization of neural dynamics, these models typically have a large number of parameters, which
323 rarely correspond to physically measurable or mechanistically interpretable quantities, and thus it can be challenging
324 to derive mechanistic insights or causal hypotheses from them (but see e.g. [115–117]).

325 Use of summary statistics

326 When fitting mechanistic models to data, it is common to target summary statistics to isolate specific behaviors,
327 rather than the full data. For example, the spike shape is known to constrain sodium and potassium conductances
328 [29, 30, 66]. When modeling population dynamics, it is often desirable to achieve realistic firing rates, rate-correlations
329 and response nonlinearities [42, 118], or specified oscillations [7]. In models of decision making, one is often interested
330 in reproducing psychometric functions or reaction-time distributions [119]. Choice of summary statistics might also
331 be guided by known limitations of either the model or the measurement approach, or necessitated by the fact
332 that published data are only available in summarized form. Several methods have been proposed to automatically
333 construct informative summary statistics [120–122]. SNPE can be applied to, and might benefit from the use of
334 summary statistics, but it also makes use of the ability of neural networks to automatically learn informative features
335 in high-dimensional data. Thus, SNPE can also be applied directly to raw data (e.g. using recurrent neural networks
336 [37]), or to high-dimensional summary statistics which are challenging for ABC approaches (Fig. 2). In all cases, care is
337 needed when interpreting models fit to summary features, as choice of features can influence the results [120–122].

338 Applicability and limitations

339 A key advantage of SNPE is its general applicability: it can be applied whenever one has a simulator that allows to
340 stochastically generate model outputs from specific parameters. Furthermore, it can be applied in a fully ‘black-box
341 manner’, i.e. does not require access to the internal workings of the simulator, likelihoods or gradients. It does not
342 impose any other limitations on the model or the summary features, and in particular does not require them to be
343 differentiable. However, it also has limitations: First, current implementations of SNPE scale well to high-dimensional
344 observations (~100s dims, also see [38]), but scaling to high-dimensional parameter spaces (>30) is challenging.
345 Second, while it is a long-term goal for these approaches to be made fully automatic, our current implementation still
346 requires choices by the user: As described in Methods, one needs to provide the architecture of the density estimation
347 network, and specify settings related to network-optimisation, and the number of simulations and inference rounds.
348 These settings depend on the complexity of the relation between summary statistics and model parameters, and the
349 number of simulations that can be afforded. In the documentation accompanying our code-package, we provide
350 examples and guidance. For small-scale problems, we have found SNPE to be robust to these settings. However, for
351 challenging, high-dimensional applications, SNPE might currently require substantial user interaction. Third, the power
352 of SNPE crucially rests on the ability of deep neural networks to perform density estimation. While deep nets have had
353 ample empirical success, we still have an incomplete understanding of their limitations, in particular in cases where

354 the mapping between data and parameters might not be smooth (e.g. near phase transitions). Fourth, when applying
355 SNPE (or any other model-identification approach), validation of the results is of crucial importance, both to assess
356 the accuracy of the inference procedure, as well as to identify possible limitations of the mechanistic model itself.
357 In the example applications, we used several procedures for assessing the quality of the inferred posteriors. One
358 common ingredient of these approaches is to sample from the inferred model, and search for systematic differences
359 between observed and simulated data, e.g. to perform *posterior predictive checks* [37, 38, 93, 123, 124] (Fig. 2G,
360 Fig. 3F,G, Fig. 4C, and Fig. 5D). There are challenges and opportunities ahead in further scaling and automating
361 simulation-based inference approaches. However, in its current form, SNPE will be a powerful tool for quantitatively
362 evaluating mechanistic hypotheses on neural data, and for designing better models of neural dynamics.

363 **Acknowledgments**

364 We thank Mahmood S. Hoseini and Michael Stryker for sharing their data for Fig. 2, and Philipp Berens, Sean Bittner,
365 Jan Boelts, John Cunningham, Richard Gao, Scott Linderman, Eve Marder, Iain Murray, George Papamakarios, Astrid
366 Prinz, Auguste Schulz and Srinivas Turaga for discussions and/or comments on the manuscript. This work was
367 supported by the German Research Foundation (DFG) through SFB 1233 'Robust Vision', (276693517), SFB 1089
368 'Synaptic Microcircuits' and SPP 2041 'Computational Connectomics', the German Federal Ministry of Education and
369 Research (BMBF, project 'ADIMEM', FKZ 01IS18052 A-D), and UK Research and Innovation, Biotechnology and Biological
370 Sciences Research Council (UKRI-BBSRC) through BB/N019512/1.

371 **Methods**

372 **Code availability**

373 Code implementing SNPE is available at <http://www.mackelab.org/delfi/>.

374 **Simulation-based inference**

375 To perform Bayesian parameter identification with SNPE, three types of input need to be specified:

- 376 1. A mechanistic model. The model only needs to be specified through a simulator, i.e. that one can generate a
377 simulation result \mathbf{x} for any parameters θ . We do not assume access to the likelihood $p(\mathbf{x}|\theta)$ or the equations
378 or internals of the code defining the model, nor do we require the model to be differentiable. This is in
379 contrast to many alternative approaches (including [42]), which require the model to be differentiable and to be
380 implemented in a software code that is amenable to automatic differentiation packages. Finally, SNPE can both
381 deal with inputs \mathbf{x} which resemble ‘raw’ outputs of the model, or summary features/statistics calculated from
382 data.
- 383 2. Observed data \mathbf{x}_o of the same form as the results \mathbf{x} produced by model simulations.
- 384 3. A prior distribution $p(\theta)$ describing the range of possible parameters. $p(\theta)$ could consist of upper and lower
385 bounds for each parameter, or a more complex distribution incorporating mechanistic first principles or
386 knowledge gained from previous inference procedures on other data.

387 For each problem, our goal was to estimate the posterior distribution $p(\theta|\mathbf{x}_o)$. To do this we used SNPE [36–38].
388 Setting up the inference procedure required three design choices:

- 389 1. A network architecture, including number of layers, units per layer, layer type (feedforward or convolutional),
390 activation function and skip connections.
- 391 2. A parametric family of probability densities $q_{\psi}(\theta)$ to represent inferred posteriors, to be used as conditional
392 density estimator. We used either a mixture of Gaussians (MoG) or a masked autoregressive flow (MAF) [41]. In
393 the former case, the number of components K must be specified; in the latter the number of *MADES* (Masked
394 Autoencoder for Distribution Estimation) n_{MADES} . Both choices are able to represent richly structured, and
395 multi-modal posterior distributions.
- 396 3. A simulation budget, i.e. number of rounds R and simulations per round N_r .

397 We emphasize that SNPE is highly modular, i.e. that the the inputs (data, the prior over parameter, the mechanistic
398 model), and algorithmic components (network architecture, probability density, optimization approach) can all be
399 modified and chosen independently. This allows neuroscientists to work with models which are designed with
400 mechanistic principles—and not convenience of inference—in mind. Furthermore, it allows SNPE to benefit from
401 advances in more flexible density estimators, more powerful network architectures, or optimization strategies.

402 With the problem and inference settings specified, SNPE adjusts the network weights ϕ based on simulation results,
403 so that $p(\theta|\mathbf{x}) \approx q_{F(\mathbf{x},\phi)}(\theta)$ for any \mathbf{x} . In the first round of SNPE simulation parameters are drawn from the prior $p(\theta)$. If
404 a single round of inference is not sufficient, SNPE can be run in multiple rounds, in which samples are drawn from the
405 version of $q_{F(\mathbf{x}_o,\phi)}(\theta)$ at the beginning of the round. After the last round, $q_{F(\mathbf{x}_o,\phi)}$ is returned as the inferred posterior on
406 parameters θ given observed data \mathbf{x}_o . If SNPE is only run for a single round, then the generated samples only depend
407 on the prior, but not on \mathbf{x}_o : In this case, the inference network is applicable to any data (covered by the prior ranges),
408 and can be used for rapid amortized inference.

409 SNPE learns the correct network weights ϕ by minimizing the objective function $\sum_j \mathcal{L}(\theta_j, \mathbf{x}_j)$ where the simulation
410 with parameters θ_j produced result \mathbf{x}_j . For the first round of SNPE $\mathcal{L}(\theta_j, \mathbf{x}_j) = -\log q_{F(\mathbf{x},\phi)}$, while in subsequent rounds
411 a different loss function accounts for the fact that simulation parameters were not sampled from the prior. Different
412 choices of the loss function for later rounds result in SNPE-A [36], SNPE-B [37] or SNPE-C algorithm [38]. To optimize
413 the networks, we used ADAM with default settings [125].

414 The details of the algorithm are below:

Algorithm 1: SNPE

Input: simulator with (implicit) density $p(\mathbf{x}|\theta)$, observed data \mathbf{x}_o , prior $p(\theta)$, density family q_ψ , neural network $F(\mathbf{x}, \phi)$, number of rounds R , simulation count for each round N_r

randomly initialize ϕ

$\tilde{p}_1(\theta) := p(\theta)$

$N := 0$

for $r = 1$ **to** R **do**

415

for $i = 1 \dots N_r$ **do**

 sample $\theta_{N+i} \sim \tilde{p}_r(\theta)$

 simulate $\mathbf{x}_{N+i} \sim p(\mathbf{x}|\theta_{N+i})$

$N \leftarrow N + N_r$

 train $\phi \leftarrow \arg \min_{\phi} \sum_{j=1}^N \mathcal{L}(\theta_j, \mathbf{x}_j)$

$\tilde{p}_r(\theta) := q_{F(\mathbf{x}_o, \phi)}(\theta)$

return $q_{F(\mathbf{x}_o, \phi)}(\theta)$

416 **Linear-nonlinear encoding models**

We used a Linear-Nonlinear (LN) encoding model (a special case of a generalized linear model, GLM, [19, 21, 44–47]) to simulate the activity of a neuron in response to a univariate time-varying stimulus. Neural activity z_i was subdivided in $T = 100$ bins and, within each bin i , spikes were generated according to a Bernoulli observation model,

$$z_i \sim \text{Bern}(\eta(\mathbf{v}_i^\top \mathbf{f} + \beta)),$$

417 where \mathbf{v}_i is a vector of white noise inputs between time bins $i - 8$ and i , \mathbf{f} a length-9 linear filter, β is the bias, and
 418 $\eta(\cdot) = \exp(\cdot)/(1 + \exp(\cdot))$ is the canonical inverse link function for a Bernoulli GLM. As summary statistics, we used
 419 the total number of spikes N and the spike-triggered average $\frac{1}{N} \mathbf{Vz}$, where $\mathbf{V} = [v_1, v_2, \dots, v_T]$ is the so-called design
 420 matrix of size $9 \times T$. We note that the spike-triggered sum \mathbf{Vz} constitutes sufficient statistics for this GLM, i.e. that
 421 selecting the STA and N together as summary statistics does not lead to loss of model relevant information over the
 422 full input-output dataset $\{\mathbf{V}, \mathbf{z}\}$. We used a Gaussian prior with zero mean and covariance matrix $\Sigma_\beta = \sigma^2(\mathbf{F}^\top \mathbf{F})^{-1}$,
 423 where \mathbf{F} encourages smoothness by penalizing the second-order differences in the vector of parameters [126].

