Simulation-based inference and the places it takes us

Jakob Macke, www.mackelab.org @mackelab
Machine Learning in Science, Tübingen University
Excellence Cluster Machine Learning & Tübingen AI Center
Bernstein Center for Computational Neuroscience Tübingen
Empirical Inference, Max Planck Institute for Intelligent Systems

Neurophysiology

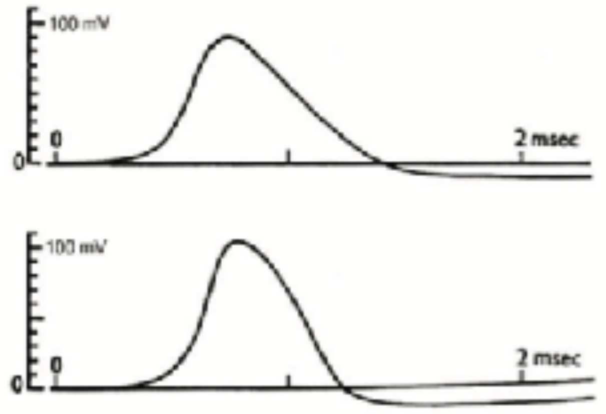

fig; Schwiening 2003

Population Dynamics

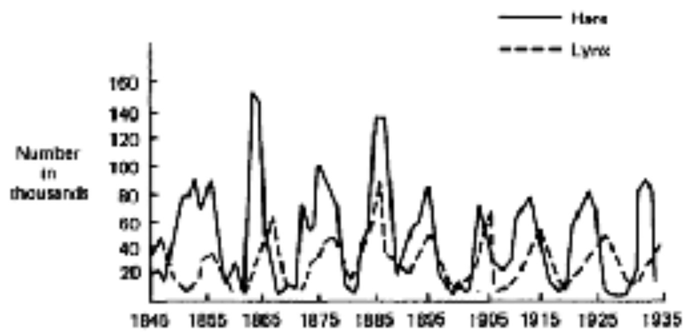

Figure 48-1 Oscillation observed in Canada of populations of lynx and hare (*data from* E. P. Odum, *Fundamentals of Ecology*, Philadelphia: W. B. Saunders, 1953).

fig; Saunders 1953

Psychophysics of Decision Making

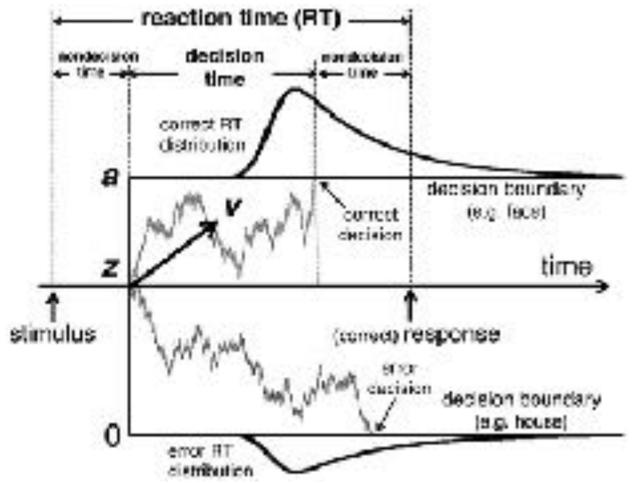

fig: Murata 2014

Particle Physics

$$\begin{aligned}\mathcal{L}_{SM} = \; & \underbrace{\frac{1}{4}\mathbf{W}_{\mu\nu}\cdot\mathbf{W}^{\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}G^{a}_{\mu\nu}G^{\mu\nu}_{a}}_{\text{kinetic energies and self-interactions of the gauge bosons}} \\ & + \underbrace{\bar{L}\gamma^{\mu}(i\partial_{\mu} - \frac{1}{2}g\tau\cdot\mathbf{W}_{\mu} - \frac{1}{2}g'YB_{\mu})L + \bar{R}\gamma^{\mu}(i\partial_{\mu} - \frac{1}{2}g'YB_{\mu})R}_{\text{kinetic energies and electroweak interactions of fermions}} \\ & + \underbrace{\frac{1}{2}\left|(i\partial_{\mu} - \frac{1}{2}g\tau\cdot\mathbf{W}_{\mu} - \frac{1}{2}g'YB_{\mu})\phi\right|^{2} - V(\phi)}_{W^{\pm}, Z, \gamma \text{ and Higgs masses and couplings}} \\ & + \underbrace{g''(\bar{q}\gamma^{\mu}T_{a}q)G^{a}_{\mu}}_{\text{interactions between quarks and gluons}} + \underbrace{(G_{1}\bar{L}\phi R + G_{2}\bar{L}\phi_{c}R + h.c.)}_{\text{fermion masses and couplings to Higgs}}\end{aligned}$$

fig: Cranmer 2017

Systems Biology

$$\dot{S} = \alpha - \gamma SI - dS, \qquad (3.10a)$$

$$\dot{I} = \gamma SI - vI - dI, \qquad (3.10b)$$

$$\dot{R} = vI - dR, \qquad (3.10c)$$

fig; Toni 2018

Cosmology

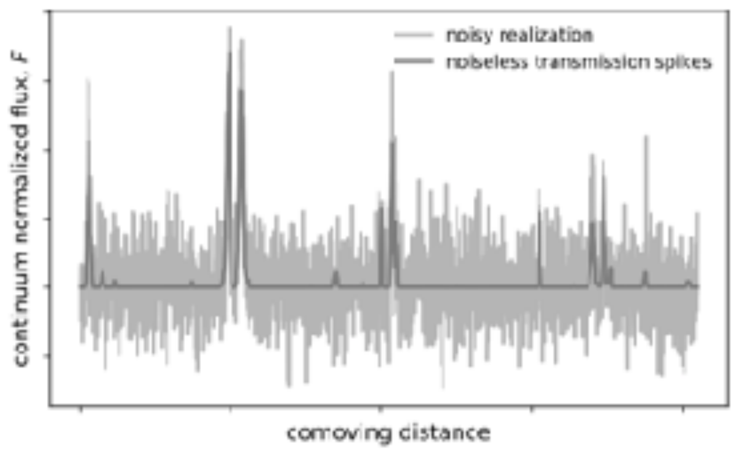

fig: Walting 2018

Molecular Dynamics

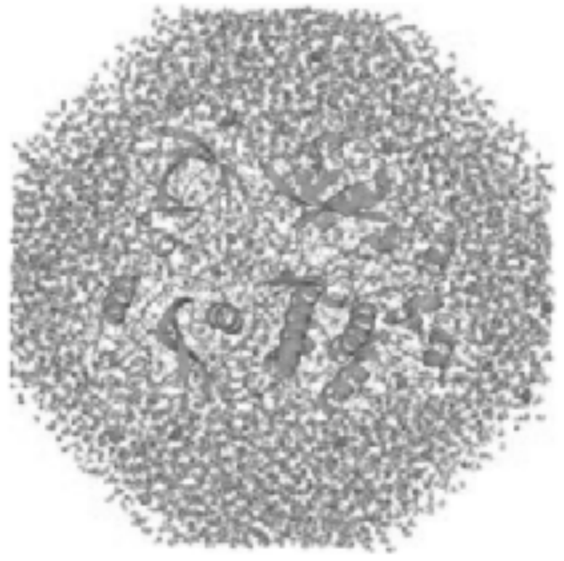

fig: Karplus 2005

...

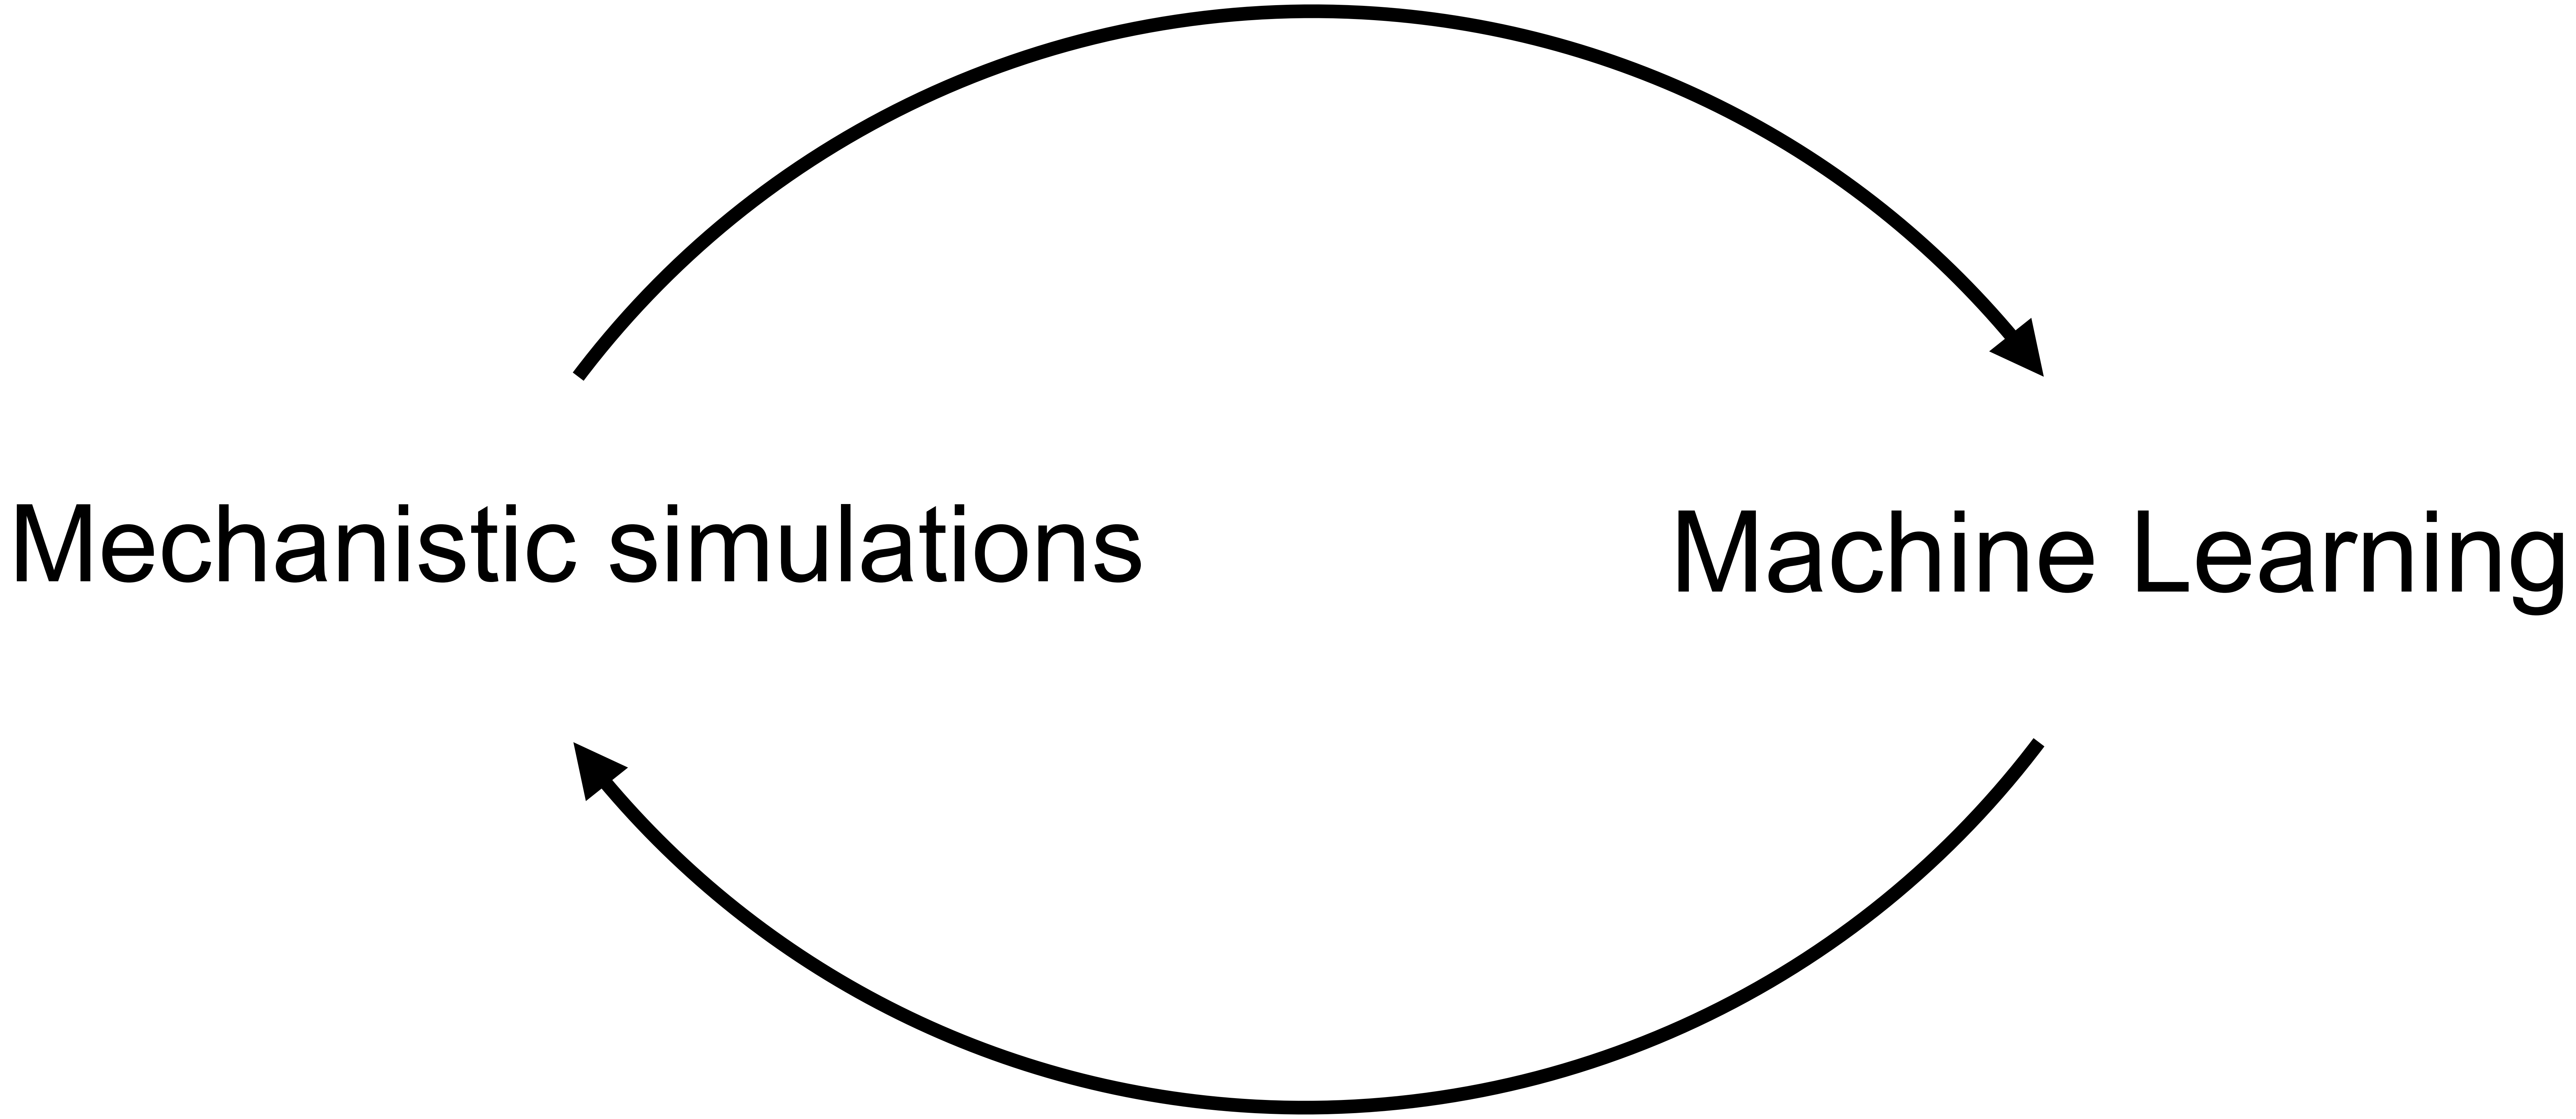

Mechanistic simulations

Machine Learning

How can we combine these two approaches
to build tools for data-driven scientific discovery?

Focus: Neuroscience

How can we find mechanistic/physical models which are compatible with observed data (and prior knowledge)?

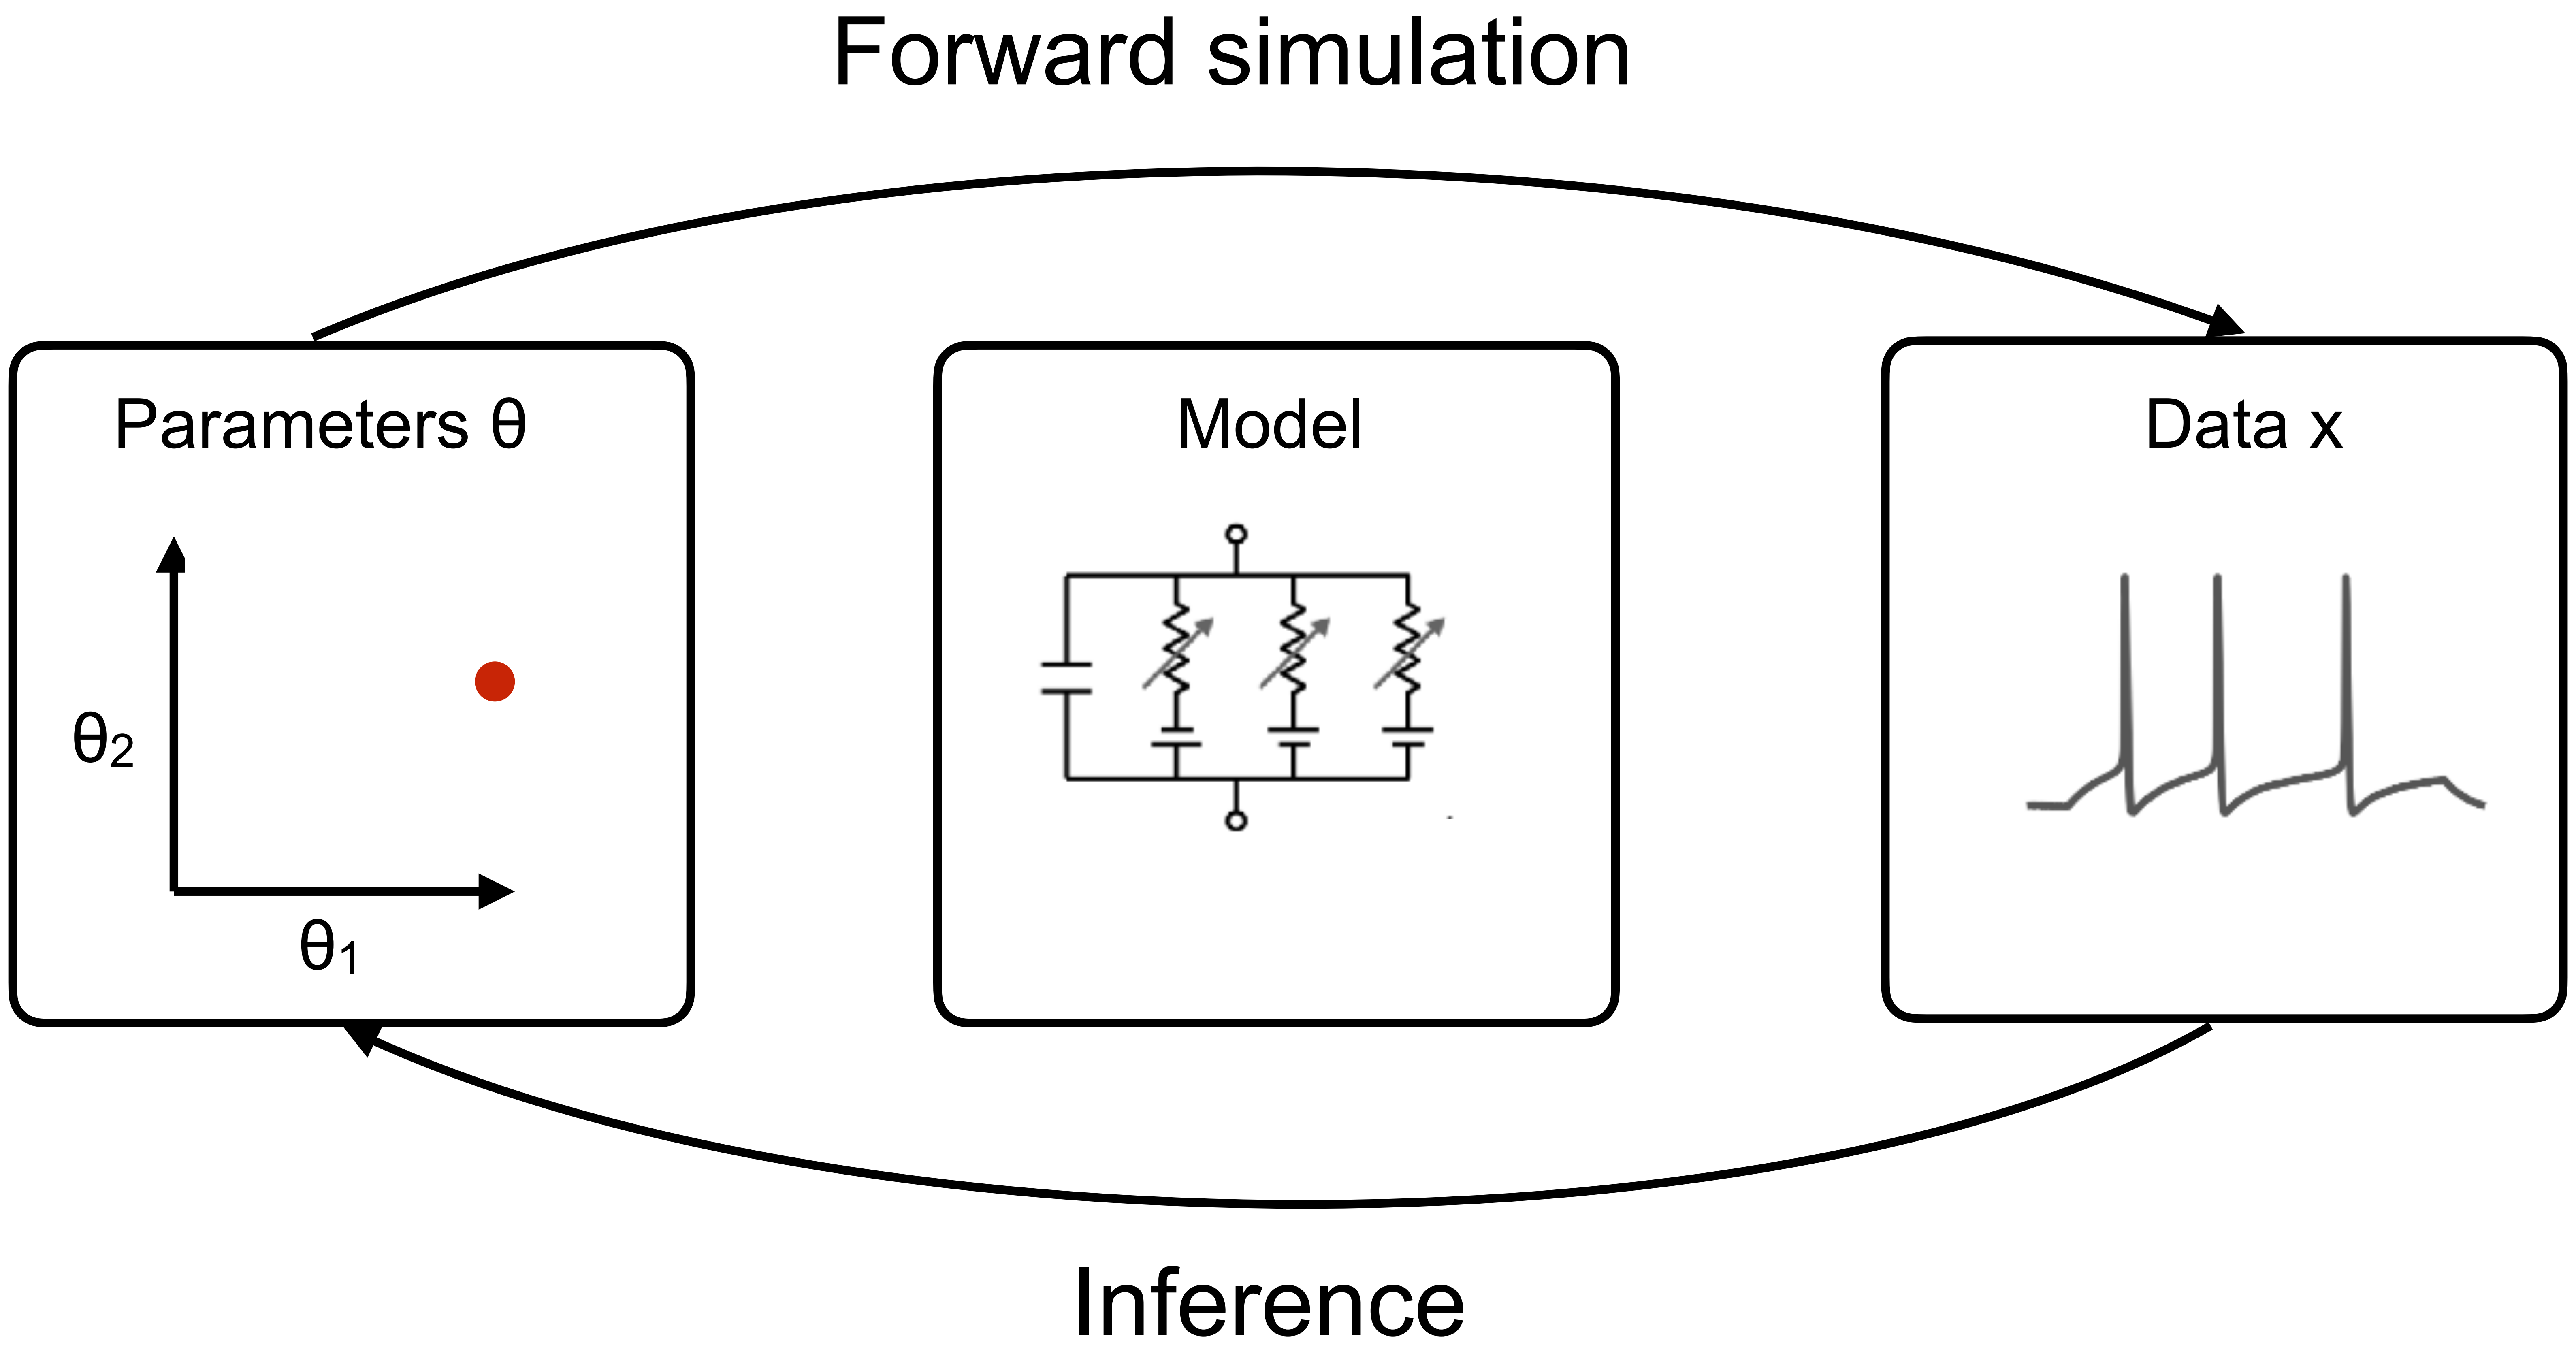

Forward simulation

Parameters θ

Model

Data x

Inference

Challenges:

1) **Forward model often given by simulator (might be complicated, black-box, non-diffable)**
2) **Inference problems vastly under-constrained**

How can we do Bayesian inference for *any* mechanistic model (including ones specified as simulators)?

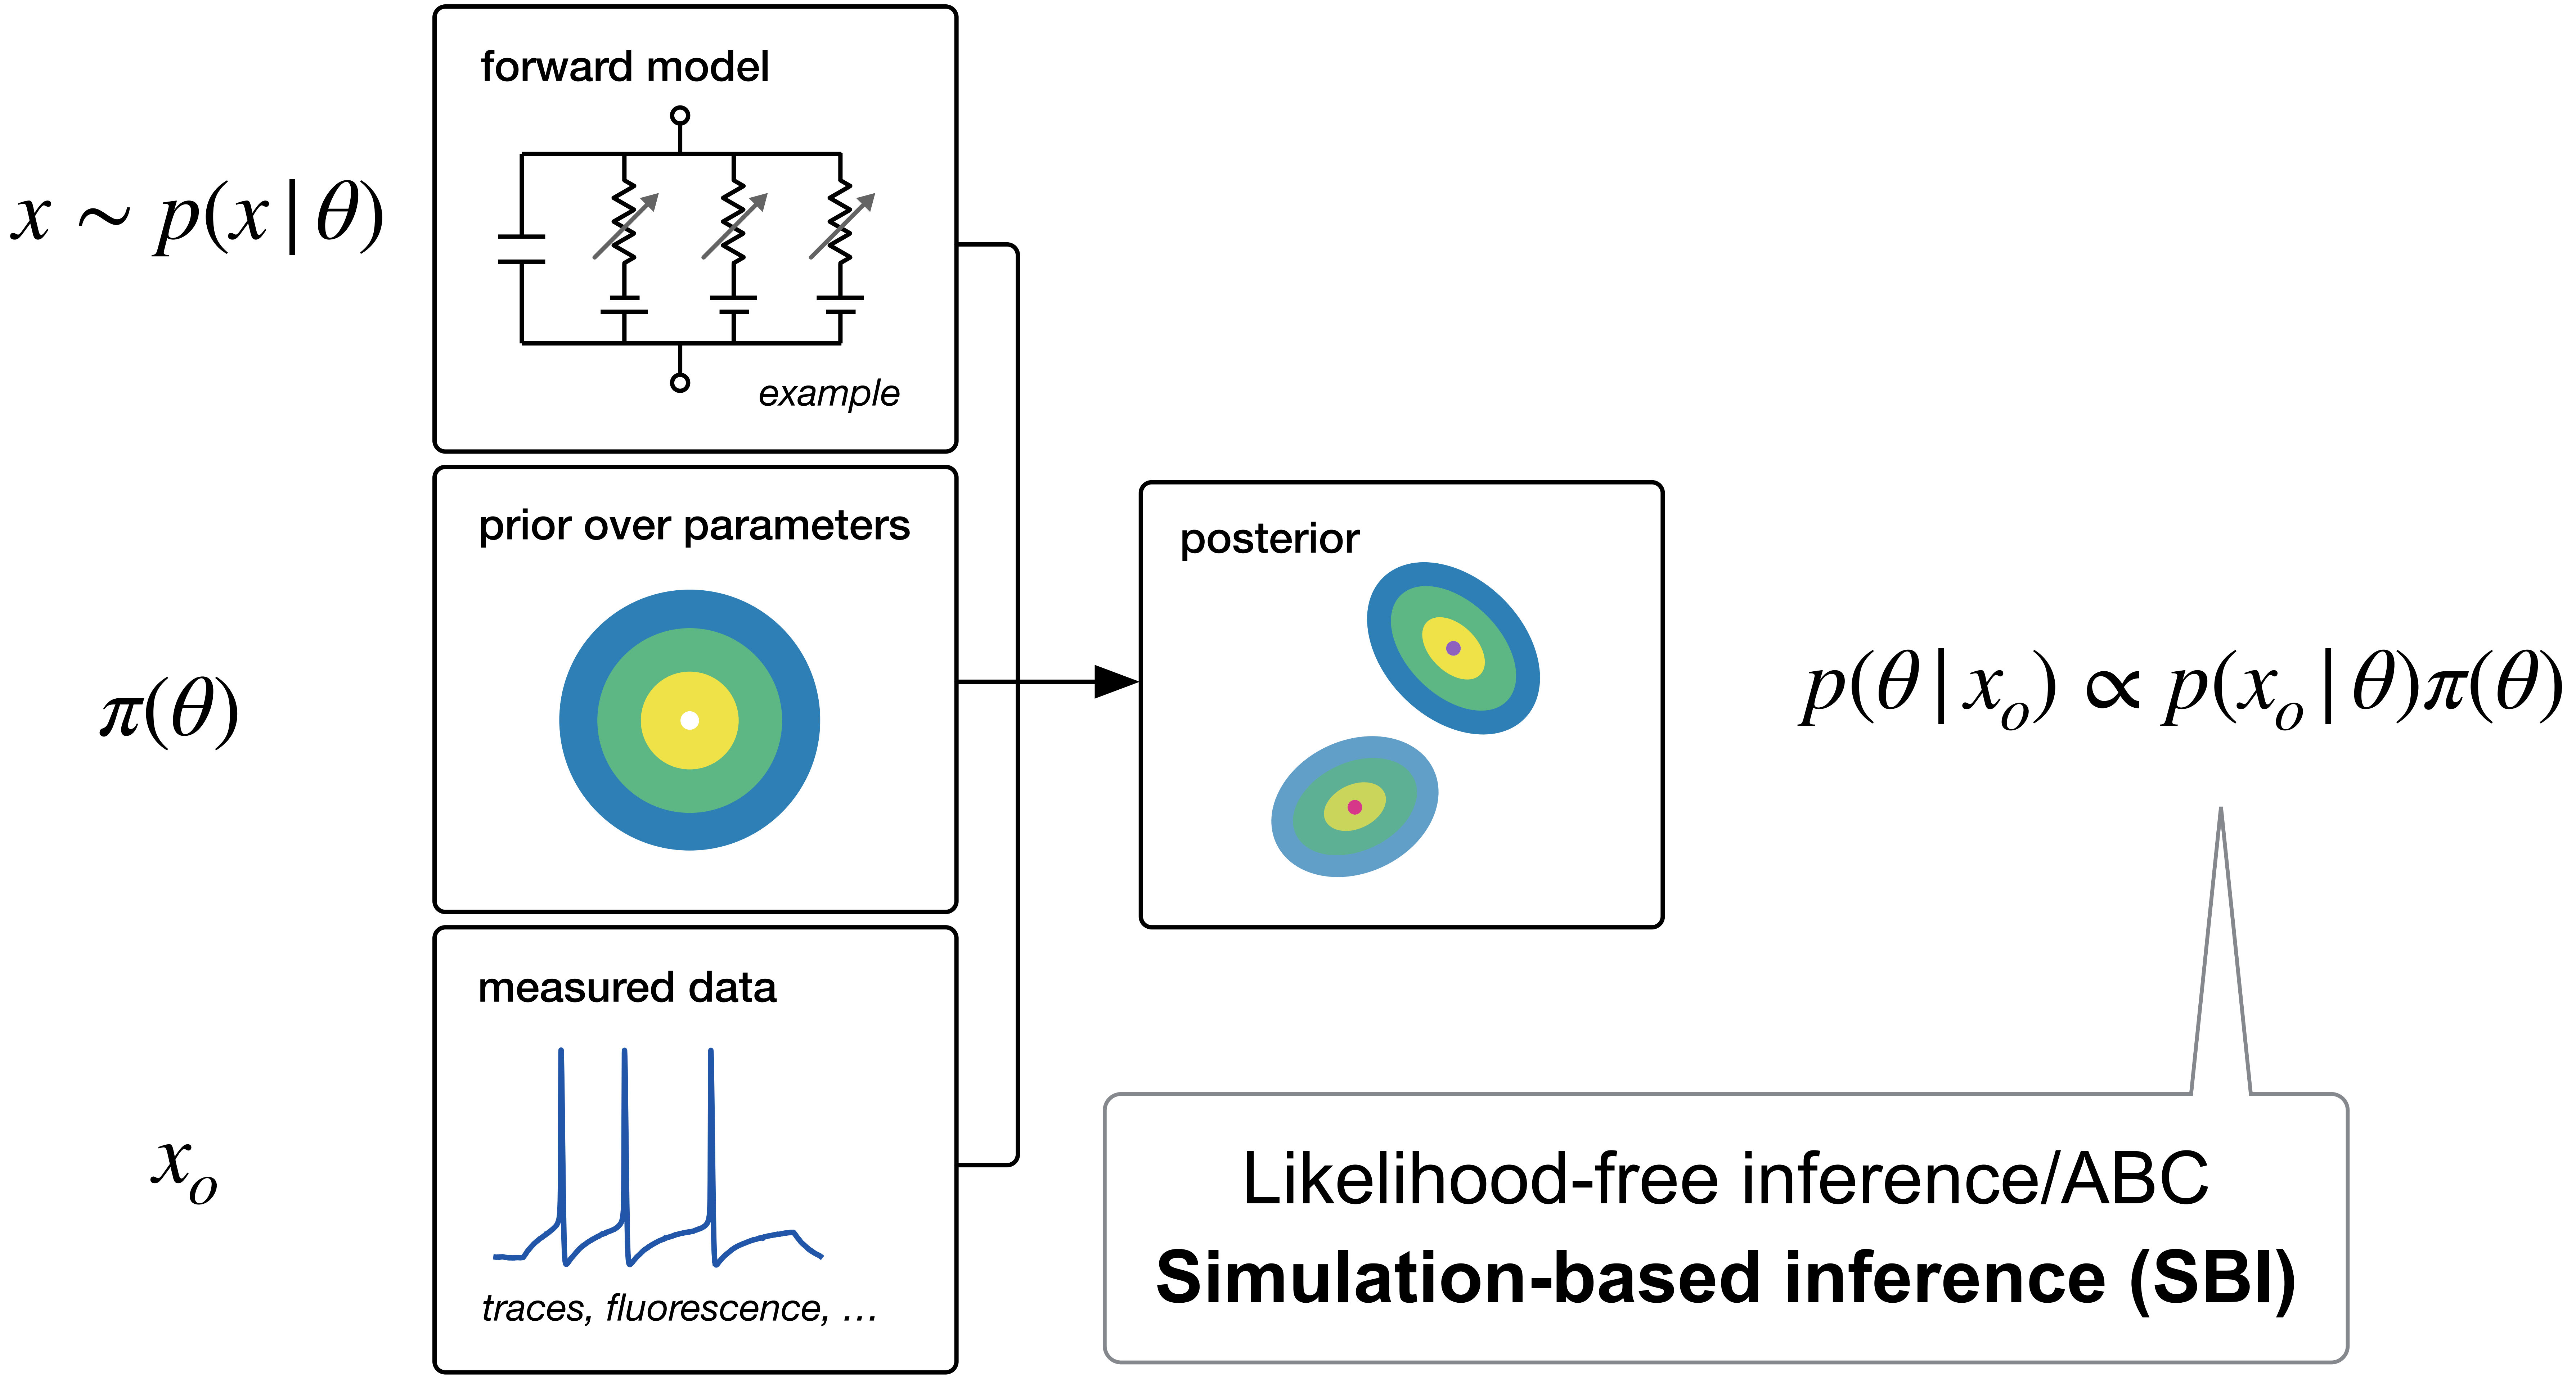

$x \sim p(x \mid \theta)$

forward model

example

$\pi(\theta)$

prior over parameters

posterior

$$p(\theta \mid x_o) \propto p(x_o \mid \theta)\pi(\theta)$$

x_o

measured data

traces, fluorescence, …

Likelihood-free inference/ABC

Simulation-based inference (SBI)

Cranmer, Brehmer, Louppe, 2020

1. Sample parameters from prior

$$\boldsymbol{\theta}_{1:N} \sim \pi(\boldsymbol{\theta})$$

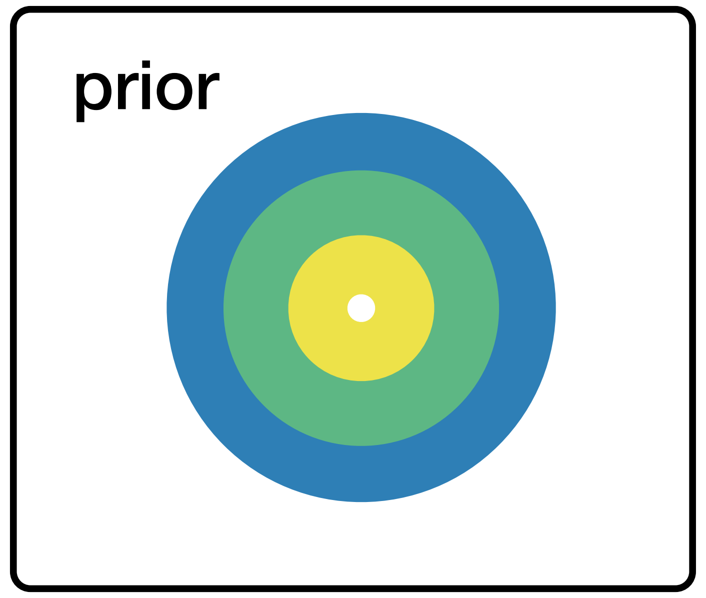

2. Simulate data from parameters

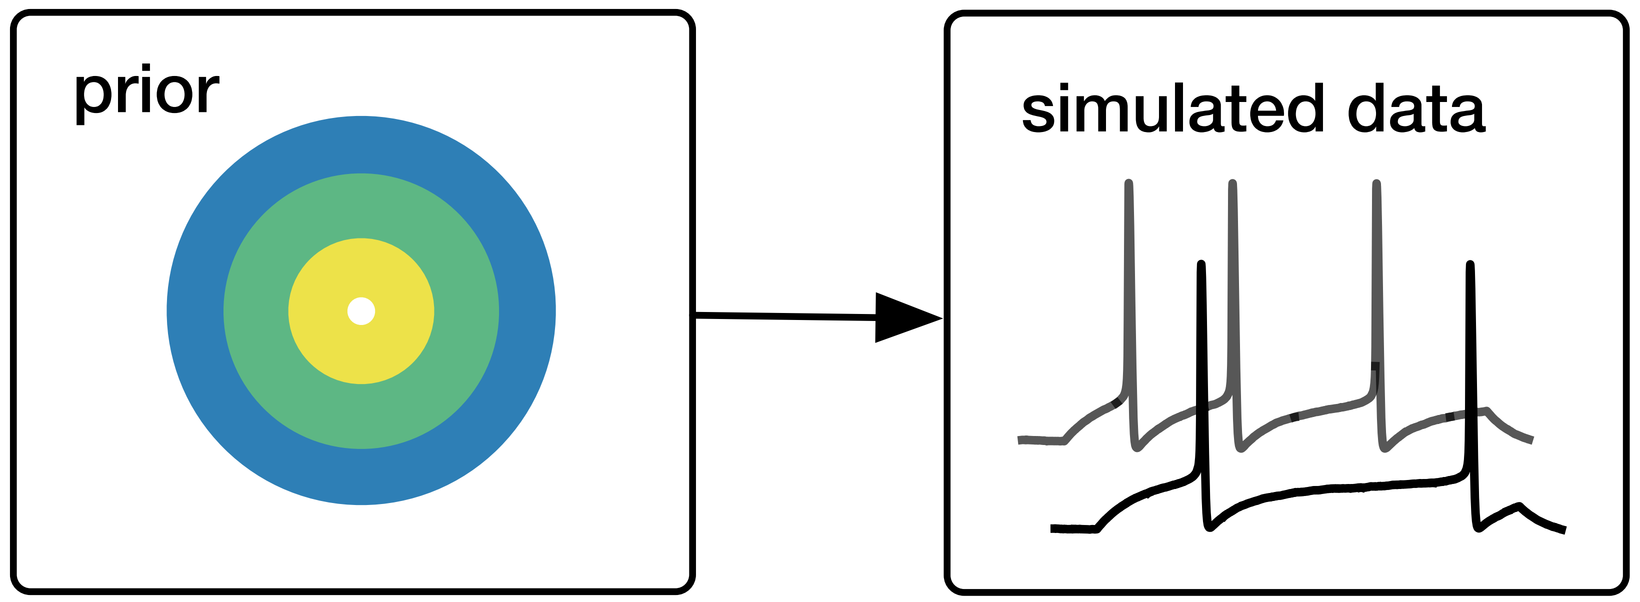

$$\boldsymbol{\theta}_{1:N} \sim \pi(\boldsymbol{\theta}) \qquad \mathbf{x}_{1:N} \sim p(\mathbf{x}|\boldsymbol{\theta}_{1:N})$$