424 For inference, we used a single round of 10000 simulations, and the posterior was approximated with a Gaussian
 425 distribution ($\theta \in \mathbb{R}^{10}$, $\mathbf{x} \in \mathbb{R}^{10}$). We used a feedforward neural network with two hidden layers of 50 units each. We
 426 used a Polya Gamma Markov Chain Monte Carlo sampling scheme [48] to estimate a reference posterior.

For the spatial receptive field model of a cell in primary visual cortex, we simulated the activity of a neuron depending on an image-valued stimulus. Neural activity was subdivided in bins of length $\Delta t = 0.025\text{s}$ and within each bin i , spikes were generated according to a Poisson observation model,

$$z_i \sim \text{Poiss}(\eta(\mathbf{v}_i^\top \mathbf{h} + \beta)),$$

where \mathbf{v}_i is the vectorized white noise stimulus at time bin i , \mathbf{h} a 41×41 linear filter, β is the bias, and $\eta(\cdot) = \exp(\cdot)$ is the canonical inverse link function for a Poisson GLM. The receptive field \mathbf{h} is constrained to be a Gabor filter:

$$h(g_x, g_y) = g \exp\left(-\frac{x'^2 + r^2 y'^2}{2\sigma^2}\right) \cos(2\pi f x' - \phi)$$

$$x' = (g_x - x) \cos \psi - (g_y - y) \sin \psi$$

$$y' = (g_x - x) \sin \psi + (g_y - y) \cos \psi$$

$$\sigma = \frac{\sqrt{2} 2^w + 1}{2\pi f 2^w - 1},$$

427 where (g_x, g_y) is a regular grid of 41×41 positions spanning the 2D image-valued stimulus. The parameters of the
428 Gabor are gain g , spatial frequency f , aspect-ratio r , width w , phase ϕ (between 0 and π), angle ψ (between 0 and
429 2π) and location x, y (assumed within the stimulated area, scaled to be between -1 and 1). Bounded parameters
430 were transformed with a log-, or logit-transform, to yield unconstrained parameters. After applying SNPE, we back-
431 transformed both the parameters and the estimated posteriors in closed form, as shown in Fig. 2. We did not
432 transform the parameters bias β and gain g .

We used a factorizing Gaussian prior for the vector of transformed Gabor parameters

$$[g, \log f, \log r, \log w, l_{0,\pi}(\phi), l_{0,2\pi}(\psi), l_{-1,1}(x), l_{-1,1}(y)],$$

433 where transforms $l_{0,\pi}(X) = \log(X/(2\pi - X))$, $l_{0,2\pi}(X) = \log(X/(\pi - X))$, $l_{-1,1}(X) = \log((X + 1)/(1 - X))$ ensured the
434 assumed ranges for the Gabor parameters ϕ, ψ, x, y . Our Gaussian prior had zero mean and standard deviations
435 $[2, 0.5, 0.5, 0.5, 1.9, 1.78, 1.78, 1.78]$. We note that a Gaussian prior on a logit-transformed random variable $\text{logit}X$ with
436 zero mean and standard deviation around 1.78 is close to a uniform prior over the original variable X . For the bias β ,
437 we used a Gaussian prior with mean -0.57 and variance 1.63, which approximately corresponds to an exponential
438 prior $\exp(\beta) \sim \text{Exp}(\lambda)$ with rate $\lambda = 1$ on the baseline firing rate $\exp(\beta)$ in absence of any stimulus.

439 The ground-truth parameters for the demonstration in Fig. 2 were chosen to give an asymptotic firing rate of 1Hz
440 for 5 minutes stimulation, resulting in 307 spikes, and a signal-to-noise ratio of -12dB .

441 As summary statistics, we used the total number of spikes N and the spike-triggered average $\frac{1}{N}\mathbf{Vz}$, where $\mathbf{V} =$
442 $[v_1, v_2, \dots, v_T]$ is the stimulation video of length $T = 300/\Delta t = 12000$. As for the GLM with a temporal filter, the
443 spike-triggered sum \mathbf{Vz} constitutes sufficient statistics for this GLM.

444 For inference, we applied SNPE-A with in total 2 rounds: an initial round serves to first roughly identify the relevant
445 region of parameter space. Here we used a Gaussian distribution to approximate the posterior from 10000 simulations
446 each. A second round then used a mixture of 8 Gaussian components to estimate the exact shape of the posterior
447 from another 100000 simulations ($\theta \in \mathbb{R}^9, \mathbf{x} \in \mathbb{R}^{1682}$). We used a convolutional network with 5 convolutional layers with
448 16 to 32 convolutional filters followed by three fully connected layers with 50 units each. The total number of spikes N
449 within a simulated experiment was passed as an additional input directly to the fully-connected layers of the network.

450 Similar to the previous GLM, this model has a tractable likelihood, so we use MCMC to obtain a reference posterior.

451 We applied this approach to extracellular recordings from primary visual cortex of alert mice obtained using silicon
452 microelectrodes in response to colored-noise visual stimulation. Experimental methods are described in Dyballa et al
453 2018 [53].

454 Comparison with Sequential Monte Carlo (SMC) ABC

455 In order to illustrate the competitive performance of SNPE, we obtained a posterior estimate with a classical ABC
456 method, Sequential Monte Carlo (SMC) ABC [35, 127]. Likelihood-free inference methods from the ABC family require
457 a distance function $d(\mathbf{x}_o, \mathbf{x})$ between observed data \mathbf{x}_o and possible simulation outputs \mathbf{x} to characterize dissimilarity
458 between simulations and data. A common choice is the (scaled) Euclidean distance $d(\mathbf{x}_o, \mathbf{x}) = \|\mathbf{x} - \mathbf{x}_o\|_2$. The Euclidean
459 distance here was computed over 1681 summary statistics given by the spike-triggered average (one per pixel) and a
460 single summary statistic given by the 'spike count'. To ensure that the distance measure was sensitive to differences in
461 both STA and spike count, we scaled the summary statistic 'spike count' to account for about 20% of the average total
462 distance (other values did not yield better results). The other 80% were computed from the remaining 1681 summary
463 statistics given by spike-triggered averages. To showcase how this situation is challenging for ABC approaches, we
464 generated 10000 input-output pairs $(\theta_i, \mathbf{x}_i) \sim p(\mathbf{x}|\theta)p(\theta)$ with the prior and simulator used above, and illustrate the 10
465 STAs and spike counts with closest $d(\mathbf{x}_o, \mathbf{x}_i)$ in Supplementary Fig. 3A. Spike counts were comparable to the observed
466 data (307 spikes), but STAs are noise-dominated and the 10 'closest' underlying receptive fields (yellow contours) show
467 substantial variability in location and shape of the receptive field. If even the 'closest' samples do not show any visible
468 receptive field, then there is little hope that even an appropriately chosen acceptance threshold will yield a good
469 approximation to the posterior. This findings were also reflected in the results from SMC-ABC with a total simulation
470 budget of 10^6 simulations (Fig. 3B). The estimated posterior marginals for 'bias' and 'gain' parameters show that the
471 parameters related to the firing rate were constrained by the data \mathbf{x}_o , but marginals of parameters related to shape
472 and location of the receptive field did not differ from the prior, highlighting that SMC-ABC was here not able to identify

473 the posterior distribution. Further comparisons of neural-density estimation approaches with ABC-methods can be
474 found in the publications describing the underlying machine-learning methodologies [36, 38, 101].

475 Ion channel models

We simulated non-inactivating potassium channel currents subject to voltage-clamp protocols as:

$$I_K = \bar{g}_K m(V - E_K),$$

where V is the membrane potential, \bar{g}_K is the density of potassium channels, E_K is the reversal potential of potassium, and m is the gating variable for potassium channel activation. m is modeled according to the first-order kinetic equation

$$\frac{dm}{dt} = \frac{m_\infty(V) - m}{\tau_m(V)},$$

where $m_\infty(V)$ is the steady-state activation, and $\tau_m(V)$ the respective time constant. We used a general formulation of $m_\infty(V)$ and $\tau_m(V)$ [59], where the steady-state activation curve has 2 parameters (slope and offset) and the time constant curve has 6 parameters, amounting to a total of 8 parameters (θ_1 to θ_8):

$$m_\infty(V) = \frac{1}{1 + e^{-\theta_1 V + \theta_2}}$$

$$\tau_m(V) = \frac{\theta_4}{e^{-[\theta_5(V - \theta_3) + \theta_6(V - \theta_3)^2]} + e^{[\theta_7(V - \theta_3) + \theta_8(V - \theta_3)^2]}}.$$

476 Since this model can be used to describe the dynamics of a wide variety of channel models, we refer to it as *Omnimodel*.

477 We modeled responses of the Omnimodel to a set of five voltage-clamp protocols described in [58]. Current
478 responses were reduced to 55 summary statistics (11 per protocol). Summary statistics were coefficients to basis
479 functions derived via Principal Components Analysis (PCA) (10 per protocol) plus a linear offset (1 per protocol) found
480 via least-squares fitting. PCA basis functions were found by simulating responses of the non-inactivating potassium
481 channel models to the five voltage-clamp protocols and reducing responses to each protocol to 10 dimensions
482 (explaining 99.9% of the variance).

483 To amortize inference on the model, we specified a wide uniform prior over the parameters: $\theta_1 \in \mathcal{U}(0, 1)$, $\theta_2 \in$
484 $\mathcal{U}(-10., 10.)$, $\theta_3 \in \mathcal{U}(-120., 120.)$, $\theta_4 \in \mathcal{U}(0., 2000)$, $\theta_5 \in \mathcal{U}(0., 0.5)$, $\theta_6 \in \mathcal{U}(0, 0.05)$, $\theta_7 \in \mathcal{U}(0., 0.5)$, $\theta_8 \in \mathcal{U}(0, 0.05)$.

485 For inference, we trained a shared inference network in a single round of 10^6 simulations generated by sampling
486 from the prior ($\theta \in \mathbb{R}^8$, $\mathbf{x} \in \mathbb{R}^{55}$). The density estimator is a masked autoregressive flow (MAF) [41] with five MADEs
487 with [250,250] hidden units each.

488 We evaluated performance on 350 non-inactivating potassium ion channels selected from IonChannelGenealogy
489 (ICG) by calculating the correlation coefficient between traces generated by the original model and traces from the
490 Omnimodel using the posterior mode.