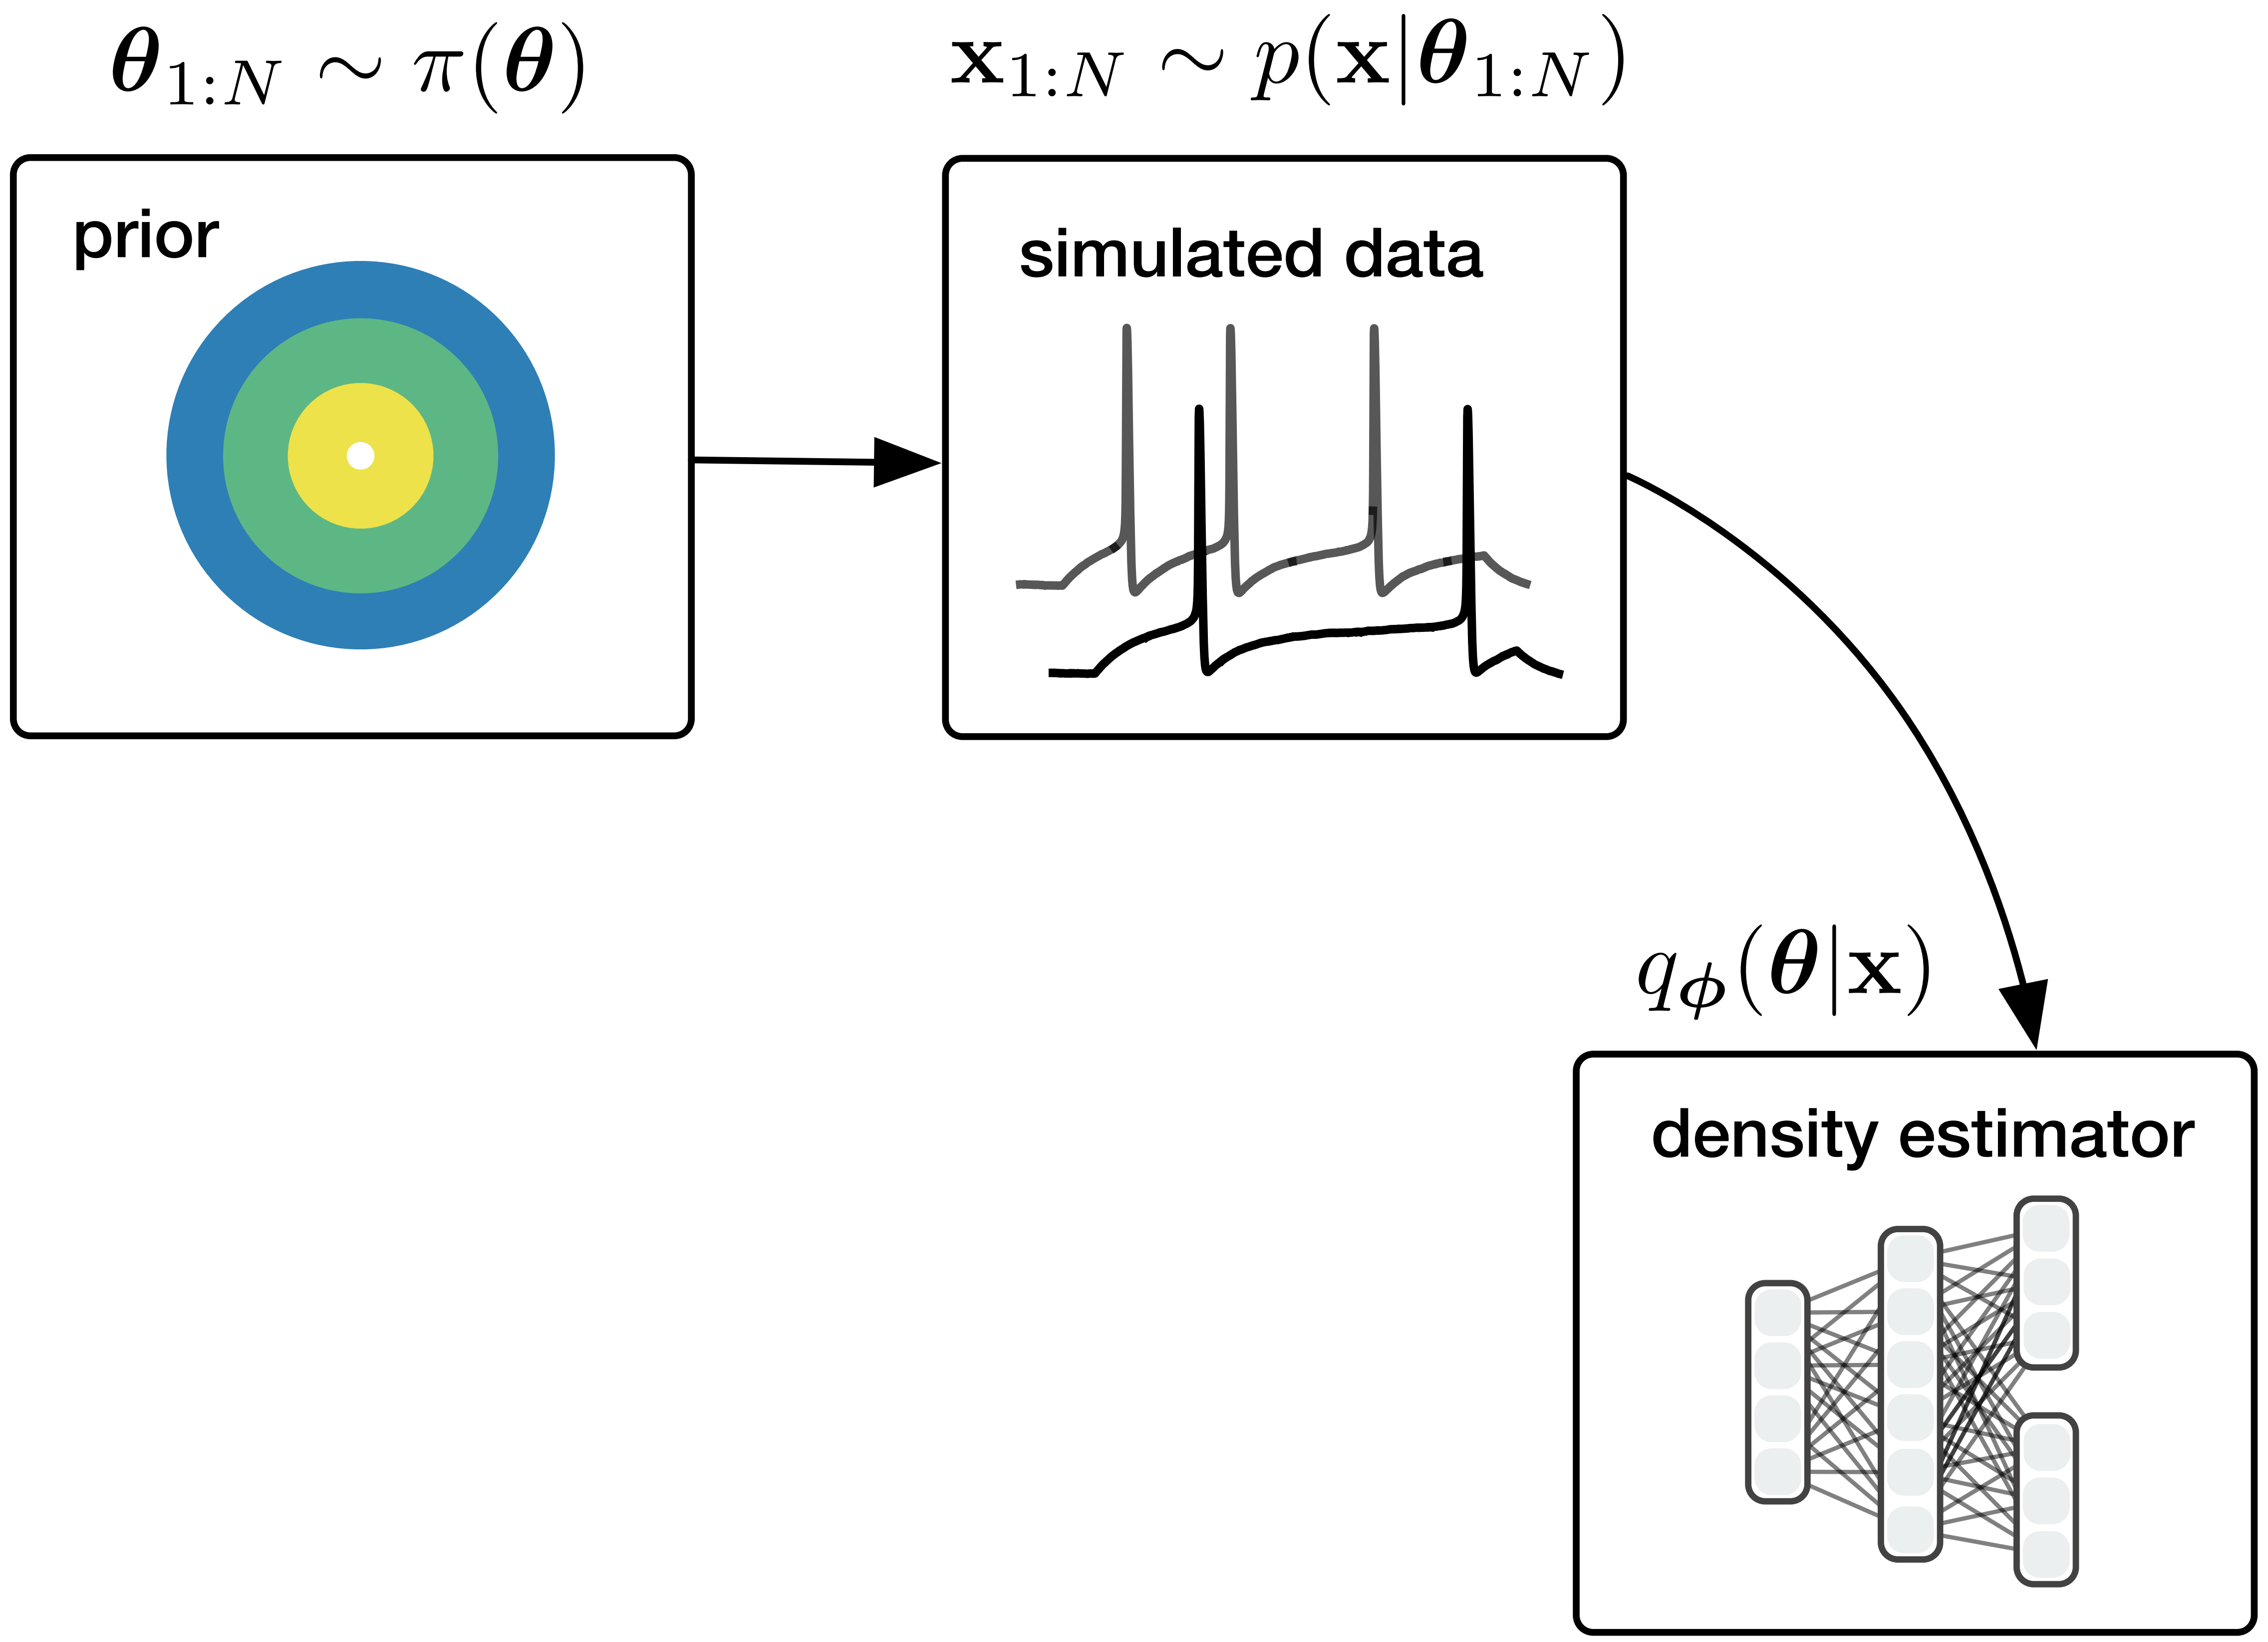

$\boldsymbol{\theta}_{1:N} \sim \pi(\boldsymbol{\theta})$

prior

$\mathbf{x}_{1:N} \sim p(\mathbf{x}|\boldsymbol{\theta}_{1:N})$

simulated data

$q_{\boldsymbol{\phi}}(\boldsymbol{\theta}|\mathbf{x})$

density estimator

3. Train conditional density estimator to predict parameters from (simulated) data

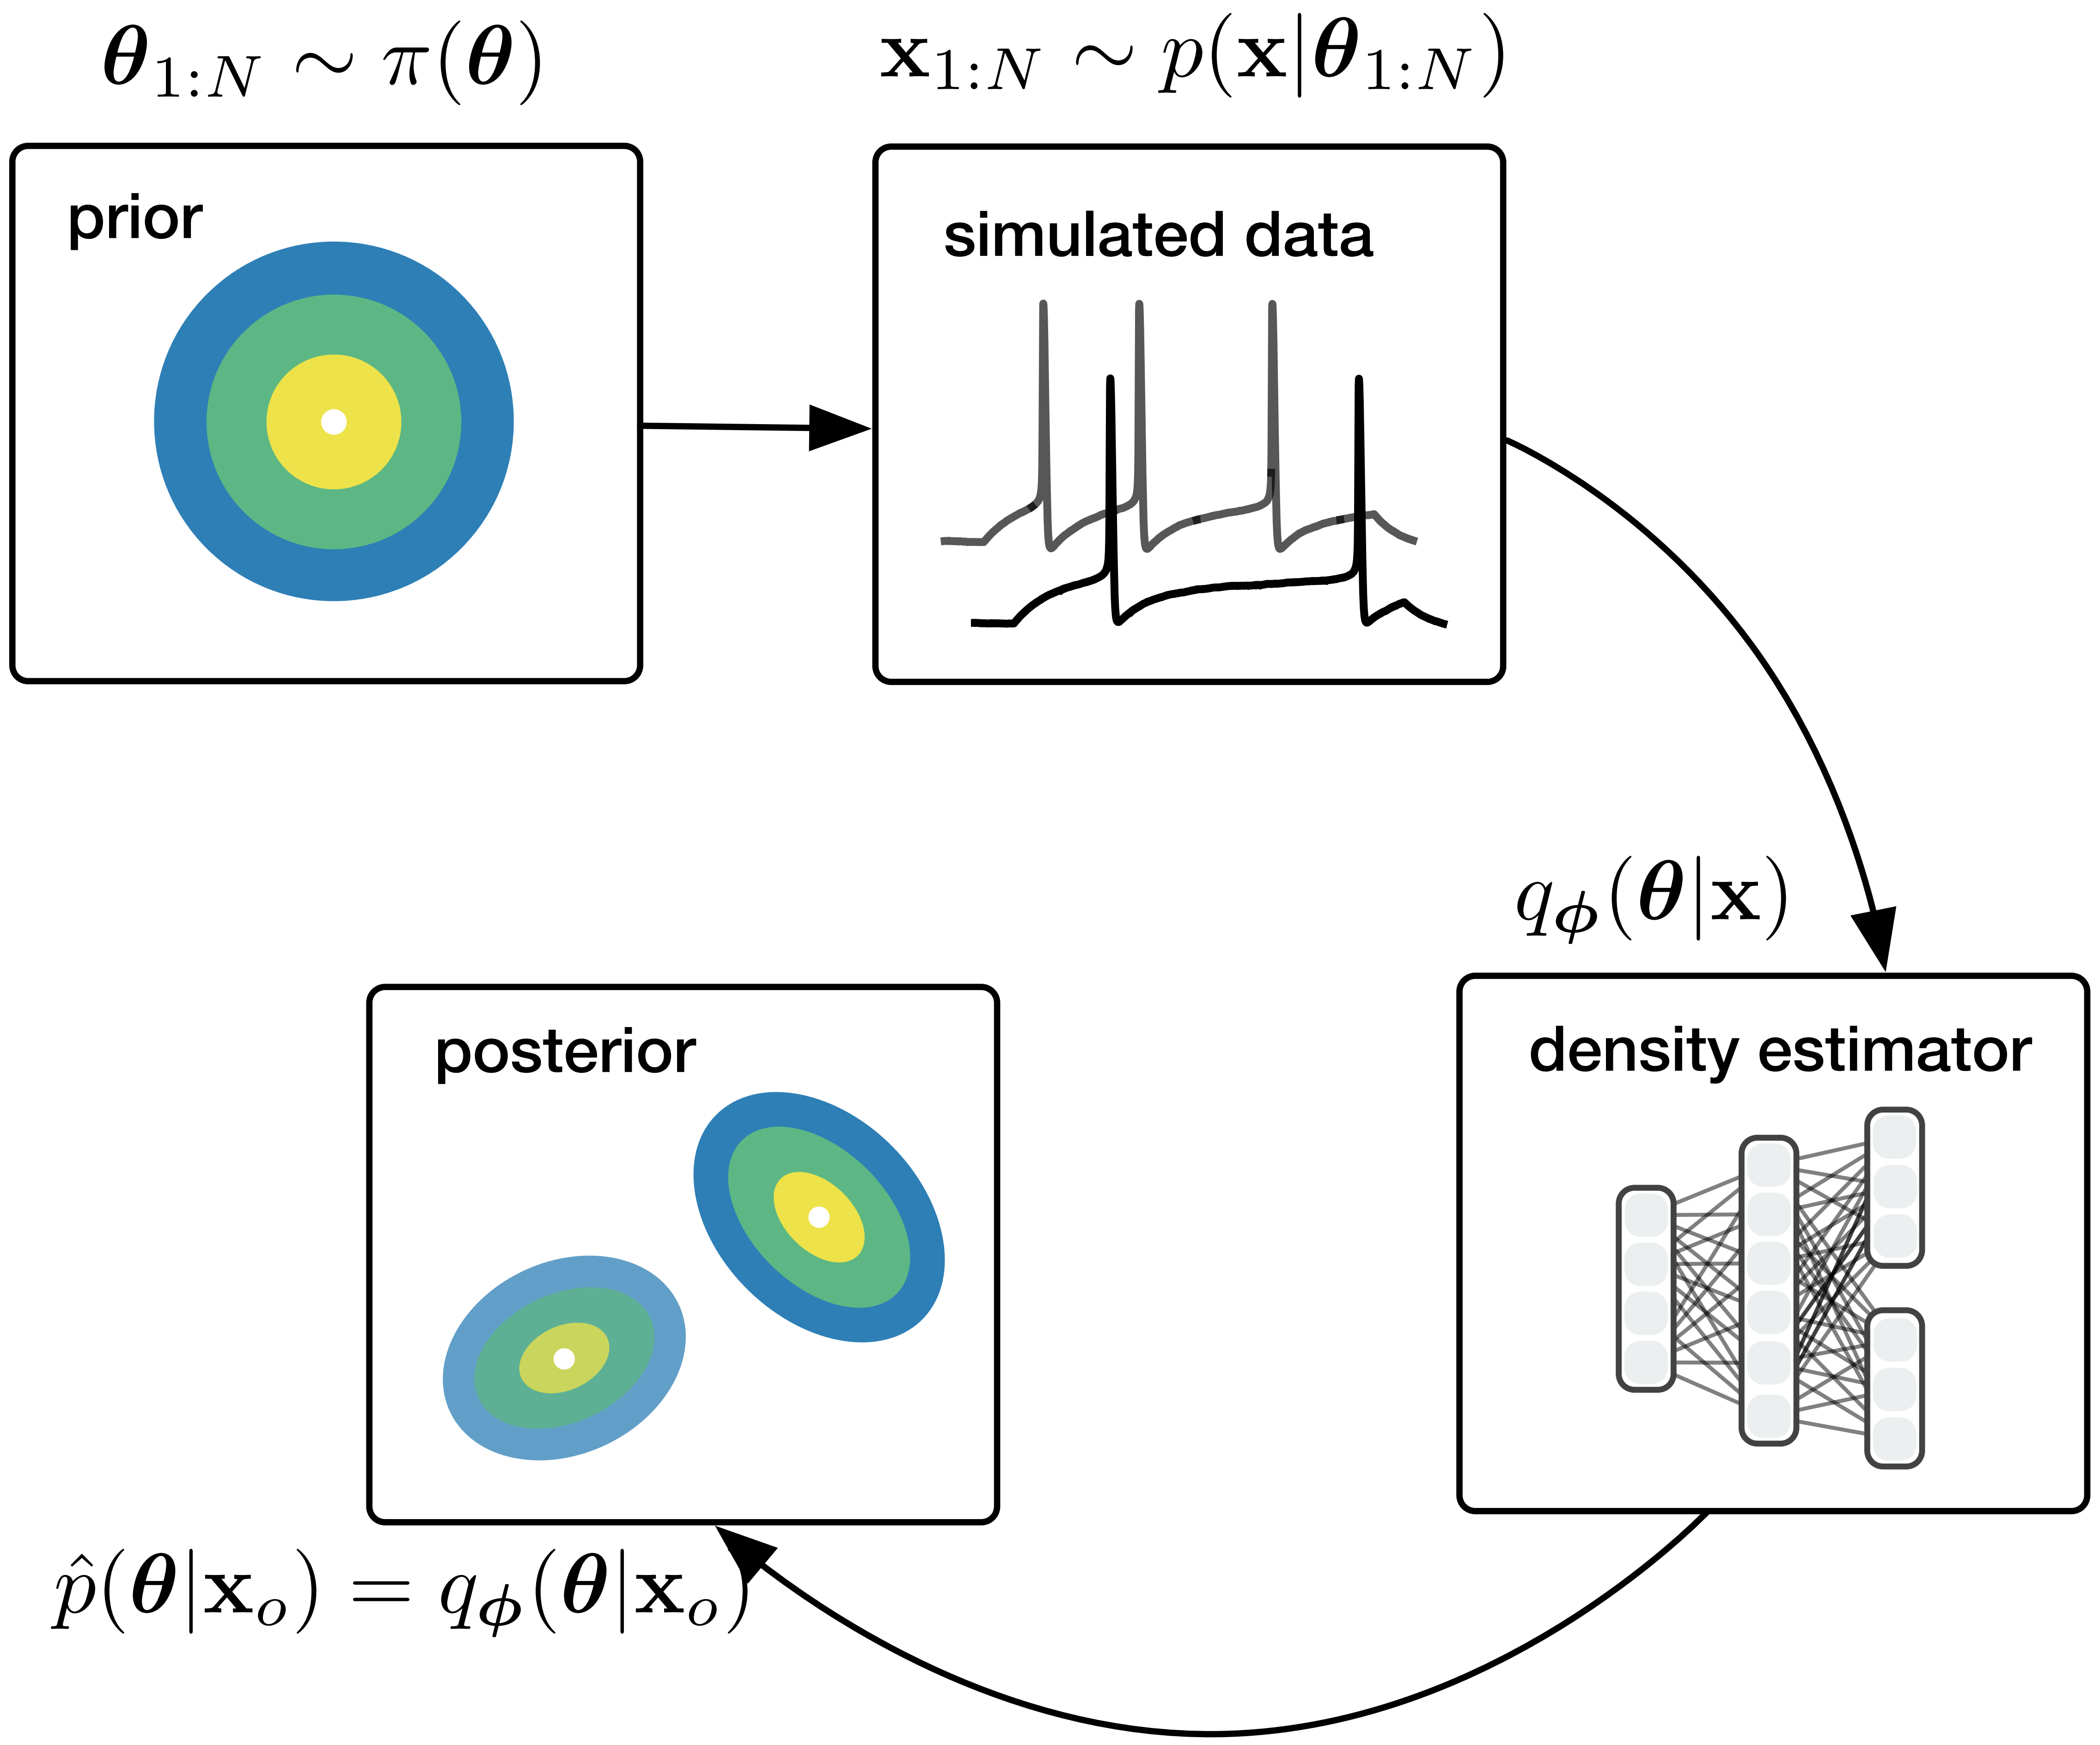

4. Plug empirical data x_o into density estimator to calculate posterior

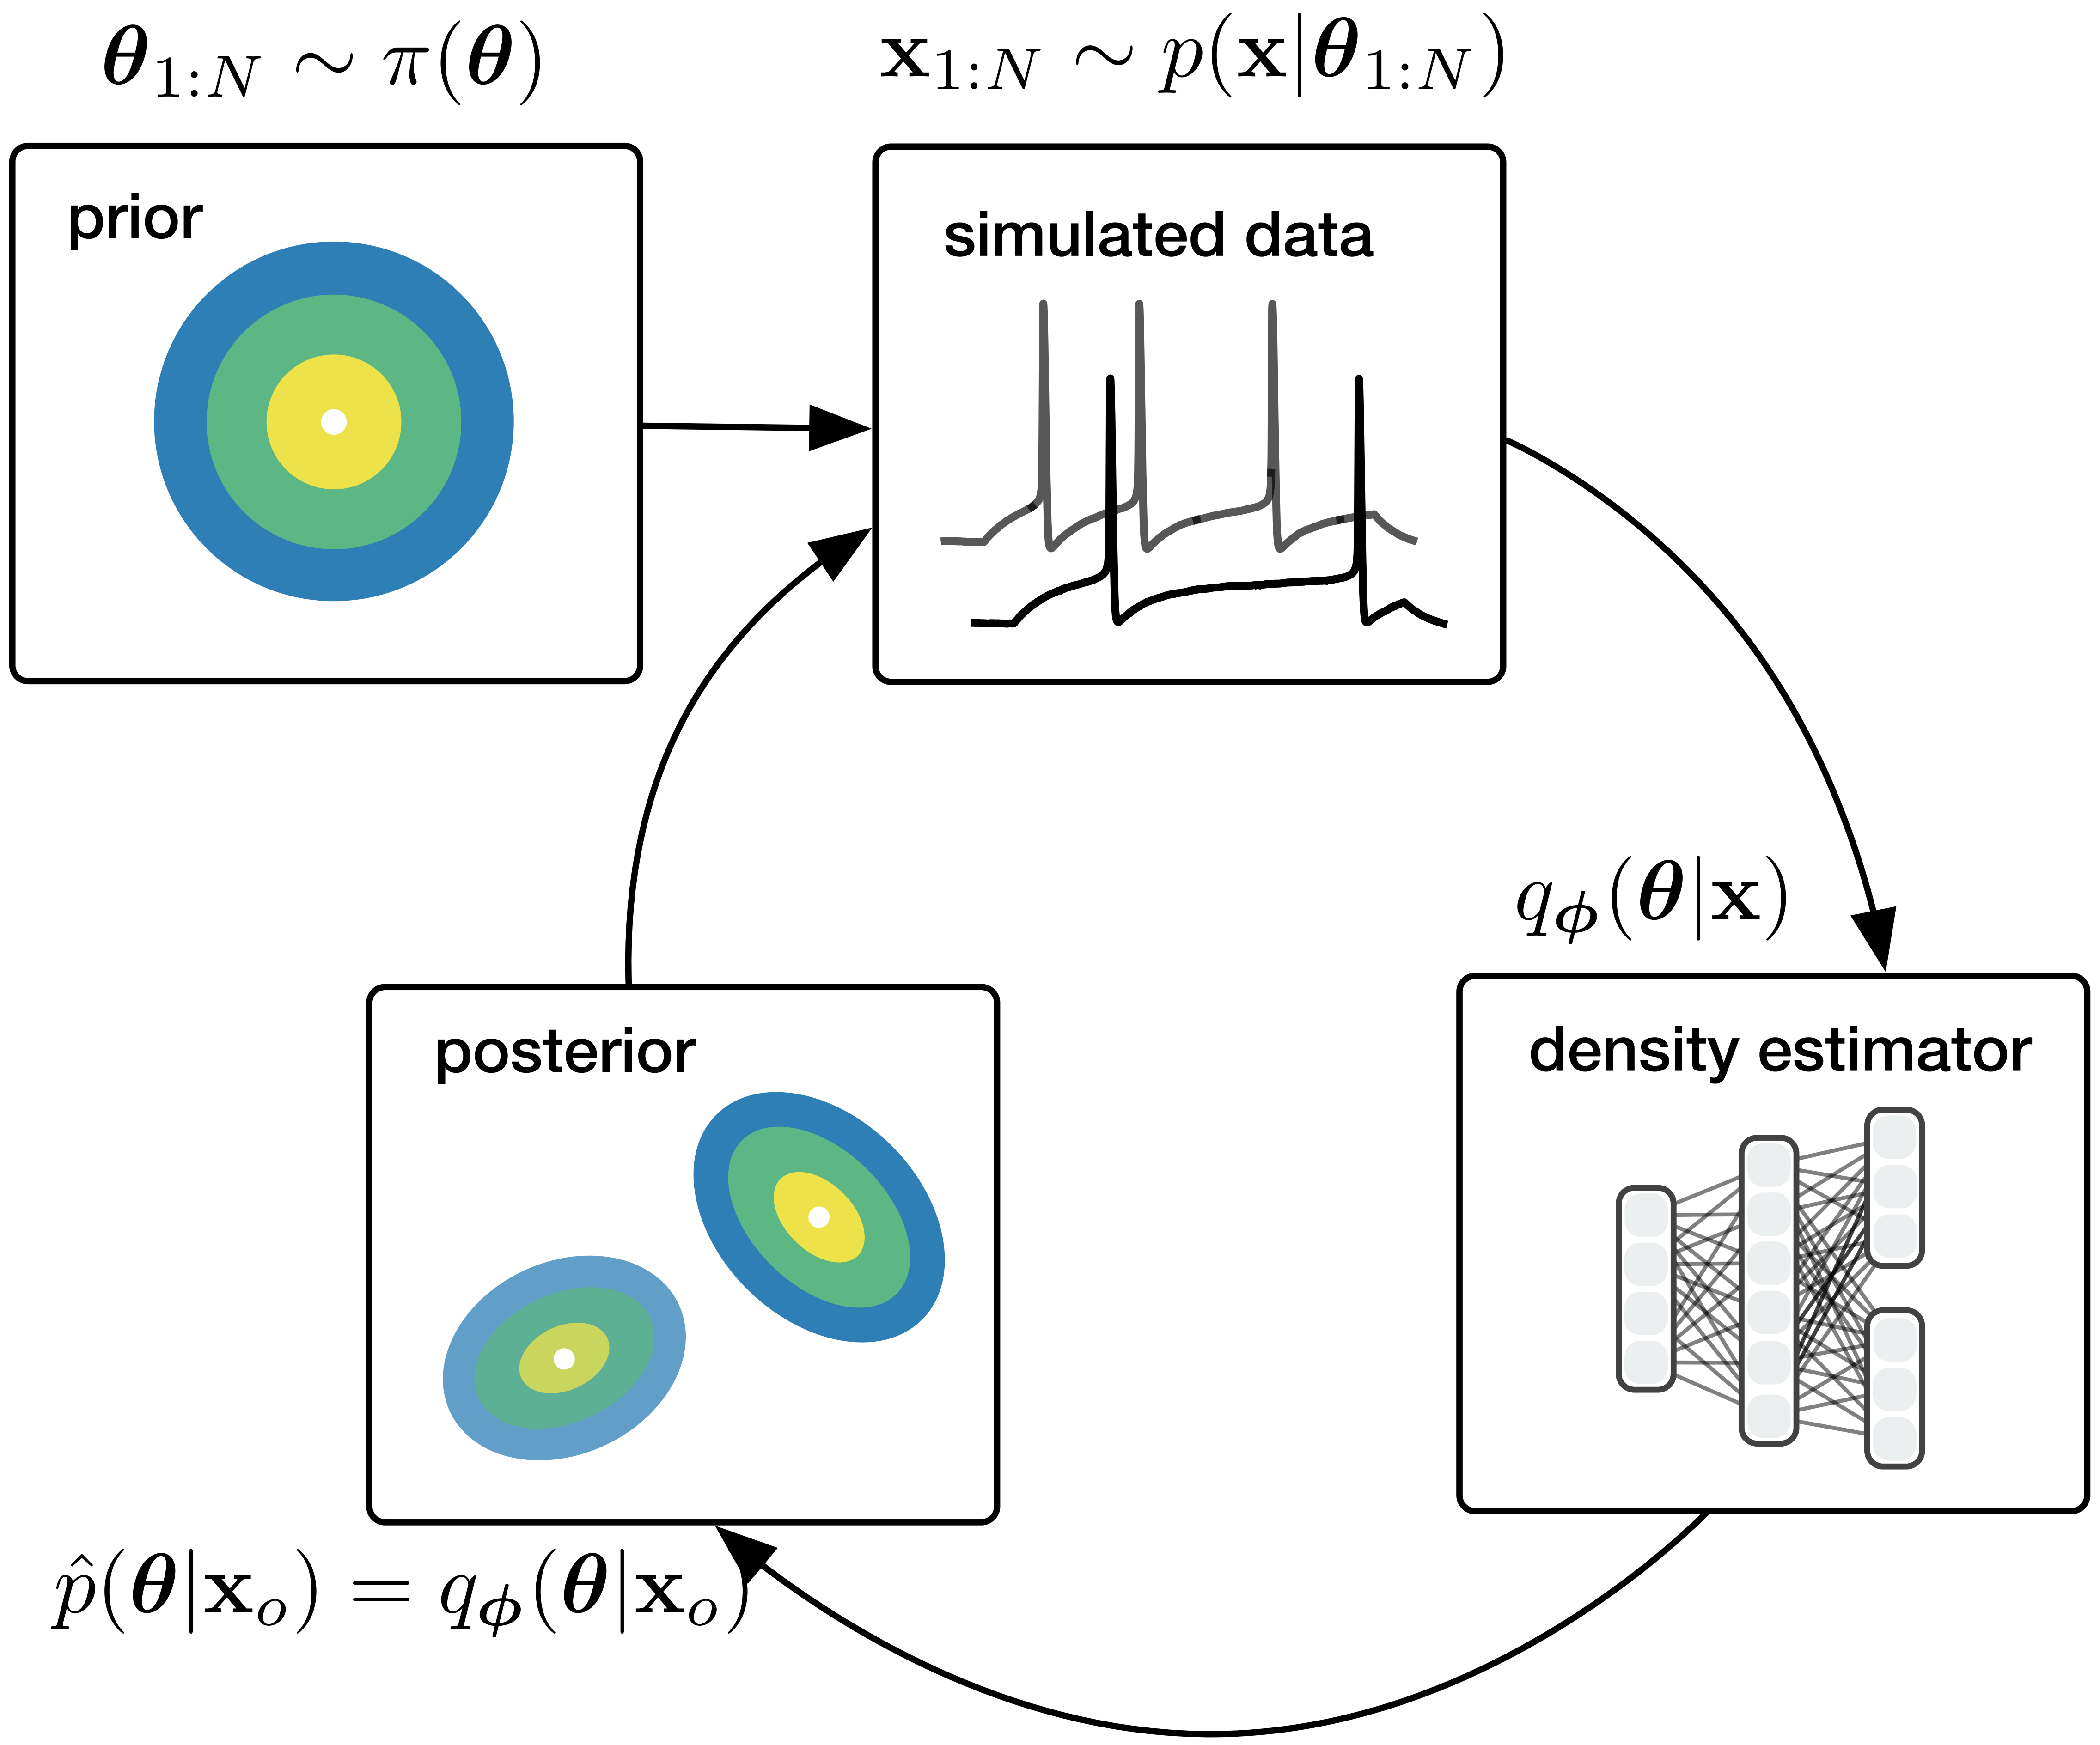

5. If needed, adaptively generate more simulations

(Sequential) Neural Posterior Estimation:
Train neural networks to perform inference on simulations

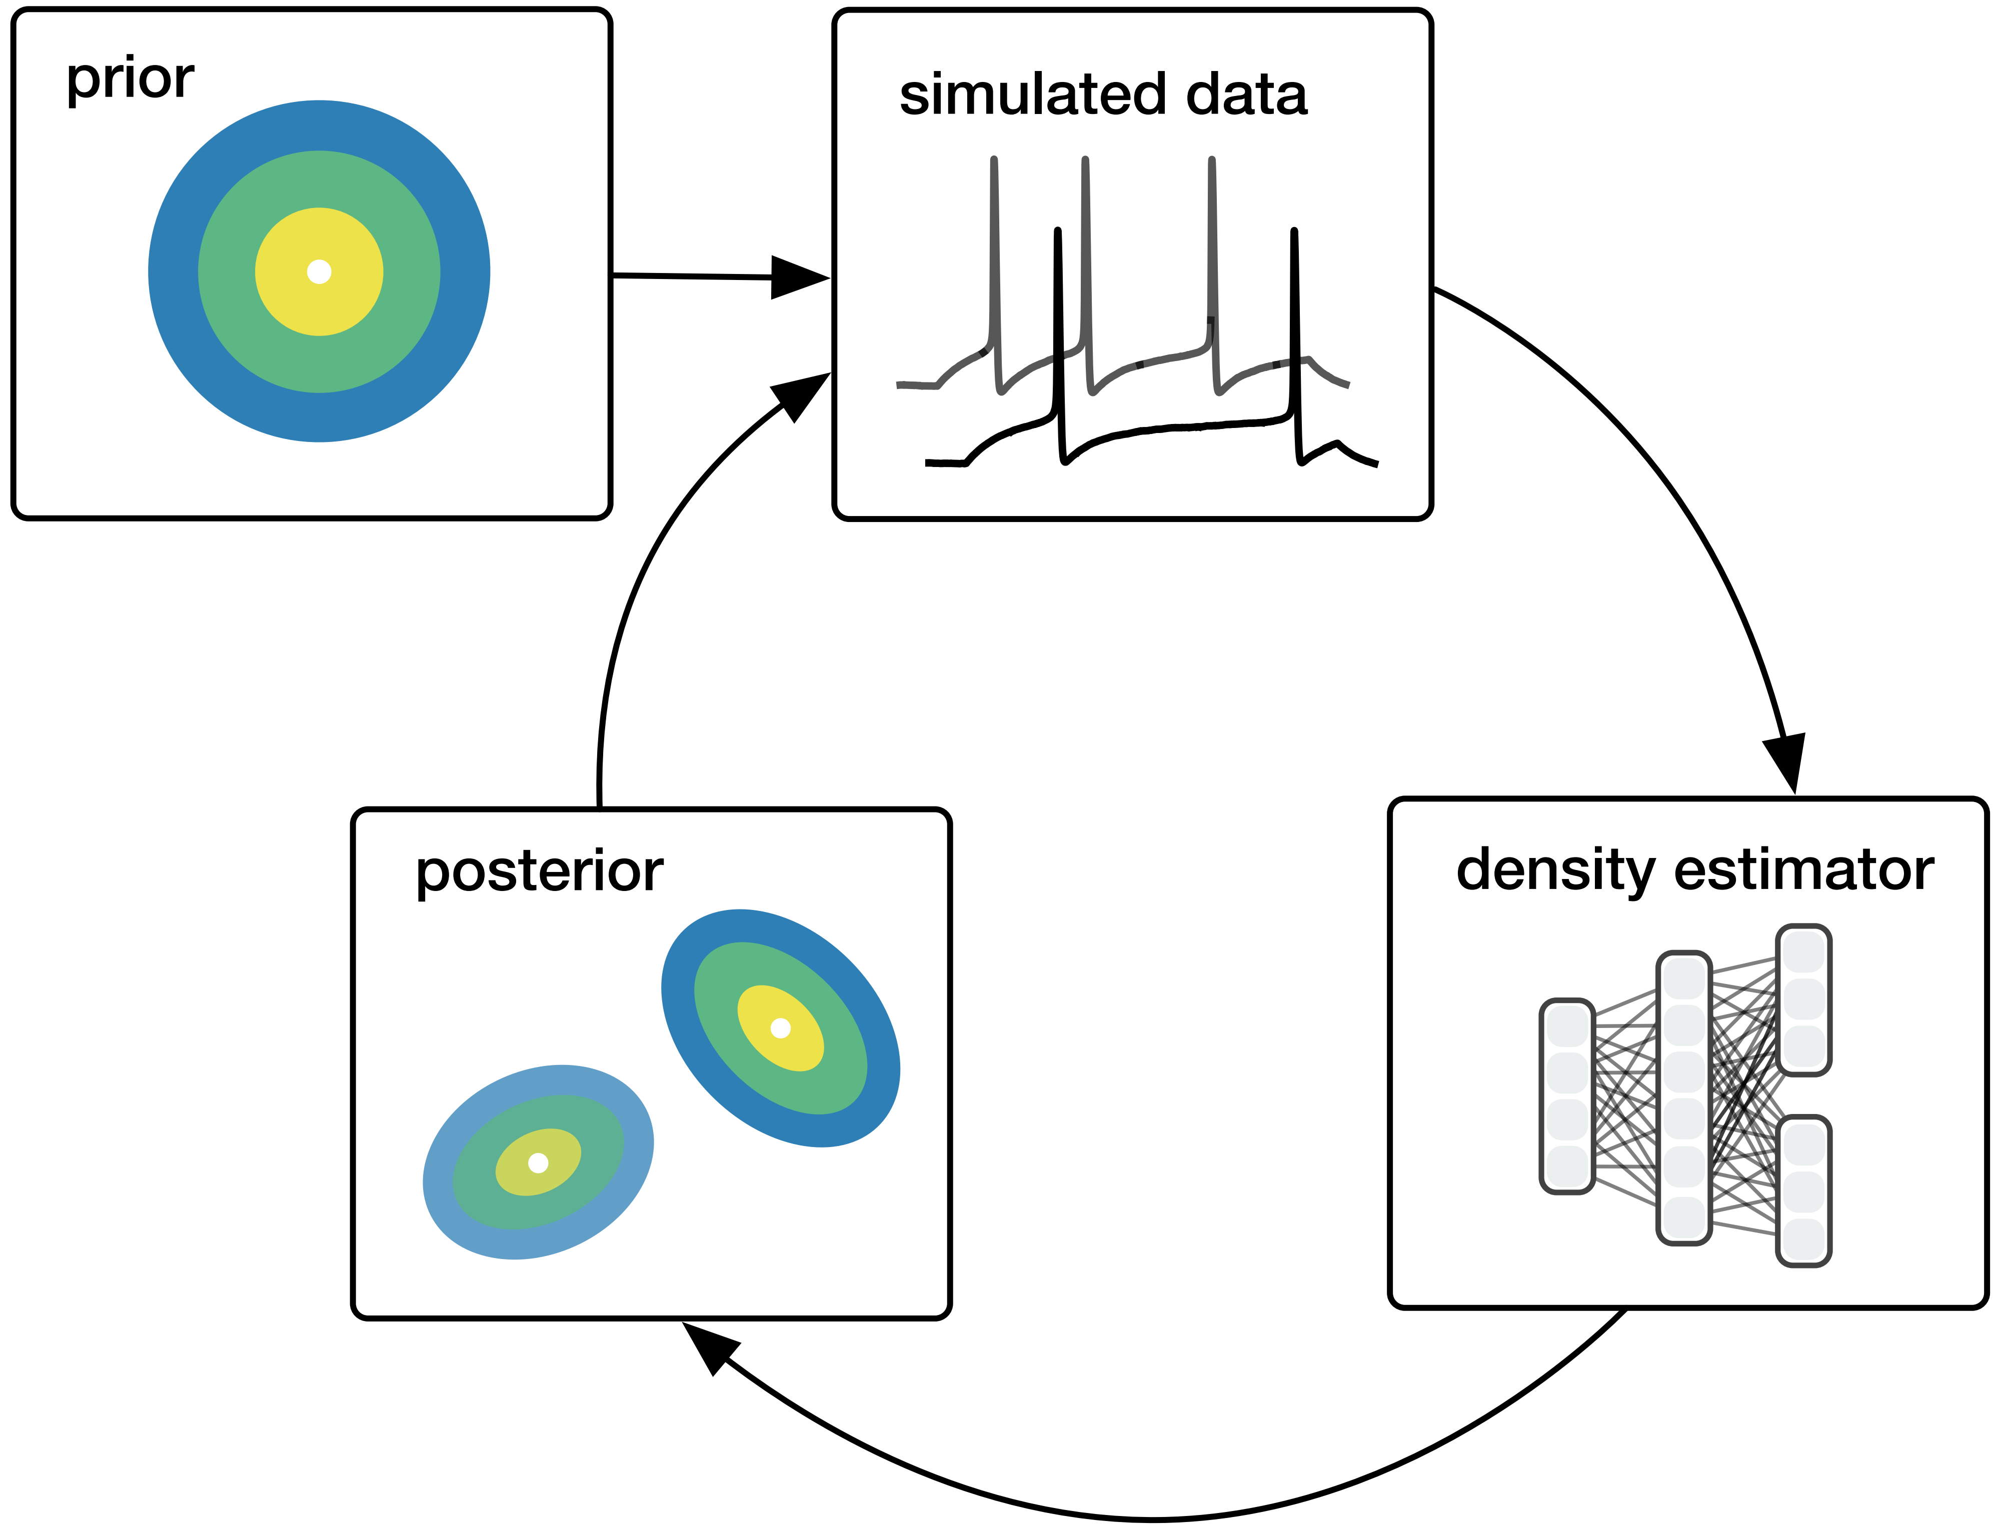

Lueckmann, Goncalves et al NeurIPS 2017
Greenberg et al ICML 2019
Lueckmann et al PRML 2019
Goncalves, Lueckmann, Deistler et al eLife 2020
Tejero-Cantero et al, JOSS 2020
Lueckmann et al AISTATS 2021
Boelts et al eLife 2022
Ramesh et al ICLR 2022
Dax et al ICLR 2022
Glöckler et al ICLR 2022
Deistler et al NeurIPS 2022
Deistler et al PNAS 2022
Beck et al NeurIPS 2022
Glöckler et al ICML 2023
Gao, Deistler, Macke NeurIPS 2023
Dax, Wildberger, Buchholz et al NeurIPS 2023
Confafreux, Ramesh et al NeurIPS 2023

Beaumont et al 2002
Blum & Francois 2010
Papamakarios & Murray NeurIPS 2016

Other SBI approaches, reviewed in Cranmer, Brehmer, Louppe 2020
Website: simulation-based-inference.org/

Code: www.mackelab.org/sbi
Benchmark: sbi-benchmark.github.io/

Alternatives to Neural Posterior Estimation: Sequential Neural Variational Inference (SNVI)