491 Single-compartment Hodgkin-Huxley neurons

We simulated a single-compartment Hodgkin-Huxley type neuron with channel kinetics as in [66],

$$C_m \frac{dV}{dt} = g_l(E_l - V) + \bar{g}_{Na} m^3 h(E_{Na} - V) + \bar{g}_K n^4(E_K - V) + \bar{g}_M p(E_K - V) + I_{inj} + \sigma \eta(t)$$

$$\frac{dq}{dt} = \frac{q_\infty(V) - q}{\tau_q(V)}, q \in \{m, h, n, p\},$$

492 where V is the membrane potential, C_m is the membrane capacitance, g_l is the leak conductance, E_l is the membrane
493 reversal potential, \bar{g}_c is the density of channels of type c (Na^+ , K^+ , M), E_c is the reversal potential of c , (m, h, n, p) are the
494 respective channel gating kinetic variables, and $\sigma \eta(t)$ is the intrinsic neural noise. The right hand side of the voltage
495 dynamics is composed of a leak current, a voltage-dependent Na^+ current, a delayed-rectifier K^+ current, a slow
496 voltage-dependent K^+ current responsible for spike-frequency adaptation, and an injected current I_{inj} . Channel gating
497 variables q have dynamics fully characterized by the neuron membrane potential V , given the respective steady-state
498 $q_\infty(V)$ and time constant $\tau_q(V)$ (details in [66]). Two additional parameters are implicit in the functions $q_\infty(V)$ and
499 $\tau_q(V)$: V_T adjusts the spike threshold through m_∞ , h_∞ , n_∞ , τ_m , τ_h and τ_n ; τ_{\max} scales the time constant of adaptation

500 through $\tau_p(V)$ (details in [66]). We set $E_{Na} = 53$ mV and $E_K = -107$ mV, similar to the values used for simulations in
 501 Allen Cell Types Database (<http://help.brain-map.org/download/attachments/8323525/BiophysModelPeri.pdf>).

502 We applied SNPE to infer the posterior over 8 parameters (\bar{g}_{Na} , \bar{g}_K , g_I , \bar{g}_M , τ_{max} , V_T , σ , E_I), given 7 voltage features
 503 (number of spikes, mean resting potential, standard deviation of the resting potential, and the first 4 voltage moments,
 504 mean, standard deviation, skewness and kurtosis).

The prior distribution over the parameters was uniform,

$$\theta \sim \mathcal{U}(p_{low}, p_{high}),$$

505 where $p_{low} = [0.5, 10^{-4}, 10^{-4}, 10^{-4}, 50, 40, 10^{-4}, 35]$ and $p_{high} = [80, 15, 0.6, 0.6, 3000, 90, 0.15, 100]$. These ranges are
 506 similar to the ones obtained in [66].

507 For inference in simulated data, we used a single round of 100000 simulations ($\theta \in \mathbb{R}^8$, $\mathbf{x} \in \mathbb{R}^{11}$). The density
 508 estimator was a masked autoregressive flow (MAF) [41] with five MADEs with [50,50] hidden units each.

509 For the inference on in vitro recordings from mouse cortex (Allen Cell Types Database, <https://celltypes.brain-map.org/data>), we selected 8 recordings corresponding to spiny neurons with at least 10 spikes during the current-
 510 clamp stimulation. The respective cell identities and sweeps are: (518290966,57), (509881736,39), (566517779,46),
 511 (567399060,38), (569469018,44), (532571720,42), (555060623,34), (534524026,29). For each recording, SNPE-B was run
 512 for 2 rounds with 125000 Hodgkin-Huxley simulations each, and the posterior was approximated by a mixture of two
 513 Gaussians. In this case, the density estimator was composed of two fully connected layers of 100 units each.
 514

515 **Circuit model of the crustacean stomatogastric ganglion**

516 We used extracellular recordings from the crab *Cancer borealis* [71]. The preparations from the stomatogastric ganglion
 517 were decentralized, i.e. the input from descending modulatory inputs was removed. The data was recorded at a
 518 temperature of 11°C. See Haddad & Marder (2018) [71] for full experimental details.

We simulated the circuit model of the crustacean stomatogastric ganglion by adapting a model described in [7].
 The model is composed of three single-compartment neurons, AB/PD, LP, and PD, where the electrically coupled AB
 and PD neurons are modeled as a single neuron. Each of the model neurons contains 8 currents, a Na^+ current I_{Na} ,
 a fast and a slow transient Ca^{2+} current I_{CaT} and I_{CaS} , a transient K^+ current I_A , a Ca^{2+} -dependent K^+ current I_{KCa} , a
 delayed rectifier K^+ current I_{Kd} , a hyperpolarization-activated inward current I_H , and a leak current I_{leak} . In addition, the
 model contains 7 synapses. As in [7], these synapses were simulated using a standard model of synaptic dynamics
 [128]. The synaptic input current into the neurons is given by $I_s = g_s s (V_{post} - E_s)$, where g_s is the maximal synapse
 conductance, V_{post} the membrane potential of the postsynaptic neuron, and E_s the reversal potential of the synapse.
 The evolution of the activation variable s is given by

$$\frac{ds}{dt} = \frac{\bar{s}(V_{pre}) - s}{\tau_s}$$

with

$$\bar{s}(V_{pre}) = \frac{1}{1 + \exp((V_{th} - V_{pre})/\delta)} \quad \text{and} \quad \tau_s = \frac{1 - \bar{s}(V_{pre})}{k_-}$$

519 Here, V_{pre} is the membrane potential of the presynaptic neuron, V_{th} is the half-activation voltage of the synapse, δ sets
 520 the slope of the activation curve, and k_- is the rate constant for transmitter-receptor dissociation rate.

521 As in [7], two types of synapses were modeled since AB, LP, and PY are glutamatergic neurons whereas PD is
 522 cholinergic. We set $E_s = -70$ mV and $k_- = 1/40$ ms for all glutamatergic synapses and $E_s = -80$ mV and $k_- = 1/100$
 523 ms for all cholinergic synapses. For both synapse types, we set $V_{th} = -35$ mV and $\delta = 5$ mV.

524 For each set of membrane and synaptic conductances, we numerically simulated the rhythm for 10 seconds with
 525 a step size of 0.025 ms. To make the model stochastic, at each time step, we added Gaussian noise with a standard
 526 deviation of 0.001 mV to the input of each neuron.

527 We applied SNPE to infer the posterior over 24 membrane parameters and 7 synaptic parameters, i.e. 31 pa-
 528 rameters in total. The 7 synaptic parameters were the maximal conductances g_s of all synapses in the circuit,
 529 each of which is varied uniformly in logarithmic domain from 0.01 nS to 1000 nS, with an exception of the synapse
 530 from AB to LP, which is varied uniformly in logarithmic domain from 0.01 nS to 10000 nS. The membrane param-
 531 eters were the maximal membrane conductances for each of the neurons. The membrane conductances were

532 varied over an extended range of previously reported values [7], which led us to the uniform prior bounds $p_{low} =$
 533 $[0, 0, 0, 0, 0, 25, 0, 0]$ mS cm⁻² and $p_{high} = [500, 7.5, 8, 60, 15, 150, 0.2, 0.01]$ mS cm⁻² for the maximal membrane conduc-
 534 tances of the AB neuron, $p_{low} = [0, 0, 2, 10, 0, 0, 0, 0.01]$ mS cm⁻² and $p_{high} = [200, 2.5, 12, 60, 10, 125, 0.06, 0.04]$ mS cm⁻²
 535 for the maximal membrane conductances of the LP neuron, and $p_{low} = [0, 0, 0, 30, 0, 50, 0, 0]$ mS cm⁻² and $p_{high} =$
 536 $[600, 12.5, 4, 60, 5, 150, 0.06, 0.04]$ mS cm⁻² for the maximal membrane conductances of the PY neuron. The order of
 537 the membrane currents was: [Na, CaT, CaS, A, KCa, Kd, H, leak].

538 We used the 15 summary statistics proposed by [7], and extended them by 3 additional features. The features
 539 proposed by [7] are 15 salient features of the pyloric rhythm, namely: cycle period T (s), AB/PD burst duration d_{AB}^b (s),
 540 LP burst duration d_{LP}^b (s), PY burst duration d_{PY}^b (s), gap AB/PD end to LP start Δt_{AB-LP}^{es} (s), gap LP end to PY start Δt_{LP-PY}^{es}
 541 (s), delay AB/PD start to LP start Δt_{AB-LP}^{ss} (s), delay LP start to PY start Δt_{LP-PY}^{ss} (s), AB/PD duty cycle d_{AB} , LP duty cycle d_{LP} ,
 542 PY duty cycle d_{PY} , phase gap AB/PD end to LP start $\Delta \phi_{AB-LP}$, phase gap LP end to PY start $\Delta \phi_{LP-PY}$, LP start phase ϕ_{LP} ,
 543 and PY start phase ϕ_{PY} . Note that several of these values are only defined if each neuron produces rhythmic bursting
 544 behavior. In addition, for each of the three neurons, we used one feature that describes the maximal duration of its
 545 voltage being above -30 mV. We did this as we observed plateaus at around -10 mV during the onset of bursts, and
 546 wanted to distinguish such activity traces from others. If the maximal duration was below 5 ms, we set this feature to 5
 547 ms. To extract the summary statistics from the observed experimental data, we first found spikes by searching for
 548 local maxima above a hand-picked voltage threshold, and then extracted the 15 above described features. We set the
 549 additional 3 features to 5 ms.

550 We used SNPE to infer the posterior distribution over the 18 summary statistics from experimental data. For
 551 inference, we used a single round with 18.5 million samples, out of which 174,000 samples contain bursts in all
 552 neurons. We therefore used these 174,000 samples with well defined summary statistics for training the inference
 553 network ($\theta \in \mathbb{R}^{31}$, $\mathbf{x} \in \mathbb{R}^{18}$). The density estimator was a masked autoregressive flow (MAF) [41] with five MADES with
 554 [200,400] hidden units each. The synaptic conductances were transformed into logarithmic space before training and
 555 for the entire analysis.

556 Finding paths in the posterior

557 In order to find directions of robust network output, we searched for a path of high posterior probability. First, as in
 558 [7], we aimed to find 2 similar model outputs with disparate parameters. To do so, we sampled from the posterior and
 559 searched for 2 parameter sets whose summary statistics were within 0.1 standard deviations of all 174,000 samples
 560 from the observed experimental data, but that had strongly disparate parameters from each other. In the following,
 561 we denote the obtained parameter sets by θ_s and θ_g .

562 Second, in order to identify whether network output can be maintained along a continuous path between these 2
 563 samples, we searched for a connection in parameter space lying in regions of high posterior probability. To do so, we
 564 considered the connection between the samples as a path and minimize the following path integral:

$$\mathcal{L}(\gamma) = \int_0^1 -\log(p_X(\gamma(s))) \|\dot{\gamma}(s)\| ds. \quad (1)$$

To minimize this term, we parameterized the path $\gamma(s)$ using sinusoidal basis-functions with coefficients $\alpha_{n,k}$:

$$\gamma(s) = \begin{bmatrix} \sum_{k=1}^K \alpha_{1,k} \cdot \sin(\pi ks) \\ \vdots \\ \sum_{k=1}^K \alpha_{N,k} \cdot \sin(\pi ks) \end{bmatrix} + \begin{bmatrix} \sum_{k=K+1}^{2K} \alpha_{N,k} \cdot \sin^2(\pi ks) \\ \vdots \\ \sum_{k=K+1}^{2K} \alpha_{N,k} \cdot \sin^2(\pi ks) \end{bmatrix} + (1-s) \cdot \theta_s + s\theta_g$$

These basis functions are defined such that, for any coefficients $\alpha_{n,k}$, the starting and end points of the path are exactly the two parameter sets defined above:

$$\gamma(0) = \theta_s \quad \gamma(1) = \theta_g$$

565 With this formulation, we have framed the problem of finding the path as an unconstrained optimization problem over
 566 the parameters $\alpha_{n,k}$. We can therefore minimize the path integral L using gradient descent over $\alpha_{n,k}$. For numerical
 567 simulations, we approximated the integral in equation 1 as a sum over 80 points along the path and use 2 basis
 568 functions for each of the 31 dimensions, i.e. $K = 2$.

In order to demonstrate the sensitivity of the pyloric network, we aimed to find a path along which the circuit output quickly breaks down. For this, we picked a starting point along the high-probability path and then minimize the posterior probability. In addition, we enforced that the orthogonal path lies within the orthogonal disk to the high-probability path, leading to the following constrained optimization problem:

$$\min_{\theta} \log(p(\theta|\mathbf{x})) \quad \text{s.t.} \quad n^T \Delta\theta = 0$$

where n is the tangent vector along the path of high probability. This optimization problem can be solved using the gradient projection method [129]:

$$\Delta\theta = -\frac{P(\nabla \log(p(\theta|\mathbf{x})))}{\sqrt{(\nabla \log(p(\theta|\mathbf{x})))^T P(\nabla \log(p(\theta|\mathbf{x})))}}$$

569 with projection matrix $P = \mathbb{1} - \frac{1}{n^T n} n n^T$ and $\mathbb{1}$ indicating the identity matrix. Each gradient update is a step along
570 the orthogonal path. We let the optimization run until the distance along the path is 1/27 of the distance along the
571 high-probability path.

572 Identifying conditional correlations

573 In order to investigate compensation mechanisms in the STG, we compared marginal and conditional correlations.
574 For the marginal correlation matrix in Fig. 6B, we calculated the Pearson correlation coefficient based on 1.26 million
575 samples from the posterior distribution $p(\theta|\mathbf{x})$. To find the 2-dimensional conditional distribution for any pair of
576 parameters, we fixed all other parameters to values taken from an arbitrary posterior sample, and varied the remaining
577 2 on an evenly spaced grid with 50 points along each dimension for figure 6C and with 20 points along each dimension
578 for figure 6D, covering the entire prior space. We evaluated the posterior distribution at every value on this grid.
579 We then calculated the conditional correlation as the Pearson correlation coefficient over this distribution. For the
580 1-dimensional conditional distribution, we varied only 1 parameter and kept all others fixed. Lastly, in Fig. 6D, we
581 sampled 500 parameter sets from the posterior, computed the respective conditional posteriors and conditional
582 correlation matrices, and took the average over the conditional correlation matrices.

583 References

- 584 [1] A. V. Herz, T. Gollisch, C. K. Machens, and D. Jaeger. Modeling single-neuron dynamics and computations: a balance of detail
585 and abstraction. *Science*, 314(5796):80–85, 2006.
- 586 [2] W. Gerstner, H. Sprekeler, and G. Deco. Theory and simulation in neuroscience. *Science*, 338(6103):60–65, 2012.
- 587 [3] T. O’Leary, A. C. Sutton, and E. Marder. Computational models in the age of large datasets. *Current Opinion in Neurobiology*, 32:
588 87–94, 2015.
- 589 [4] R. E. Baker, J.-M. Pena, J. Jayamohan, and A. Jérusalem. Mechanistic models versus machine learning, a fight worth fighting for
590 the biological community? *Biology Letters*, 14(5), 2018.
- 591 [5] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation
592 in nerve. *The Journal of Physiology*, 117(4):500–544, 1952.
- 593 [6] C. van Vreeswijk and H. Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274
594 (5293), 1996.
- 595 [7] A. A. Prinz, D. Bucher, and E. Marder. Similar network activity from disparate circuit parameters. *Nature Neuroscience*, 7(12):
596 1345, 2004.
- 597 [8] T. P. Vogels, K. Rajan, and L. F. Abbott. Neural network dynamics. *Annual Review of Neuroscience*, 28:357–376, 2005.
- 598 [9] T. C. Potjans and M. Diesmann. The cell-type specific cortical microcircuit: relating structure and activity in a full-scale spiking
599 network model. *Cerebral Cortex*, 24(3):785–806, 2012.
- 600 [10] A. Litwin-Kumar and B. Doiron. Slow dynamics and high variability in balanced cortical networks with clustered connections.
601 *Nature Neuroscience*, 15(11):1498, 2012.
- 602 [11] O. Sporns. Contributions and challenges for network models in cognitive neuroscience. *Nature Neuroscience*, 17(5):652, 2014.

- 603 [12] D. S. Bassett, P. Zurn, and J. I. Gold. On the nature and use of models in network neuroscience. *Nature Reviews Neuroscience*, 19
604 (9):566, 2018.
- 605 [13] J. I. Gold and M. N. Shadlen. The neural basis of decision making. *Annual Review of Neuroscience*, 30, 2007.
- 606 [14] X.-J. Wang. Decision making in recurrent neuronal circuits. *Neuron*, 60(2):215–234, 2008.
- 607 [15] J. Benda, T. Gollisch, C. K. Machens, and A. V. Herz. From response to stimulus: adaptive sampling in sensory physiology.
608 *Current Opinion in Neurobiology*, 17(4):430–436, 2007.
- 609 [16] R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, and J. P. Sethna. Universally sloppy parameter sensitivities
610 in systems biology models. *PLoS Computational Biology*, 3(10):e189, 2007.
- 611 [17] P. Achard and E. De Schutter. Complex parameter landscape for a complex neuron model. *PLoS Computational Biology*, 2(7):
612 e94, 2006.
- 613 [18] L. M. Alonso and E. Marder. Visualization of currents in neural models with similar behavior and different conductance
614 densities. *eLife*, 8:e42722, 2019.
- 615 [19] W. Truccolo, U. T. Eden, M. R. Fellows, J. P. Donoghue, and E. N. Brown. A point process framework for relating neural spiking
616 activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*, 93(2):1074–1089, 2005.
- 617 [20] E. Schneidman, M. J. Berry II, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a
618 neural population. *Nature*, 440(7087):1007, 2006.
- 619 [21] J. W. Pillow, J. Shlens, L. Paninski, A. Sher, A. M. Litke, E. J. Chichilnisky, and E. P. Simoncelli. Spatio-temporal correlations and
620 visual signalling in a complete neuronal population. *Nature*, 454(7207), 2008.
- 621 [22] B. M. Yu, J. P. Cunningham, G. Santhanam, S. I. Ryu, K. V. Shenoy, and M. Sahani. Gaussian-process factor analysis for
622 low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, 102(1):614–35, Jul 2009.
623 doi:10.1152/jn.90941.2008.
- 624 [23] J. H. Macke, L. Buesing, J. P. Cunningham, M. Y. Byron, K. V. Shenoy, and M. Sahani. Empirical models of spiking in neural
625 populations. In *Advances in Neural Information Processing Systems*, pages 1350–1358, 2011.
- 626 [24] J. P. Cunningham and M. Y. Byron. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500,
627 2014.
- 628 [25] C. Pandarinath, D. J. O’Shea, J. Collins, R. Jozefowicz, S. D. Stavisky, J. C. Kao, E. M. Trautmann, M. T. Kaufman, S. I. Ryu, L. R.
629 Hochberg, et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, page 1,
630 2018.
- 631 [26] A. A. Prinz, C. P. Billimoria, and E. Marder. Alternative to hand-tuning conductance-based models: construction and analysis of
632 databases of model neurons. *Journal of Neurophysiology*, 90(6):3998–4015, 2003.
- 633 [27] C. Tomm, M. Avermann, T. Vogels, W. Gerstner, and C. Petersen. The influence of structure on the response properties of
634 biologically plausible neural network models. *BMC neuroscience*, 12(1):P30, 2011.
- 635 [28] C. Stringer, M. Pachitariu, N. A. Steinmetz, M. Okun, P. Bartho, K. D. Harris, M. Sahani, and N. A. Lesica. Inhibitory control of
636 correlated intrinsic variability in cortical networks. *eLife*, 5, 2016.
- 637 [29] S. Druckmann, Y. Banitt, A. A. Gidon, F. Schürmann, H. Markram, and I. Segev. A novel multiple objective optimization framework
638 for constraining conductance-based neuron models by experimental data. *Frontiers in Neuroscience*, 1:1, 2007.
- 639 [30] E. Hay, S. Hill, F. Schürmann, H. Markram, and I. Segev. Models of neocortical layer 5b pyramidal cells capturing a wide range of
640 dendritic and perisomatic active properties. *PLoS Computational Biology*, 7(7), 2011.
- 641 [31] C. Rossant, D. F. M. Goodman, B. Fontaine, J. Platkiewicz, A. K. Magnusson, and R. Brette. Fitting neuron models to spike trains.
642 *Frontiers in Neuroscience*, 5:9, 2011.
- 643 [32] W. Van Geit, M. Gevaert, G. Chindemi, C. Rössert, J. Courcol, E. B. Muller, F. Schürmann, I. Segev, and H. Markram. Bluepyopt:
644 Leveraging open source software and cloud infrastructure to optimise model parameters in neuroscience. *Frontiers in
645 Neuroinformatics*, 10:17, 2016.
- 646 [33] M. Beaumont, W. Zhang, and D. J. Balding. Approximate bayesian computation in population genetics. *Genetics*, 162(4), 2002.
- 647 [34] P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov chain monte carlo without likelihoods. *Proceedings of the National
648 Academy of Sciences*, 100(26), 2003.