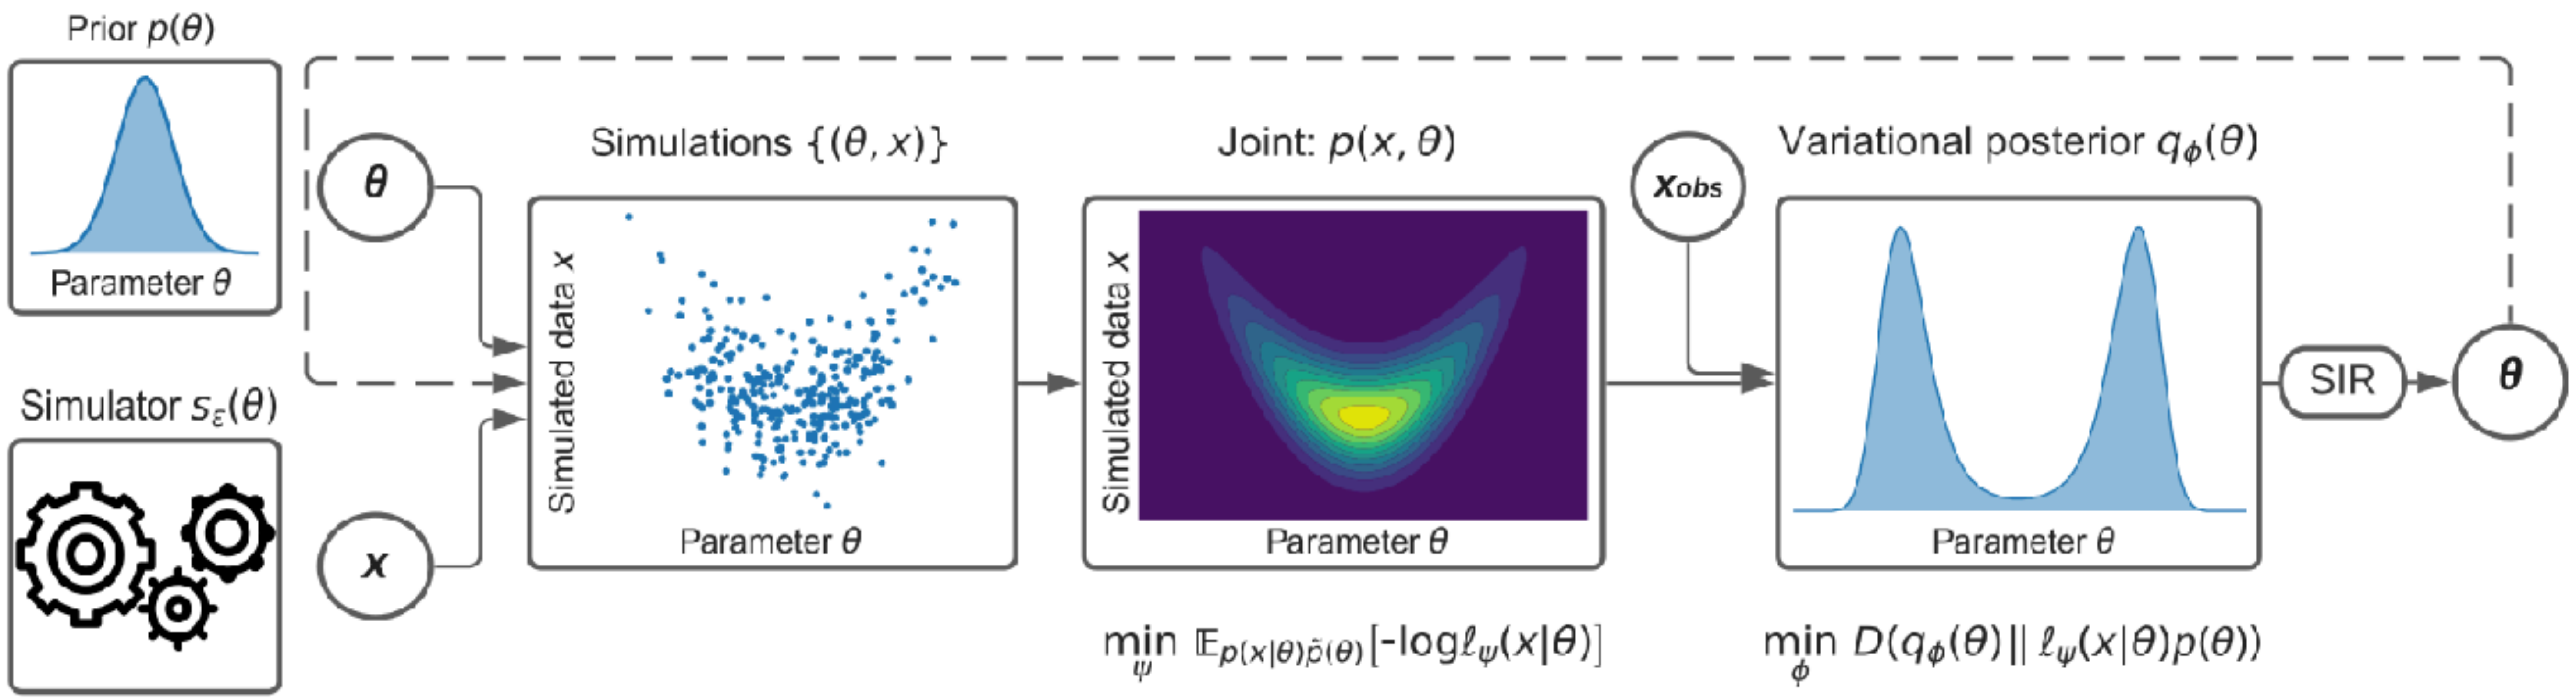

- Sequential Neural Likelihood (SNL, Papamakarios et al 2018): Estimate likelihood from simulations
- Requires additional MCMC
- Solution: Combine SNL with variational inference, learn likelihood and posterior in parallel

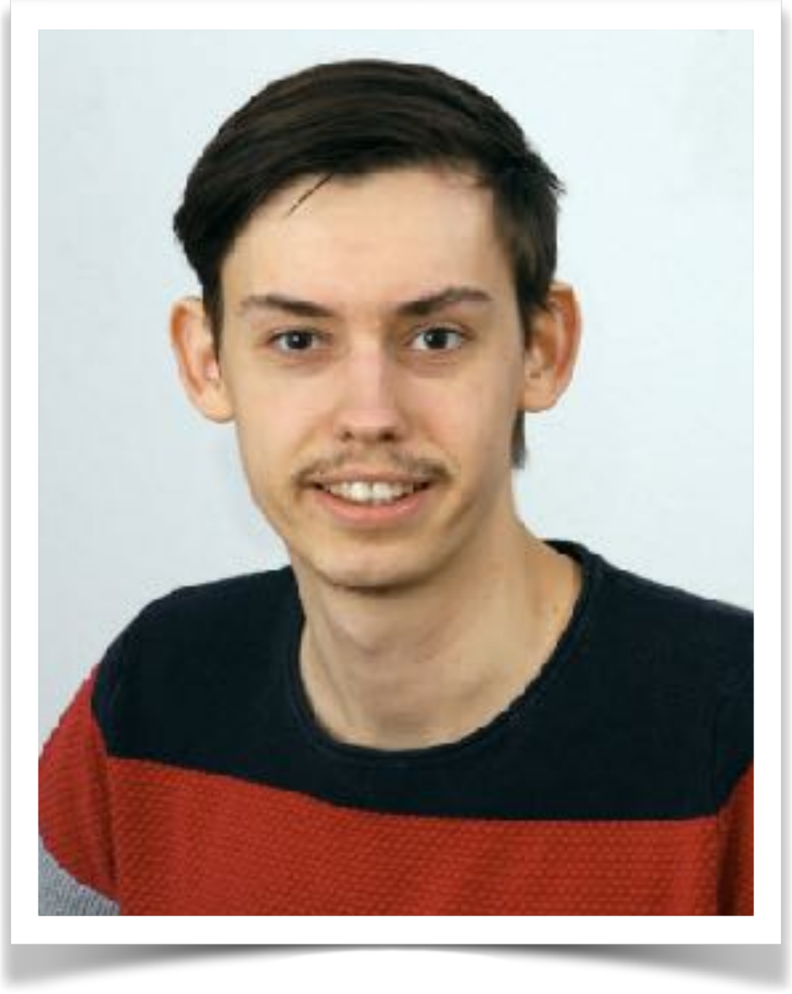

Manuel
Glöckler

Glöckler, Deistler, Macke, ICLR 2022
also see Wiqvist et al 2021

Alternatives to Neural Posterior Estimation:
GATSBI: Generative Adversarial Training for SBI

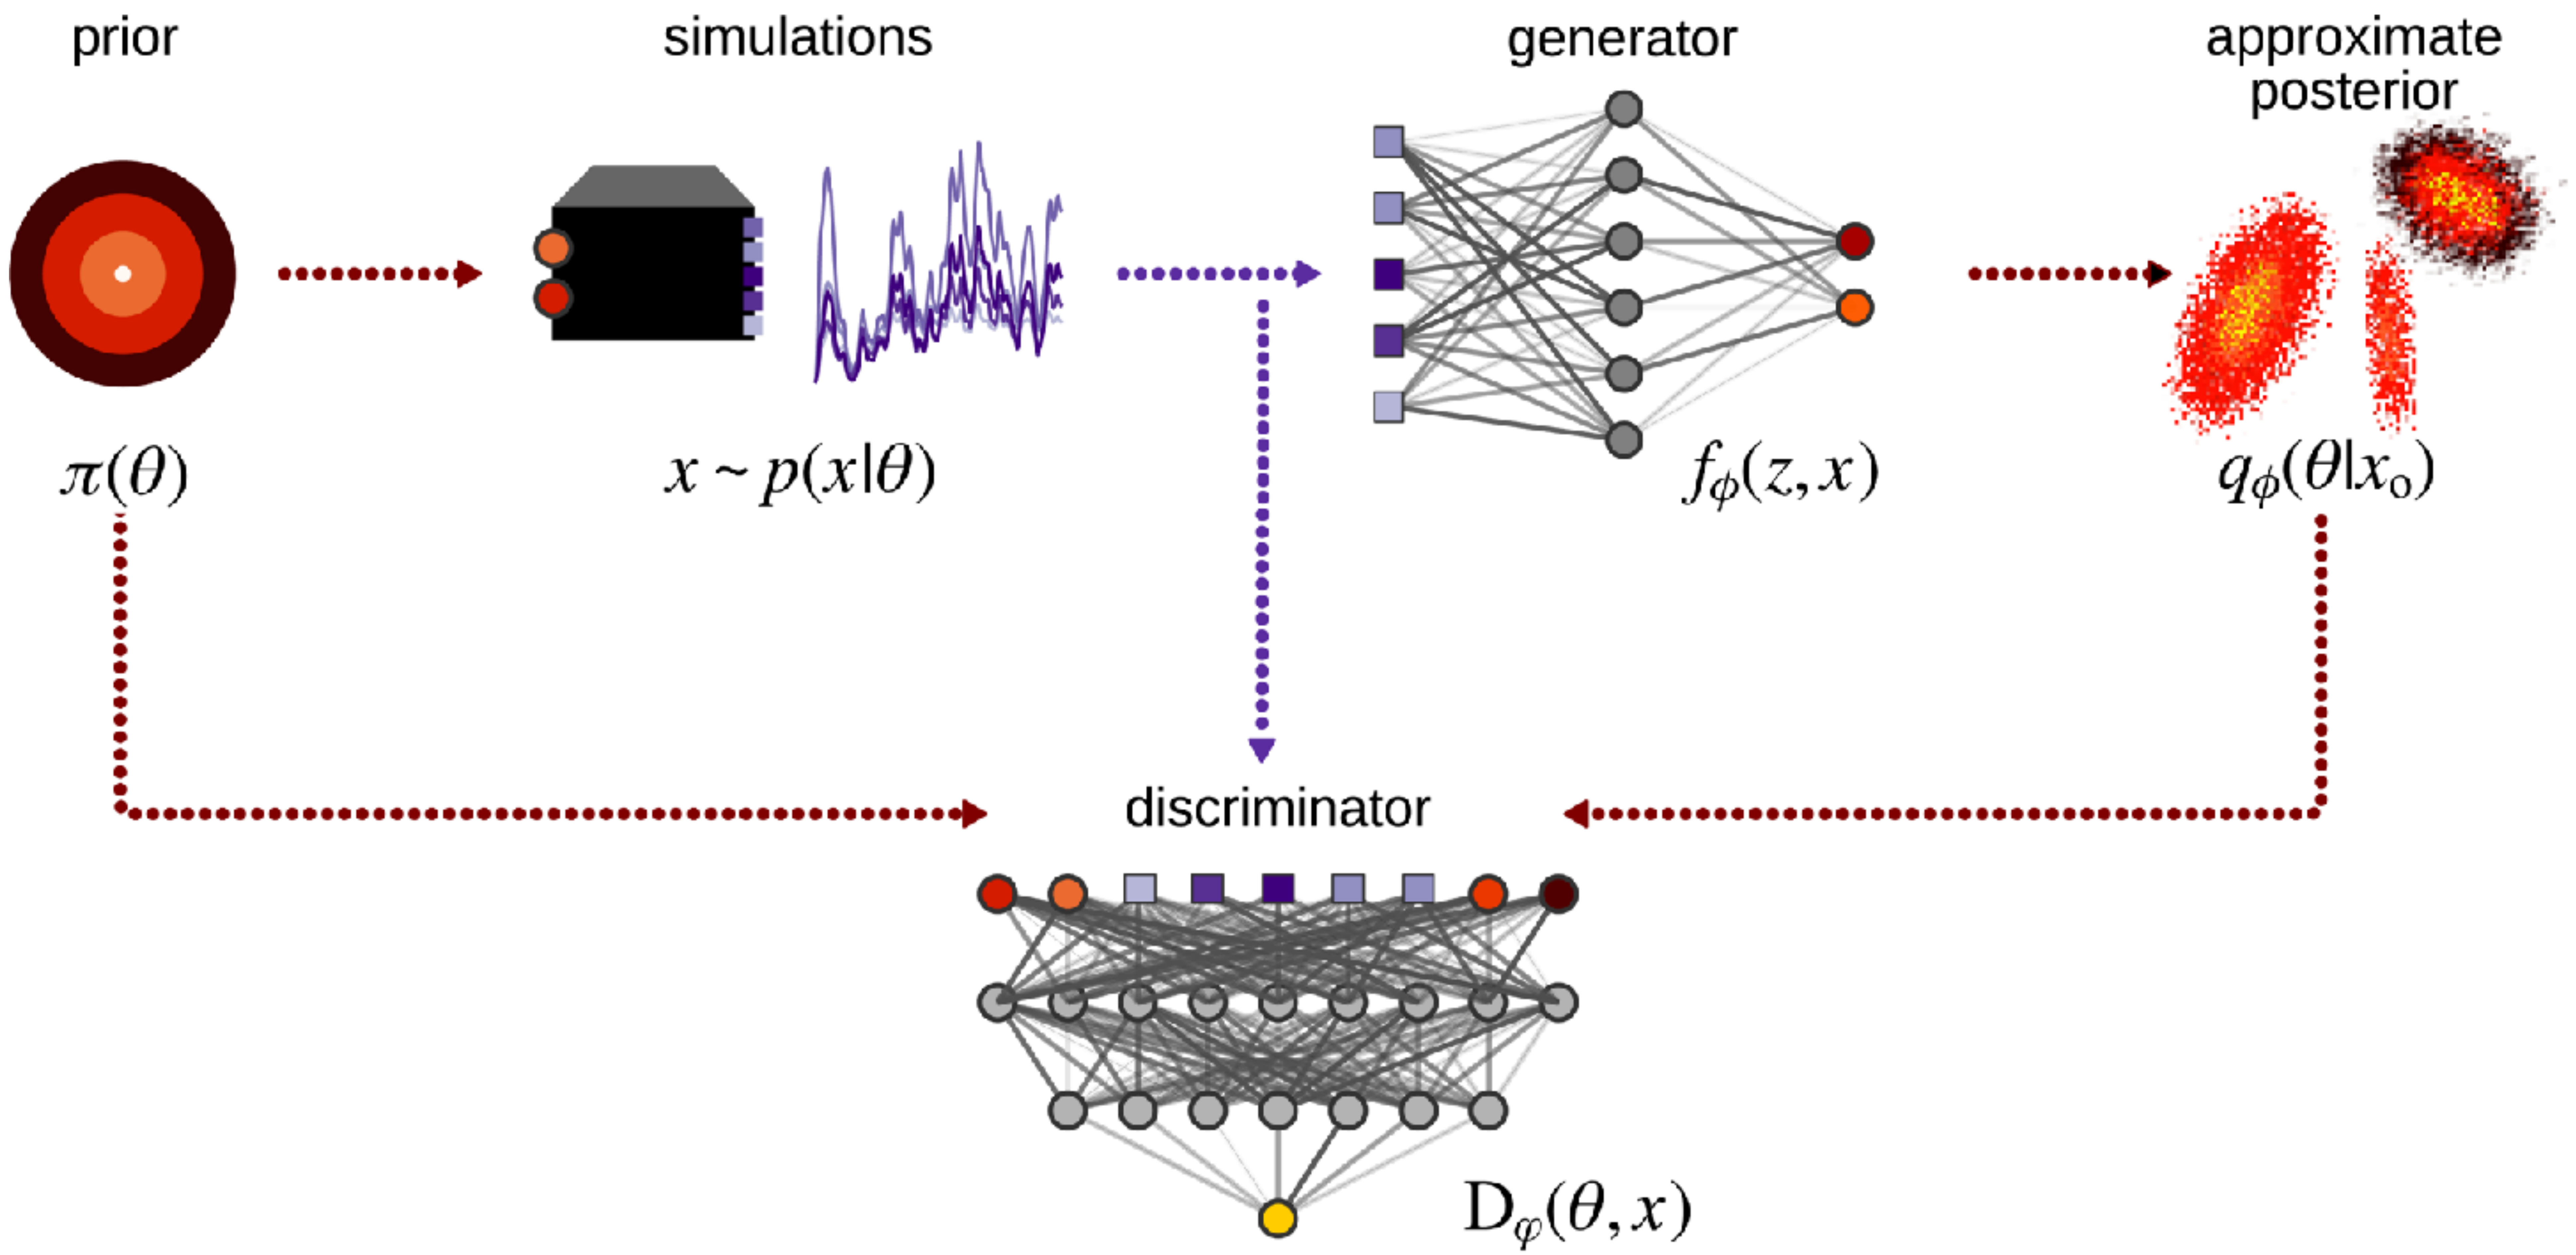

Poornima Ramesh

Ramesh et al, ICLR 2022

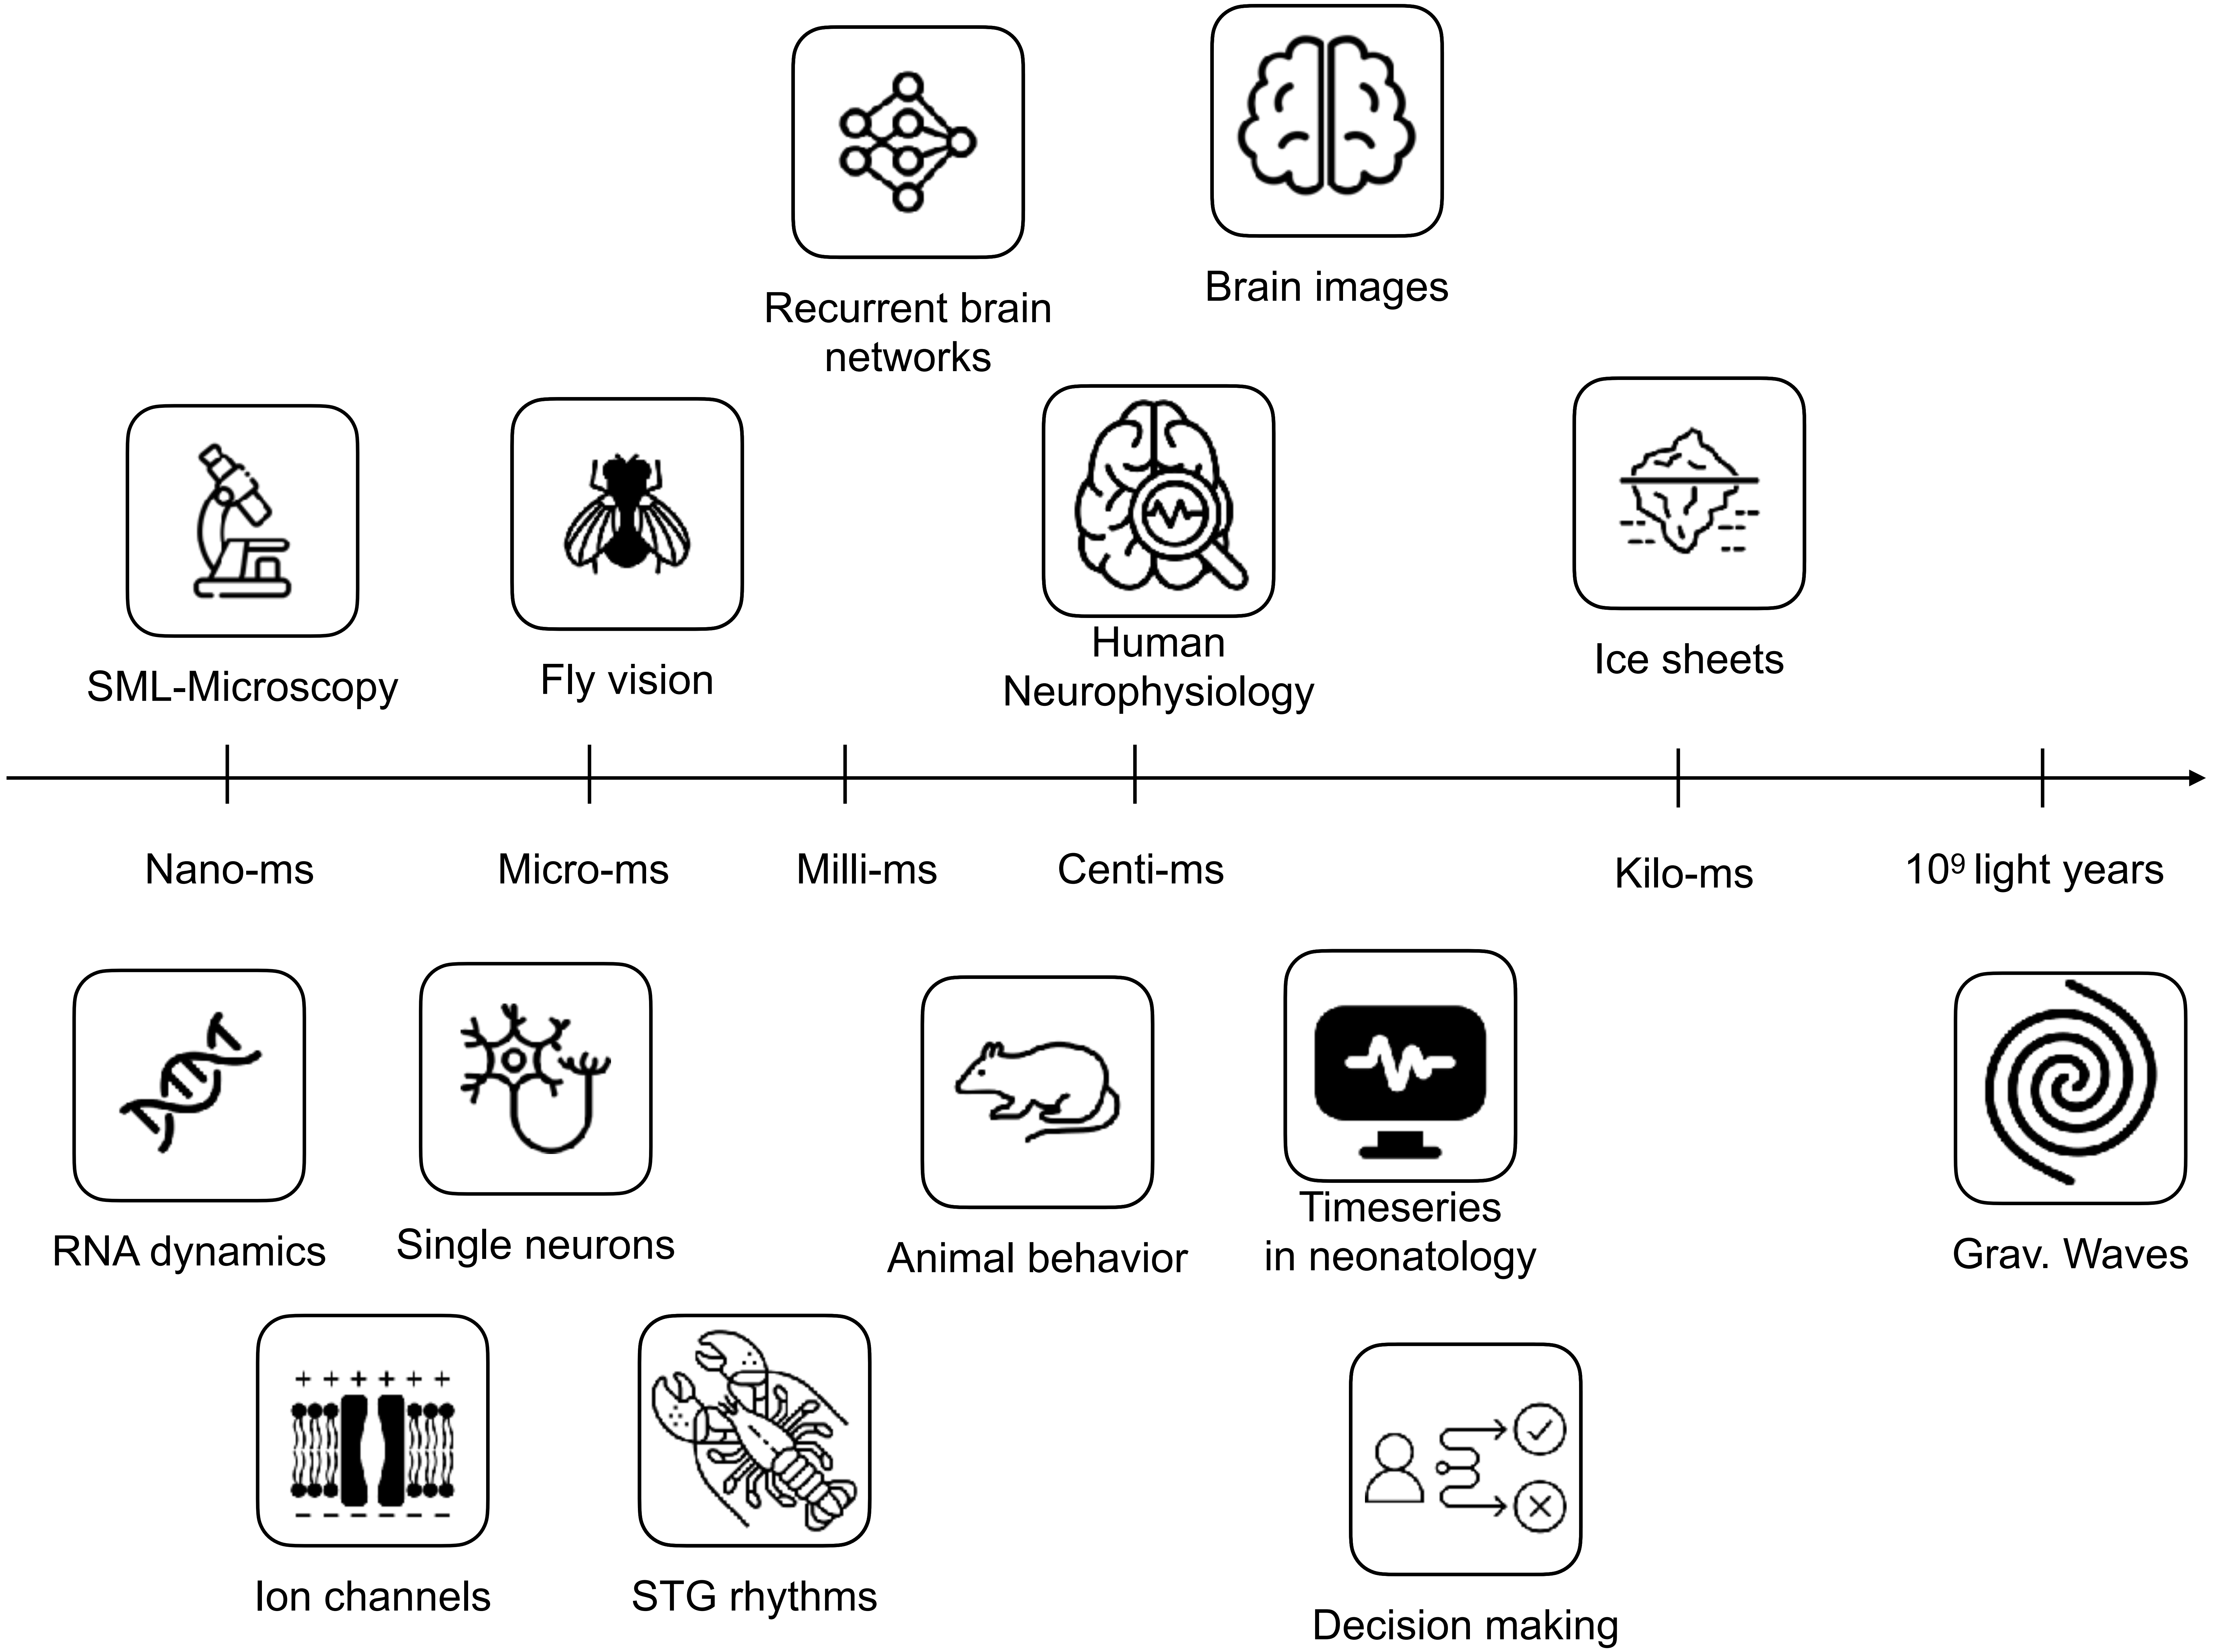

Recurrent brain networks

Brain images

SML-Microscopy

Fly vision

Human Neurophysiology

Ice sheets

Nano-ms

Micro-ms

Milli-ms

Centi-ms

Kilo-ms

10^9 light years

RNA dynamics

Single neurons

Animal behavior

Timeseries in neonatology

Grav. Waves

Ion channels

STG rhythms

Decision making

Application: Faster inference of gravitational wave models

Bernhard Schölkopf (MPI IS Tübingen)
Alessandra Buonanno (MPI Potsdam)

Image Credit: SXS project

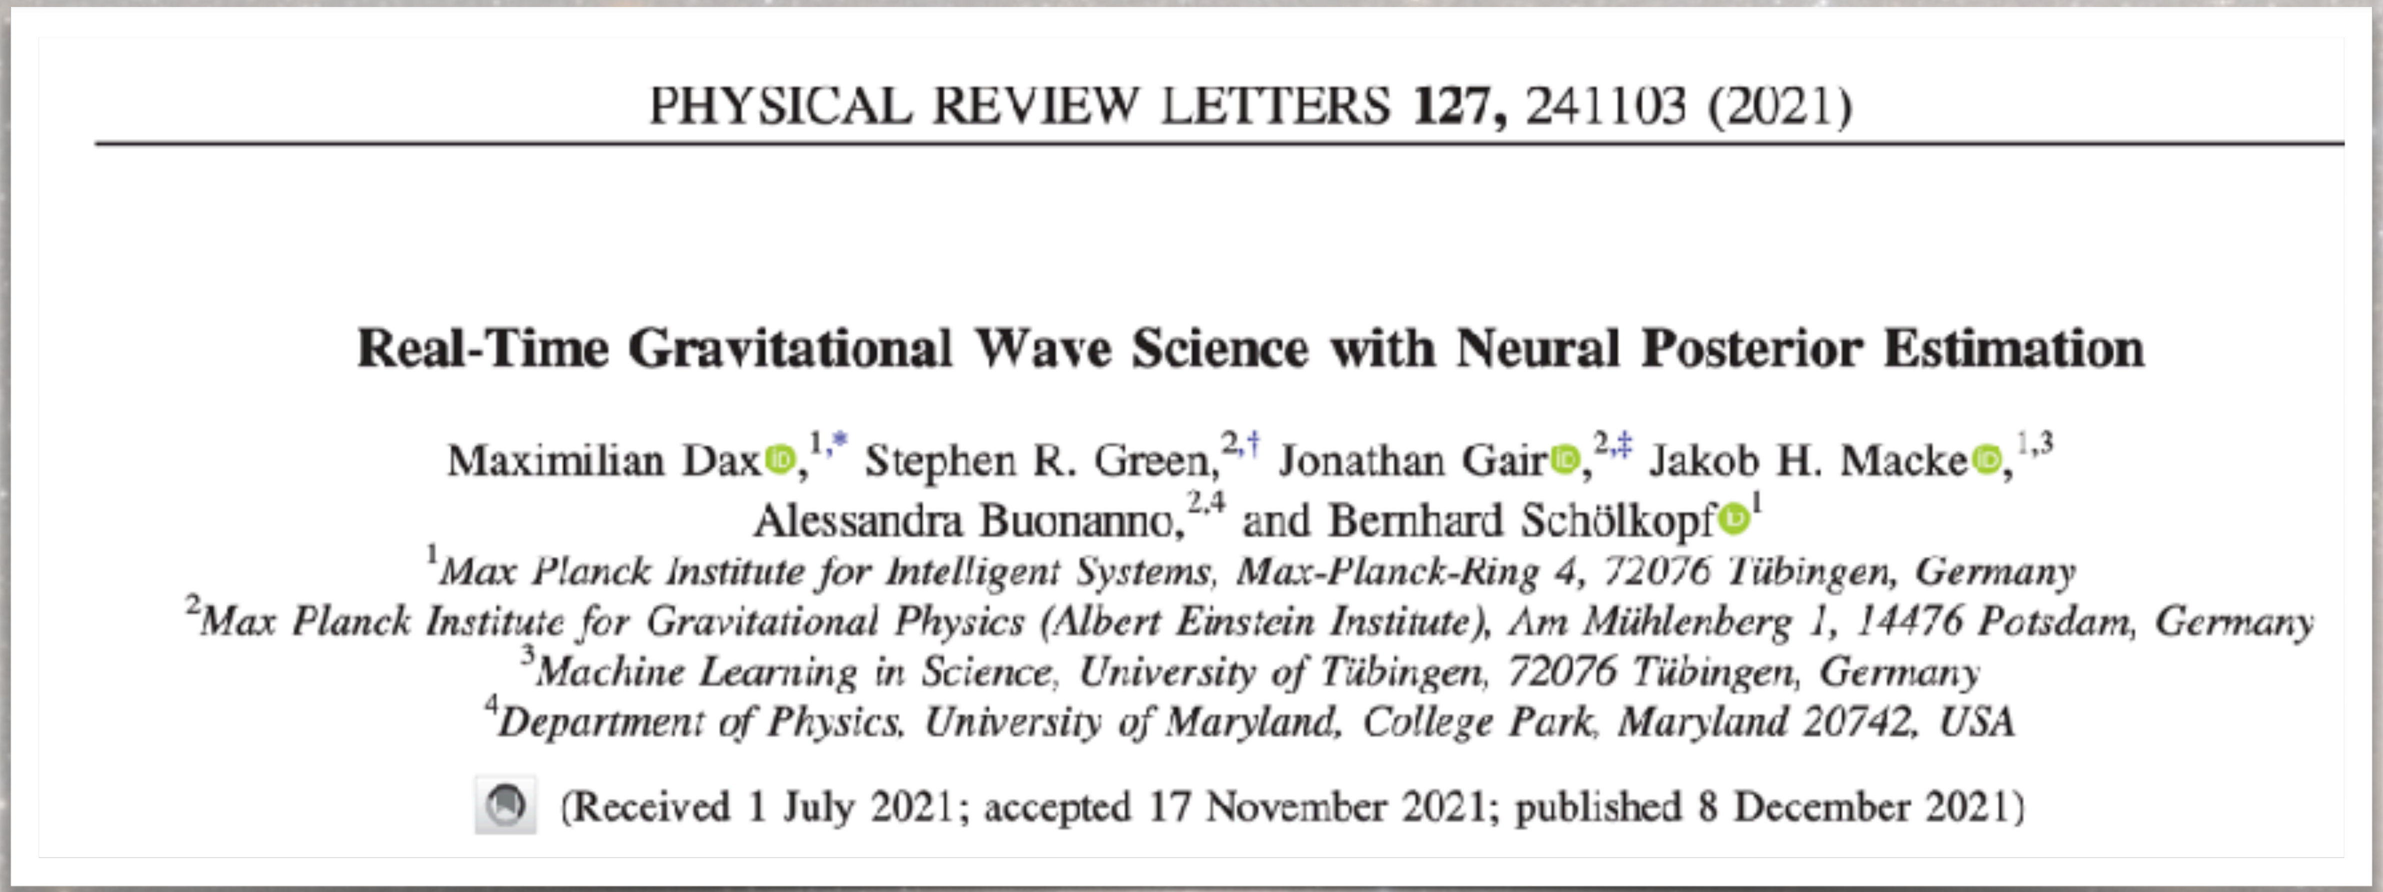

PHYSICAL REVIEW LETTERS **127,** 241103 (2021)

Real-Time Gravitational Wave Science with Neural Posterior Estimation

Maximilian Dax,[1,*] Stephen R. Green,[2,†] Jonathan Gair,[2,‡] Jakob H. Macke,[1,3]
Alessandra Buonanno,[2,4] and Bernhard Schölkopf[1]
[1]*Max Planck Institute for Intelligent Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany*
[2]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, 14476 Potsdam, Germany*
[3]*Machine Learning in Science, University of Tübingen, 72076 Tübingen, Germany*
[4]*Department of Physics, University of Maryland, College Park, Maryland 20742, USA*

(Received 1 July 2021; accepted 17 November 2021; published 8 December 2021)

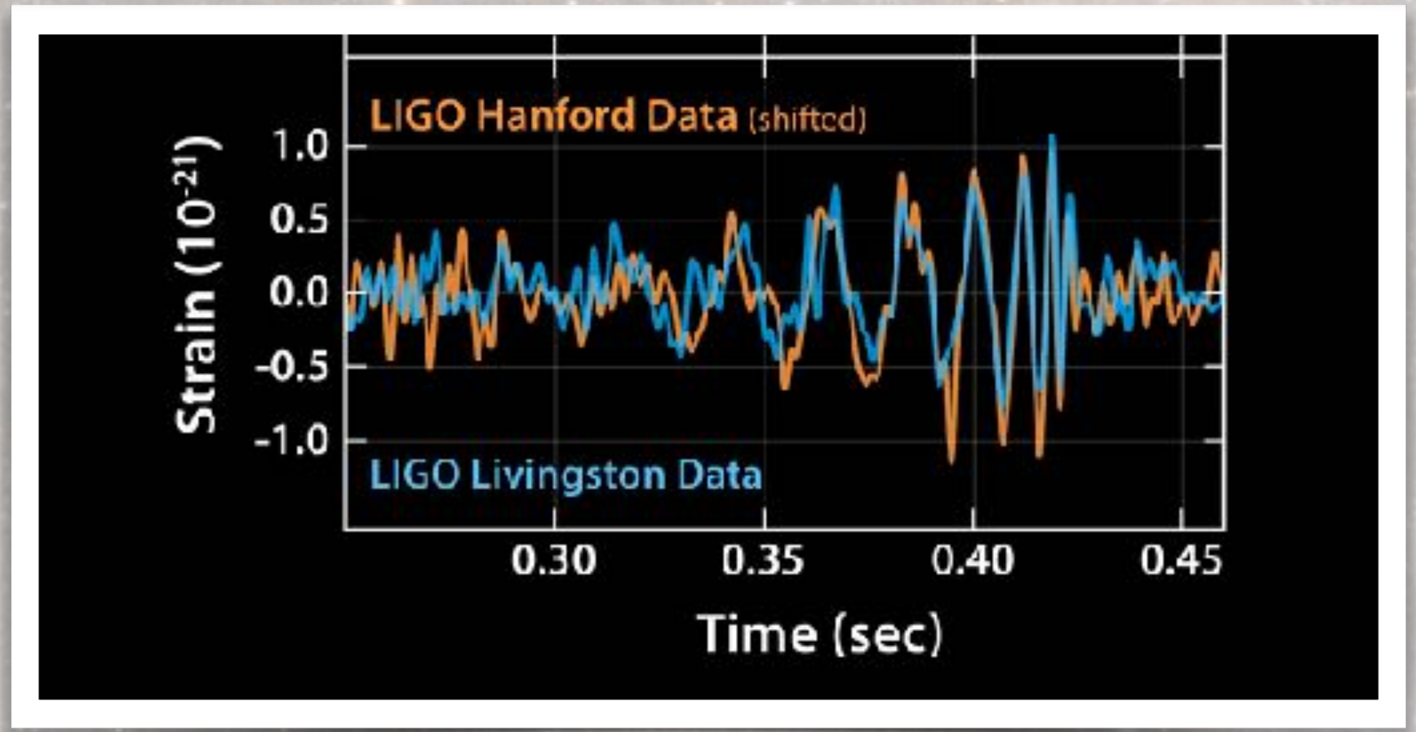

infer

Location?
Masses?
Spins?
…

Max Dax

Stephen Green

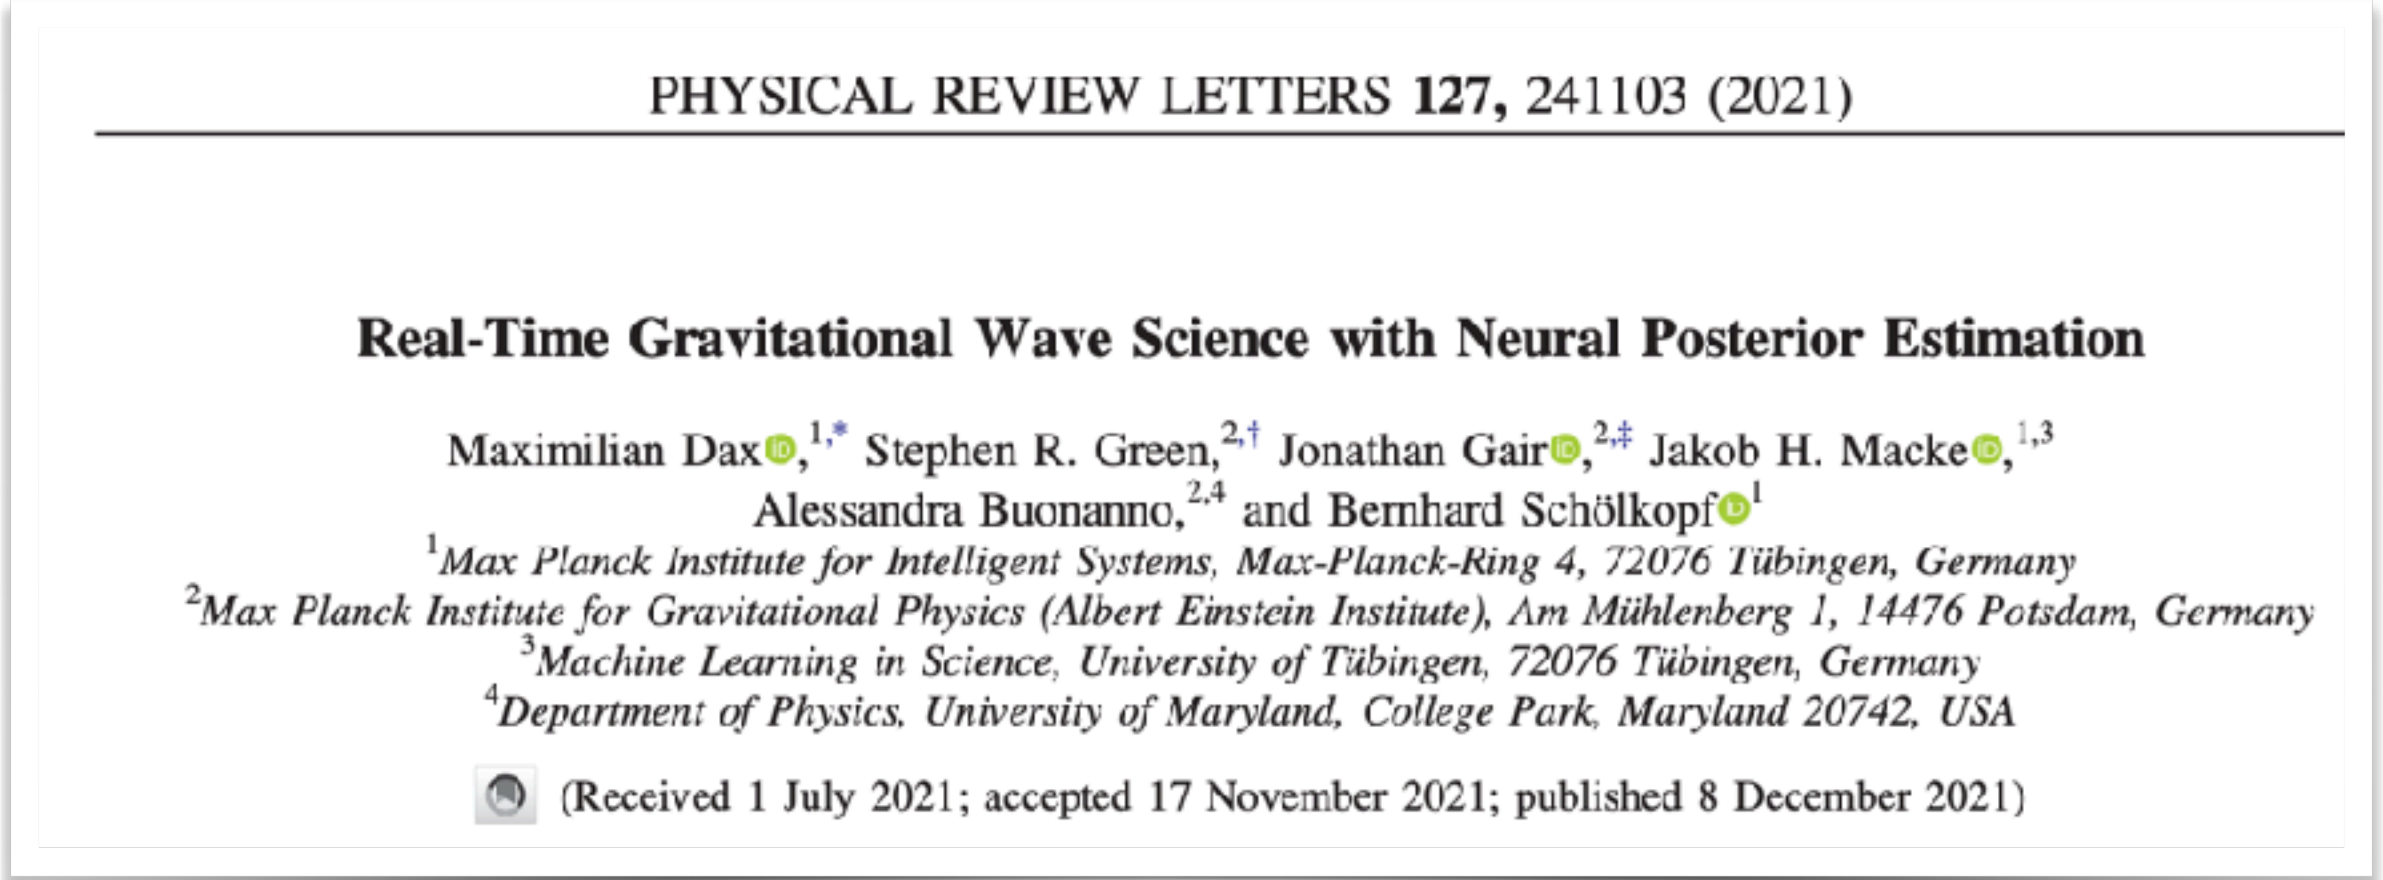

PHYSICAL REVIEW LETTERS **127,** 241103 (2021)