- 649 [35] S. A. Sisson, Y. Fan, and M. M. Tanaka. Sequential monte carlo without likelihoods. *Proceedings of the National Academy of*
650 *Sciences*, 104(6):1760–1765, 2007.
- 651 [36] G. Papamakarios and I. Murray. Fast ϵ -free inference of simulation models with bayesian conditional density estimation. In
652 *Advances in Neural Information Processing Systems*, pages 1028–1036, 2016.
- 653 [37] J.-M. Lueckmann, P. J. Goncalves, G. Bassetto, K. Öcal, M. Nonnenmacher, and J. H. Macke. Flexible statistical inference for
654 mechanistic models of neural dynamics. In *Advances in Neural Information Processing Systems*, pages 1289–1299, 2017.
- 655 [38] D. Greenberg, M. Nonnenmacher, and J. Macke. Automatic posterior transformation for likelihood-free inference. In *International*
656 *Conference on Machine Learning*, pages 2404–2414, 2019.
- 657 [39] K. Cranmer, J. Brehmer, and G. Louppe. The frontier of simulation-based inference. *arXiv preprint arXiv:1911.01429*, 2019.
- 658 [40] D. J. Rezende and S. Mohamed. Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference*
659 *on International Conference on Machine Learning-Volume 37*, pages 1530–1538. JMLR. org, 2015.
- 660 [41] G. Papamakarios, T. Pavlakou, and I. Murray. Masked autoregressive flow for density estimation. In *Advances in Neural*
661 *Information Processing Systems*, pages 2338–2347, 2017.
- 662 [42] S. R. Bittner, A. Palmigiano, K. D. Miller, and J. P. Cunningham. Degenerate solution networks for theoretical neuroscience. In
663 *Computational and Systems Neuroscience Meeting (COSYNE) 2019, Lisbon, Portugal*, 2019.
- 664 [43] G. Loaiza-Ganem, Y. Gao, and J. P. Cunningham. Maximum entropy flow networks. In *5th International Conference on Learning*
665 *Representations, ICLR*, 2017.
- 666 [44] E. N. Brown, L. M. Frank, D. Tang, M. C. Quirk, and M. A. Wilson. A statistical paradigm for neural spike train decoding applied to
667 position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience*, 18(18):7411–7425,
668 1998.
- 669 [45] L. Paninski. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural*
670 *Systems*, 15(4):243–262, 2004.
- 671 [46] J. Pillow. Likelihood-based approaches to modeling the neural code. *Bayesian Brain: Probabilistic Approaches to Neural Coding*,
672 pages 53–70, 2007.
- 673 [47] S. Gerwinn, J. H. Macke, and M. Bethge. Bayesian inference for generalized linear models for spiking neurons. *Frontiers in*
674 *Computational Neuroscience*, 4:12, 2010.
- 675 [48] N. G. Polson, J. G. Scott, and J. Windle. Bayesian inference for logistic models using pólya-gamma latent variables. *Journal of*
676 *the American statistical Association*, 108(504):1339–1349, 2013.
- 677 [49] J. W. Pillow and J. Scott. Fully bayesian inference for neural models with negative-binomial spiking. In *Advances in Neural*
678 *Information Processing Systems*, pages 1898–1906, 2012.
- 679 [50] J. W. Pillow, L. Paninski, V. J. Uzzell, E. P. Simoncelli, and E. Chichilnisky. Prediction and decoding of retinal ganglion cell responses
680 with a probabilistic spiking model. *Journal of Neuroscience*, 25(47):11003–11013, 2005.
- 681 [51] E. Chichilnisky. A simple white noise analysis of neuronal light responses. *Network: Computation in Neural Systems*, 12(2):
682 199–213, 2001.
- 683 [52] C. M. Niell and M. P. Stryker. Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience*, 28(30):7520–7536,
684 2008. ISSN 0270-6474. doi:10.1523/JNEUROSCI.0623-08.2008.
- 685 [53] L. Dyballa, M. S. Hoseini, M. C. Dadarlat, S. W. Zucker, and M. P. Stryker. Flow stimuli reveal ecologically appropriate responses
686 in mouse visual cortex. *Proceedings of the National Academy of Sciences*, 115(44):11304–11309, 2018. ISSN 0027-8424.
687 doi:10.1073/pnas.1811265115.
- 688 [54] J. P. Jones and L. A. Palmer. An evaluation of the two-dimensional gabor filter model of simple receptive fields in cat striate
689 cortex. *Journal of Neurophysiology*, 58(6):1233–1258, 1987.
- 690 [55] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in*
691 *Neural Information Processing Systems*, pages 1097–1105, 2012.
- 692 [56] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference*
693 *on Learning Representations*, 2015.
- 694 [57] C. M. Bishop. Mixture density networks. *Technical Report. Aston University, Birmingham*, 1994.

- 695 [58] W. F. Podlaski, A. Seeholzer, L. N. Groschner, G. Miesenböck, R. Ranjan, and T. P. Vogels. Mapping the function of neuronal ion
696 channels in model and experiment. *eLife*, 6:e22152, March 2017. ISSN 2050-084X. doi:10.7554/eLife.22152.
- 697 [59] A. Destexhe and J. R. Huguenard. Nonlinear Thermodynamic Models of Voltage-Dependent Currents. *Journal of Computational*
698 *Neuroscience*, 9(3):259–270, November 2000. ISSN 1573-6873. doi:10.1023/A:1026535704537.
- 699 [60] A. I. for Brain Science. Allen institute for brain science. allen cell types database. <http://celltypes.brain-map.org/>, 2016.
- 700 [61] J. Dunlop, M. Bowlby, R. Peri, D. Vasilyev, and R. Arias. High-throughput electrophysiology: an emerging paradigm for
701 ion-channel screening and physiology. *Nature Reviews Drug Discovery*, 7(4):358, 2008.
- 702 [62] H.-J. Suk, E. S. Boyden, and I. van Welie. Advances in the automation of whole-cell patch clamp technology. *Journal of*
703 *Neuroscience Methods*, 326:108357, 2019. ISSN 0165-0270. doi:<https://doi.org/10.1016/j.jneumeth.2019.108357>.
- 704 [63] R. Ranjan, E. Logette, M. Marani, M. Herzog, V. Tache, and H. Markram. A kinetic map of the homomeric voltage-gated
705 potassium channel (kv) family. *Frontiers in Cellular Neuroscience*, 13:358, 2019.
- 706 [64] A. Speiser, J. Yan, E. W. Archer, L. Buesing, S. C. Turaga, and J. H. Macke. Fast amortized inference of neural activity from calcium
707 imaging data with variational autoencoders. In *Advances in Neural Information Processing Systems*, pages 4024–4034, 2017.
- 708 [65] S. Webb, A. Golinski, R. Zinkov, S. Narayanaswamy, T. Rainforth, Y. W. Teh, and F. Wood. Faithful inversion of generative models
709 for effective amortized inference. In *Advances in Neural Information Processing Systems*, pages 3070–3080, 2018.
- 710 [66] M. Pospischil, M. Toledo-Rodriguez, C. Monier, Z. Piwkowska, T. Bal, Y. Frégnac, H. Markram, and A. Destexhe. Minimal
711 hodgkin-huxley type models for different classes of cortical and thalamic neurons. *Biological Cybernetics*, 99(4-5), 2008.
- 712 [67] C. D. Meliza, M. Kostuk, H. Huang, A. Nogaret, D. Margoliash, and H. D. Abarbanel. Estimating parameters and predicting
713 membrane voltages with conductance-based neuron models. *Biological Cybernetics*, 108(4):495–516, 2014.
- 714 [68] Q. J. M. Huys, M. B. Ahrens, and L. Paninski. Efficient estimation of detailed single-neuron models. *Journal of Neurophysiology*,
715 96(2), 2006.
- 716 [69] A. C. Daly, D. J. Gavaghan, C. Holmes, and J. Cooper. Hodgkin-huxley revisited: reparametrization and identifiability analysis of
717 the classic action potential model with approximate bayesian methods. *Royal Society Open Science*, 2(12):150499, 2015.
- 718 [70] C. Teeter, R. Iyer, V. Menon, N. Gouwens, D. Feng, J. Berg, A. Szafer, N. Cain, H. Zeng, M. Hawrylycz, et al. Generalized leaky
719 integrate-and-fire models classify multiple neuron types. *Nature Communications*, 9(1):709, 2018.
- 720 [71] S. A. Haddad and E. Marder. Circuit robustness to temperature perturbation is altered by neuromodulators. *Neuron*, 100(3):
721 609–623, 2018.
- 722 [72] E. Marder and J.-M. Goaillard. Variability, compensation and homeostasis in neuron and network function. *Nature Reviews*
723 *Neuroscience*, 7(7):563, 2006.
- 724 [73] G. J. Gutierrez, T. O’Leary, and E. Marder. Multiple mechanisms switch an electrically coupled, synaptically inhibited neuron
725 between competing rhythmic oscillators. *Neuron*, 77(5):845–858, 2013.
- 726 [74] D. Fisher, I. Olasagasti, D. W. Tank, E. R. Aksay, and M. S. Goldman. A modeling framework for deriving the structural and
727 functional architecture of a short-term memory microcircuit. *Neuron*, 79(5):987–1000, 2013.
- 728 [75] E. Marder, M. L. Goeritz, and A. G. Otopalik. Robust circuit rhythms in small circuits arise from variable circuit components and
729 mechanisms. *Current Opinion in Neurobiology*, 31:156–163, 2015.
- 730 [76] T. O’Leary, A. H. Williams, A. Franci, and E. Marder. Cell types, network homeostasis, and pathological compensation from a
731 biologically plausible ion channel expression model. *Neuron*, 82(4):809–821, 2014.
- 732 [77] T. O’Leary and E. Marder. Temperature-robust neural function from activity-dependent ion channel regulation. *Current Biology*,
733 26(21):2935–2941, 2016.
- 734 [78] B. B. Machta, R. Chachra, M. K. Transtrum, and J. P. Sethna. Parameter space compression underlies emergent theories and
735 predictive models. *Science*, 342(6158):604–607, 2013.
- 736 [79] A. L. Taylor, T. J. Hickey, A. A. Prinz, and E. Marder. Structure and visualization of high-dimensional conductance spaces. *Journal*
737 *of Neurophysiology*, 96(2):891–905, 2006.
- 738 [80] J. N. MacLean, Y. Zhang, M. L. Goeritz, R. Casey, R. Oliva, J. Guckenheimer, and R. M. Harris-Warrick. Activity-independent
739 coregulation of ia and ih in rhythmically active neurons. *Journal of Neurophysiology*, 94(5):3601–3617, 2005.