Real-Time Gravitational Wave Science with Neural Posterior Estimation

Maximilian Dax,[1,*] Stephen R. Green,[2,†] Jonathan Gair,[2,‡] Jakob H. Macke,[1,3]
Alessandra Buonanno,[2,4] and Bernhard Schölkopf[1]
[1]*Max Planck Institute for Intelligent Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany*
[2]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, 14476 Potsdam, Germany*
[3]*Machine Learning in Science, University of Tübingen, 72076 Tübingen, Germany*
[4]*Department of Physics, University of Maryland, College Park, Maryland 20742, USA*

(Received 1 July 2021; accepted 17 November 2021; published 8 December 2021)

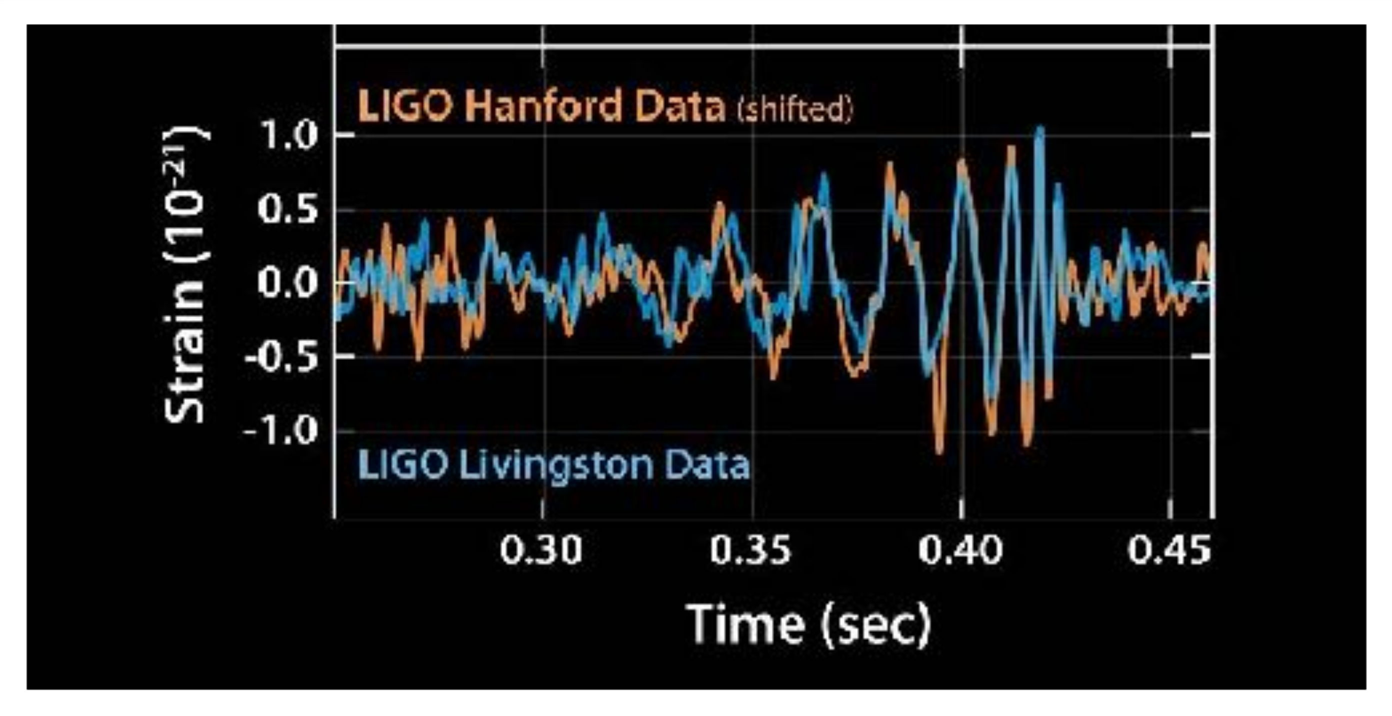

infer

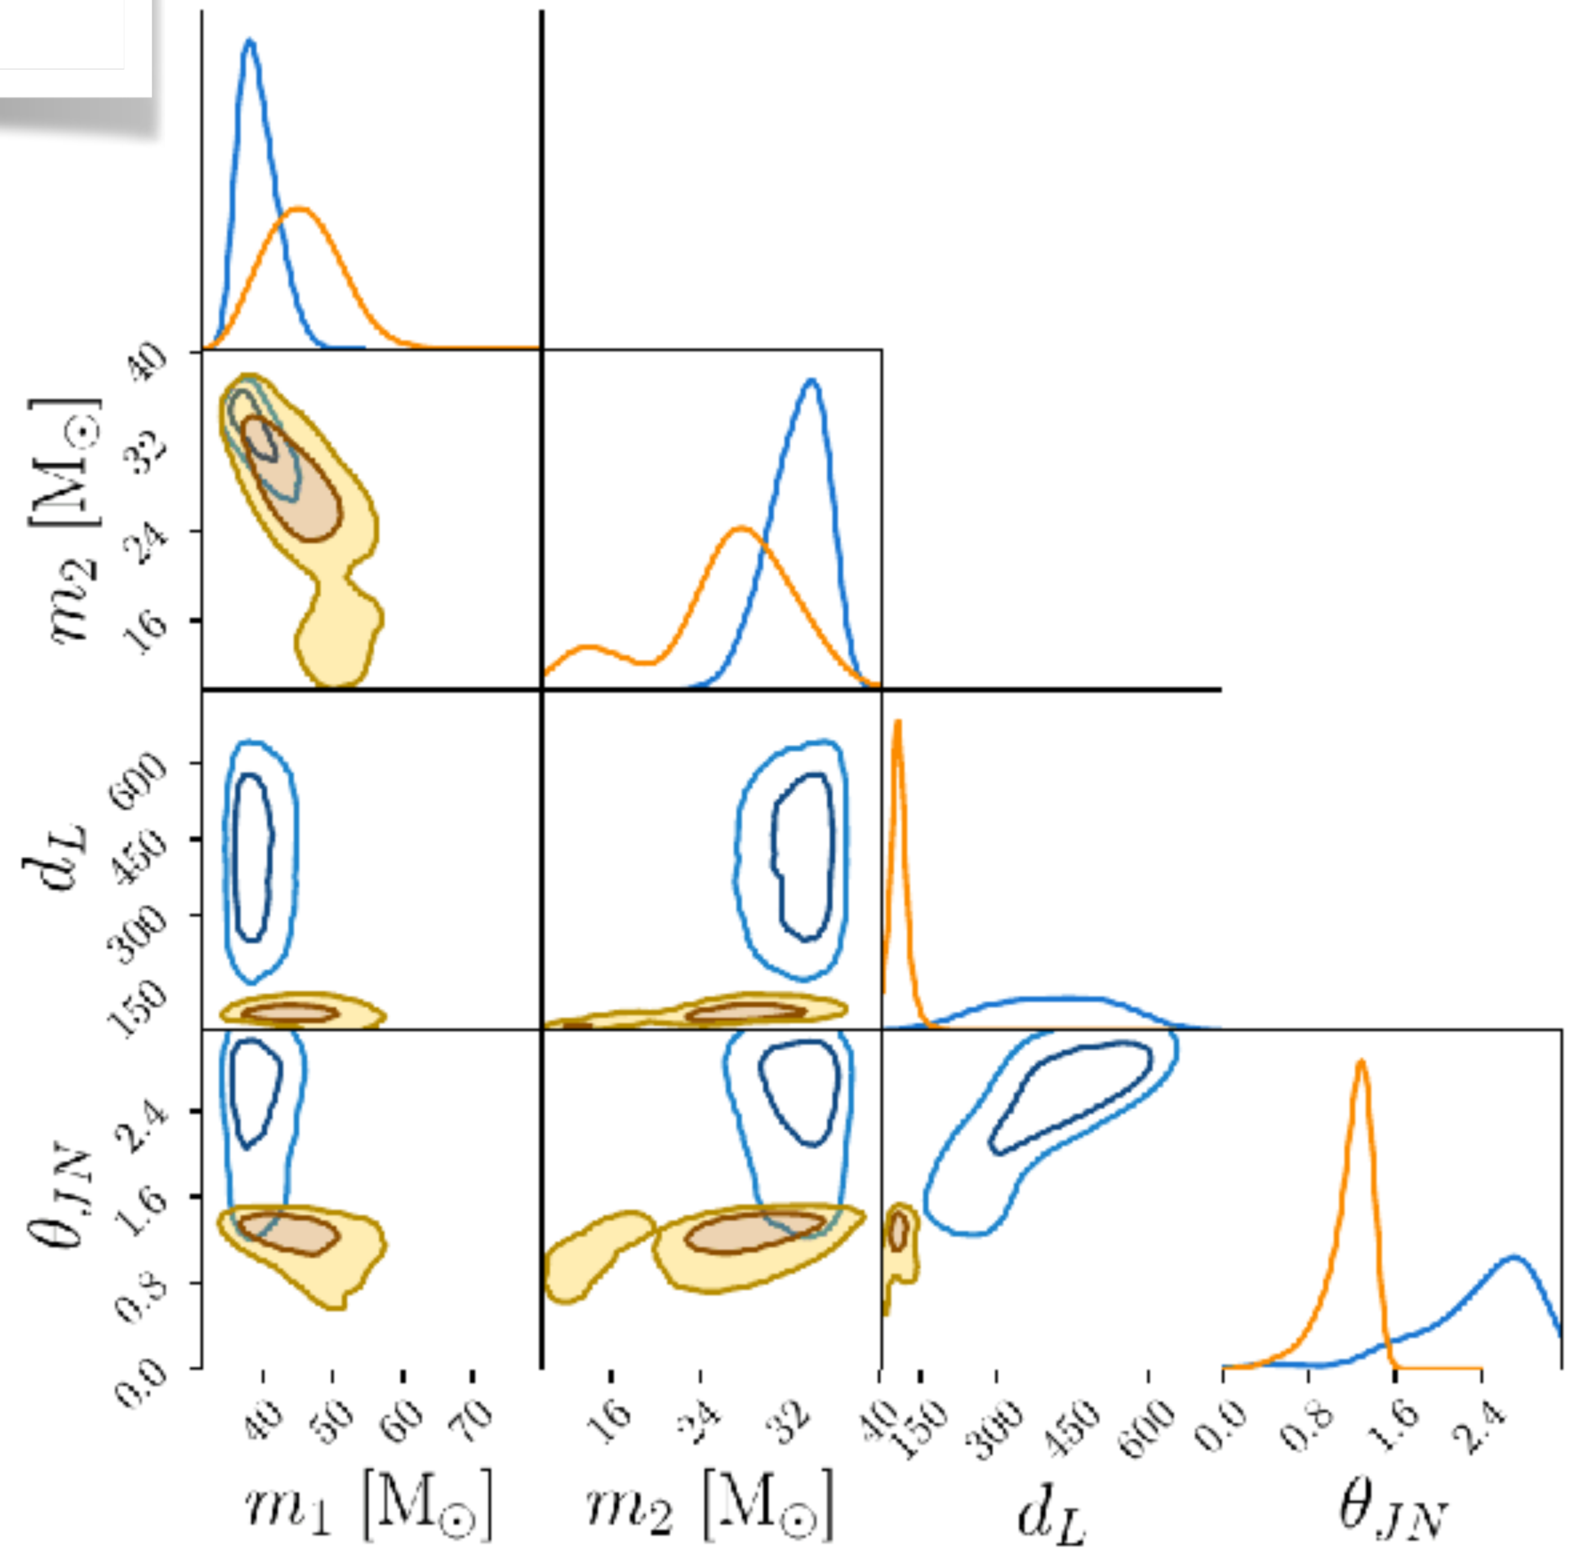

Standard method: **O(days to hours)**
Our SBI method(DINGO): **O(seconds)**

Max Dax

Stephen Green

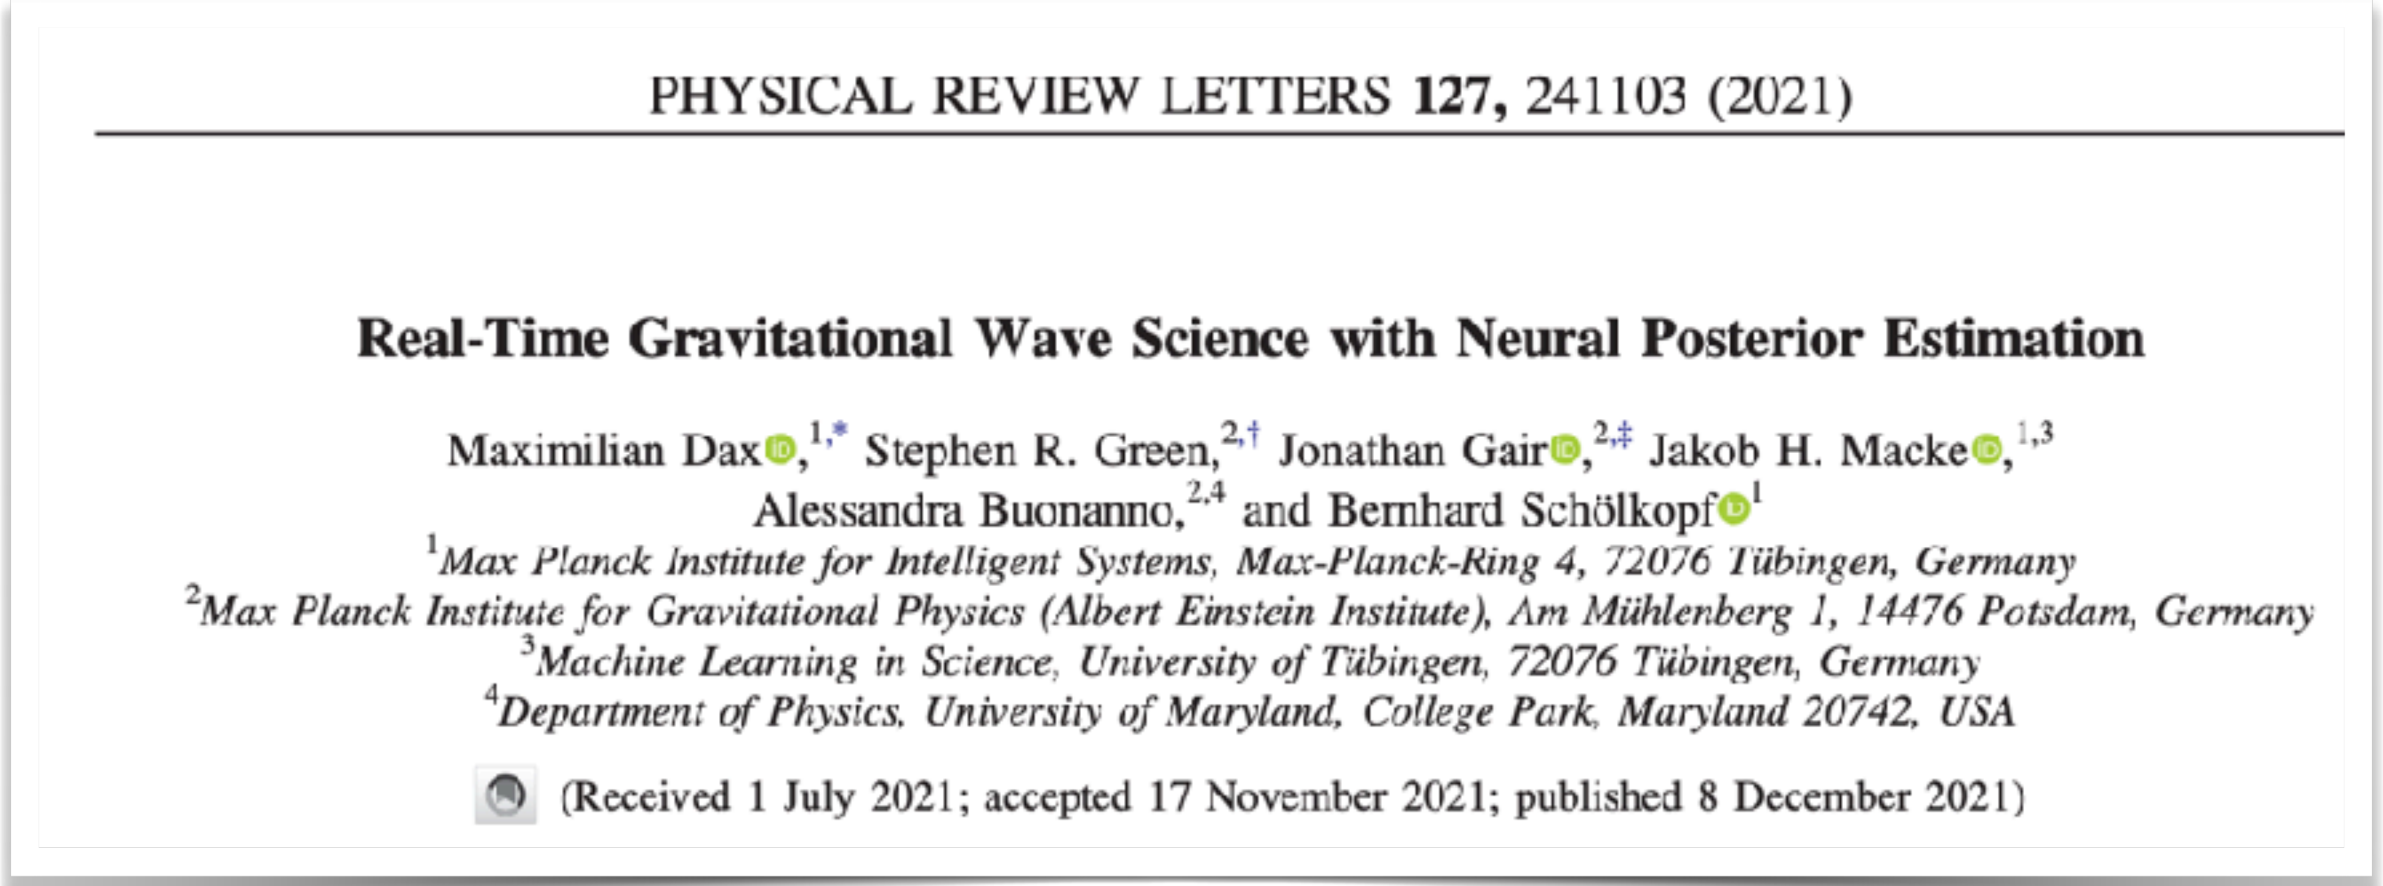

PHYSICAL REVIEW LETTERS **127,** 241103 (2021)

Real-Time Gravitational Wave Science with Neural Posterior Estimation

Maximilian Dax,[1,*] Stephen R. Green,[2,†] Jonathan Gair,[2,‡] Jakob H. Macke,[1,3]
Alessandra Buonanno,[2,4] and Bernhard Schölkopf[1]
[1]*Max Planck Institute for Intelligent Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany*
[2]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, 14476 Potsdam, Germany*
[3]*Machine Learning in Science, University of Tübingen, 72076 Tübingen, Germany*
[4]*Department of Physics, University of Maryland, College Park, Maryland 20742, USA*

(Received 1 July 2021; accepted 17 November 2021; published 8 December 2021)

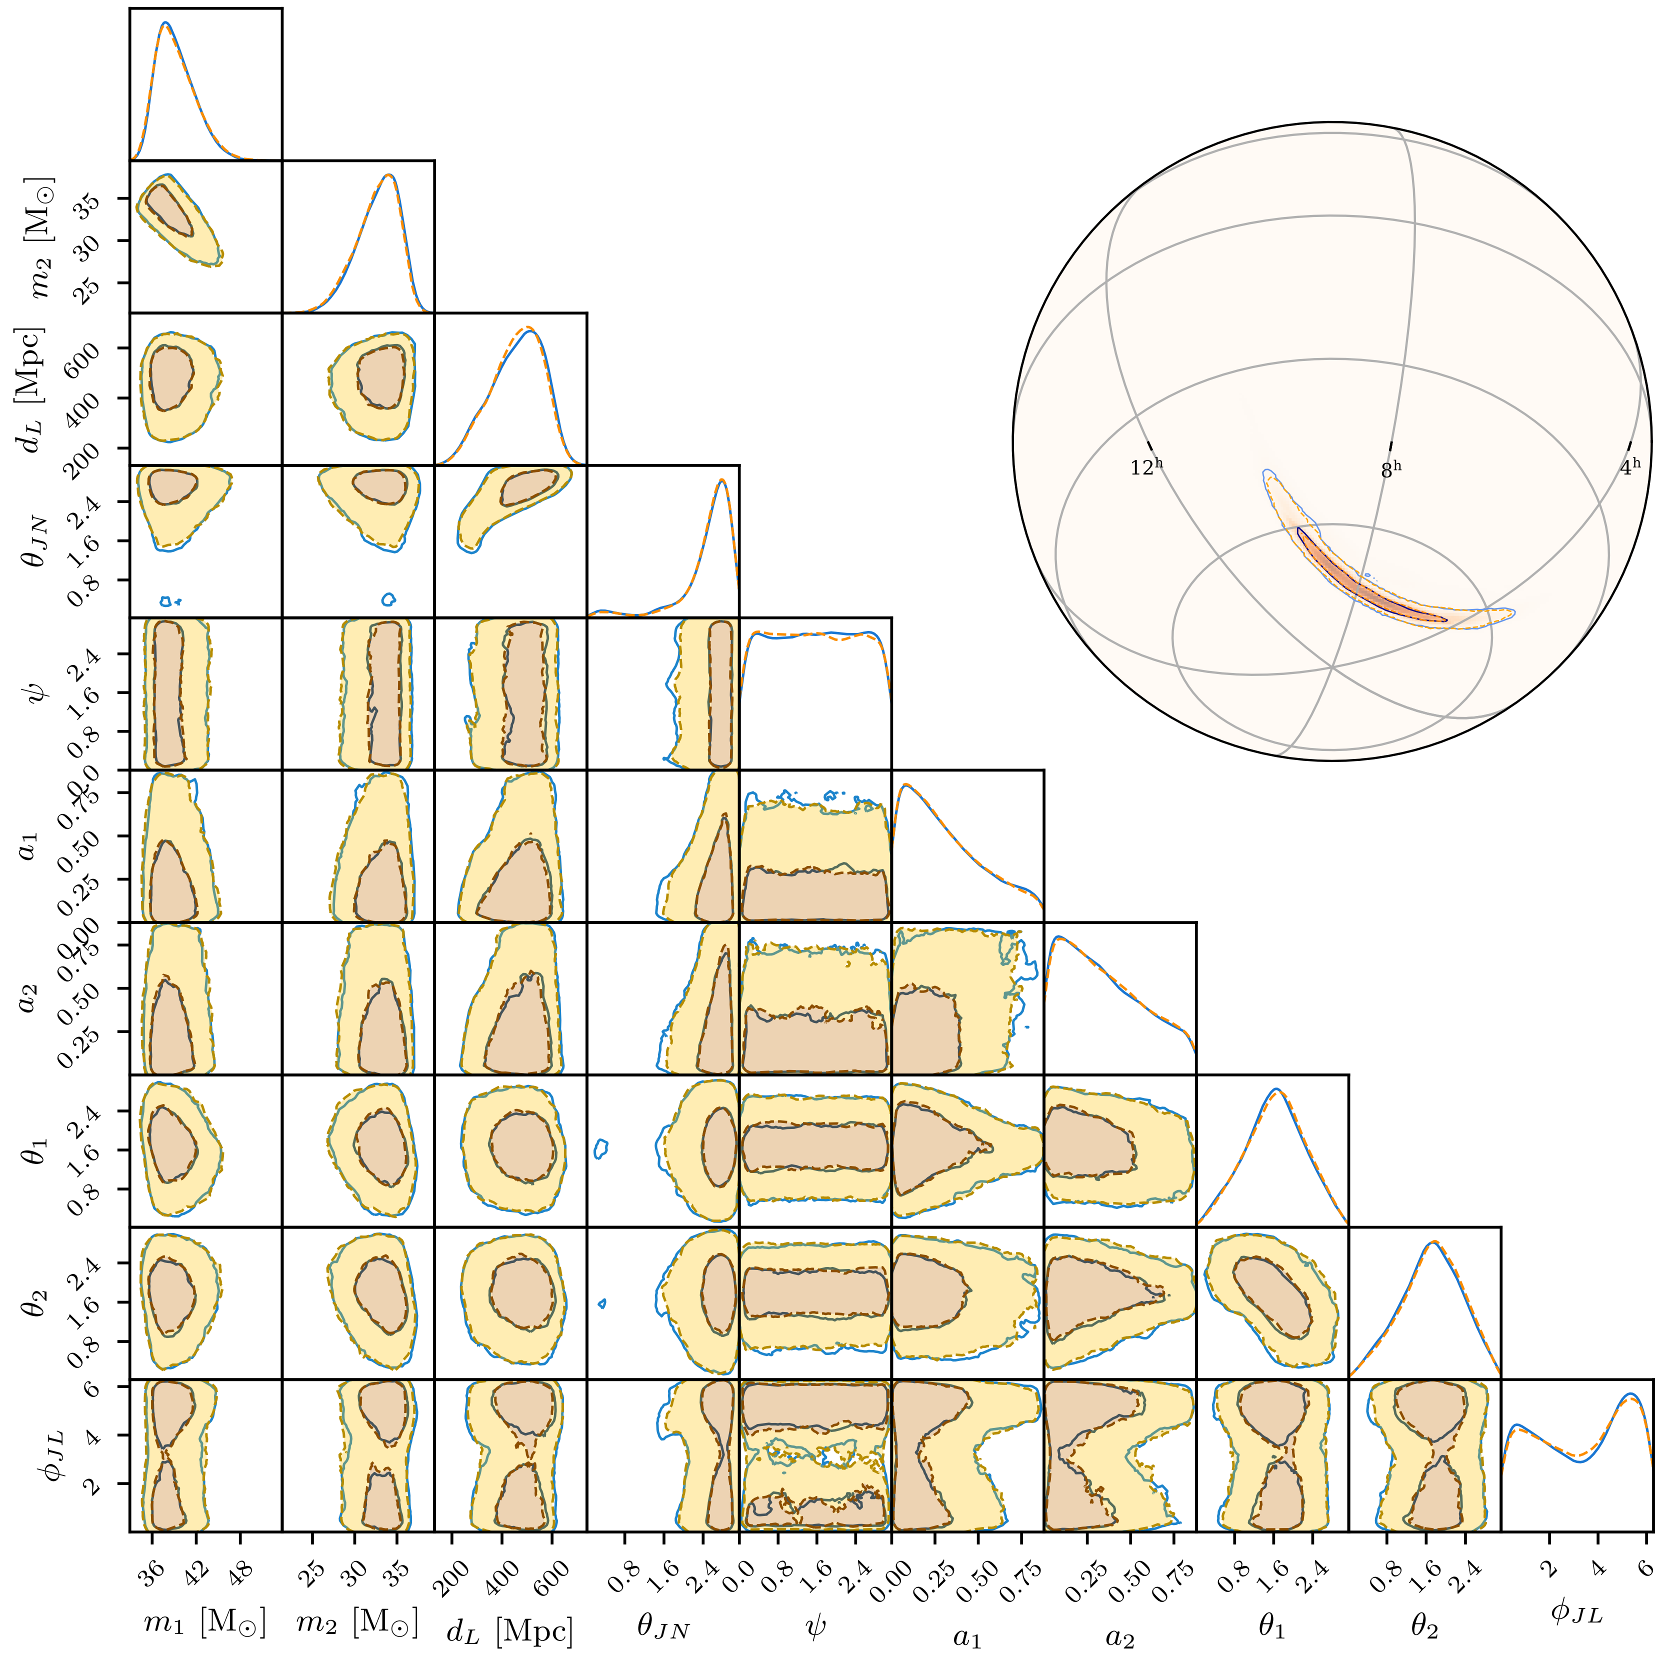

GW150914

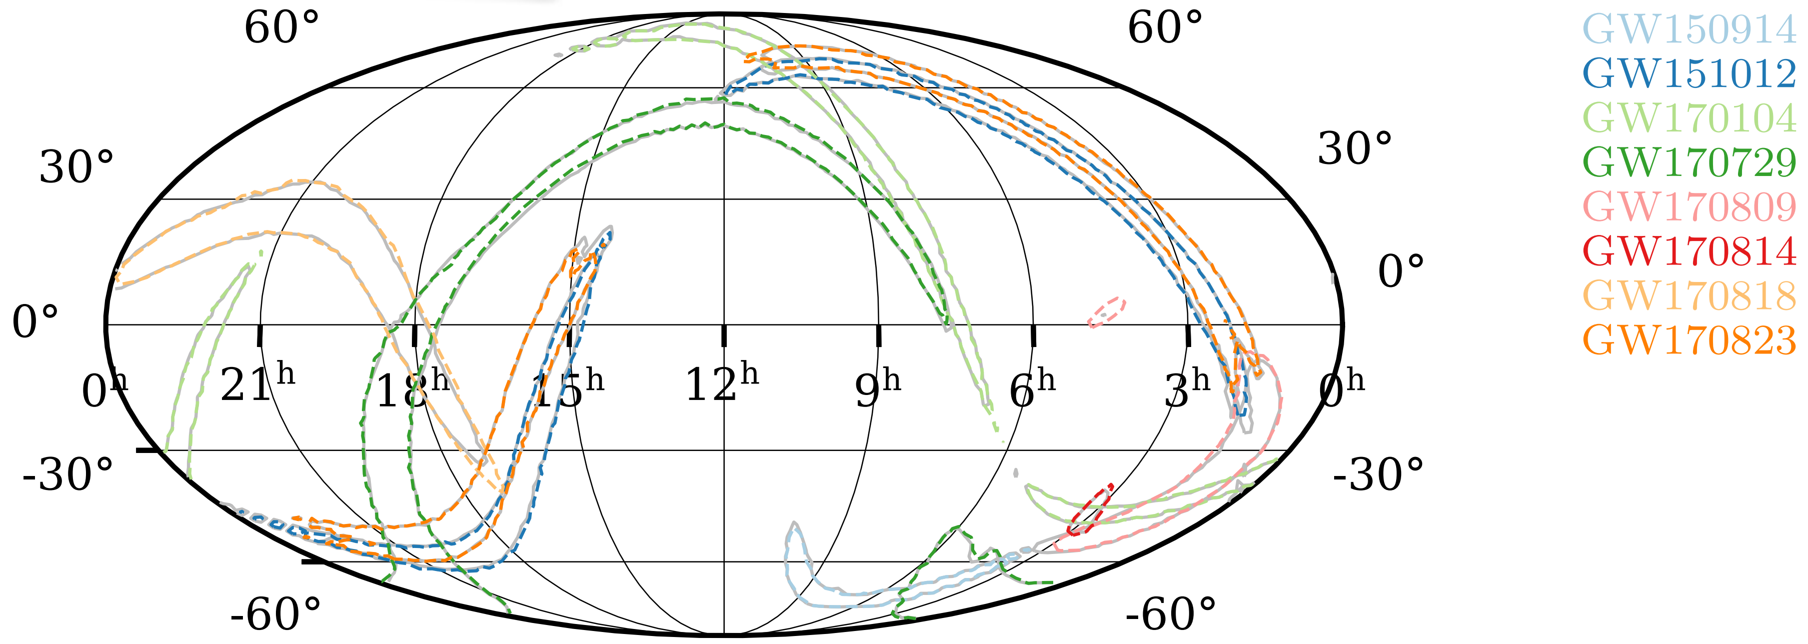

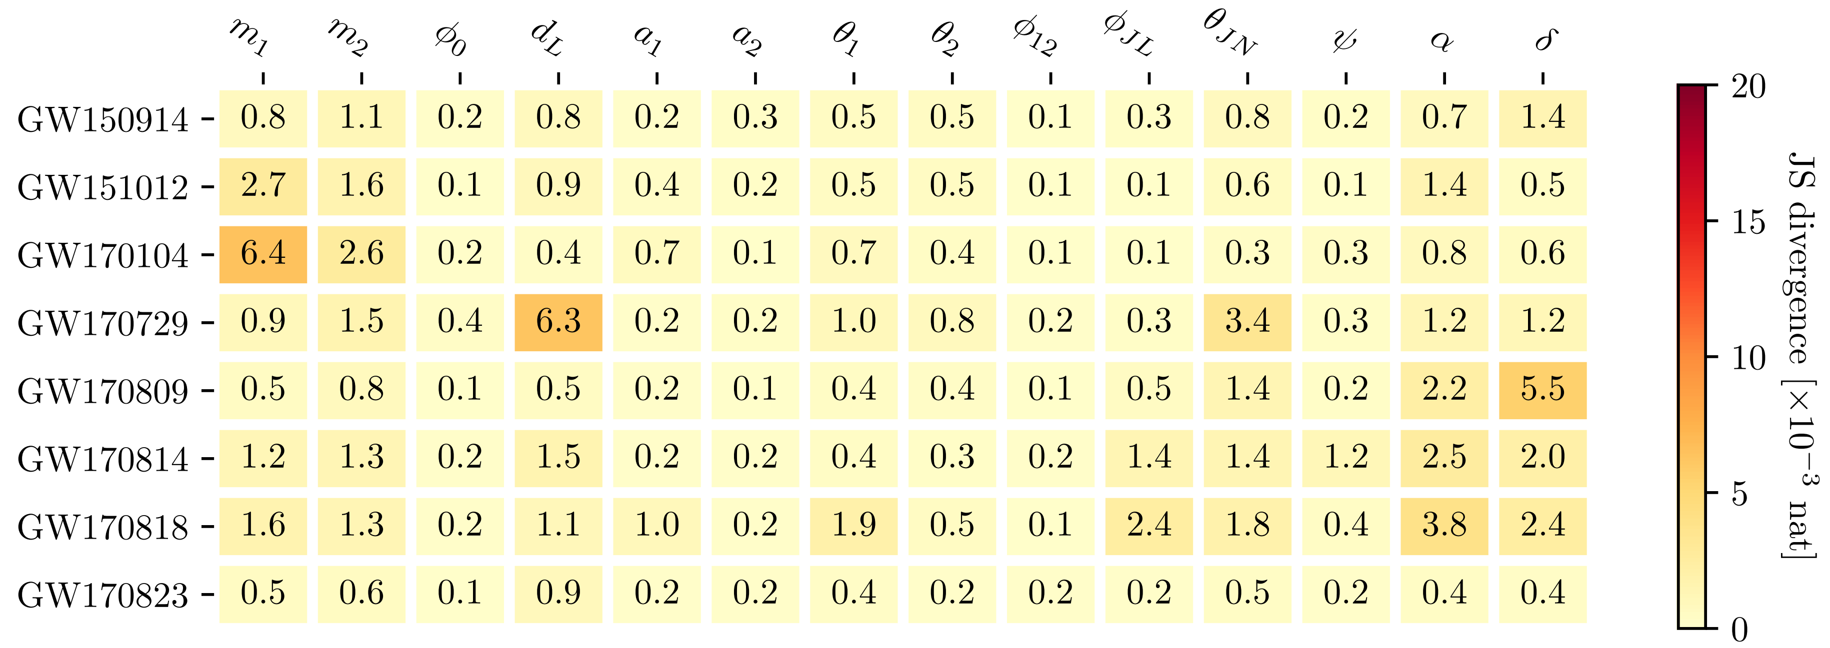

	m_1	m_2	ϕ_0	d_L	a_1	a_2	θ_1	θ_2	ϕ_{12}	ϕ_{JL}	θ_{JN}	ψ	α	δ
GW150914	0.8	1.1	0.2	0.8	0.2	0.3	0.5	0.5	0.1	0.3	0.8	0.2	0.7	1.4
GW151012	2.7	1.6	0.1	0.9	0.4	0.2	0.5	0.5	0.1	0.1	0.6	0.1	1.4	0.5
GW170104	6.4	2.6	0.2	0.4	0.7	0.1	0.7	0.4	0.1	0.1	0.3	0.3	0.8	0.6
GW170729	0.9	1.5	0.4	6.3	0.2	0.2	1.0	0.8	0.2	0.3	3.4	0.3	1.2	1.2
GW170809	0.5	0.8	0.1	0.5	0.2	0.1	0.4	0.4	0.1	0.5	1.4	0.2	2.2	5.5
GW170814	1.2	1.3	0.2	1.5	0.2	0.2	0.4	0.3	0.2	1.4	1.4	1.2	2.5	2.0
GW170818	1.6	1.3	0.2	1.1	1.0	0.2	1.9	0.5	0.1	2.4	1.8	0.4	3.8	2.4
GW170823	0.5	0.6	0.1	0.9	0.2	0.2	0.4	0.2	0.2	0.2	0.5	0.2	0.4	0.4

Jensen-Shannon divergence between Dingo and MCMC.
Average: 0.0009 nat (MCMC-MCMC: 0.0007 nat)

Exploiting problem structure: Symmetries + Noise

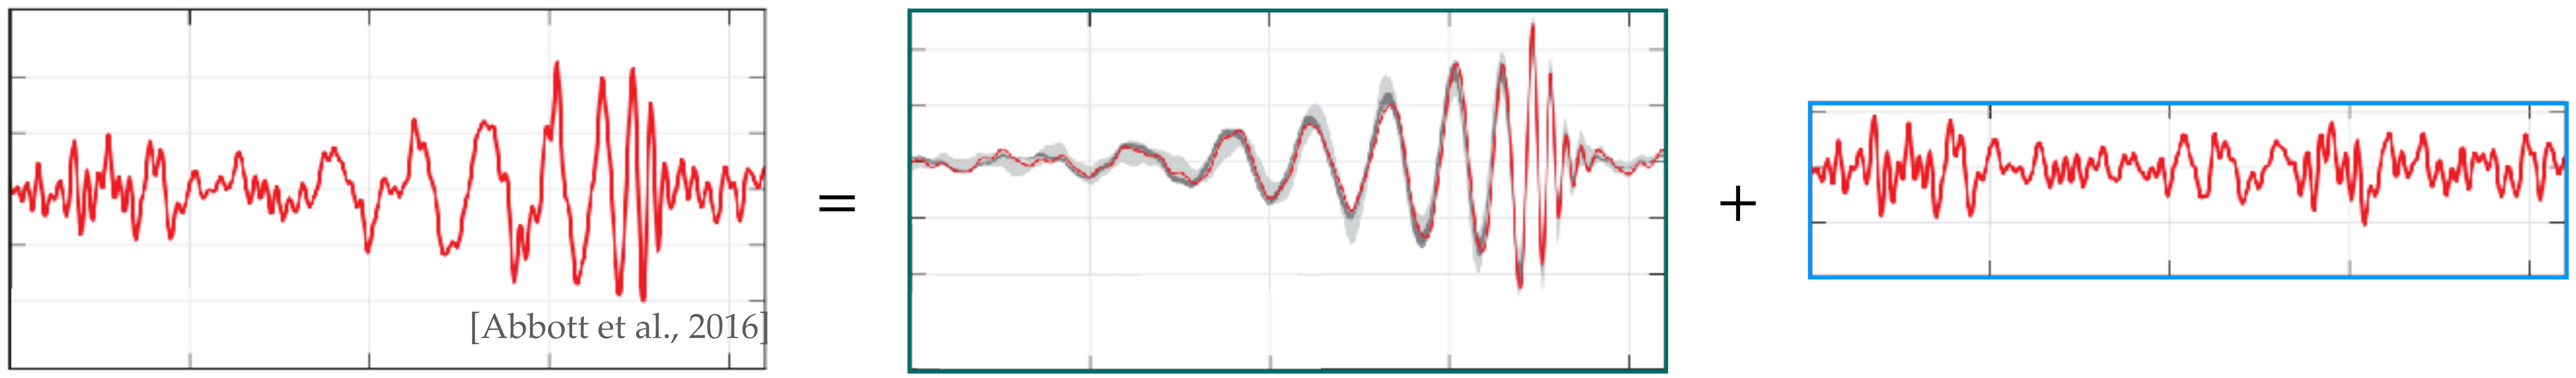

$$d \sim p(d\,|\,\theta) = h(\theta) + n \sim \mathcal{N}(0,S_n)$$

[Abbott et al., 2016]

deterministic GW signal from general relativity model

Noise: stationary Gaussian with PSD Sn

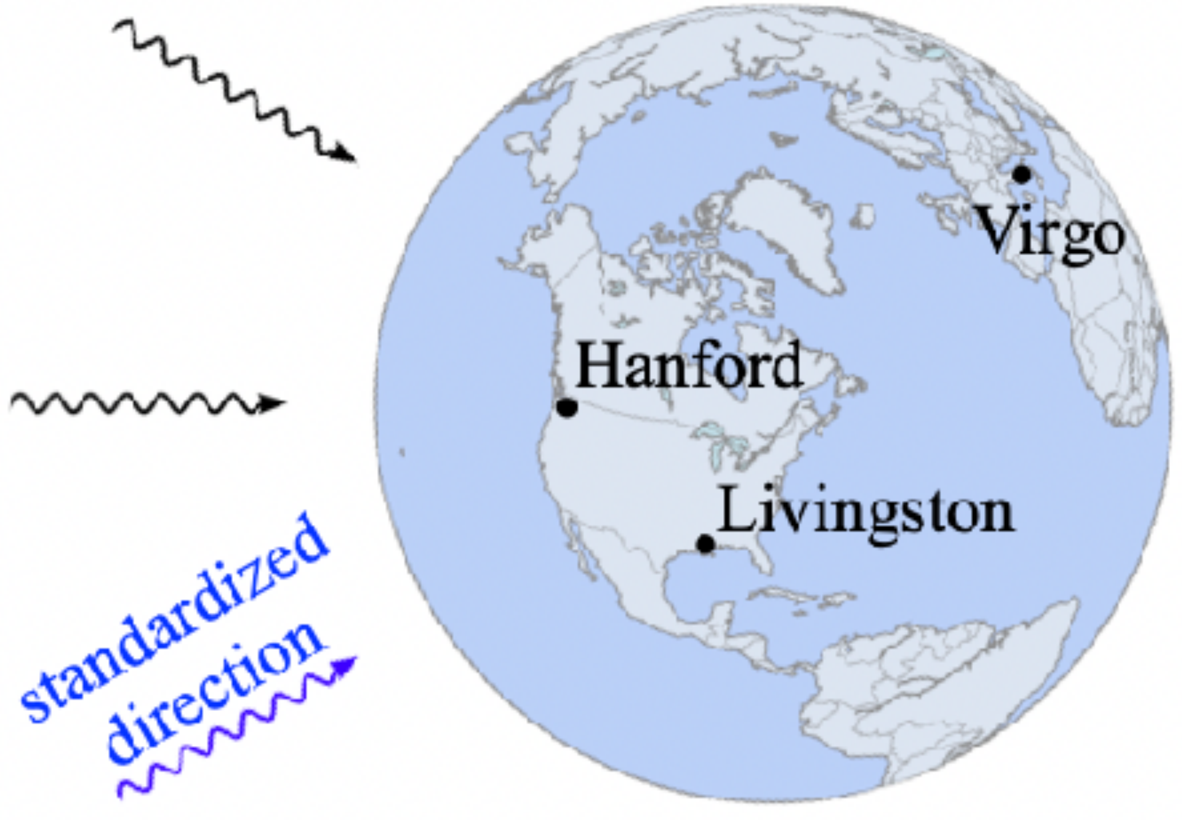

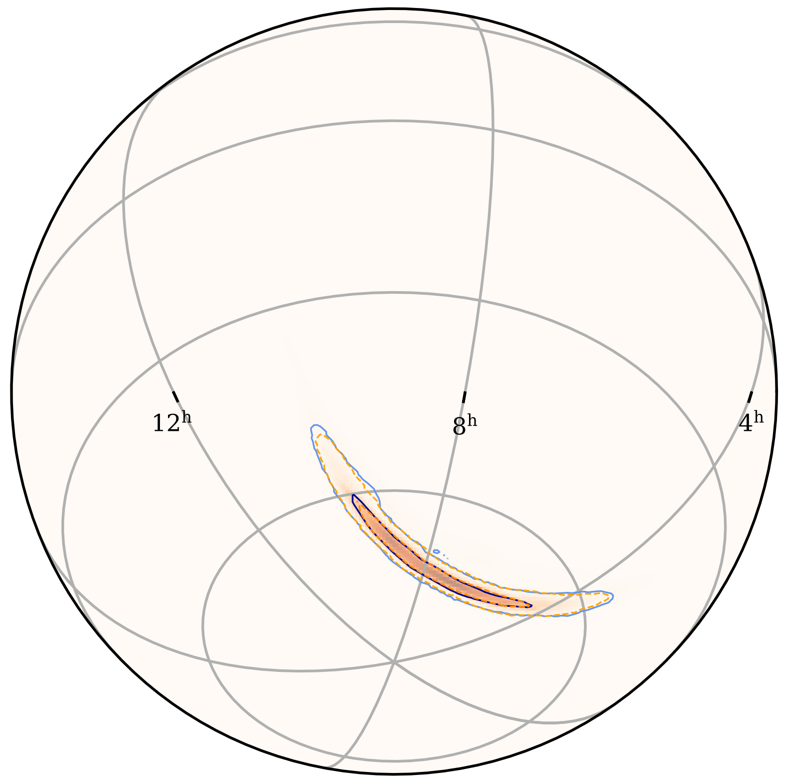

GW simulation:

- **exactly equivariant** under shifts of time $t \in \theta$
- **approx. equivariant** under rotations $(\alpha, \delta) \in \theta$

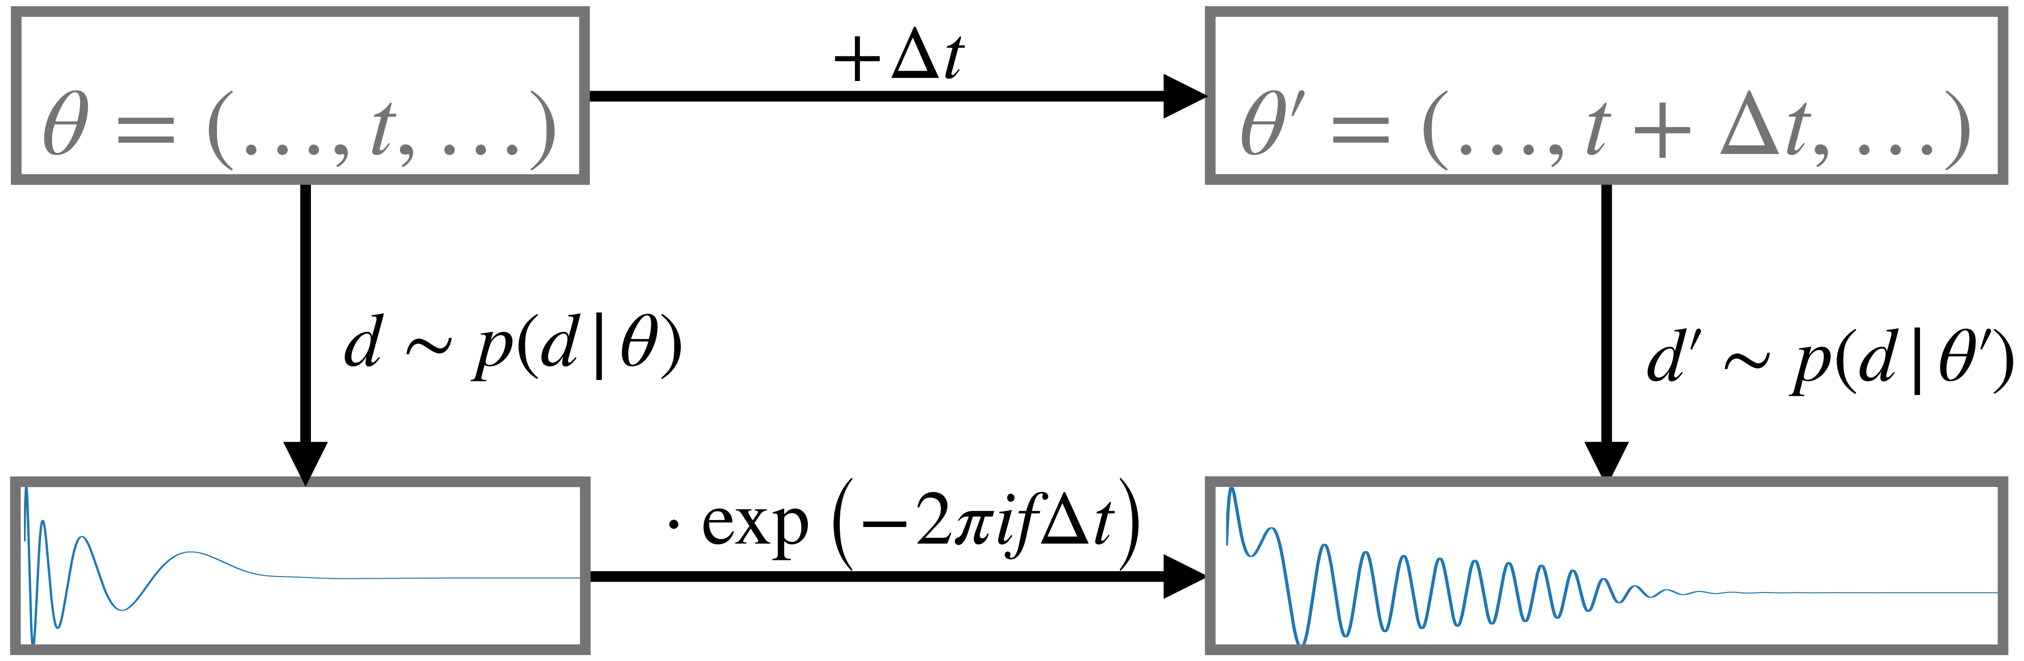

$$\begin{tikzcd} \theta = (\ldots, t, \ldots) \arrow[r, "+\Delta t"] \arrow[d, "d \sim p(d\,|\,\theta)"] & \theta' = (\ldots, t + \Delta t, \ldots) \arrow[d, "d' \sim p(d\,|\,\theta')"] \\ {} \arrow[r, "\cdot \exp\left(-2\pi i f \Delta t\right)"] & {} \end{tikzcd}$$

Exploiting problem structure: Symmetries + Noise

$$d \sim p(d\,|\,\theta) = h(\theta) + n \sim \mathcal{N}(0,S_n)$$

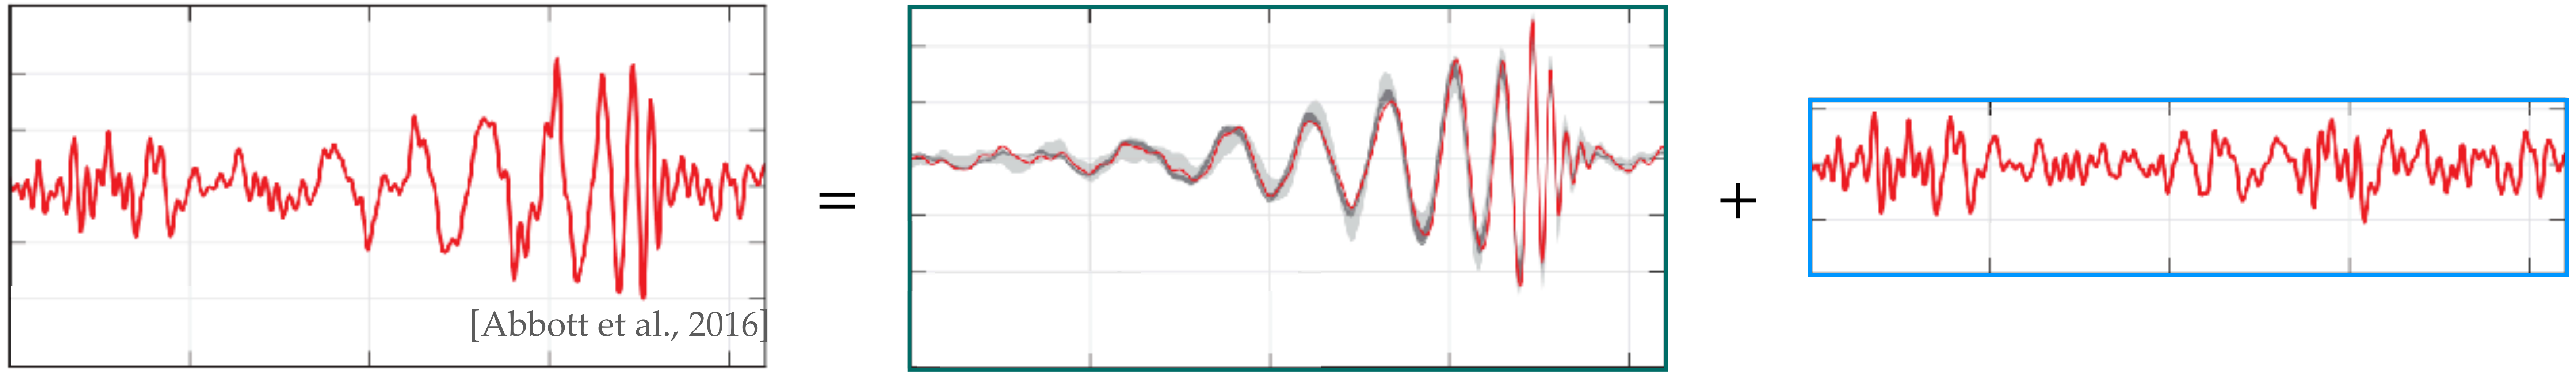
[Abbott et al., 2016]

deterministic GW signal from general relativity model

Noise: stationary Gaussian with PSD Sn

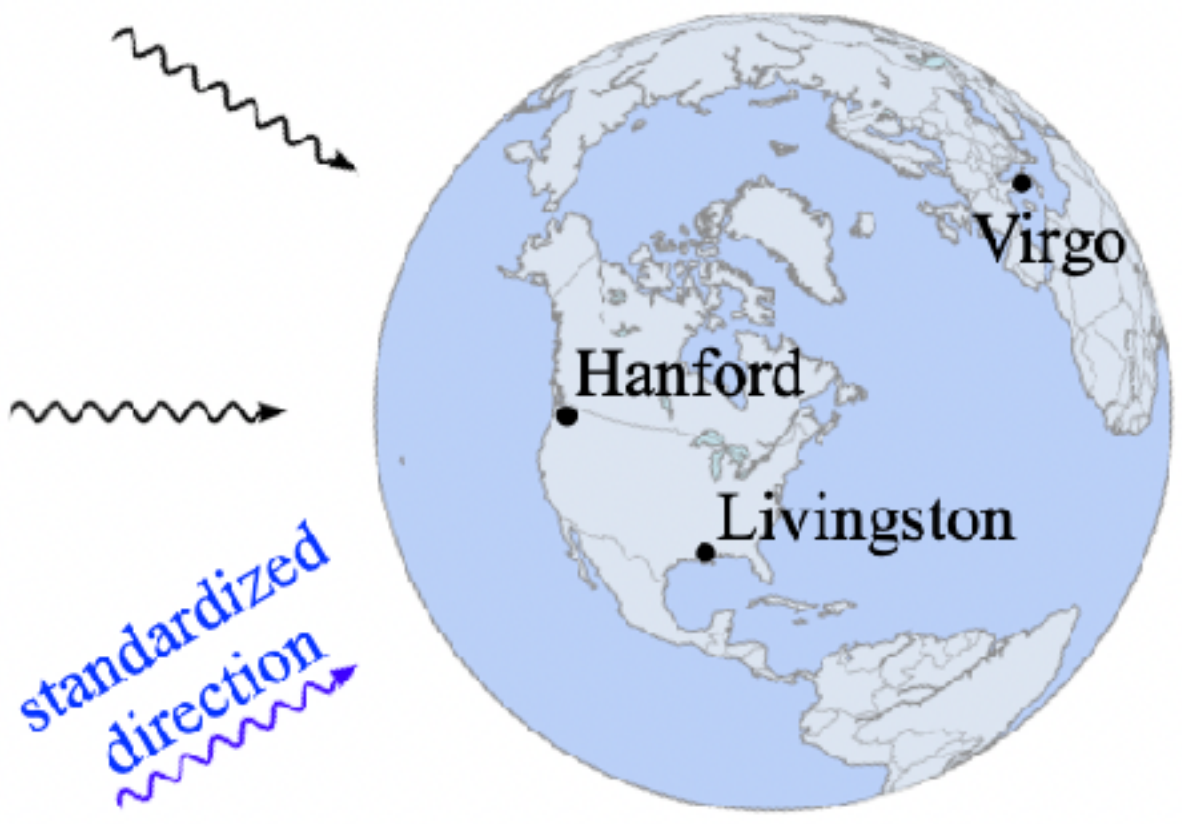

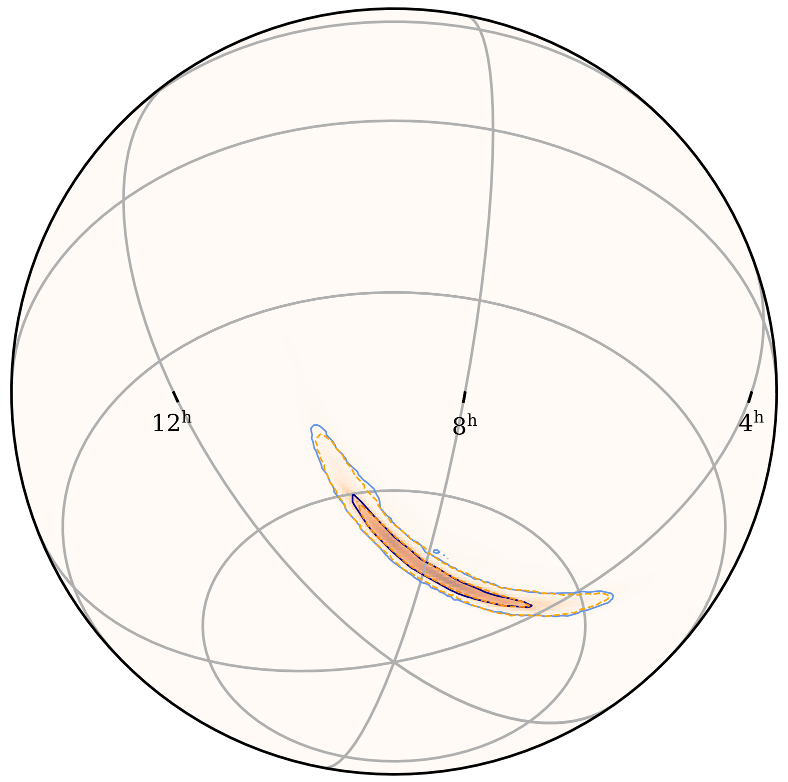

- Chicken-and-egg problem:
    - Once we align on standardized direction, inferring other parameters is easy
    - But dont know direction and time …
- Solution: Gibbs Sampling
- Group-equivariant neural posterior estimation: Dax, Green, Gair, Schölkopf, Macke, ICLR 2022

Exploiting problem structure: Symmetries + Noise

$d \sim p(d\,|\,\theta)$

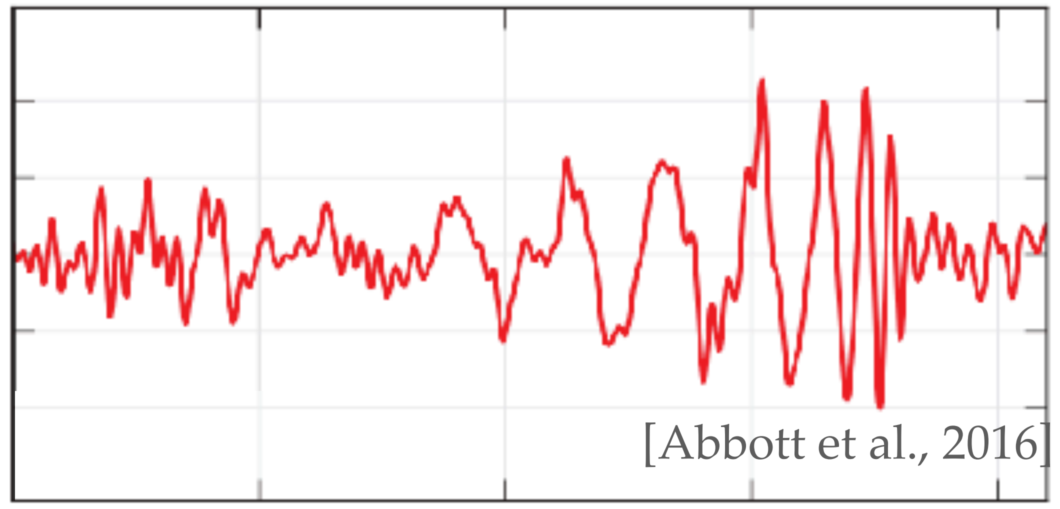

[Abbott et al., 2016]

=

$h(\theta)$

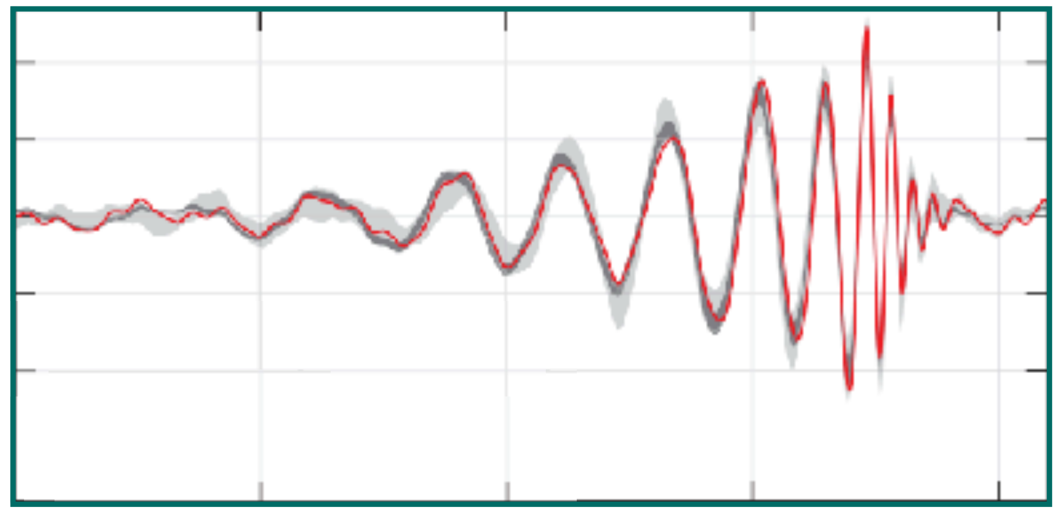

+

$n \sim \mathcal{N}(0, S_n)$

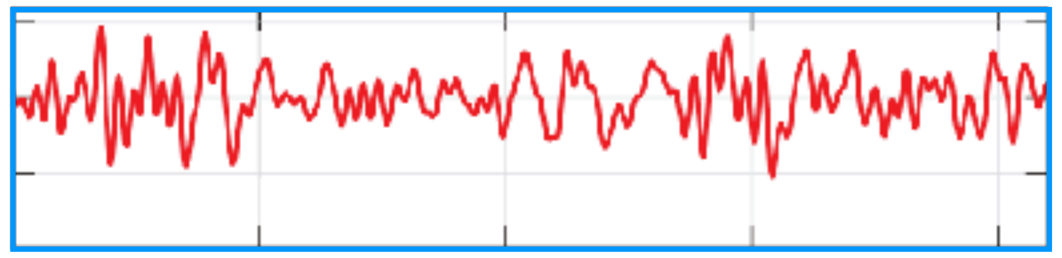

deterministic GW signal from general relativity model

Noise: stationary Gaussian with PSD Sn

- But noise is non-stationary?

  Use probabilistic model of the PSD Sn!

  Wildberger et al PRD 2023
- But noise is Gaussian?

  Use Importance Sampling to fine-tune!!