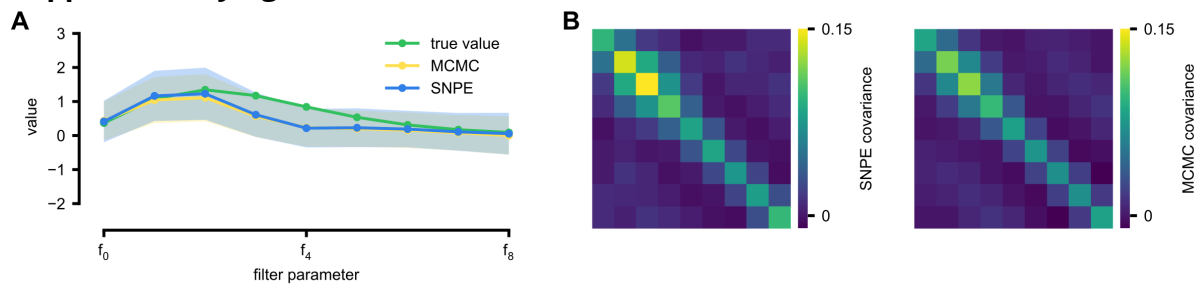
- 740 [81] R. Grashow, T. Brookings, and E. Marder. Compensation for variable intrinsic neuronal excitability by circuit-synaptic interactions.
741 *Journal of Neuroscience*, 30(27):9145–9156, 2010.
- 742 [82] E. Marder and A. L. Taylor. Multiple models to capture the variability in biological neurons and networks. *Nature Neuroscience*,
743 14(2):133, 2011.
- 744 [83] T. O’Leary. Homeostasis, failure of homeostasis and degenerate ion channel regulation. *Current Opinion in Physiology*, 2:
745 129–138, 2018.
- 746 [84] A. L. Taylor, J.-M. Goaillard, and E. Marder. How multiple conductances determine electrophysiological properties in a
747 multicompartiment model. *Journal of Neuroscience*, 29(17):5573–5586, 2009.
- 748 [85] J. N. MacLean, Y. Zhang, B. R. Johnson, and R. M. Harris-Warrick. Activity-independent homeostasis in rhythmically active
749 neurons. *Neuron*, 37(1):109–120, 2003.
- 750 [86] A. V. M. Herz, T. Gollisch, C. K. Machens, and D. Jaeger. Modeling single-neuron dynamics and computations: a balance of
751 detail and abstraction. *Science*, 314(5796), 2006.
- 752 [87] R. Brette. What is the most realistic single-compartment model of spike initiation? *PLoS Computational Biology*, 11(4):e1004114,
753 2015.
- 754 [88] T. A. Le, A. G. Baydin, R. Zinkov, and F. Wood. Using synthetic data to train neural networks is model-based reasoning. In *2017*
755 *International Joint Conference on Neural Networks (IJCNN)*, pages 3514–3521. IEEE, 2017.
- 756 [89] J. Chan, V. Perrone, J. Spence, P. Jenkins, S. Mathieson, and Y. Song. A likelihood-free inference framework for population
757 genetic data using exchangeable neural networks. In *Advances in Neural Information Processing Systems*, pages 8594–8605,
758 2018.
- 759 [90] M. G. Blum and O. François. Non-linear regression models for approximate bayesian computation. *Statistics and Computing*, 20
760 (1):63–73, 2010.
- 761 [91] D. B. Rubin et al. Bayesianly justifiable and relevant frequency calculations for the applied statistician. *The Annals of Statistics*,
762 12(4):1151–1172, 1984.
- 763 [92] J. K. Pritchard, M. T. Seielstad, A. Perez-Lezaun, and M. W. Feldman. Population growth of human y chromosomes: a study of y
764 chromosome microsatellites. *Molecular Biology and Evolution*, 16(12):1791–1798, 1999.
- 765 [93] J. Liepe, P. Kirk, S. Filippi, T. Toni, C. P. Barnes, and M. P. Stumpf. A framework for parameter estimation and model selection
766 from experimental data in systems biology using approximate bayesian computation. *Nature Protocols*, 9(2):439, 2014.
- 767 [94] M. U. Gutmann and J. Corander. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *The*
768 *Journal of Machine Learning Research*, 17(1):4256–4302, 2016.
- 769 [95] O. J. Britton, A. Bueno-Orovio, K. Van Ammel, H. R. Lu, R. Towart, D. J. Gallacher, and B. Rodriguez. Experimentally calibrated
770 population of models predicts and explains intersubject variability in cardiac cellular electrophysiology. *Proceedings of the*
771 *National Academy of Sciences*, 110(23):E2098–E2105, 2013.
- 772 [96] B. A. Lawson, C. C. Drovandi, N. Cusimano, P. Burrage, B. Rodriguez, and K. Burrage. Unlocking data sets by calibrating
773 populations of models to data density: A study in atrial electrophysiology. *Science Advances*, 4(1):e1701676, 2018.
- 774 [97] S. N. Wood. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310), 2010.
- 775 [98] R. Wilkinson. Accelerating abc methods using gaussian processes. In *AISTATS*, 2014.
- 776 [99] E. Meeds and M. Welling. Gps-abc: Gaussian process surrogate approximate bayesian computation. In *Conference on*
777 *Uncertainty in Artificial Intelligence*, 2014.
- 778 [100] G. Papamakarios, D. Sterratt, and I. Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive
779 flows. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 837–848, 2019.
- 780 [101] J.-M. Lueckmann, G. Bassetto, T. Karaletsos, and J. H. Macke. Likelihood-free inference with emulator networks. In F. Ruiz,
781 C. Zhang, D. Liang, and T. Bui, editors, *Proceedings of The 1st Symposium on Advances in Approximate Bayesian Inference*, volume 96
782 of *Proceedings of Machine Learning Research*, pages 32–53, 2019.
- 783 [102] C. Durkan, G. Papamakarios, and I. Murray. Sequential neural methods for likelihood-free inference. *NeurIPS Bayesian Deep*
784 *Learning Workshop*, 2018.

- 785 [103] S. Barthelmé and N. Chopin. Expectation propagation for likelihood-free inference. *Journal of the American Statistical Association*,
786 109(505):315–333, 2014.
- 787 [104] C. Schröder, L. Lagnado, B. James, and P. Berens. Approximate bayesian inference for a mechanistic model of vesicle release at
788 a ribbon synapse. *BioRxiv*, page 669218, 2019.
- 789 [105] T. A. Le, A. G. Baydin, and F. Wood. Inference compilation and universal probabilistic programming. In *Artificial Intelligence and*
790 *Statistics*, pages 1338–1348, 2017.
- 791 [106] M. L. Casado, A. G. Baydin, D. M. Rubio, T. A. Le, F. Wood, L. Heinrich, G. Louppe, K. Cranmer, K. Ng, W. Bhimji, et al.
792 Improvements to inference compilation for probabilistic programming in large-scale scientific simulators. *NeurIPS Workshop on*
793 *Deep Learning for Physical Sciences*, 2017.
- 794 [107] J. Hermans, V. Begy, and G. Louppe. Likelihood-free mcmc with approximate likelihood ratios. *arXiv preprint arXiv:1903.04057*,
795 2019.
- 796 [108] Z. Chen et al. Bayesian filtering: From kalman filters to particle filters, and beyond. *Statistics*, 182(1):1–69, 2003.
- 797 [109] Q. J. M. Huys and L. Paninski. Smoothing of, and parameter estimation from, noisy biophysical recordings. *PLoS Computational*
798 *Biology*, 5(5), 2009.
- 799 [110] L. Hertäg, J. Hass, T. Golovko, and D. Durstewitz. An approximation to the adaptive exponential integrate-and-fire neuron
800 model allows fast and predictive fitting to physiological data. *Frontiers in Computational Neuroscience*, 6:62, 2012.
- 801 [111] C. Pozzorini, S. Mensi, O. Hagens, R. Naud, C. Koch, and W. Gerstner. Automated high-throughput characterization of single
802 neurons by means of simplified spiking models. *PLoS Computational Biology*, 11(6):e1004275, 2015.
- 803 [112] J. Ladenbauer, S. McKenzie, D. F. English, O. Hagens, and S. Ostojic. Inferring and validating mechanistic models of neural
804 microcircuits based on spike-train data. *bioRxiv*, page 261016, 2018.
- 805 [113] A. René, A. Longtin, and J. H. Macke. Inference of a mesoscopic population model from population spike trains. *arXiv preprint*
806 *arXiv:1910.01618*, 2019.
- 807 [114] D. Sussillo and L. F. Abbott. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–557,
808 2009.
- 809 [115] V. Mante, D. Sussillo, K. V. Shenoy, and W. T. Newsome. Context-dependent computation by recurrent dynamics in prefrontal
810 cortex. *Nature*, 503(7474):78, 2013.
- 811 [116] D. Sussillo and O. Barak. Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks.
812 *Neural Computation*, 25(3):626–649, 2013.
- 813 [117] N. Maheswaranathan, A. Williams, M. D. Golub, S. Ganguli, and D. Sussillo. Reverse engineering recurrent networks for
814 sentiment classification reveals line attractor dynamics. *arXiv preprint arXiv:1906.10720*, 2019.
- 815 [118] D. B. Rubin, S. D. Van Hooser, and K. D. Miller. The stabilized supralinear network: a unifying circuit motif underlying multi-input
816 integration in sensory cortex. *Neuron*, 85(2):402–417, 2015.
- 817 [119] R. Ratcliff and G. McKoon. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20
818 (4):873–922, 2008.
- 819 [120] M. G. B. Blum, M. A. Nunes, D. Prangle, S. A. Sisson, et al. A comparative review of dimension reduction methods in approximate
820 bayesian computation. *Statistical Science*, 28(2), 2013.
- 821 [121] B. Jiang, T.-y. Wu, C. Zheng, and W. H. Wong. Learning summary statistic for approximate bayesian computation via deep
822 neural network. *Statistica Sinica*, pages 1595–1618, 2017.
- 823 [122] R. Izbicki, A. B. Lee, and T. Pospisil. Abc-cde: Toward approximate bayesian computation with complex high-dimensional data
824 and limited simulations. *Journal of Computational and Graphical Statistics*, pages 1–20, 2019.
- 825 [123] S. R. Cook, A. Gelman, and D. B. Rubin. Validation of software for bayesian models using posterior quantiles. *Journal of*
826 *Computational and Graphical Statistics*, 15(3):675–692, 2006.
- 827 [124] S. Talts, M. Betancourt, D. Simpson, A. Vehtari, and A. Gelman. Validating bayesian inference algorithms with simulation-based
828 calibration. *arXiv preprint arXiv:1804.06788*, 2018.
- 829 [125] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*,
830 2014.