  Dax et al PRL 2022

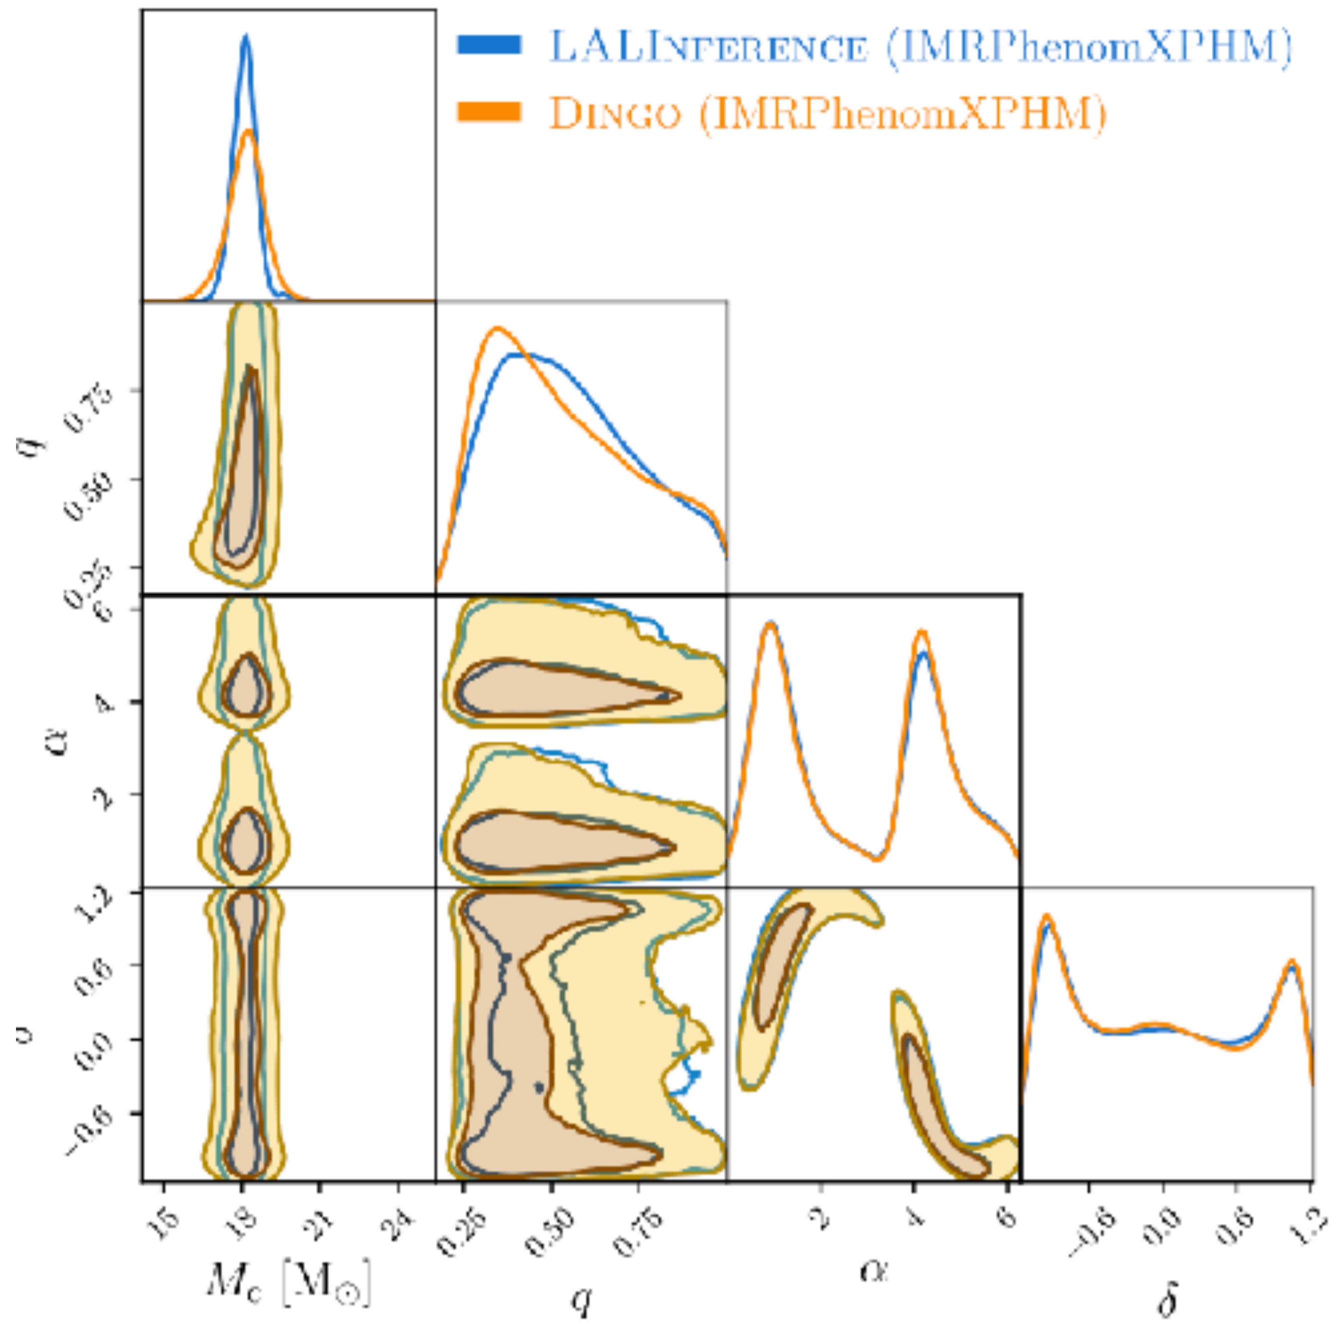

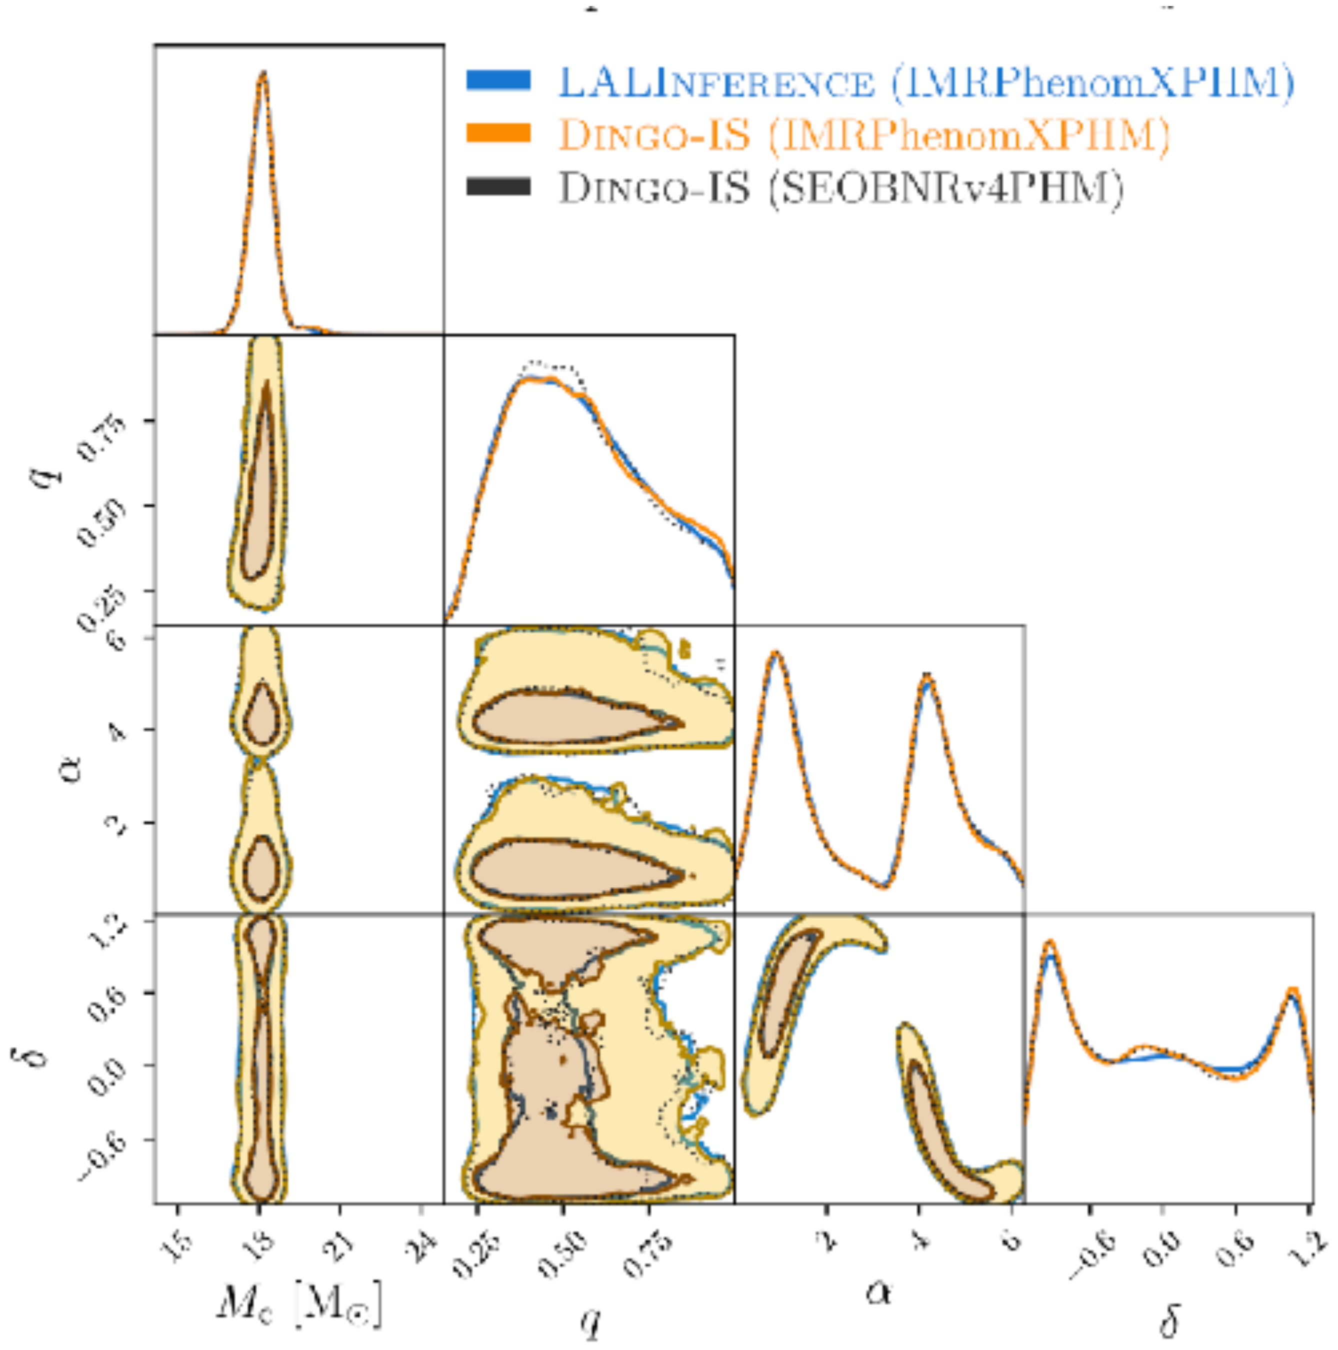

Scaling up: Flow-Matching Neural Posterior Estimation

NPE

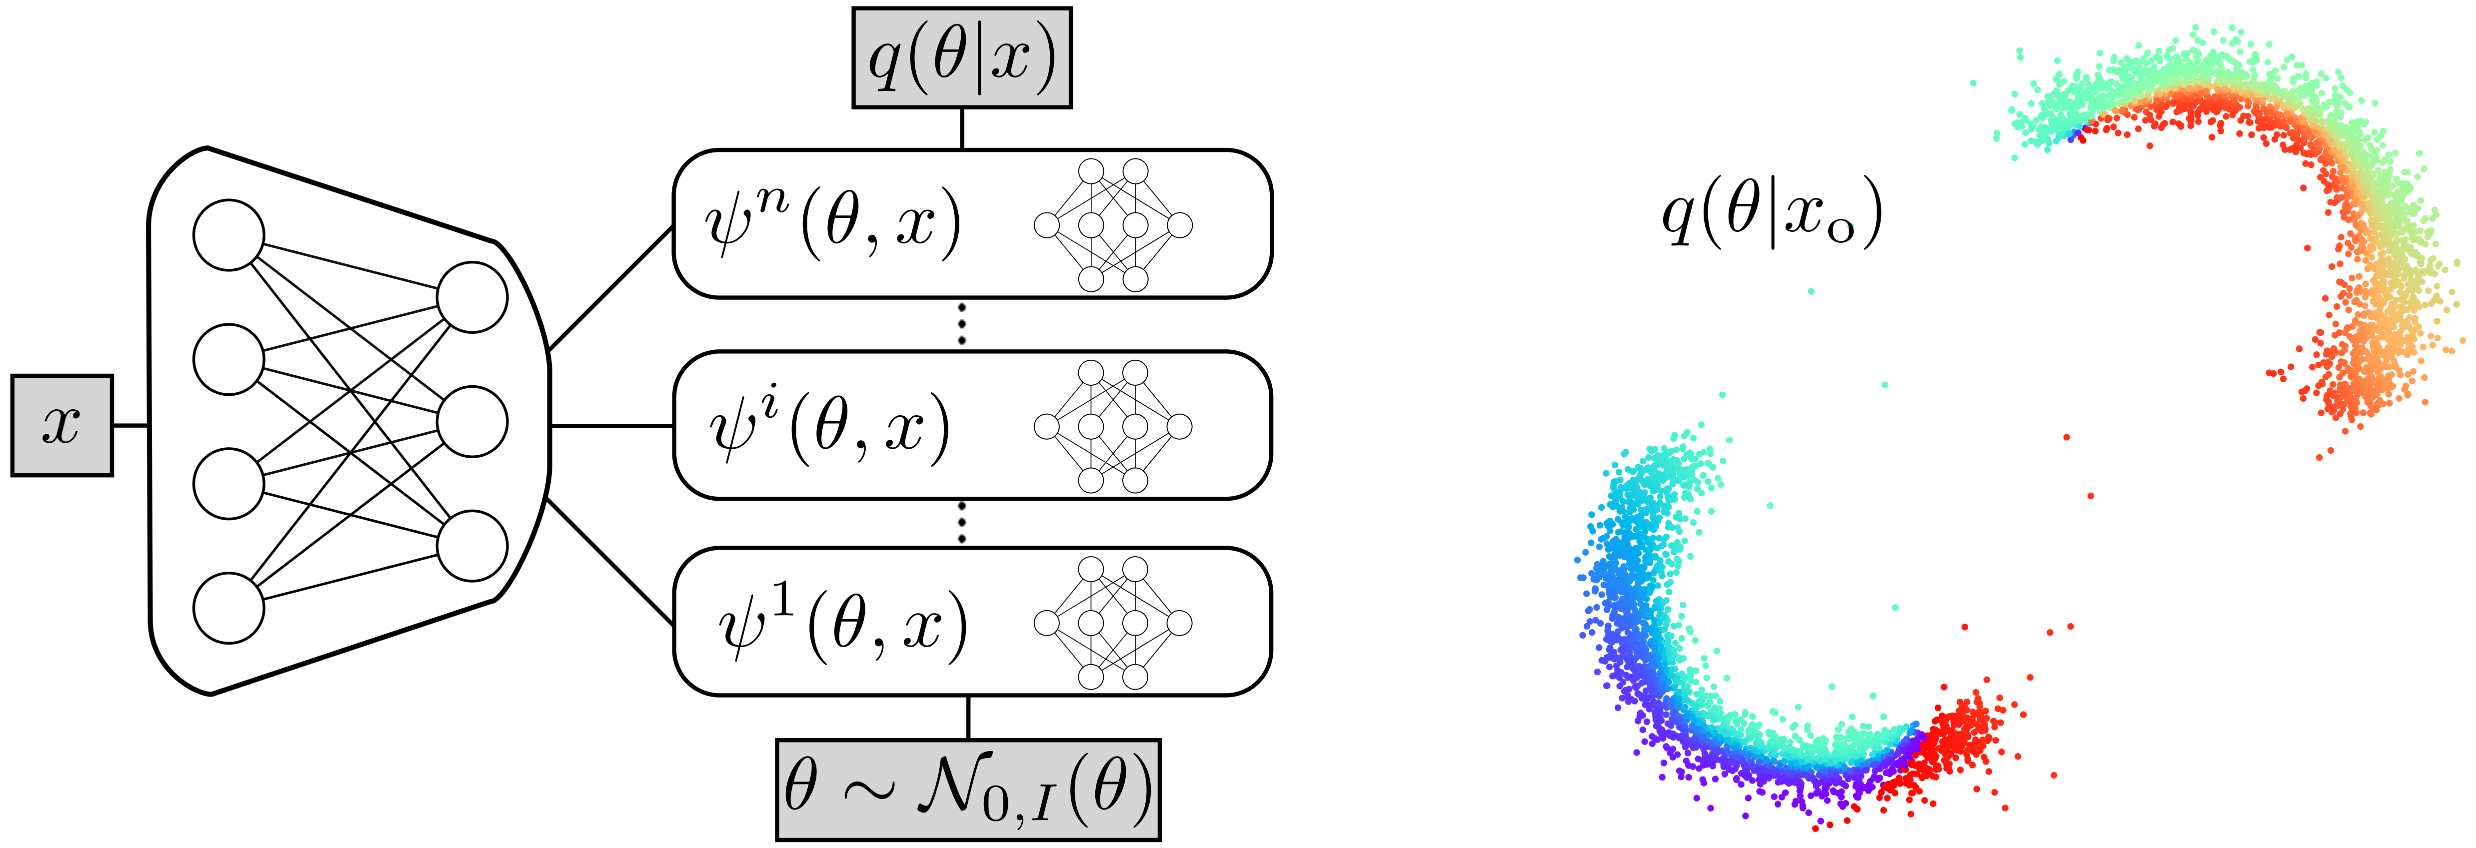

FMPE

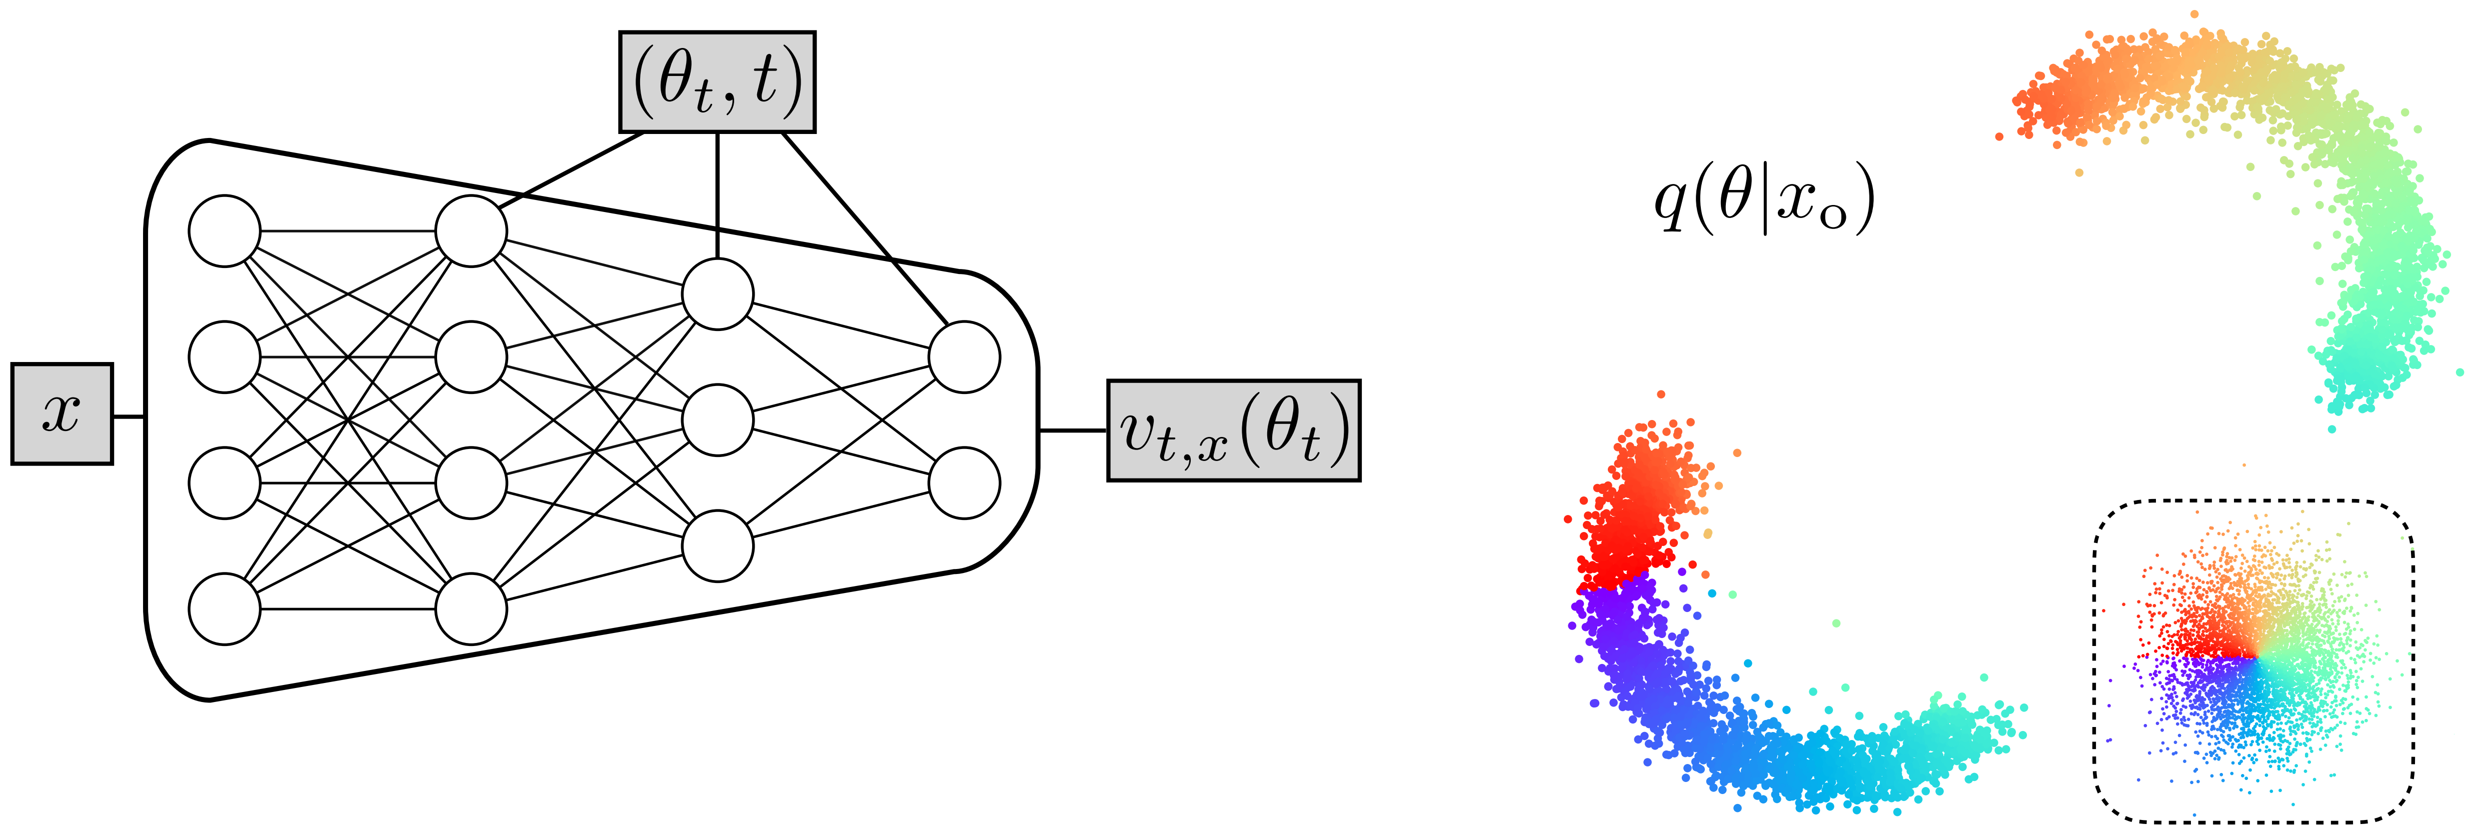

Dax*, Wildberger*, Buchholz* et al, NeurIPS 2023

Jonas Wildberger

Max Dax

Simon Buchholz
all at MPI-IS

Scaling up: Flow-Matching Neural Posterior Estimation

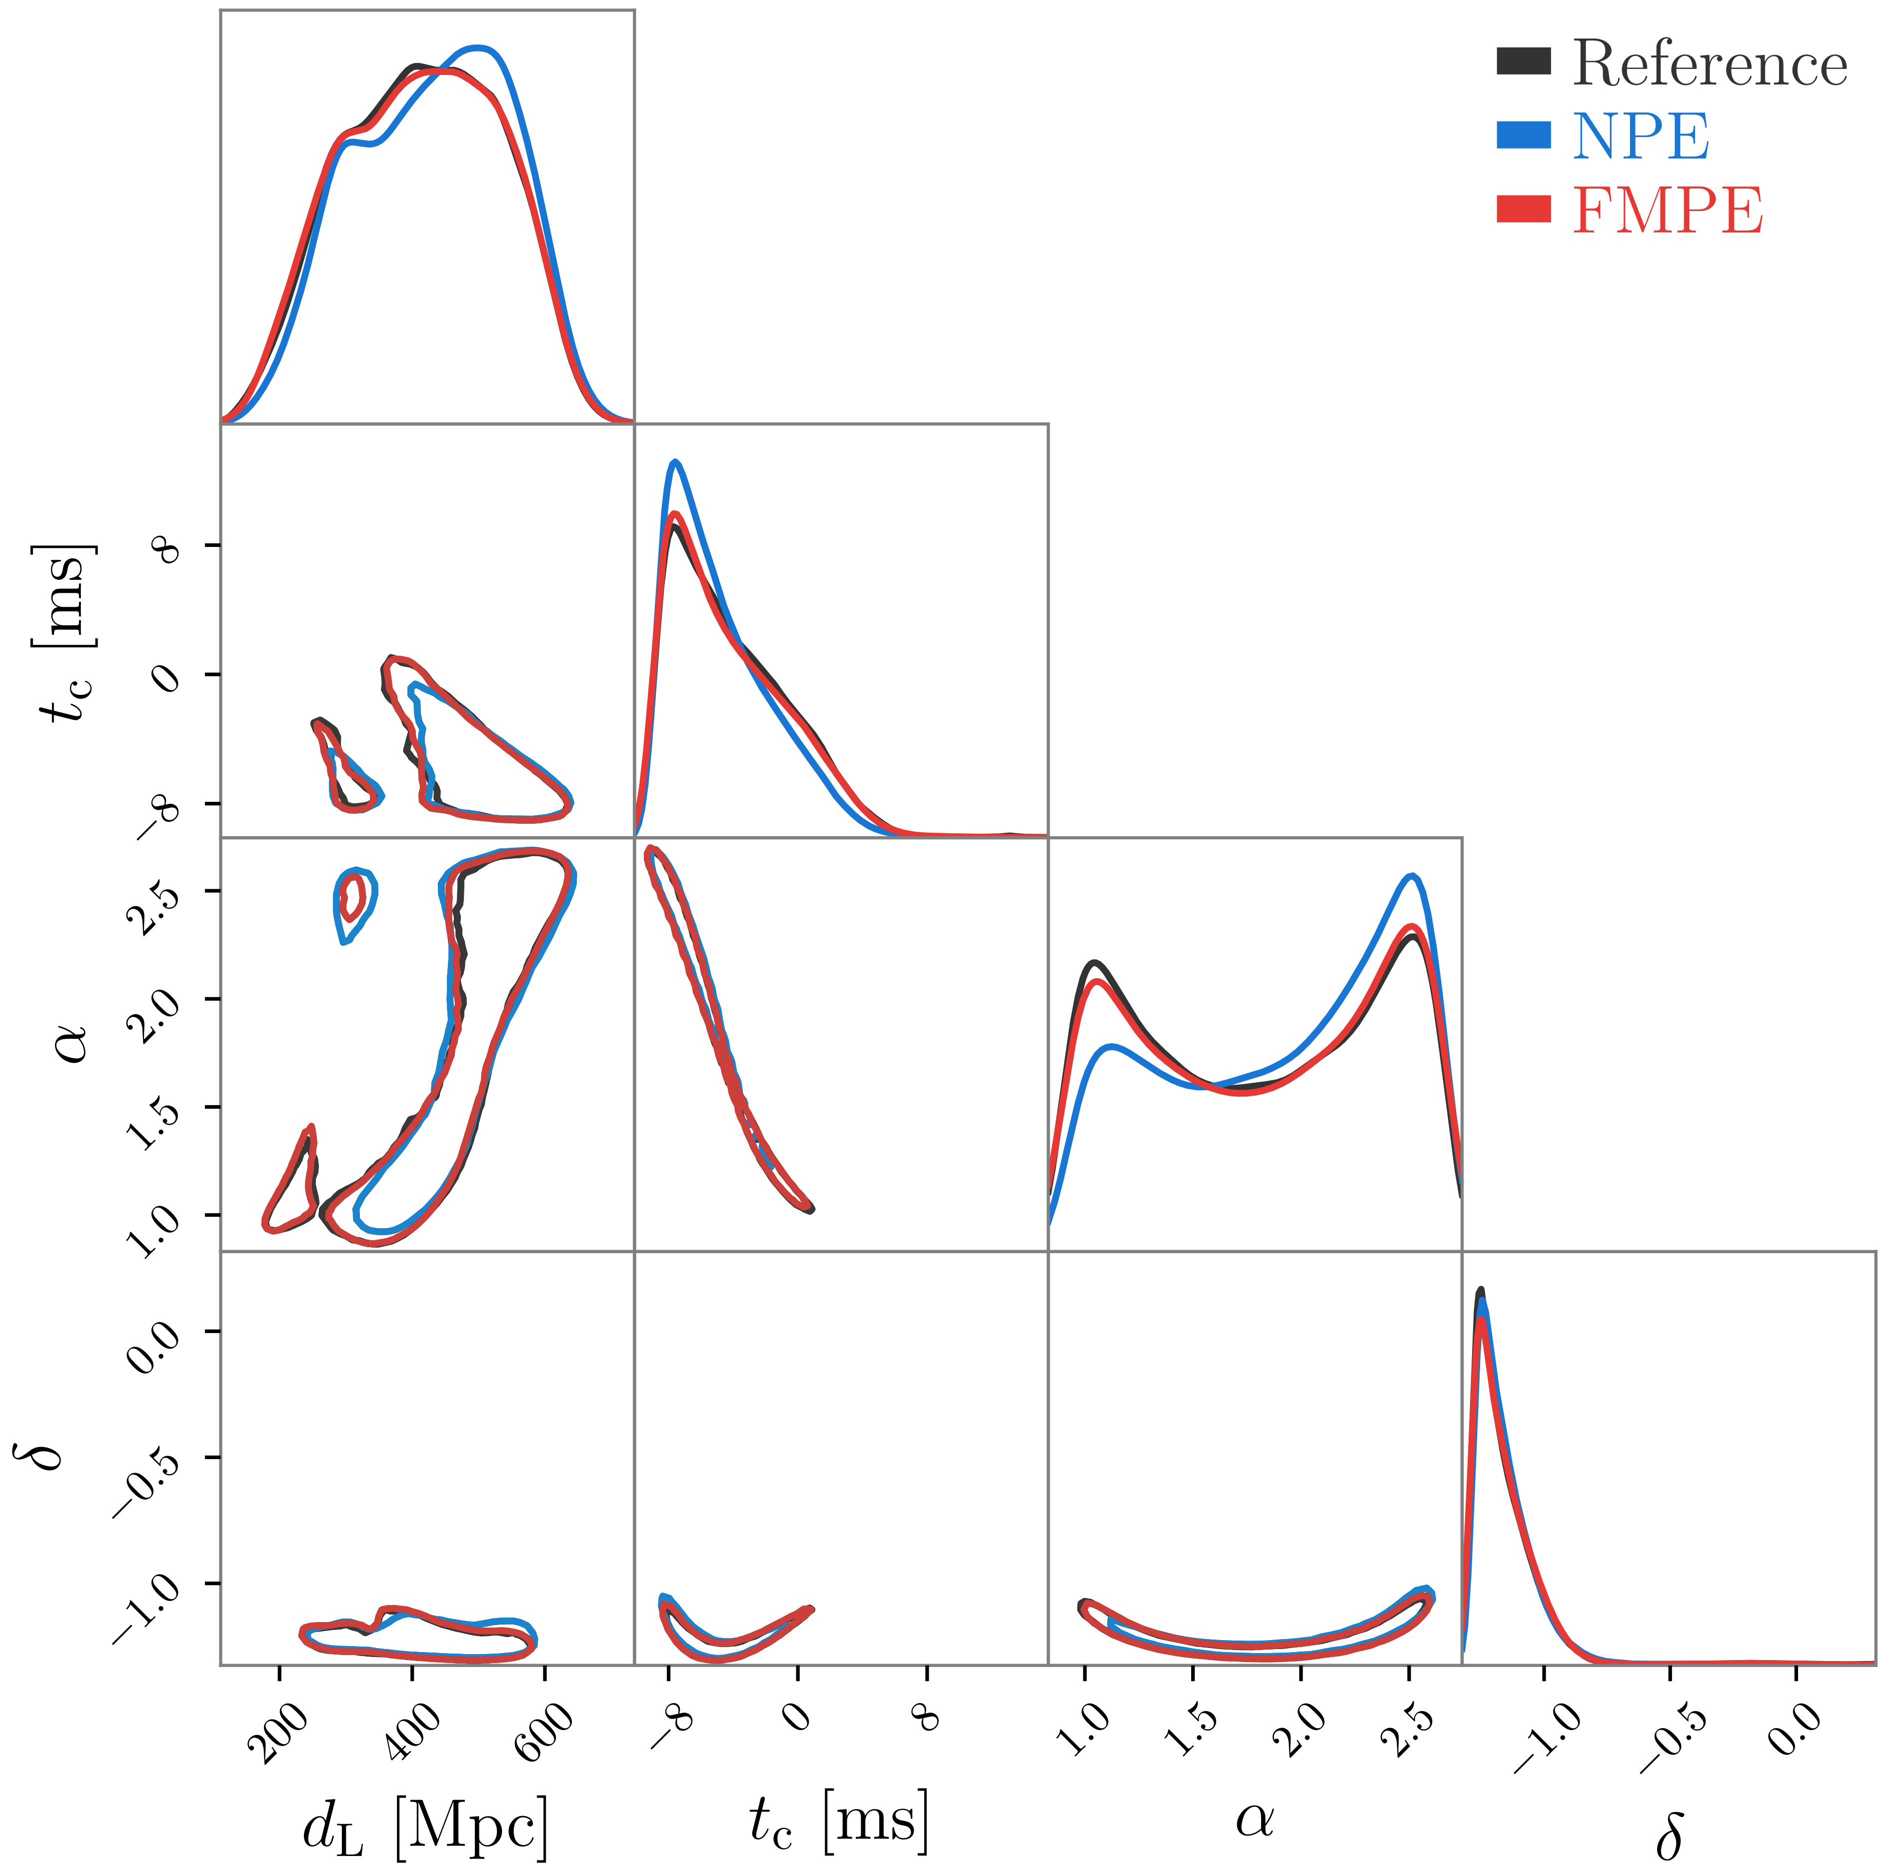

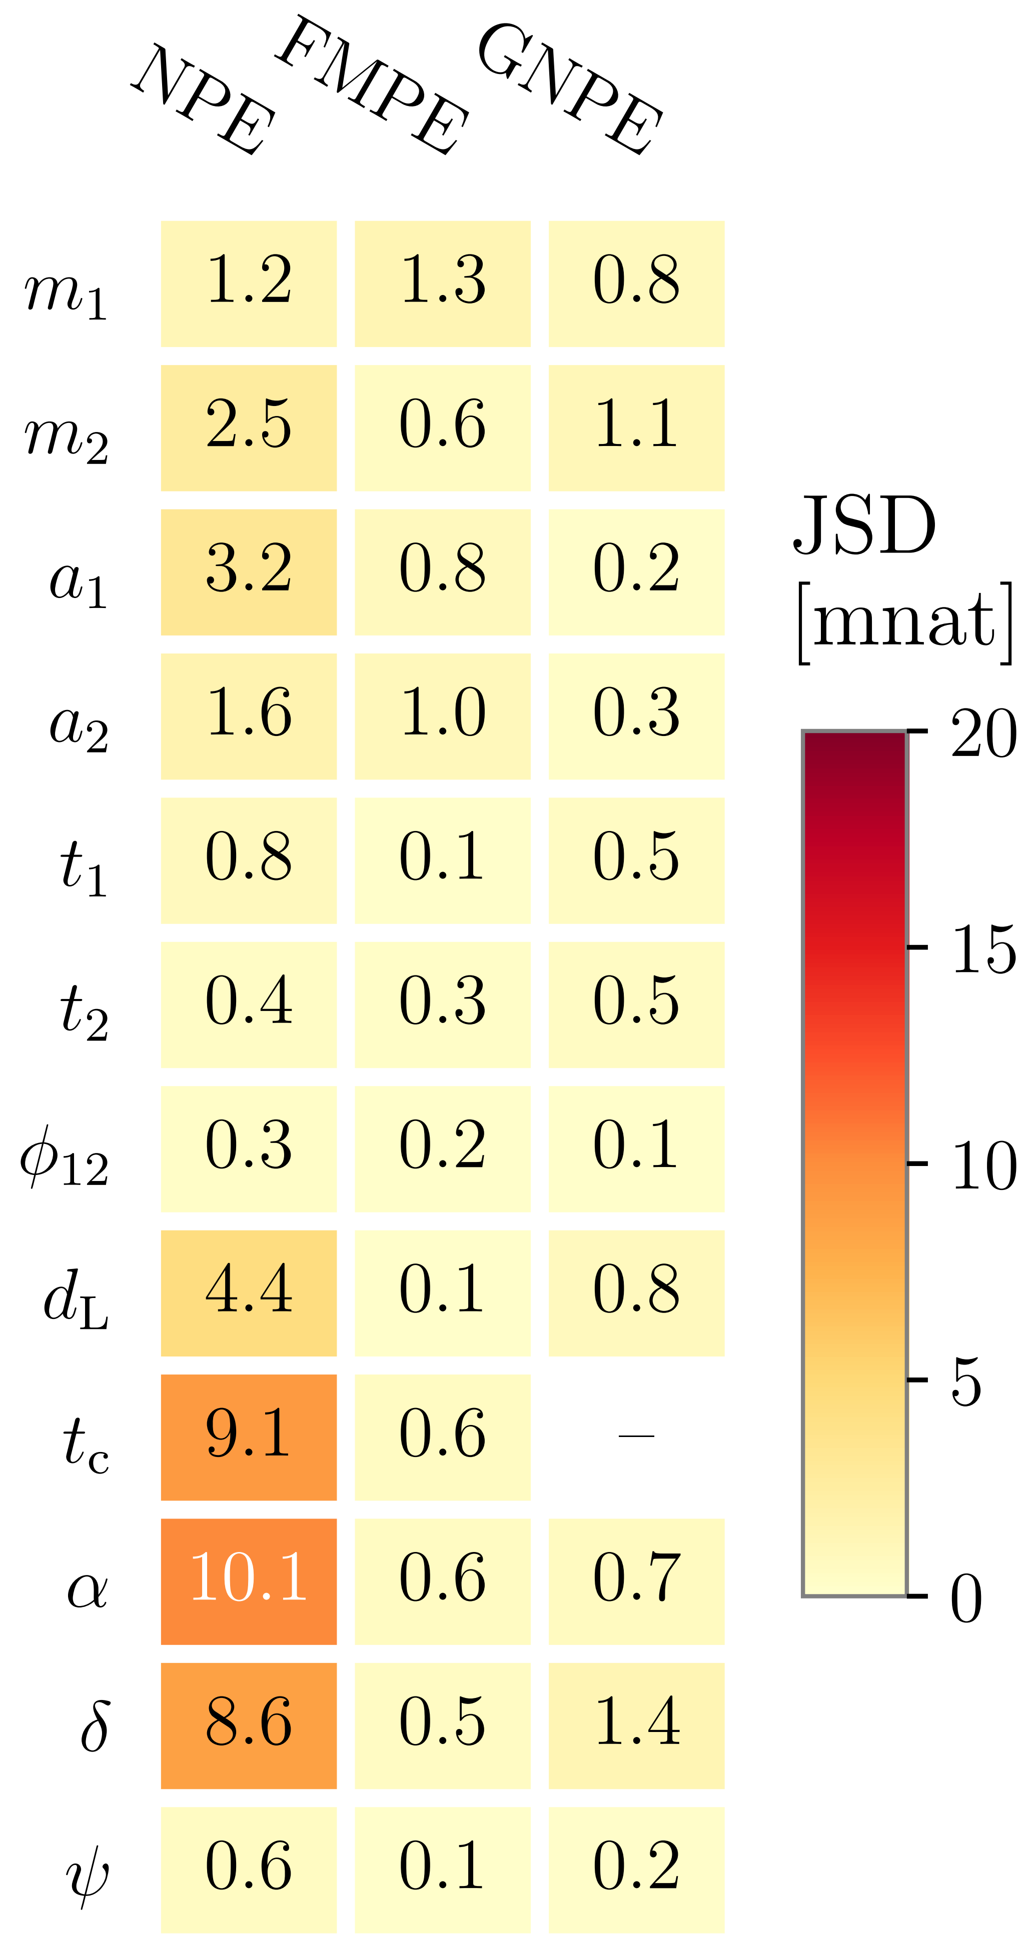

	NPE	FMPE	GNPE
m_1	1.2	1.3	0.8
m_2	2.5	0.6	1.1
a_1	3.2	0.8	0.2
a_2	1.6	1.0	0.3
t_1	0.8	0.1	0.5
t_2	0.4	0.3	0.5
ϕ_{12}	0.3	0.2	0.1
d_L	4.4	0.1	0.8
t_c	9.1	0.6	–
α	10.1	0.6	0.7
δ	8.6	0.5	1.4
ψ	0.6	0.1	0.2

- Gravitational Wave parameters $\theta \in \mathbb{R}^{15}$ (masses, spins etc.)
- Gravitational waves $x \in \mathbb{R}^{15,744}$ (frequency domain data)

- FMPE network $1.9 \cdot 10^8$ learnable parameters, training around **2 days**
- NPE network $1.3 \cdot 10^8$ parameters, training around **3 days**

- FMPE performance **en par with GNPE**

Dax*, Wildberger*, Buchholz* et al, NeurIPS 2023

Simulation-based inference:

New perspectives and opportunities across multiple scientific disciplines …

… but …

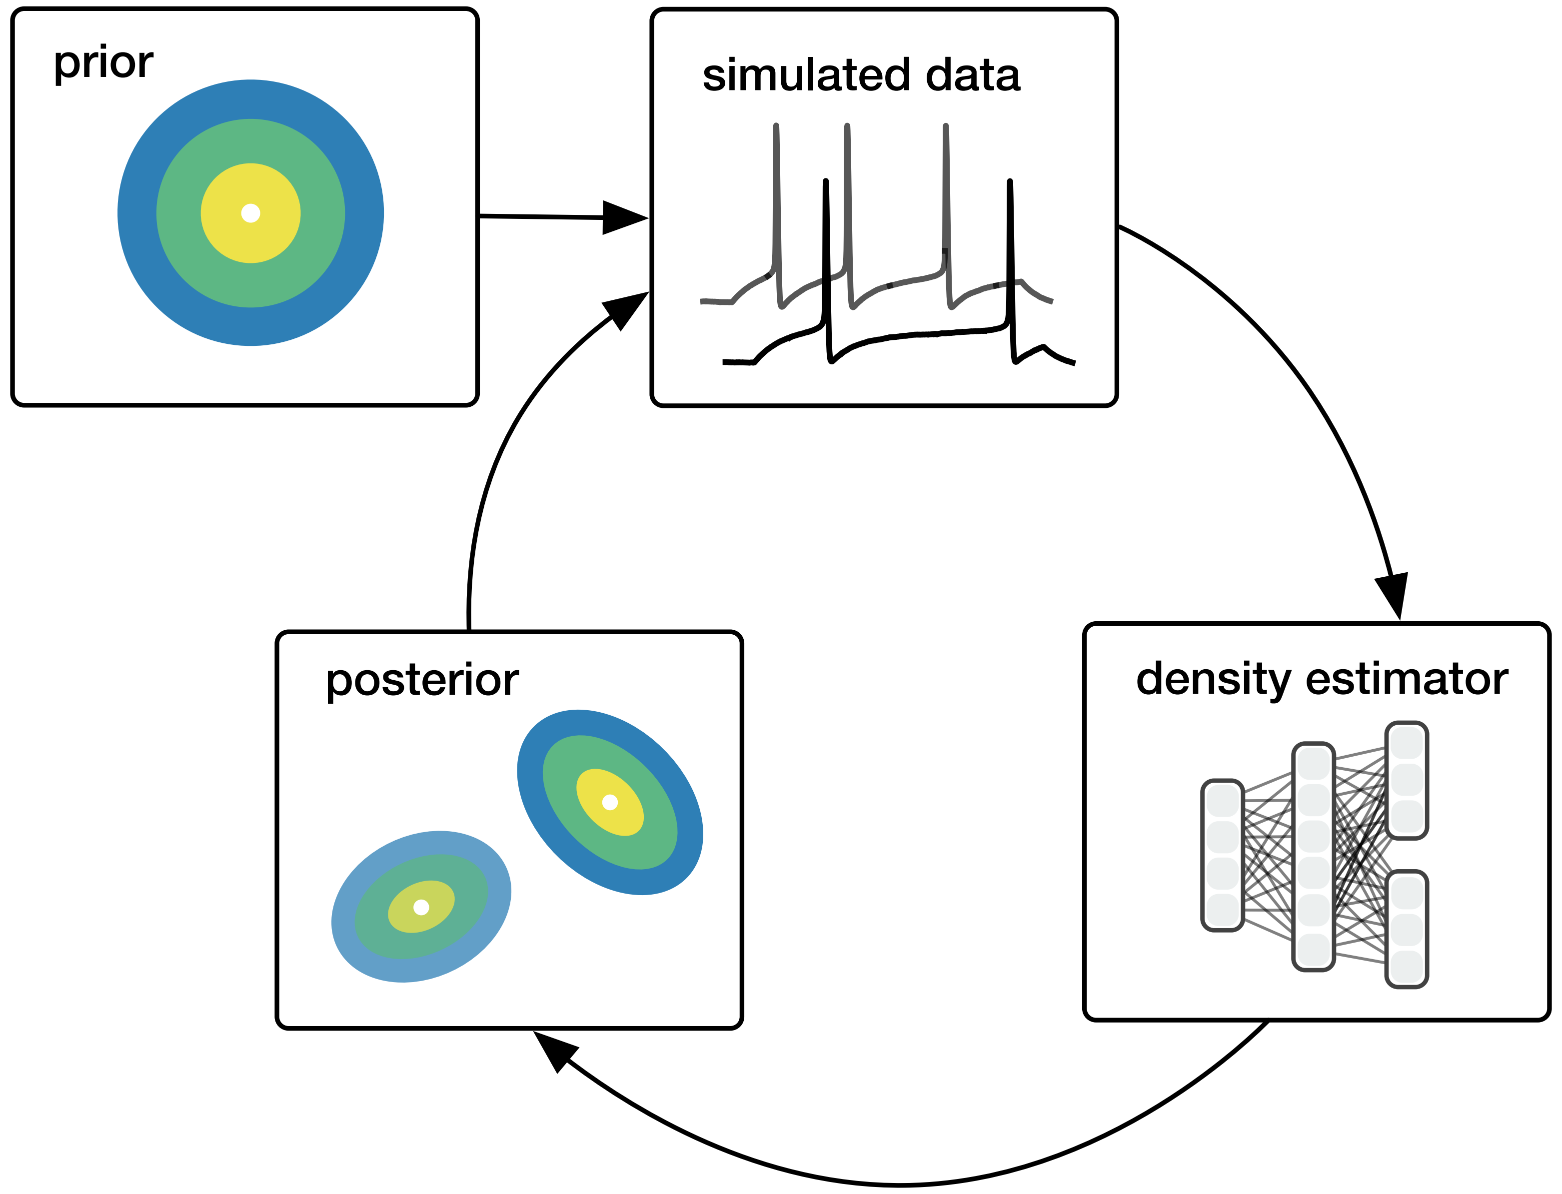

- Simulation efficiency
- Mis-specification
  Gao, Deistler, Macke, NeurIPS 2023
- Infer models, not just parameters?
  Schröder, Macke Arxive 2023
- Robustness?
  Glöckler, Deistler, Macke, ICML 2023

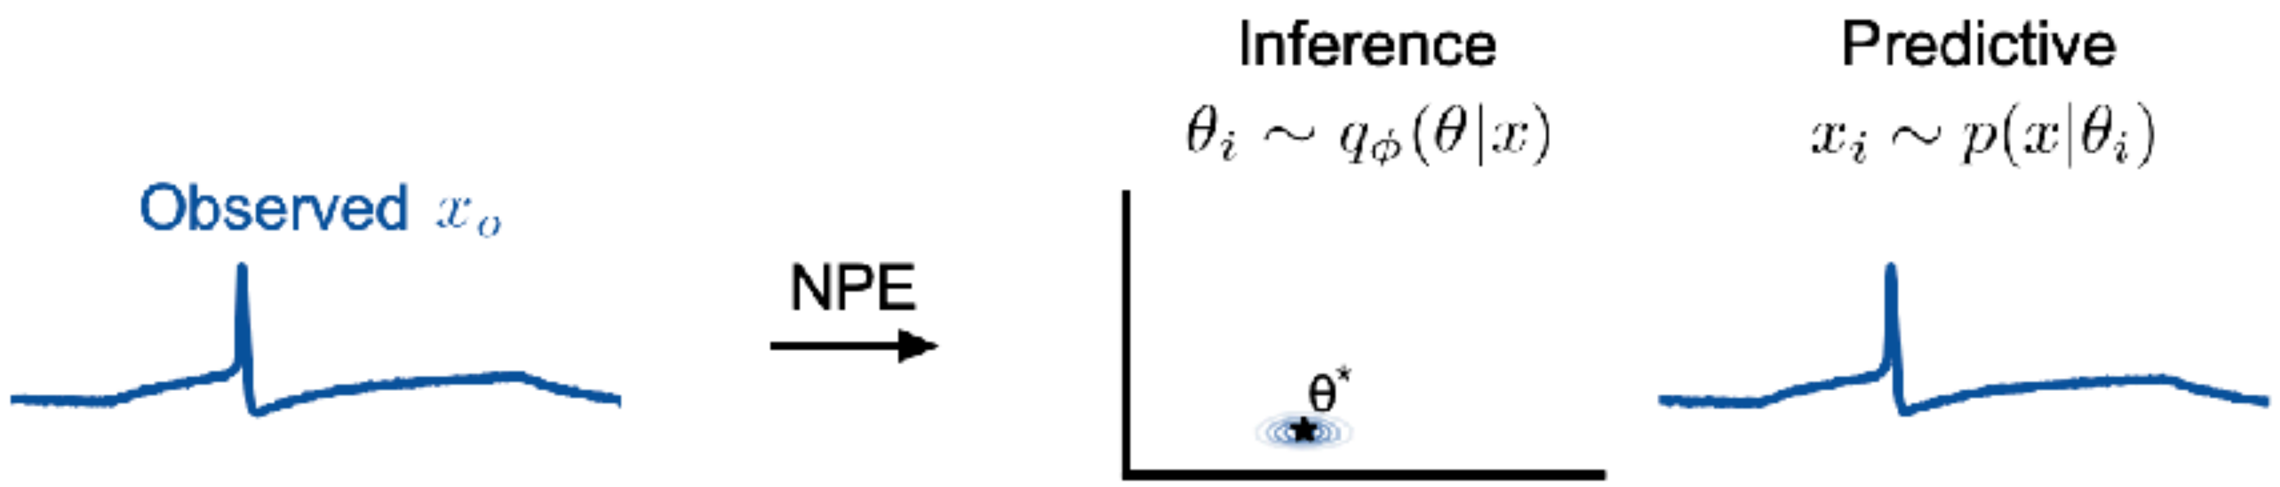

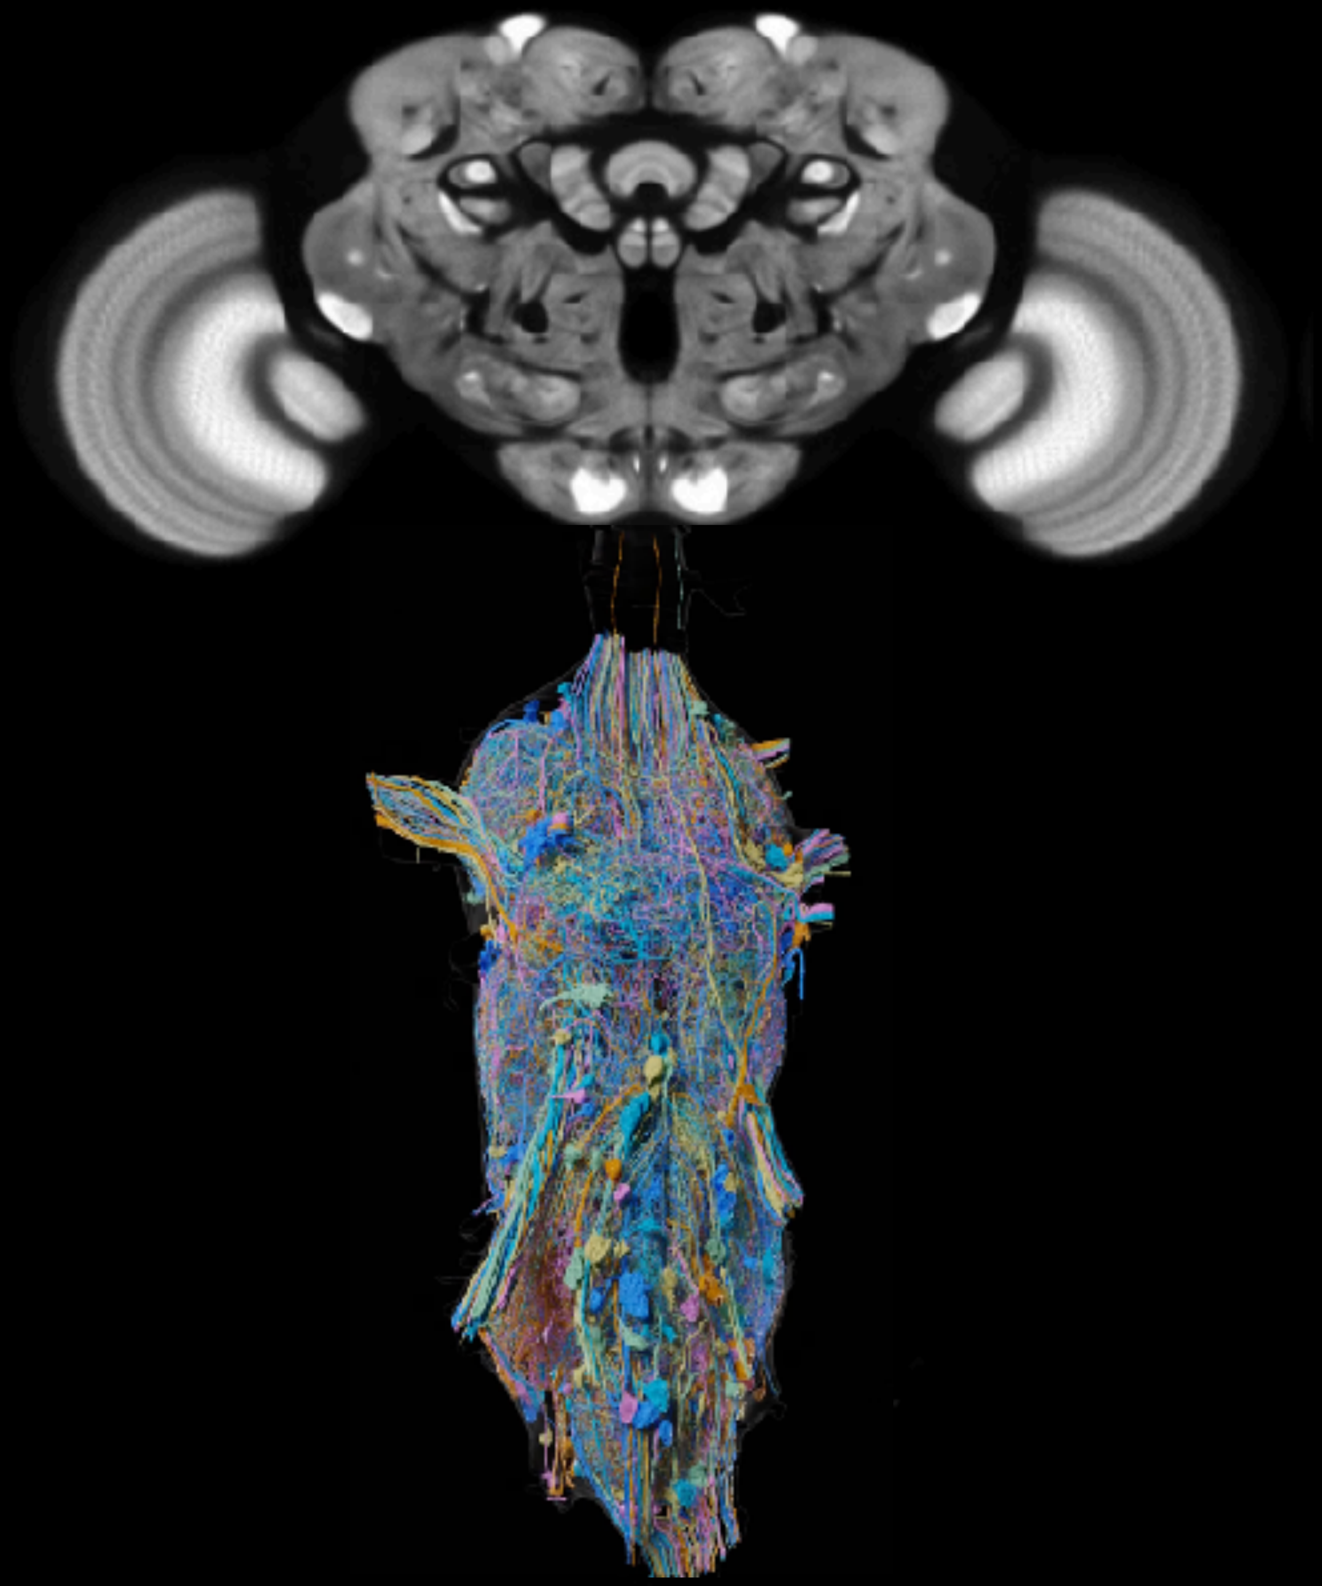

Bogovic et al 2020
Takemura et al 2023

Srinivas Turaga

(HHMI Janelia)

Janne Lappalainen

Optic lobe of the *Drosophila*

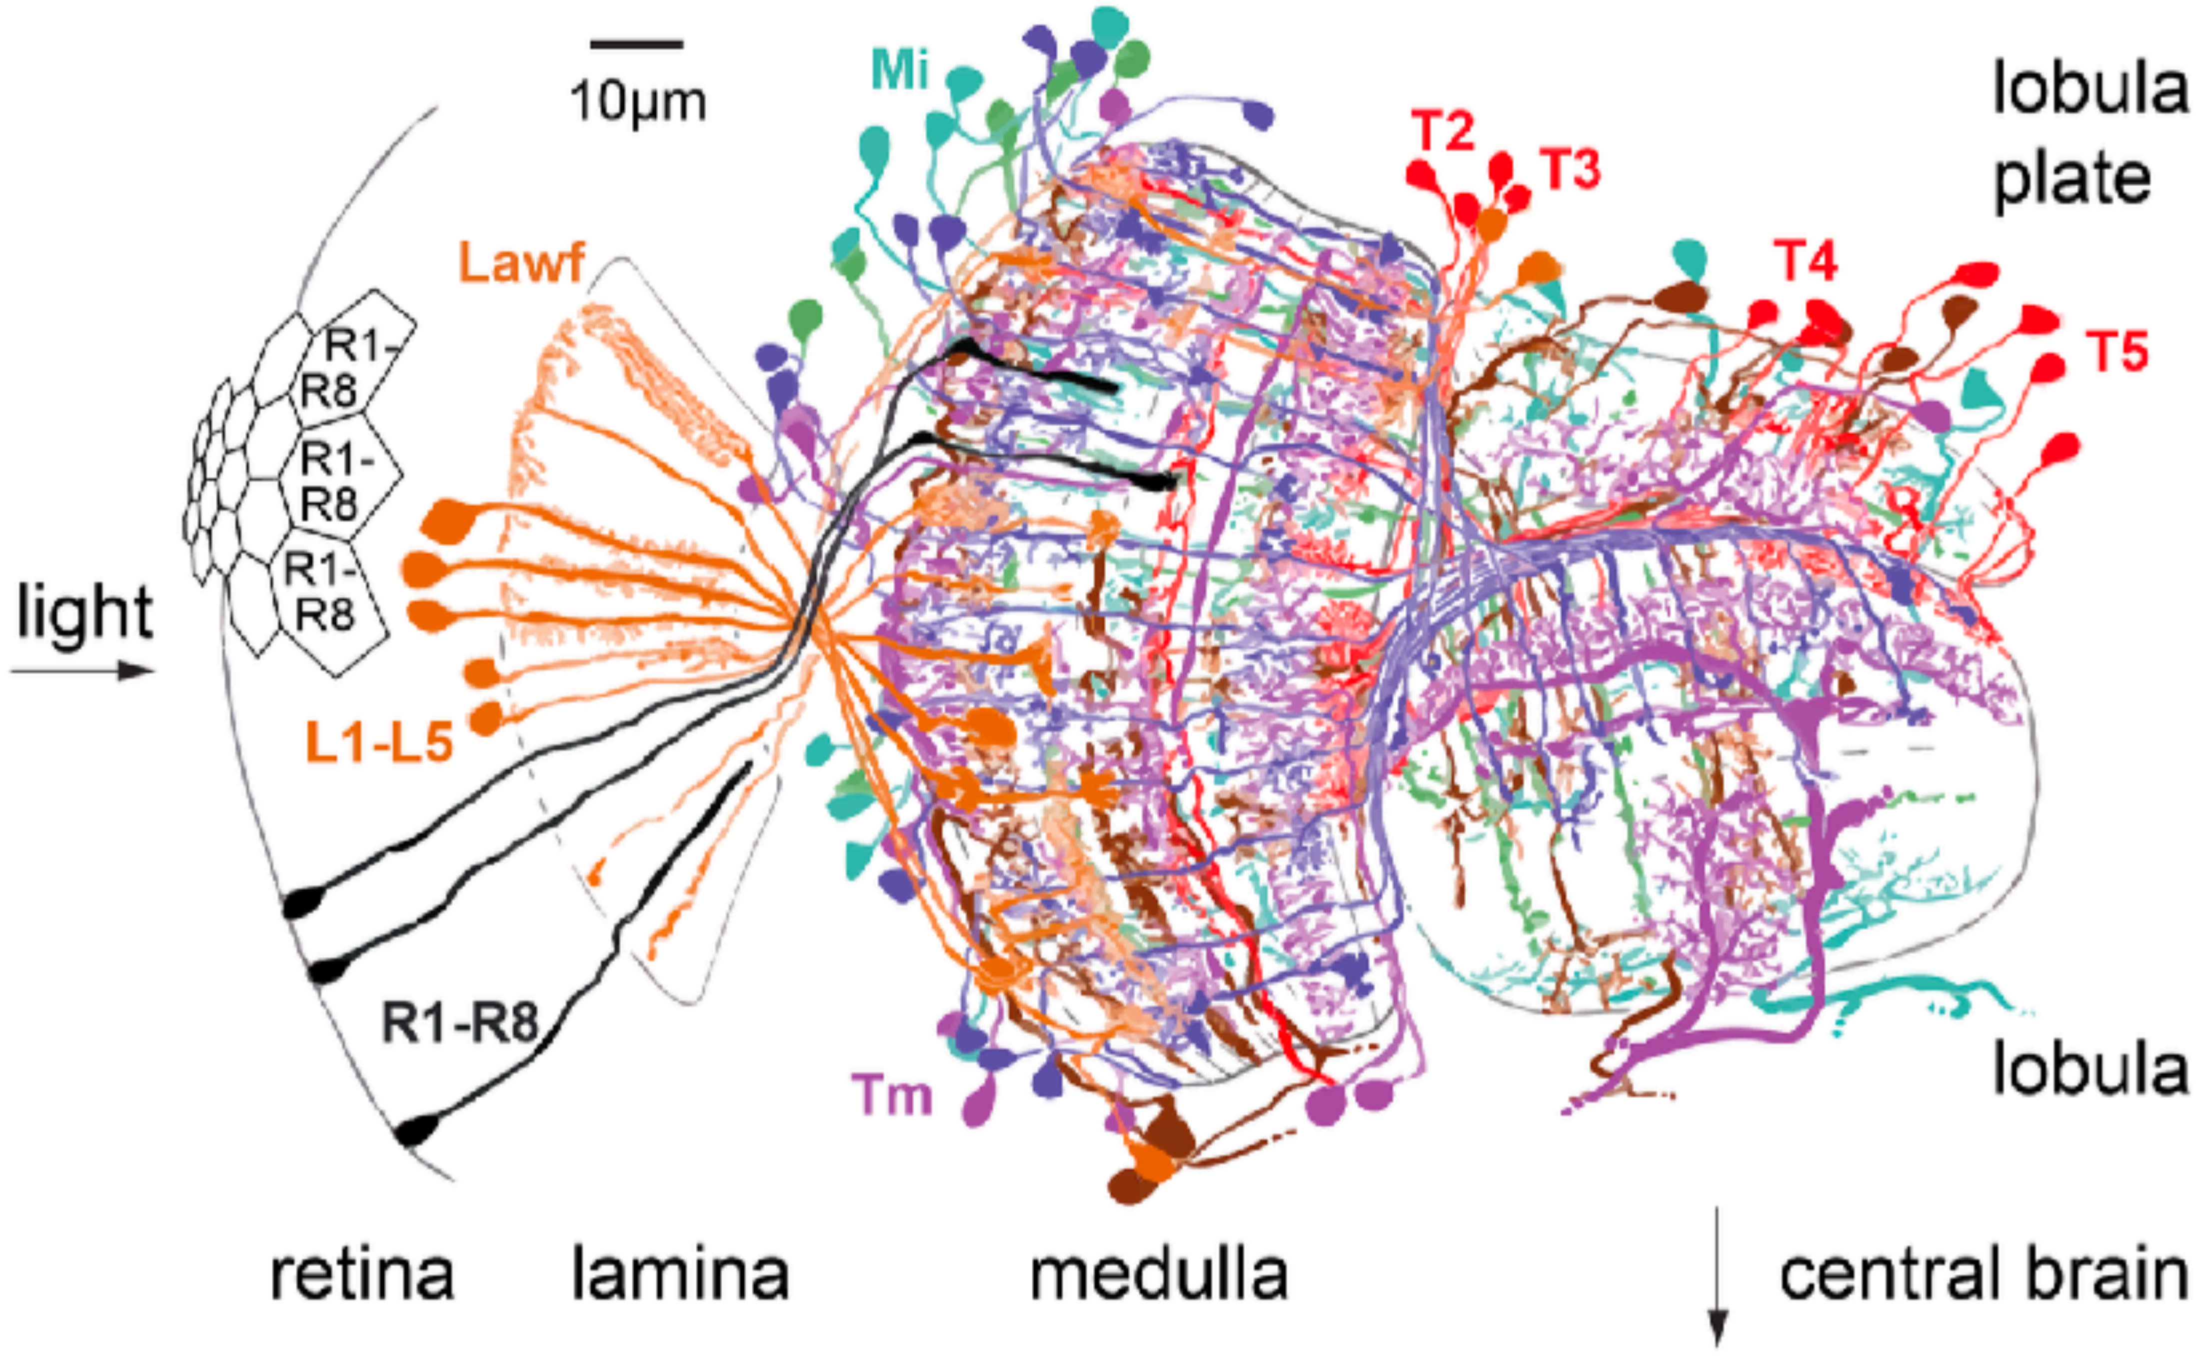

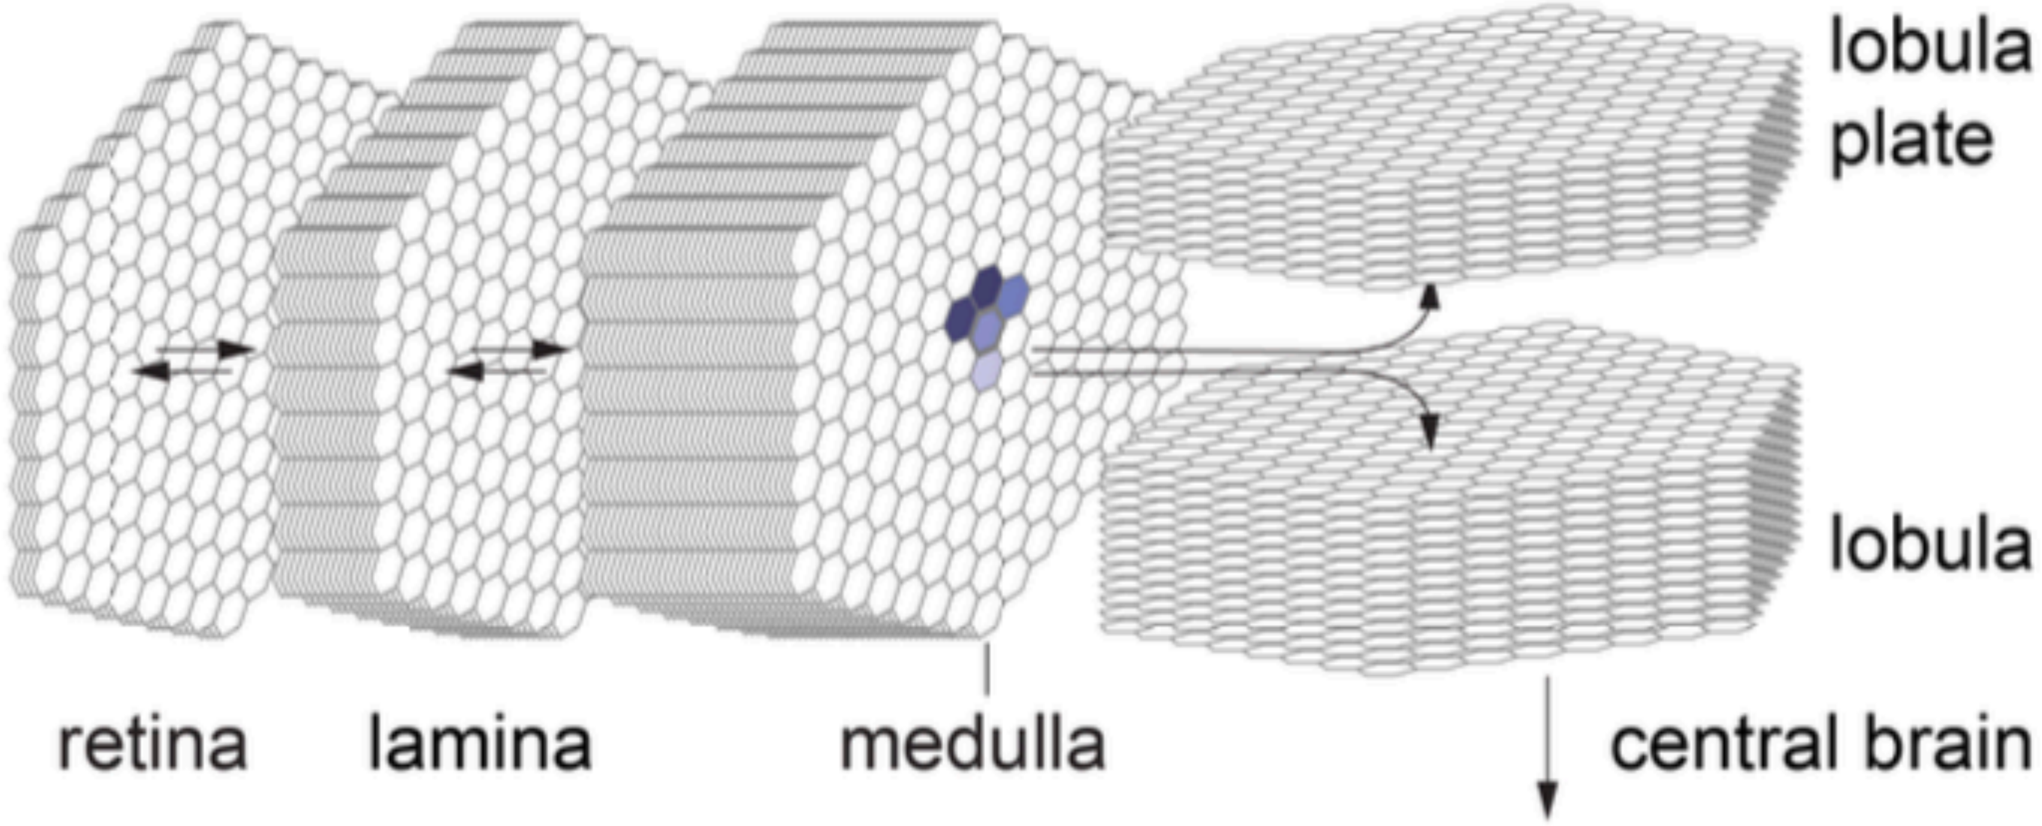

We built a recurrent, hexagonally convolutional neural network constrained by the fruit fly optic lobe— a “Deep Mechanistic Network” (DMN)

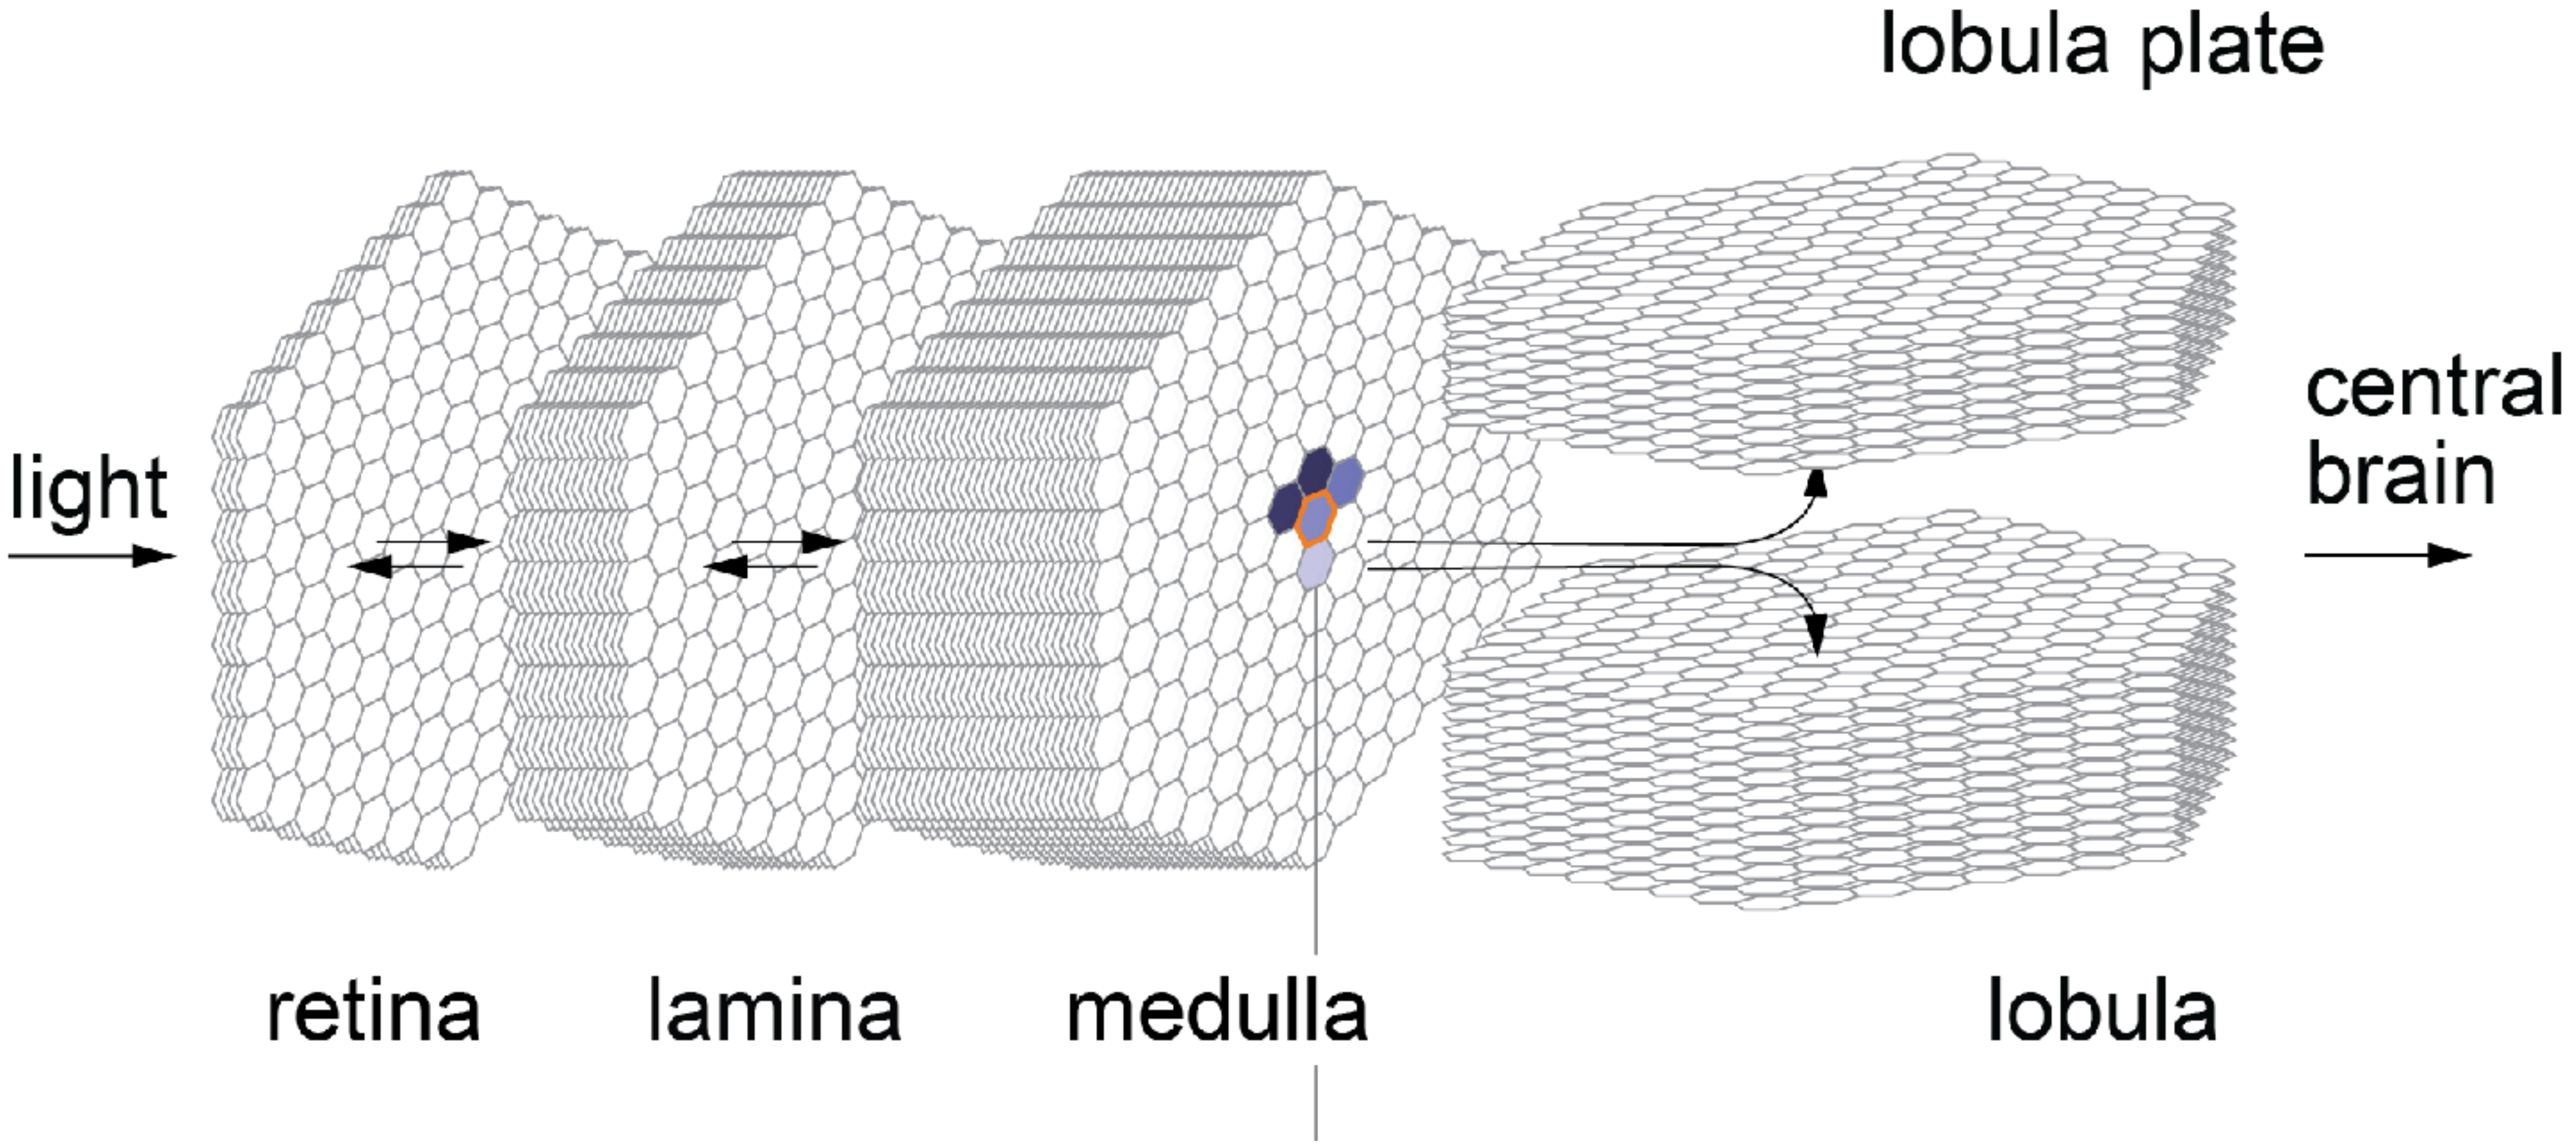

64 cell types, 44k cells, 1.5Mio connections

We use simple models of single-cell and synaptic dynamics

Non-spiking point neurons:

$$\tau_{t_i}\dot{V}_i = -V_i + \sum_j s_{ij} + V_{t_i}^{\mathrm{rest}}$$

time constants

resting potentials

Current-based synapses:

scaling coefficient of convolutional filters

synapse counts

$$s_{ij} = w_{ij} f(V_j) = \alpha_{t_i t_j} \sigma_{t_i t_j} N_{t_i t_j \Delta u \Delta v} f(V_j)$$