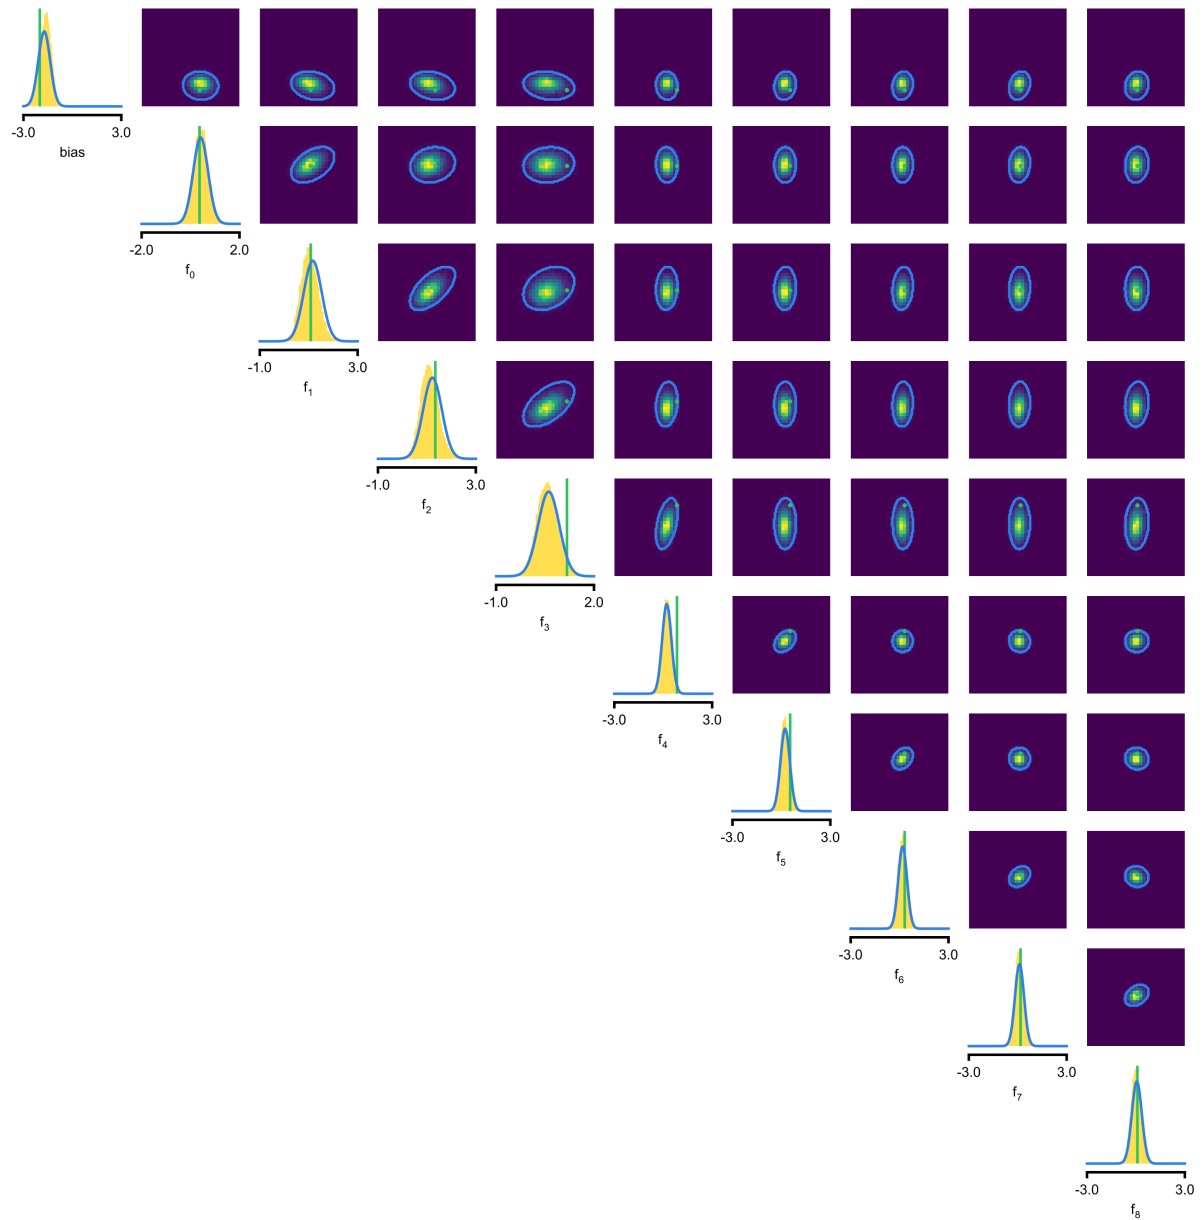
- 831 [126] G. De Nicolao, G. Sparacino, and C. Cobelli. Nonparametric input estimation in physiological systems: problems, methods, and
832 case studies. *Automatica*, 33(5), 1997.
- 833 [127] M. A. Beaumont, J. Cornuet, J. Marin, and C. P. Robert. Adaptive approximate bayesian computation. *Biometrika*, 2009.
- 834 [128] L. Abbott and E. Marder. Modeling small networks, 1998.
- 835 [129] J. B. Rosen. The gradient projection method for nonlinear programming. part i. linear constraints. *Journal of the Society for*
836 *Industrial and Applied Mathematics*, 8(1):181–217, 1960.

837 **Supplementary material**

838 **Supplementary figures**

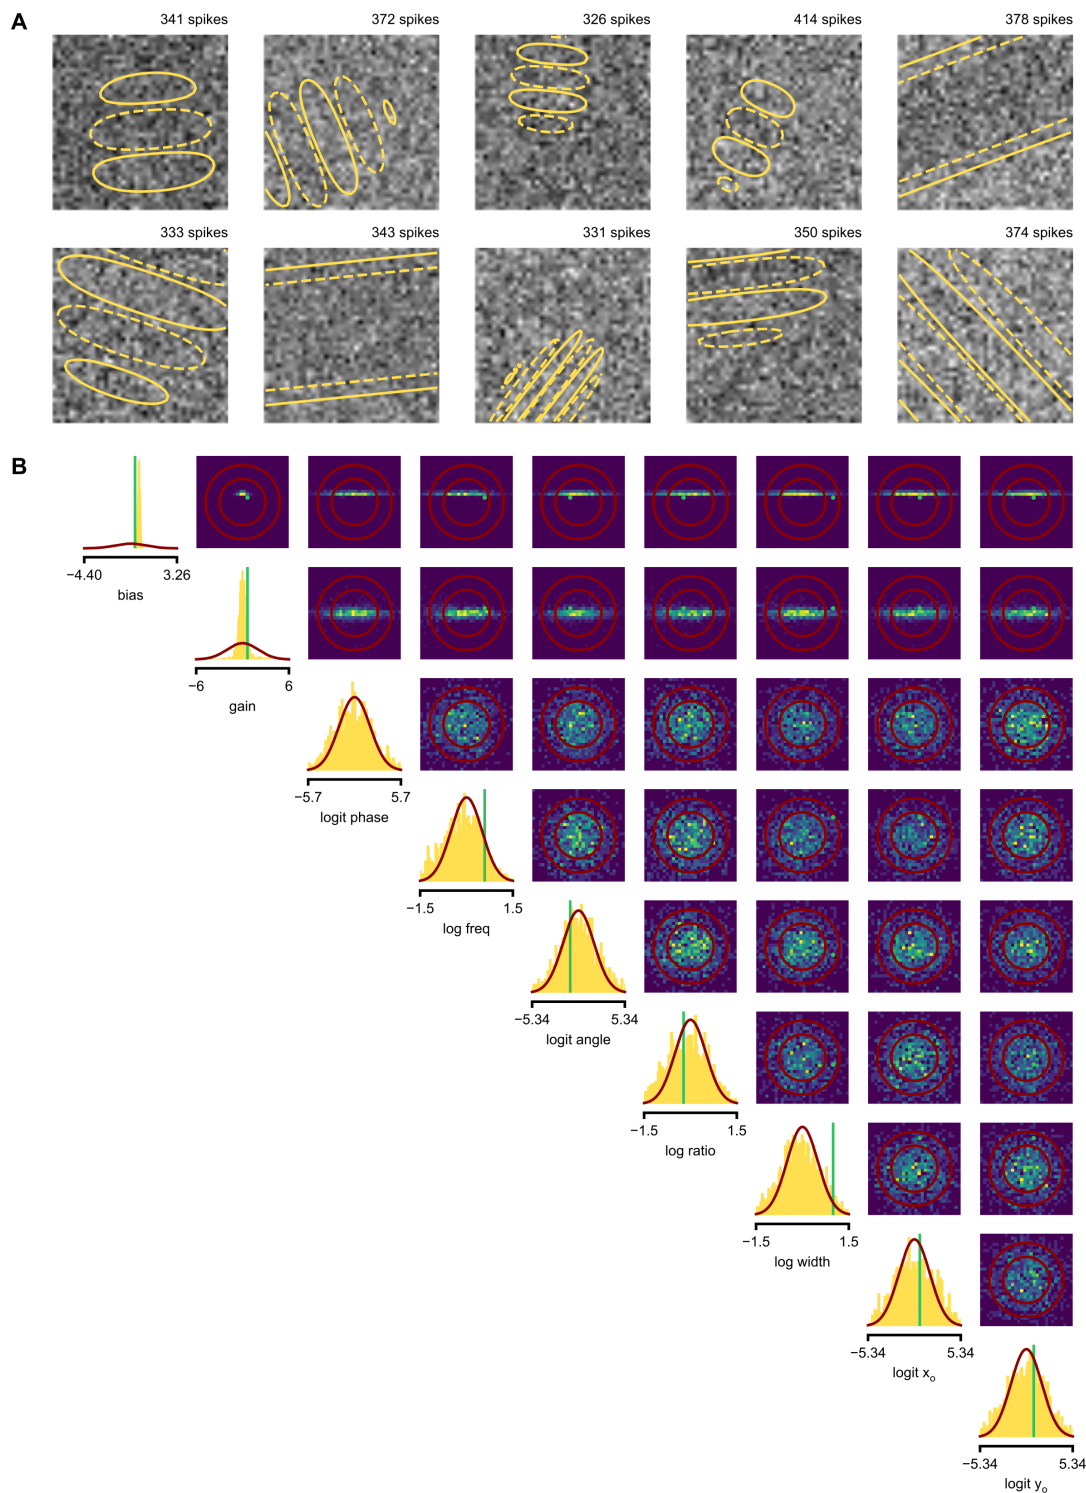


840 **Supplementary Figure 1. Comparison between SNPE and MCMC results on LN model.** (A) Posterior mean \pm standard
841 deviation of temporal filter (receptive field) from SNPE posterior (SNPE, blue) and reference posterior (MCMC, yellow). (B) Full
842 covariance matrices from SNPE posterior (left) and reference (MCMC, right).




845
846


Supplementary Figure 2. Full posterior for LN model. In green, ground-truth parameters. Marginals (blue lines) and 2D marginals for SNPE (contour lines correspond to 95% of the mass) and MCMC (yellow histograms).



849 **Supplementary Figure 3. SMC-ABC posterior estimate for Gabor GLM receptive field model.** (A) Spike-triggered
 850 averages (STAs) and spike counts with closest distance $d(x_o, x_i)$ to the observed data x_o out of 10000 simulations with θ_i
 851 sampled from the prior. Spike counts are comparable to the observed data (x_o : 307 spikes), but receptive fields (contours)
 852 are not well constrained. (B) Results for SMC-ABC with 10^6 simulations total. Histograms of 1000 particles (yellow) returned in
 853 the final iteration of SMC-ABC, compared to prior (red contour lines) and ground-truth parameters (green). Distributions over
 854 (log-/logit-)transformed parameters, axis limits scaled to mean ± 3 standard deviations of the prior.