synapse signs

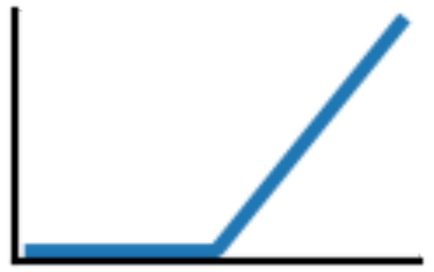

We train the DMN to perform motion-estimation on naturalistic inputs

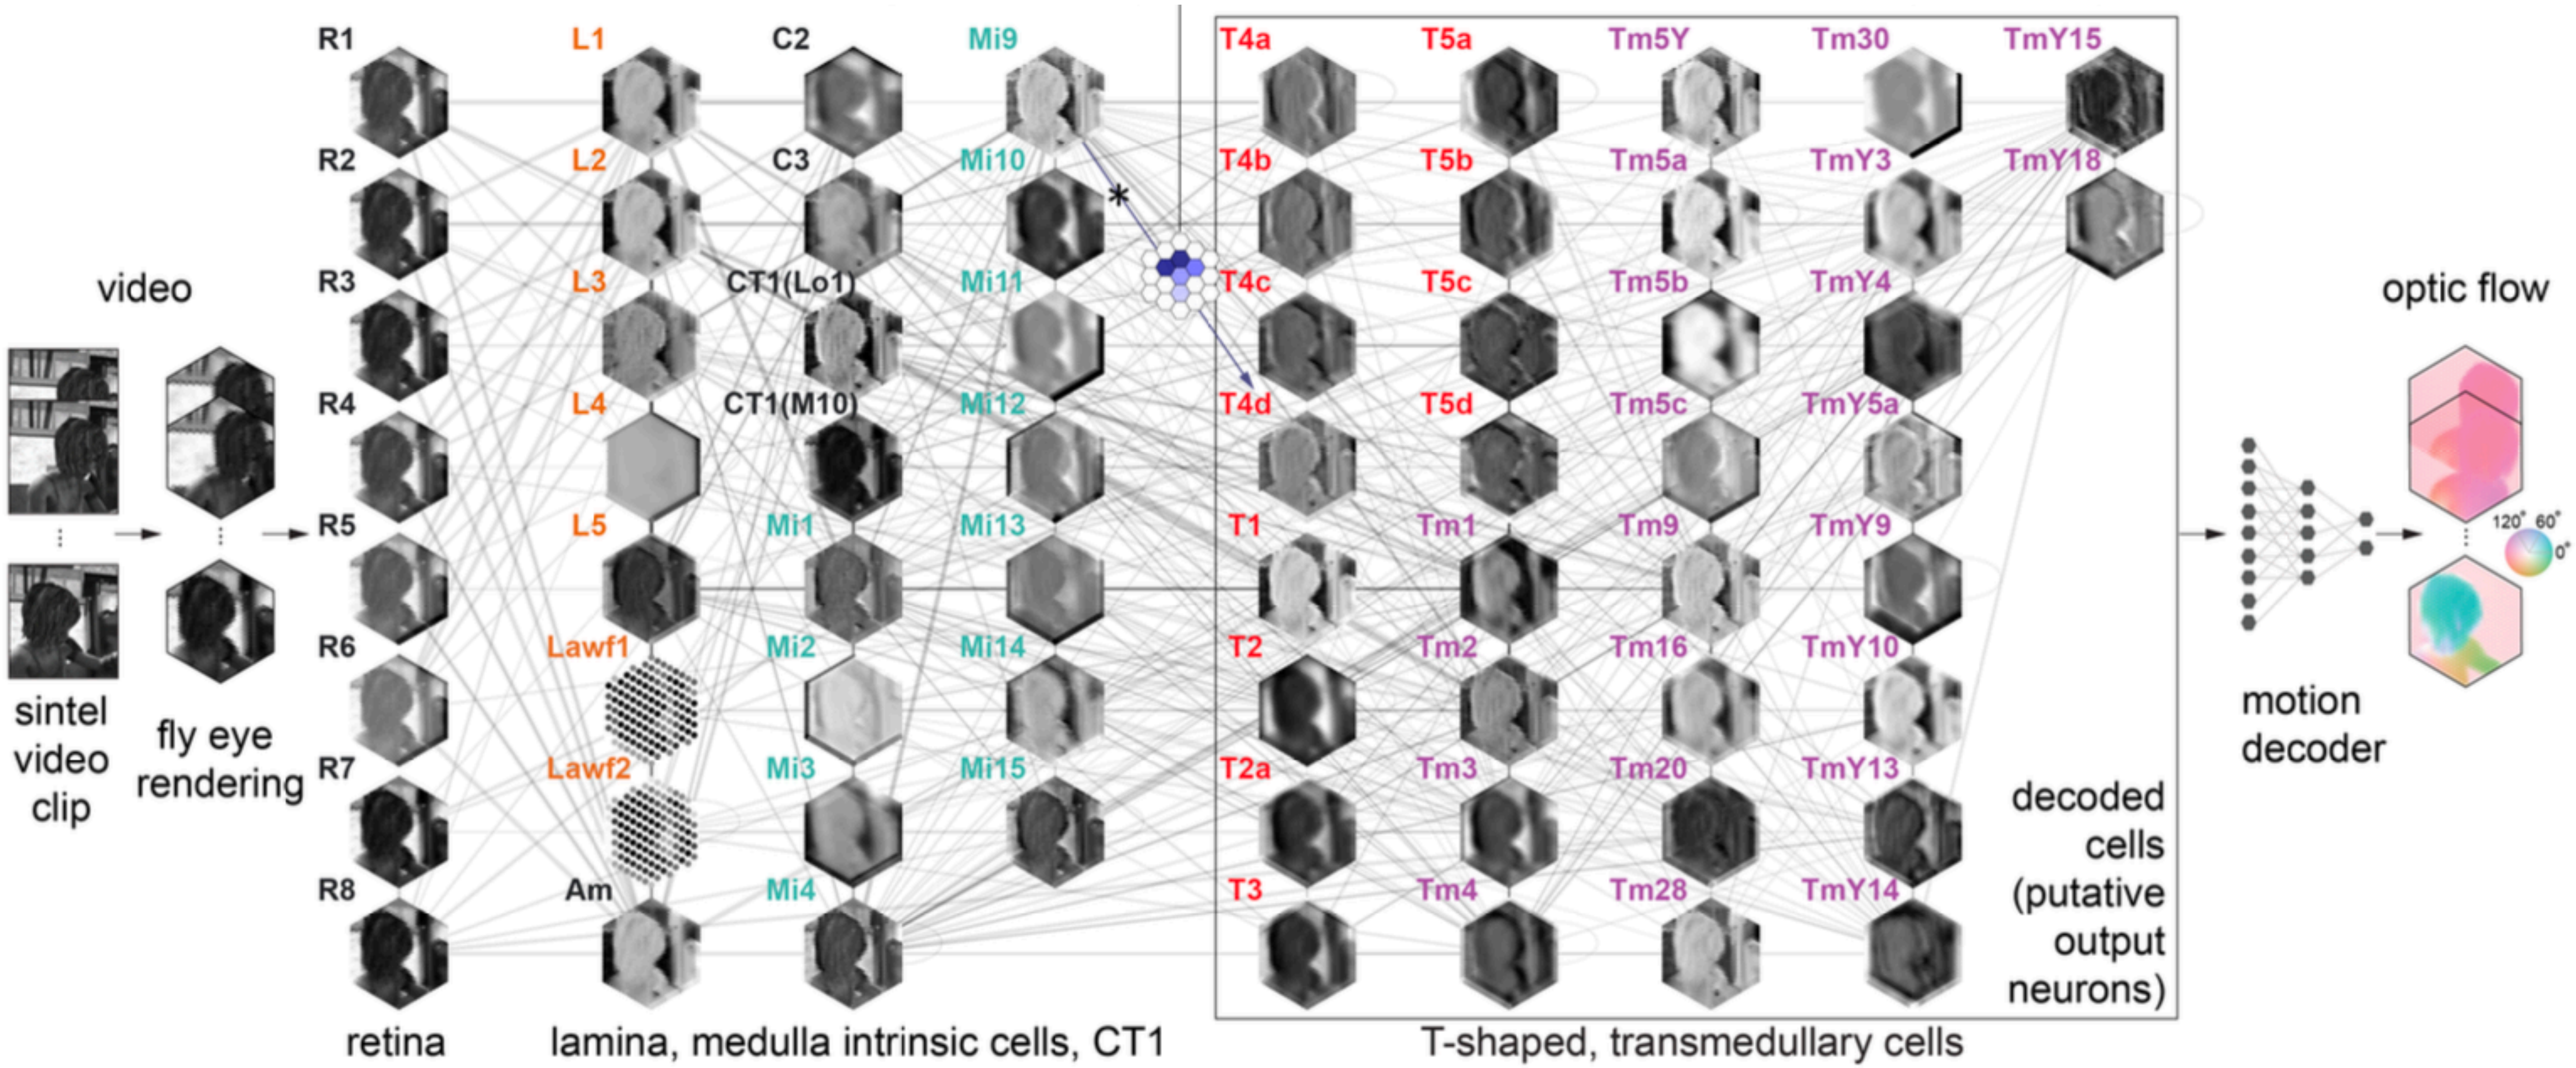

In-silico neurophysiology on the model: Flash responses

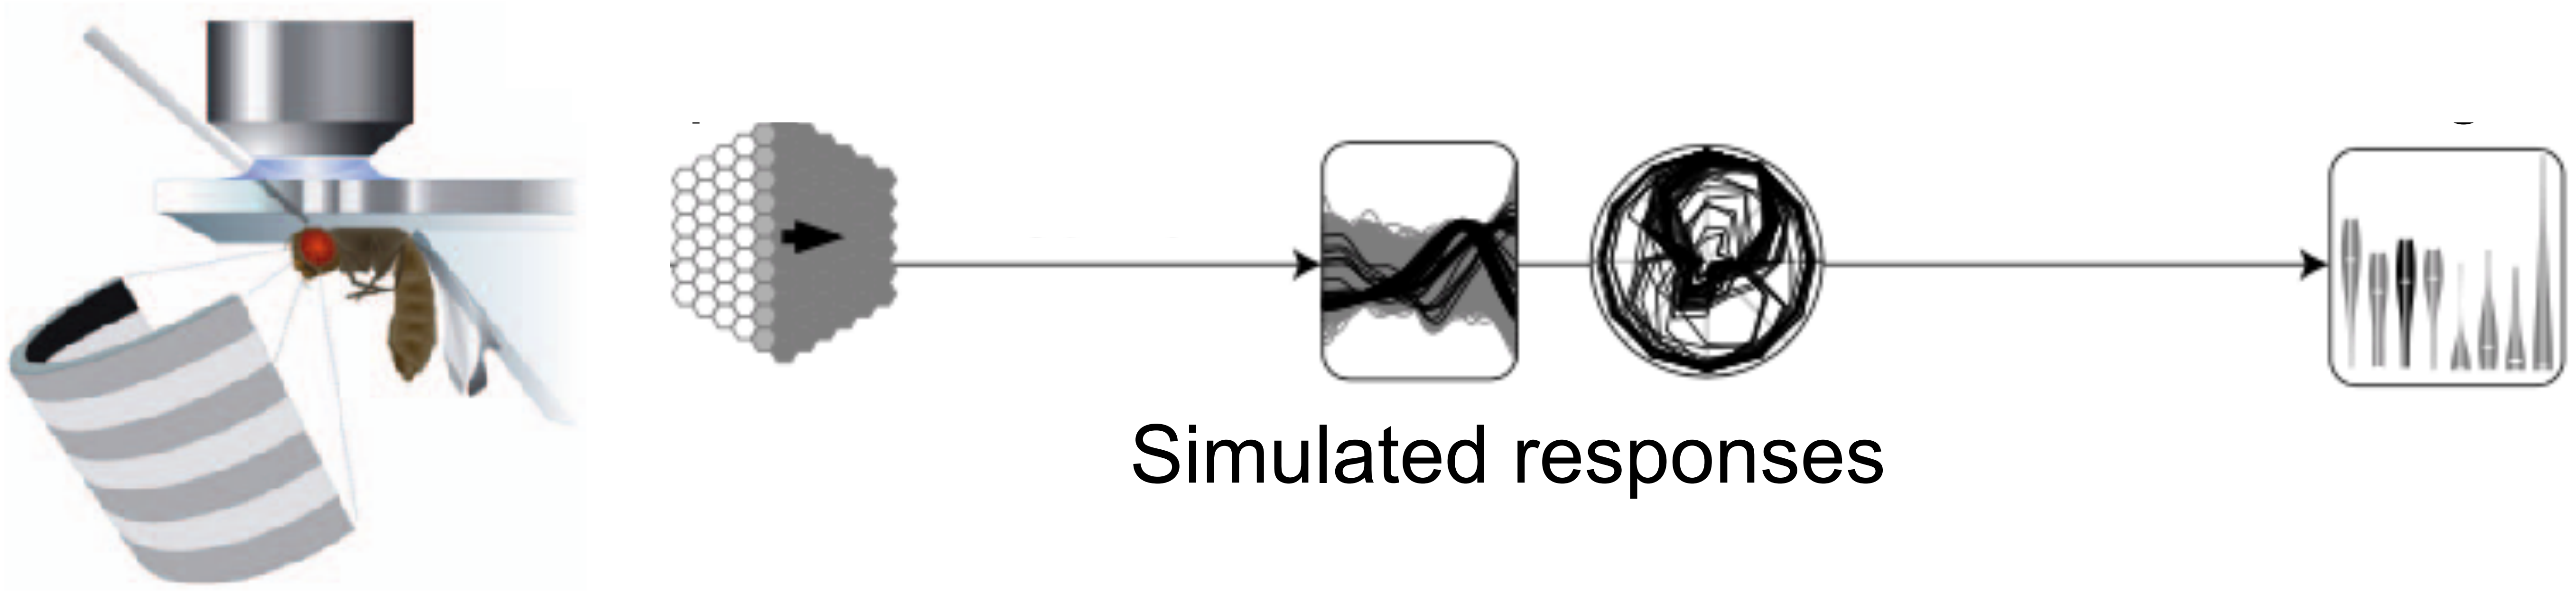

not yet established

In-silico neurophysiology on the model: Direction selectivity

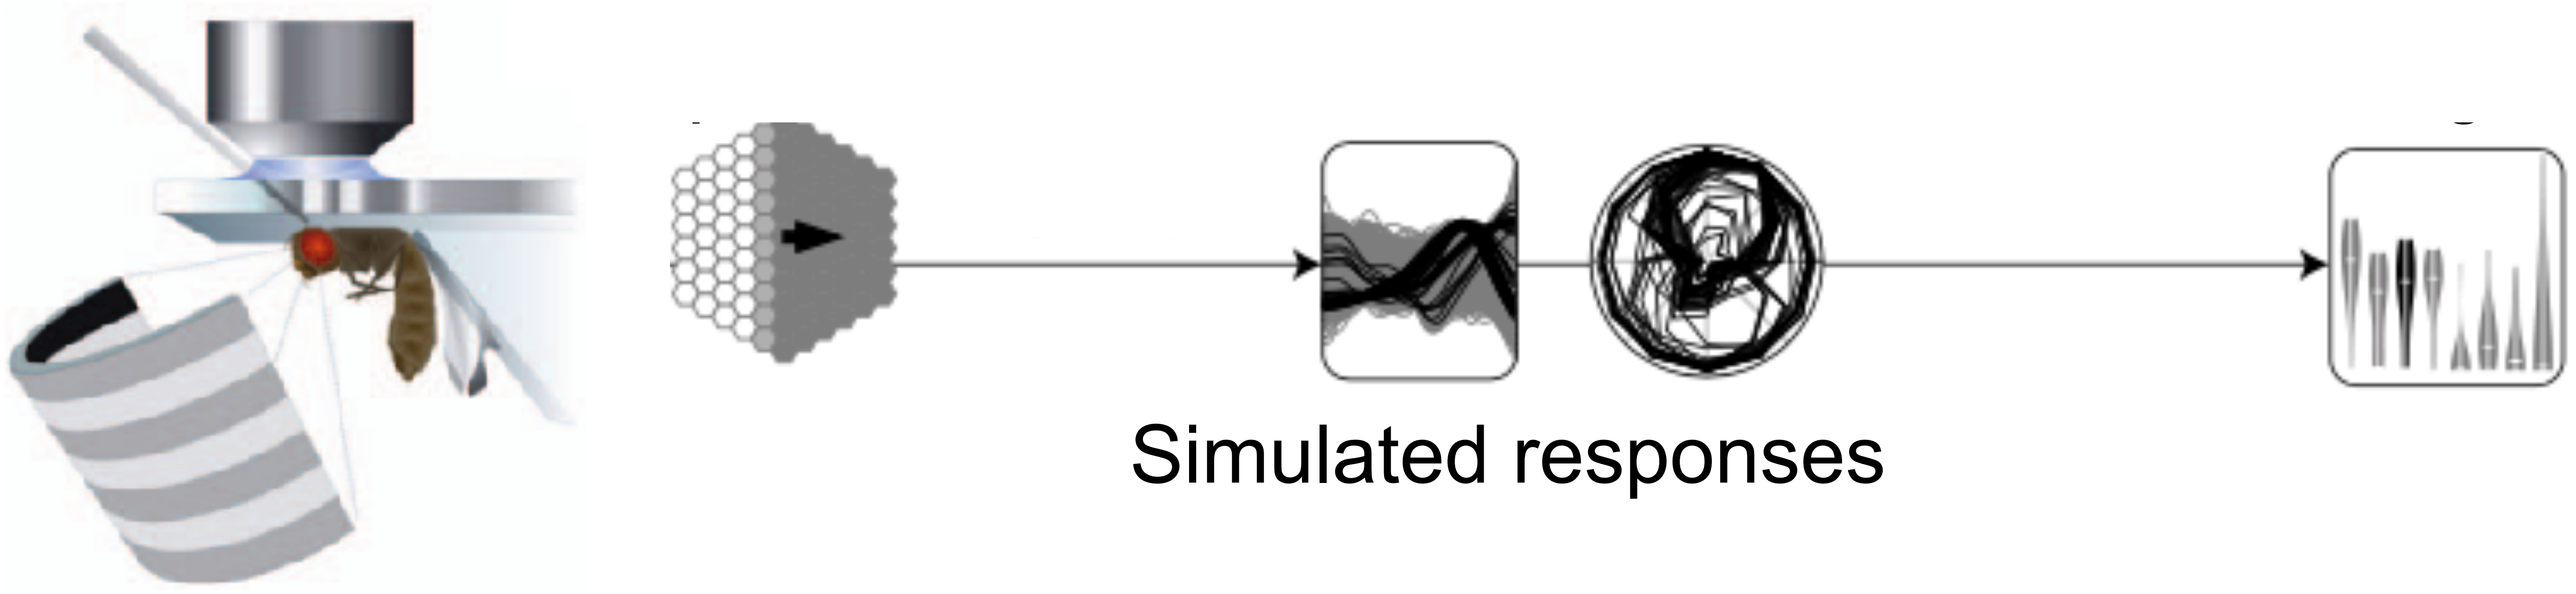

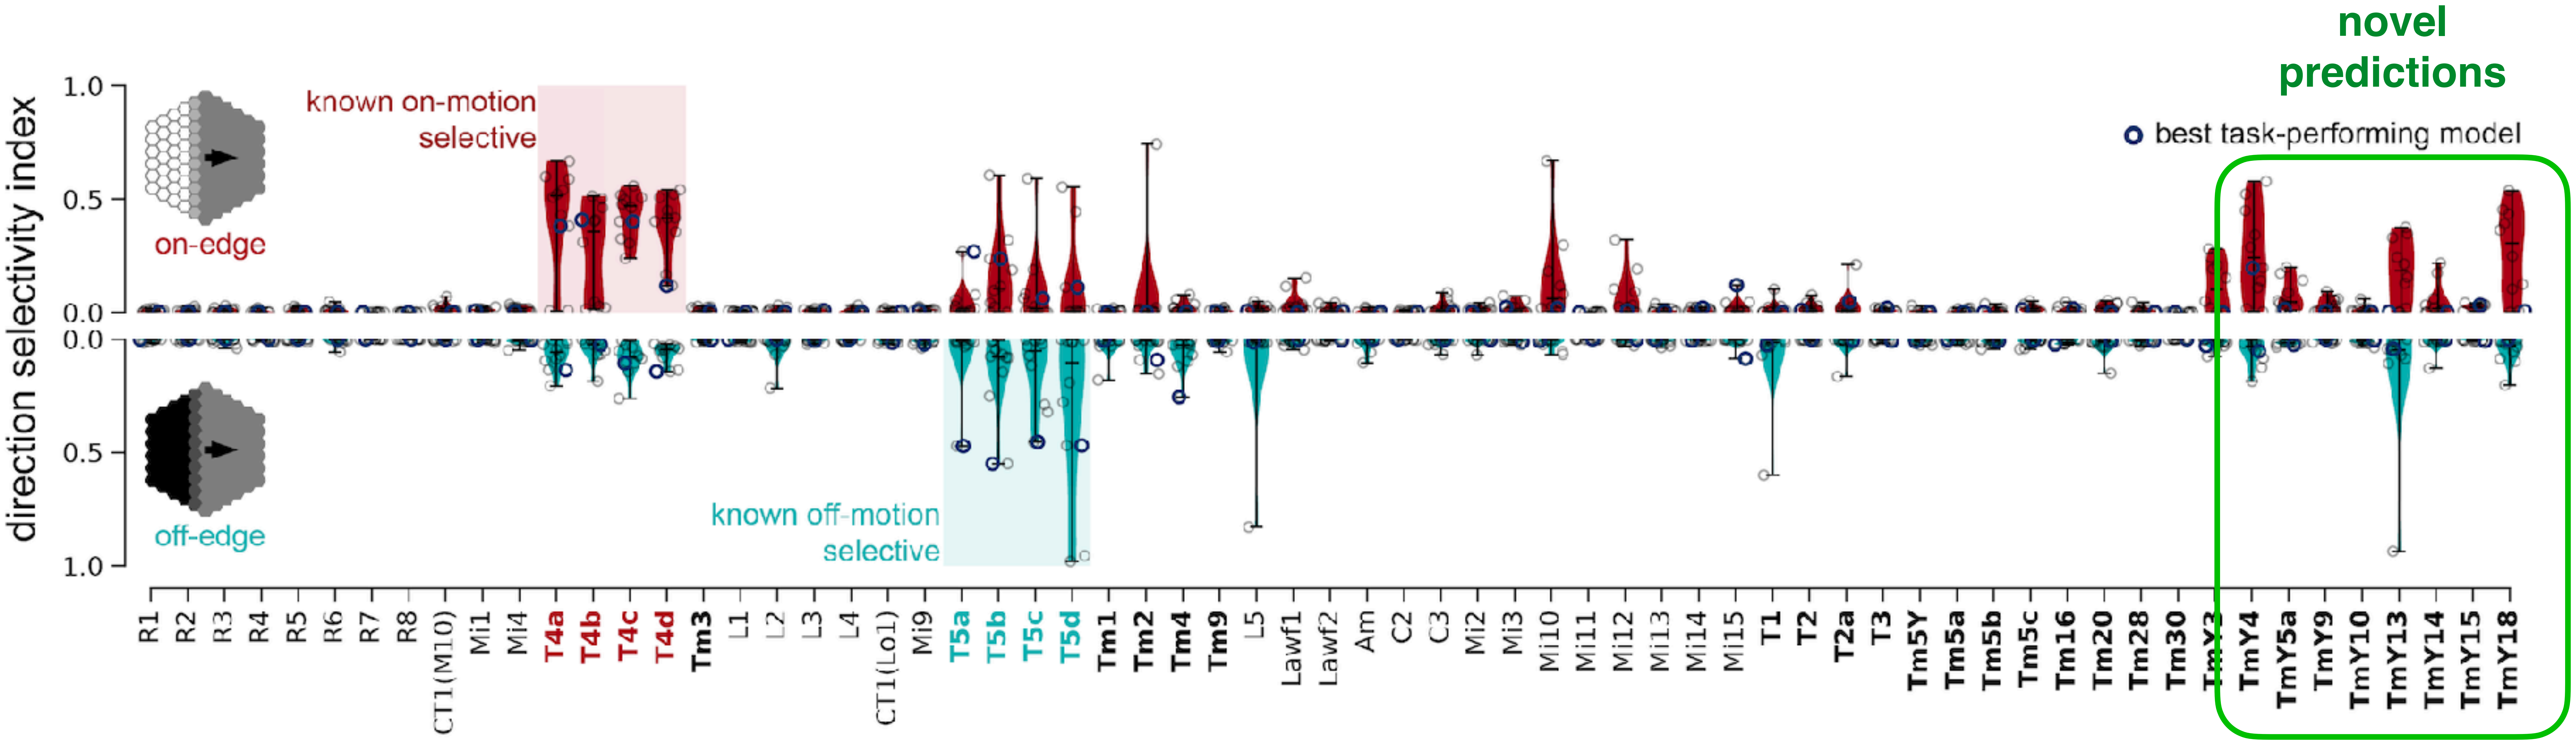

In-silico neurophysiology on the model: T4 and T5 cells

Simulated responses

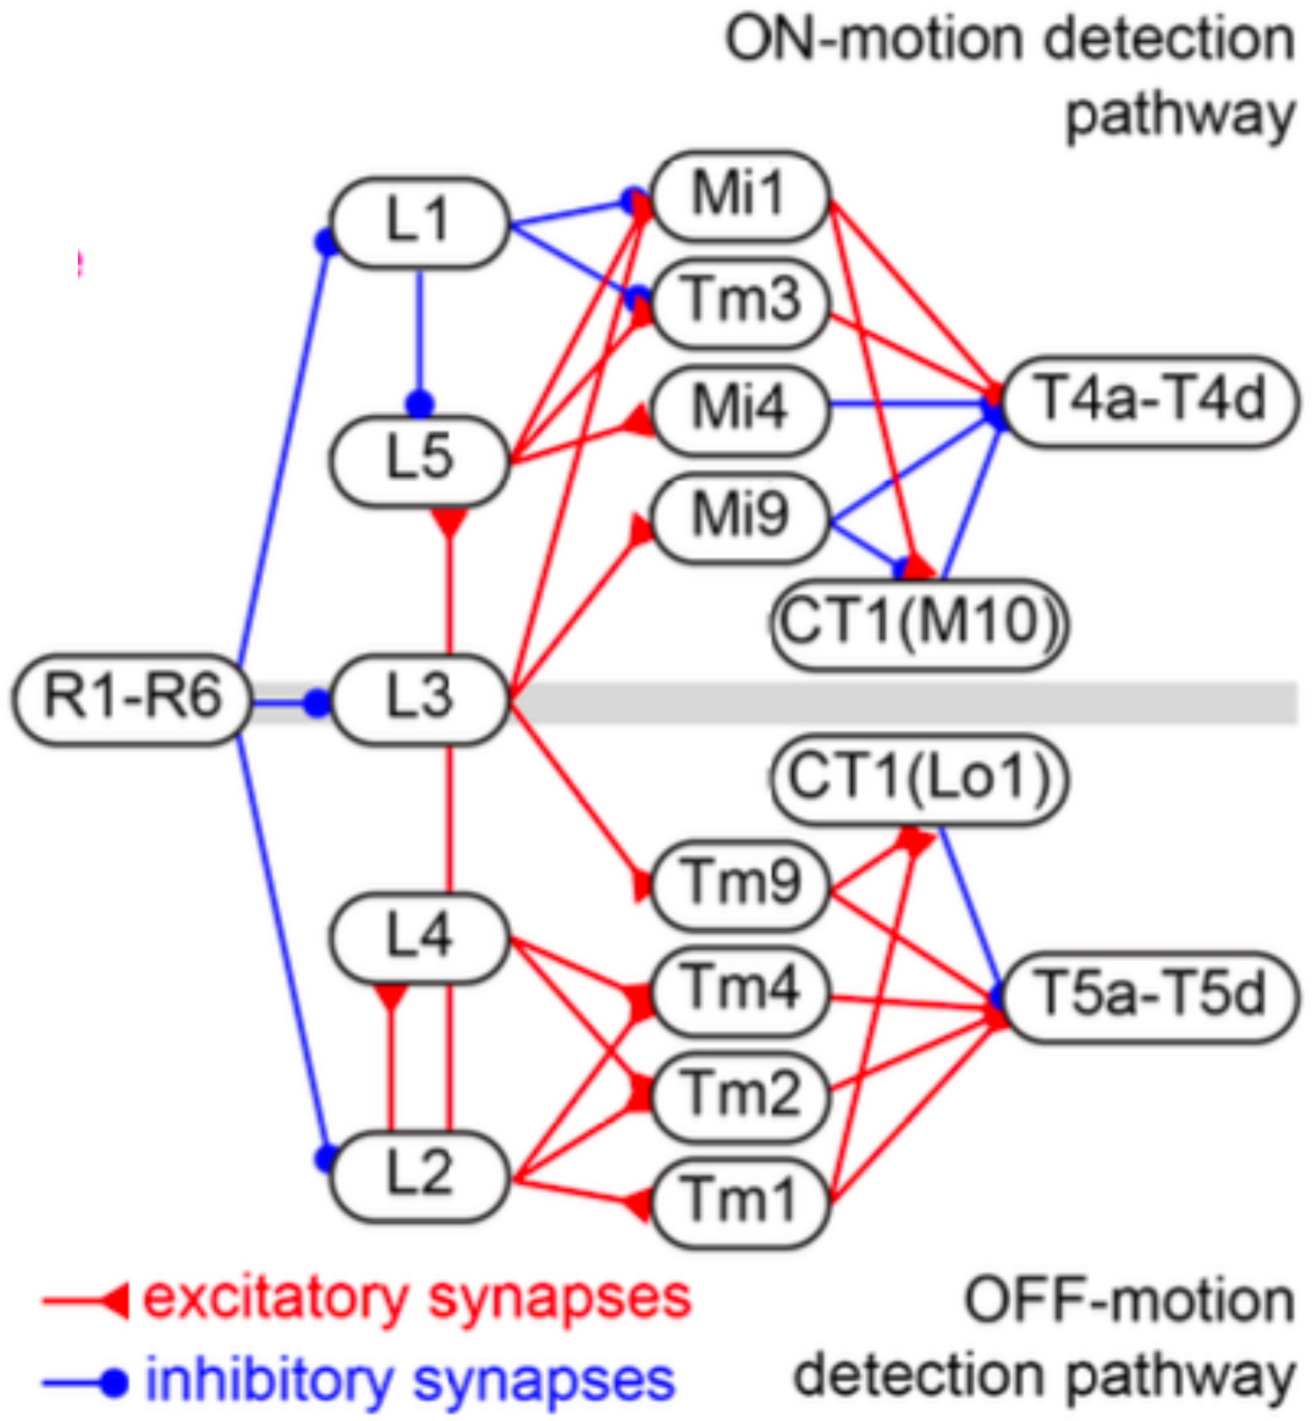

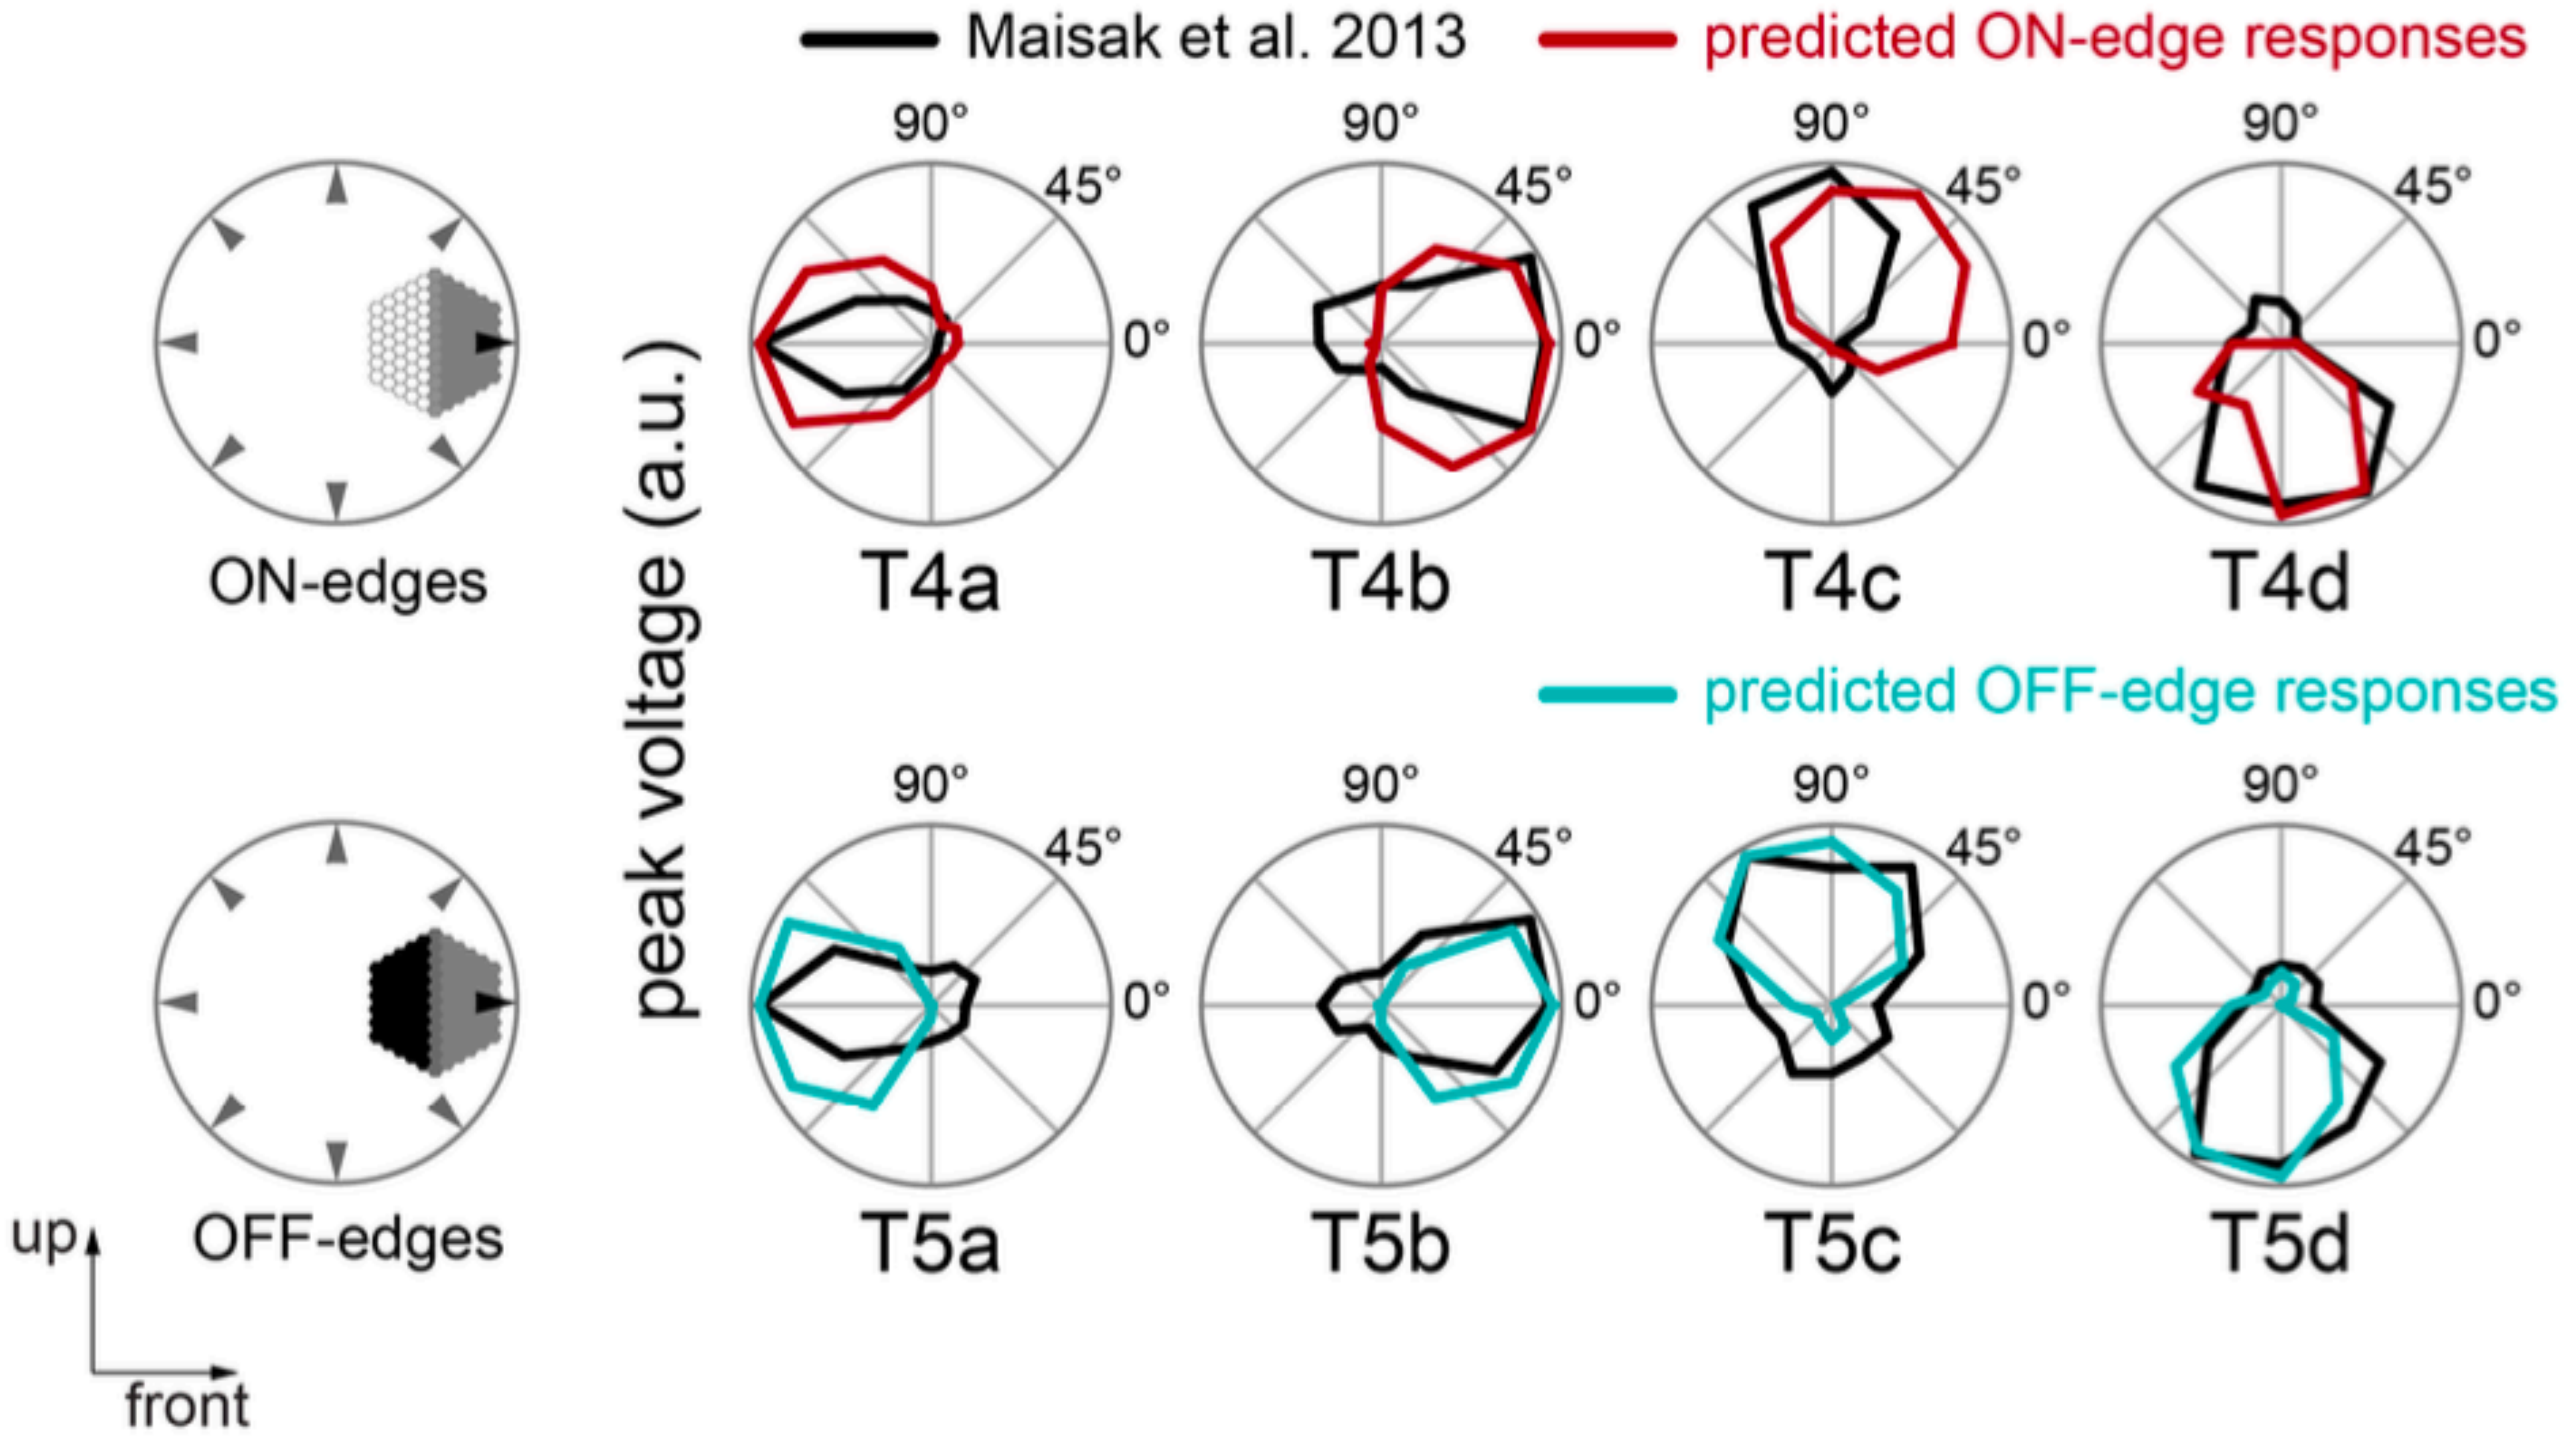

Maisak et al 2013, Fisher & Clandinin 2015, Gruntmann et al 2018., Grundmann et al 2019.

Application III: Learning to invert a (model of a) microscope to make better images, faster

Data from Legant. et al, 2016

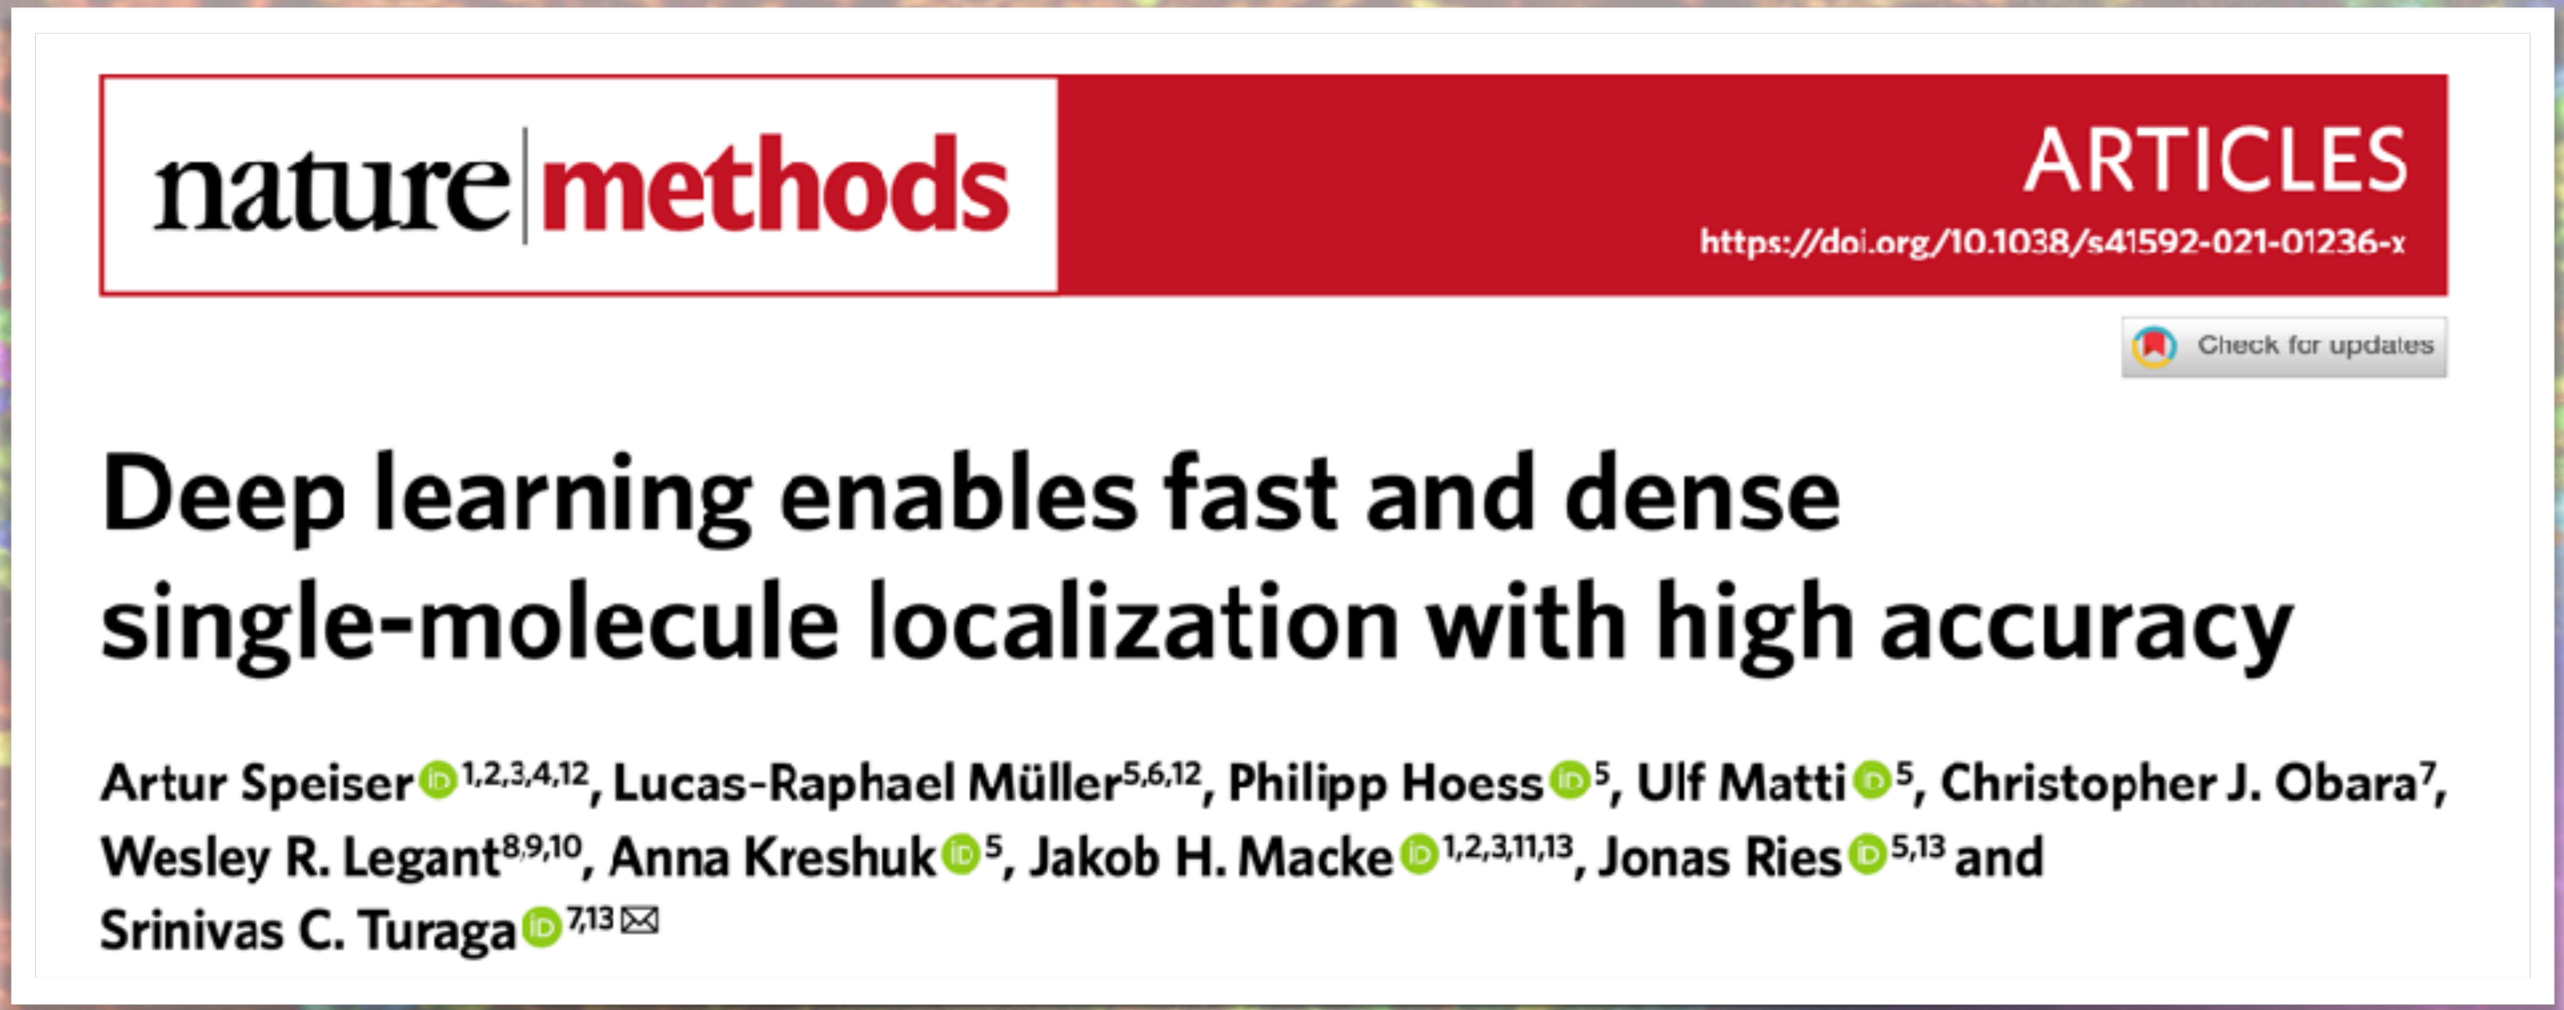

nature methods

ARTICLES

https://doi.org/10.1038/s41592-021-01236-x

Check for updates

Deep learning enables fast and dense single-molecule localization with high accuracy

Artur Speiser[1,2,3,4,12], Lucas-Raphael Müller[5,6,12], Philipp Hoess[5], Ulf Matti[5], Christopher J. Obara[7], Wesley R. Legant[8,9,10], Anna Kreshuk[5], Jakob H. Macke[1,2,3,11,13], Jonas Ries[5,13] and Srinivas C. Turaga[7,13] ✉

Forward model of imaging process

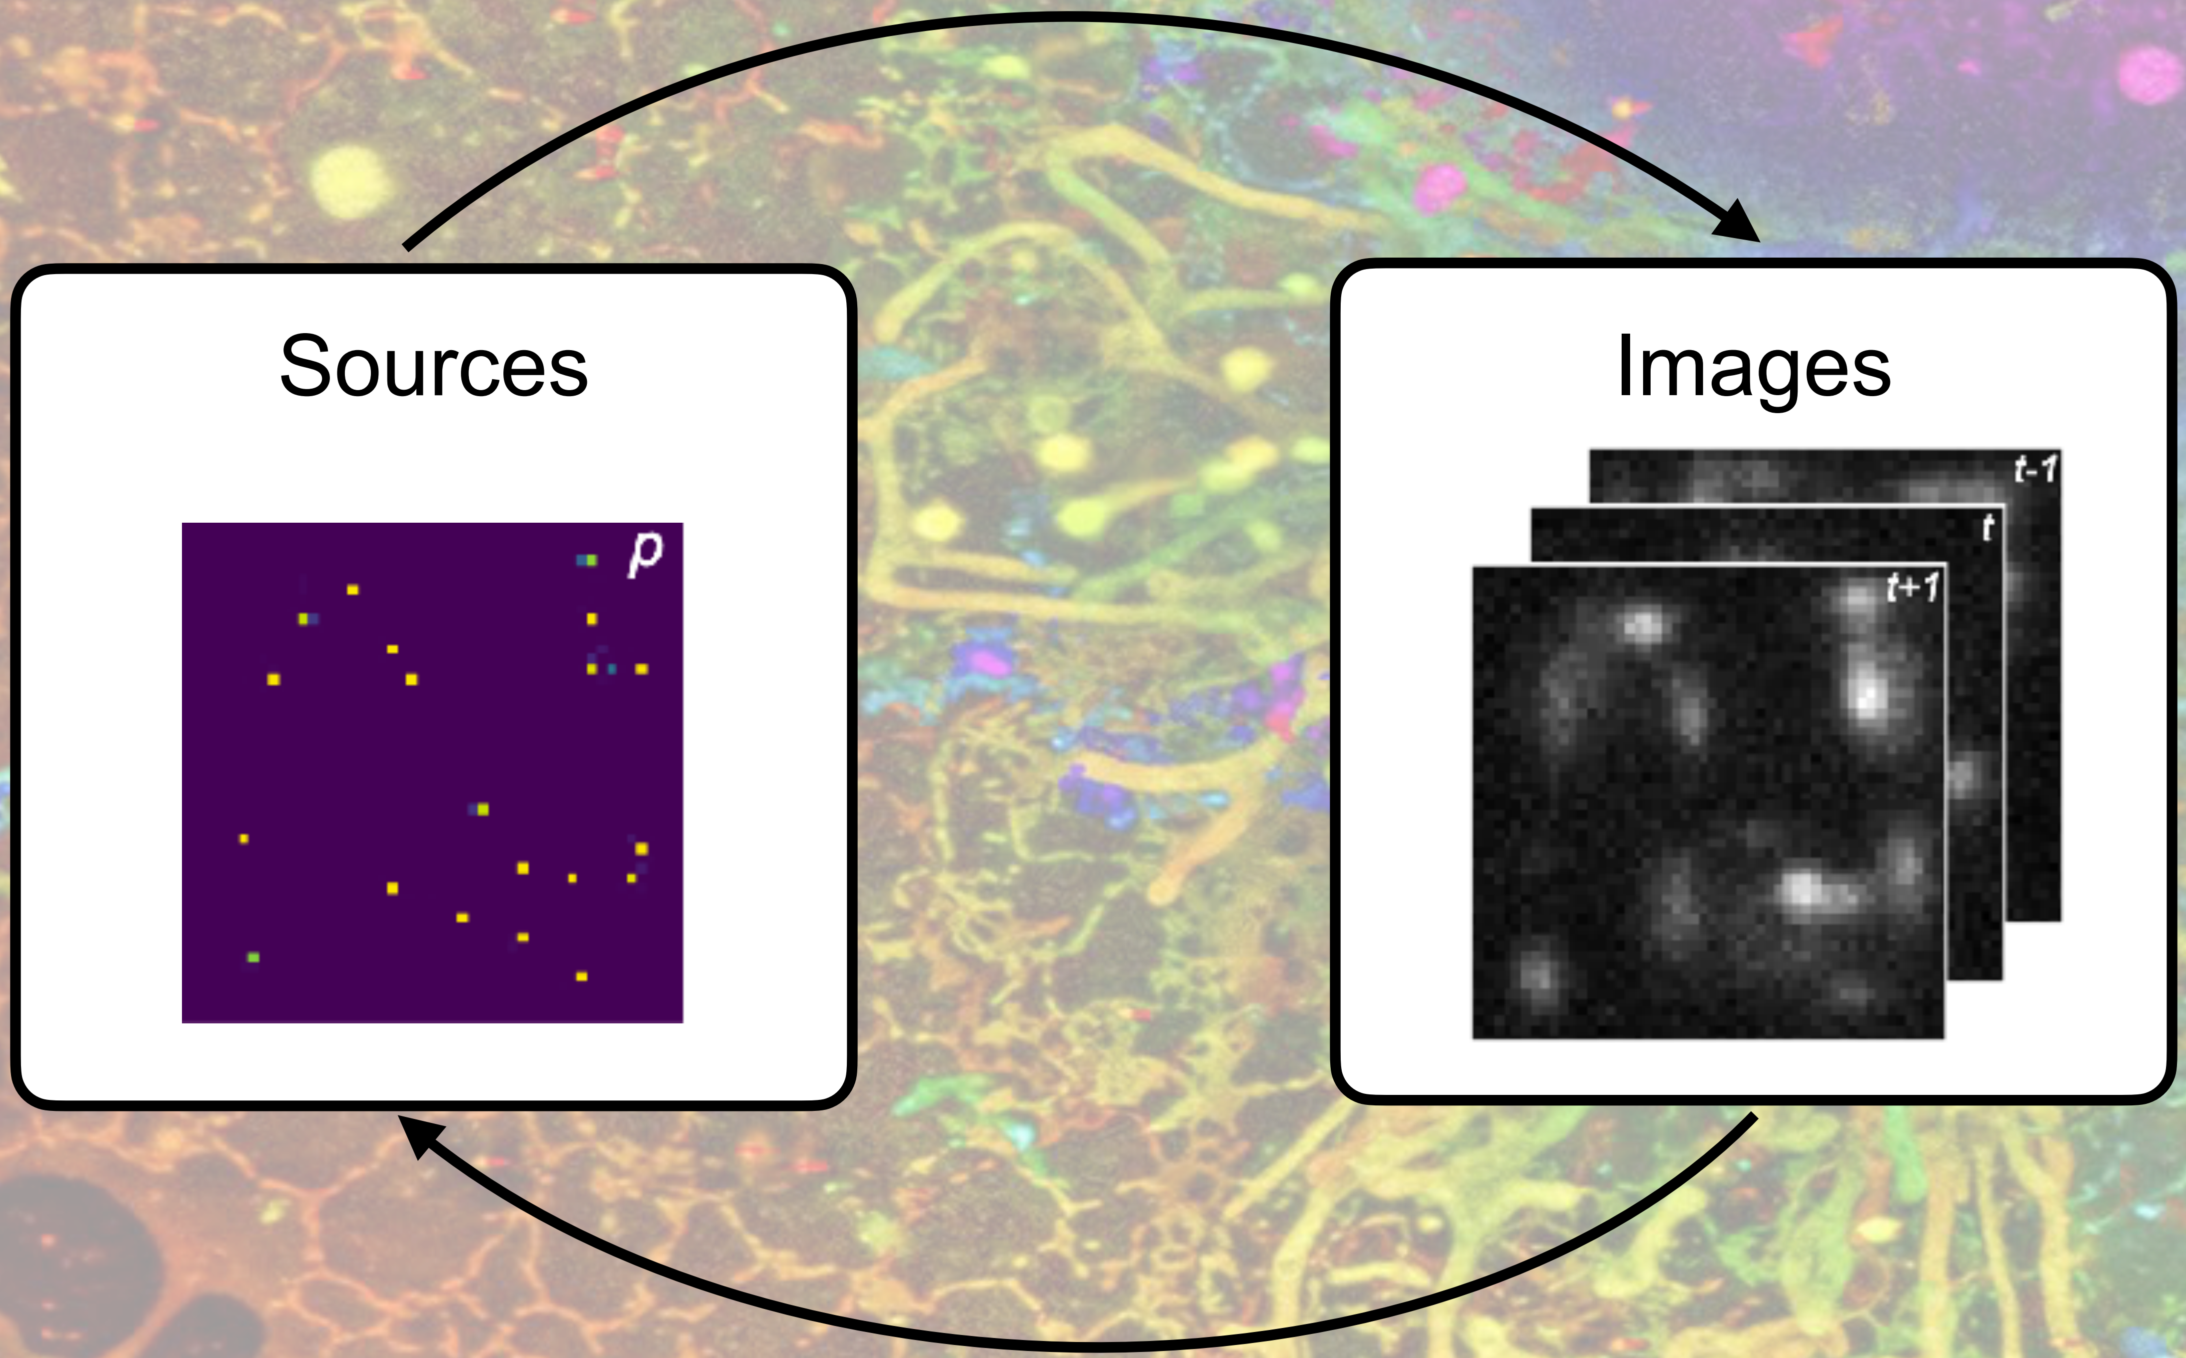

Sources

Images

Inference network ('DECODE')

A Speiser

L Müller

S Turaga

J Ries

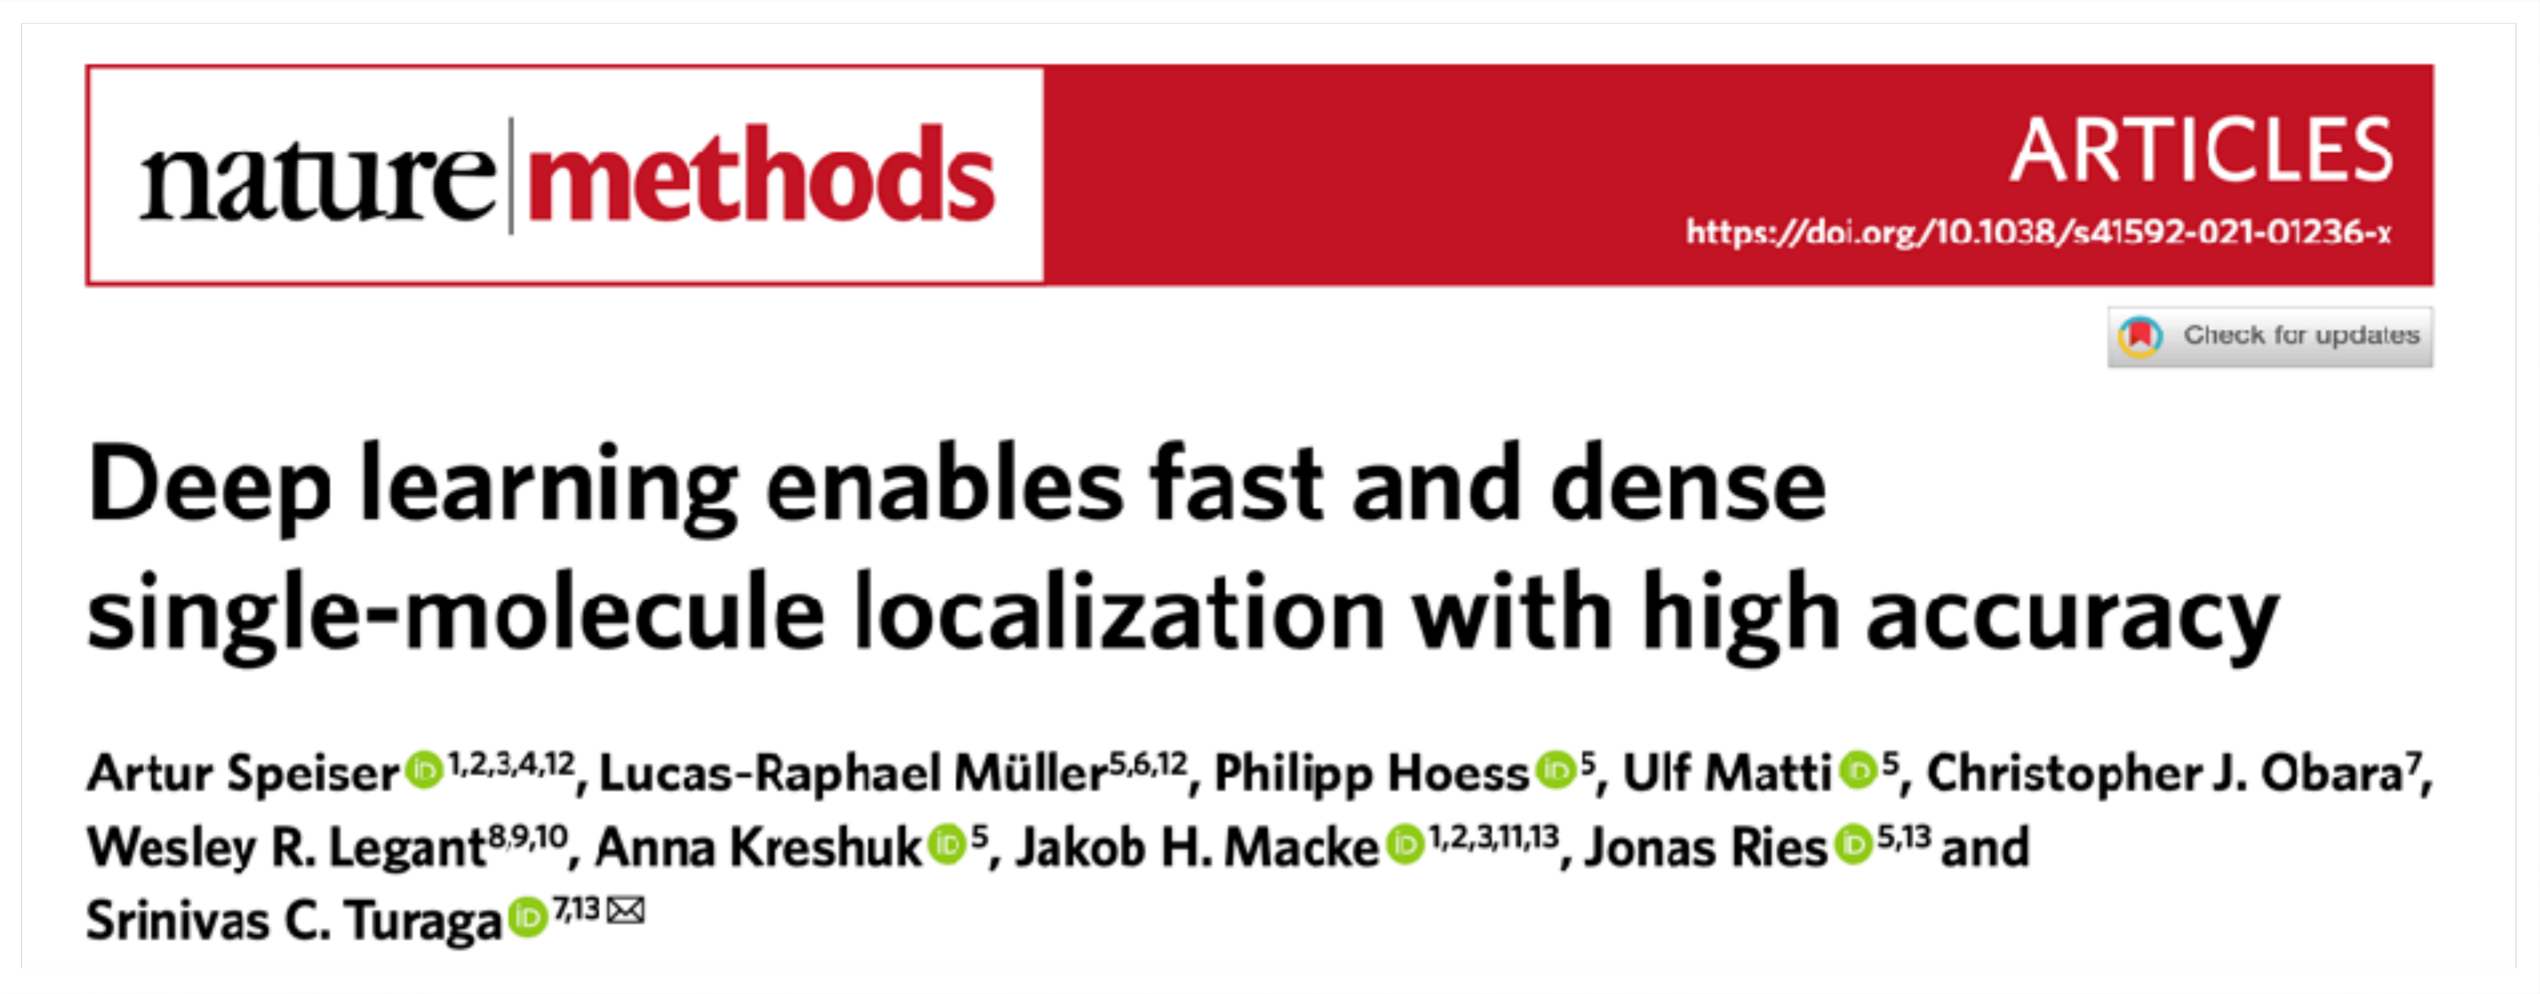

nature methods ARTICLES

https://doi.org/10.1038/s41592-021-01236-x

Check for updates

Deep learning enables fast and dense single-molecule localization with high accuracy

Artur Speiser[1,2,3,4,12], Lucas-Raphael Müller[5,6,12], Philipp Hoess[5], Ulf Matti[5], Christopher J. Obara[7], Wesley R. Legant[8,9,10], Anna Kreshuk[5], Jakob H. Macke[1,2,3,11,13], Jonas Ries[5,13] and Srinivas C. Turaga[7,13] ✉

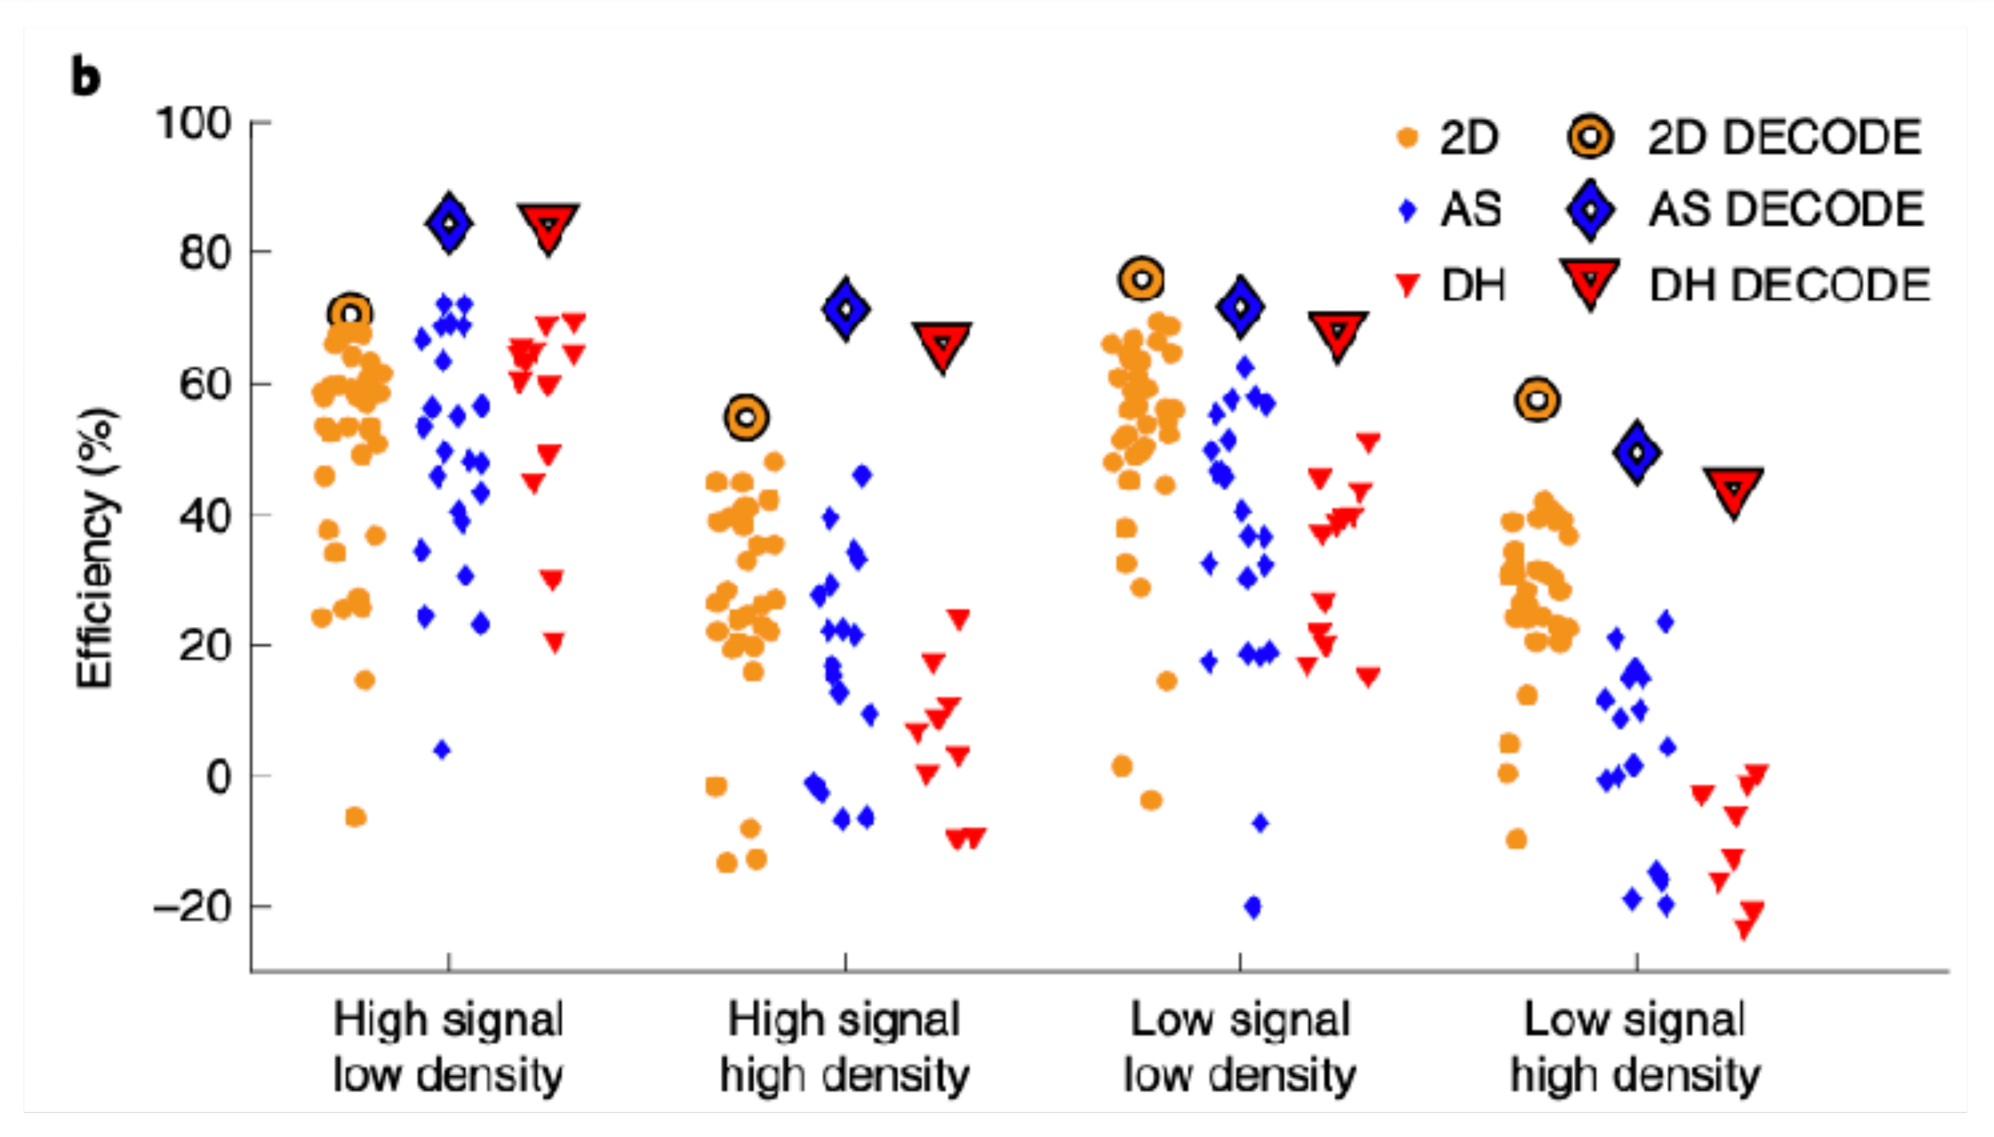

SMLM Challenge, Sage et al 2019

Problem: How to represent posterior distributions over an unknown number of points?

A Speiser

L Müller

S Turaga

J Ries

Michael Deistler

Maximilian Dax
(MPI-IS, supervised by
Bernhard Schölkopf)

Thanks to ...

Giacomo Bassetto
Sebastian Bischoff
Jan Boelts
Michael Deistler
Dr. Richard Gao
Manuel Glöckler
Dr. Pedro Goncalves
Jaivardhan Kapoor
Janne Lappalainen
Guy Moss
Lucas Raphael Müller
Matthijs Pals
Rachel Rapp
Dr. Cornelius Schröder
Zinovia Stefanidi
Auguste Schulz
Julius Vetter

Franziska Weiler
Eszter Stuber
Tharanika Thevururasa
Alana Darcher
Max Dax
Franziska Gerken
Dr. Stefanie Liebe
Philipp von Bachmann

machine learning
new perspectives
for science

DFG

Tübingen AI
Center

Bundesministerium
für Bildung
und Forschung

FENS KAVLI
NETWORK OF EXCELLENCE

erc

HFSP
Human Frontier
Science Program

SFB 1089

SPP 2041
Computational
Connectomics

ROBUST VISION

Open Positions!
(Postdoc + PhD Students)

16 October 2023 | Tübingen

6 Principal Investigators (m/f/d) as Hector Endowed Fellows of the ELLIS Institute Tübingen

The new ELLIS Institute in Tübingen, Germany, invites applications for positions of up to 6 Principal Investigators (m/f/d) as Hector Endowed Fellows of the ELLIS Institute Tübingen.

2022 heinlewischer

Tübingen AI Center

Multiple positions for (full) professors watch out for calls!

ellis

INSTITUTE TÜBINGEN

machine learning
new perspectives
for science