857 **Supplementary Figure 4. Full posterior for Gabor GLM receptive field model.** SNPE posterior estimate (blue lines)
858 compared to MCMC samples from ground-truth posterior (histograms). Ground-truth parameters used to simulate the data
859 in green. We depict the distributions over the original receptive field parameters, whereas we estimate the posterior as a
860 Gaussian mixture over transformed parameters, see Methods for details. We find that a (back-transformed) Gaussian
861 mixture with four components approximates the posterior well in this case.




864
865
866

Supplementary Figure 5. Full posterior for Gabor LN receptive field model on V1 recordings. We depict the distributions over the receptive field parameters, derived from the Gaussian mixture over transformed-parameters (see Methods for details).

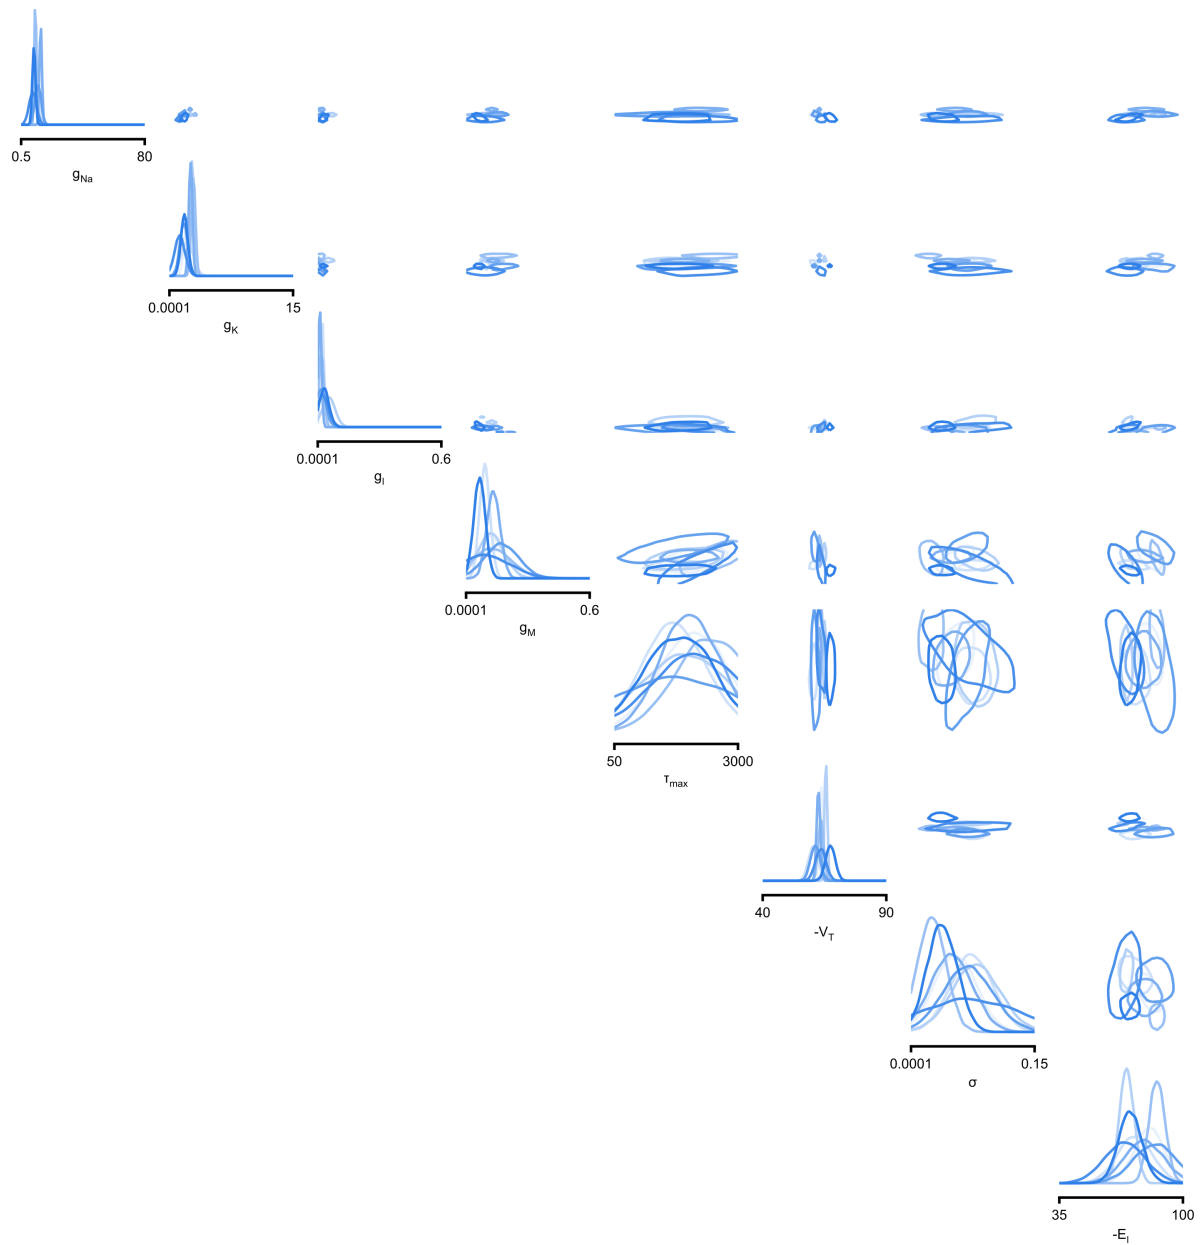


869 **Supplementary Figure 6. Summary results on ICG channel models, and comparison with direct fits.** We generate
870 predictions either with the posterior mode (blue) or with parameters obtained by directly fitting steady-state activation and
871 time-constant curves (yellow). We calculate the correlation coefficient (CC) between observation and prediction. The
872 distribution of CCs is similar for both approaches.



875
876
878

Supplementary Figure 7. Full posteriors for Hodgkin-Huxley model for 1, 4 and 7 features. Images show the pairwise marginals for 7 features. Each contour line corresponds to 68% density mass for a different inferred posterior. Light blue corresponds to 1 feature and dark blue to 7 features.




880

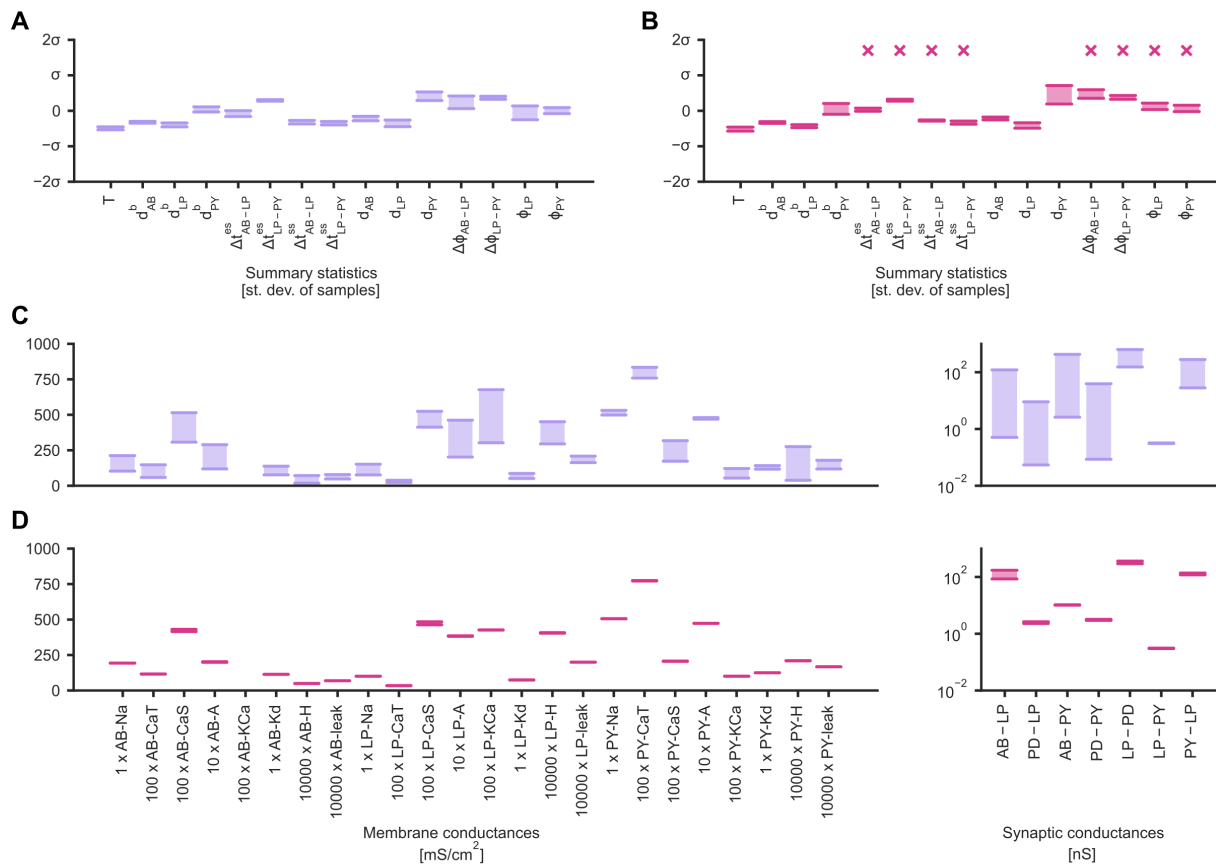
881

882

Supplementary Figure 8. Full posteriors for Hodgkin-Huxley model on 8 different recordings from Allen Cell Type Database. Images show the pairwise marginals for 7 features. Each contour line corresponds to 68% density mass for a different inferred posterior.



885 **Supplementary Figure 9. Full posterior for the stomatogastric ganglion over 24 membrane and 7 synaptic**
 886 **conductances.** The first 24 dimensions depict membrane conductances (top left), the last 7 depict synaptic conductances
 887 (bottom right). All synaptic conductances are logarithmically spaced. Between two samples from the posterior with high
 888 posterior probability (purple dots), there is a path of high posterior probability (white).



891 **Supplementary Figure 10. Identifying directions of sloppiness and stiffness in the pyloric network of the**
 892 **crustacean stomatogastric ganglion.** (A) Minimal and maximal values of all summary statistics along the path lying in
 893 regions of high posterior probability, sampled at 20 evenly spaced points. Summary statistics change only little. The
 894 summary statistics are scaled with the standard deviation of the 170,000 bursting samples in the created dataset. (B)
 895 Summary statistics sampled at 20 evenly spaced points along the orthogonal path. The summary statistics show stronger
 896 changes than in panel A and, in particular, often could not be defined because neurons bursted irregularly, as indicated by
 897 an 'x' above barplots. (C) Minimal and maximal values of the circuit parameters along the path lying in regions of high
 898 posterior probability. Both membrane conductances (left) and synaptic conductances (right) vary over large ranges. Axes as
 899 in panel D. (D) Circuit parameters along the orthogonal path. The difference between the minimal and maximal value is
 900 much smaller than in panel C.