Learning interpretable continuous-time models of latent stochastic dynamical systems

Lea Duncker, Gergő Bohner, Julien Boussard, Maneesh Sahani

Gatsby Computational Neuroscience Unit
University College London

ICML

June 12, 2019

nonlinear stochastic dynamical system

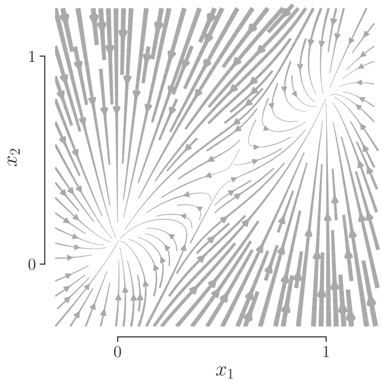

nonlinear stochastic dynamical system

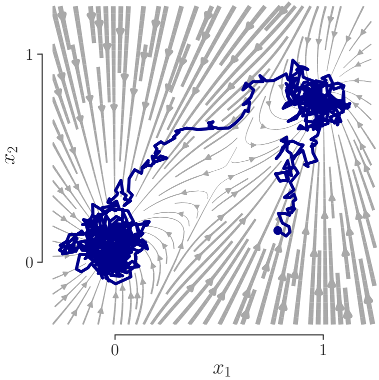

$$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma}d\boldsymbol{W}$$

nonlinear stochastic dynamical system

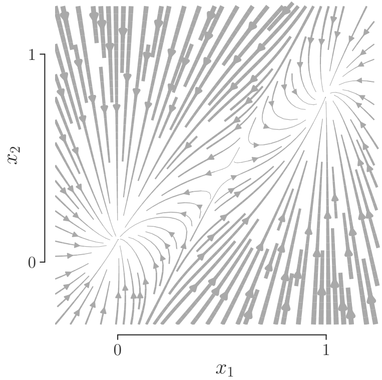

$$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma}d\boldsymbol{W}$$

nonlinear stochastic dynamical system

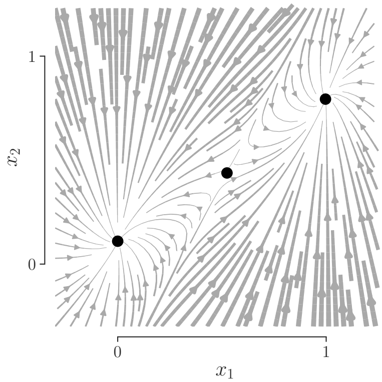

$$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma}d\boldsymbol{W}$$

$$\boldsymbol{f}(\boldsymbol{s}) = \boldsymbol{0}$$

fixed point

nonlinear stochastic dynamical system

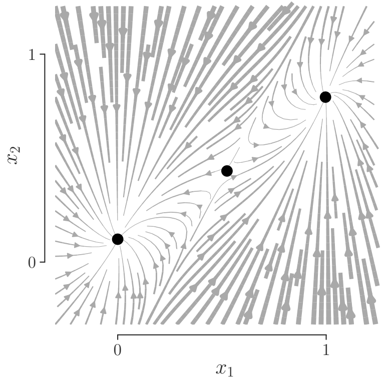

$$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma}d\boldsymbol{W}$$

$\boldsymbol{f}(\boldsymbol{s}) = \boldsymbol{0}$ **fixed point**

$$\boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{s}) + \nabla_x \boldsymbol{f}(x)|_{x=s}(\boldsymbol{x} - \boldsymbol{s}) + \ldots$$

nonlinear stochastic dynamical system

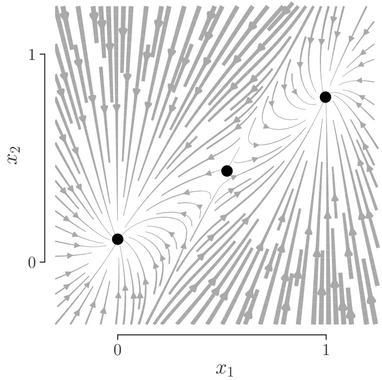

$$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma}d\boldsymbol{W}$$

$$\boldsymbol{f}(\boldsymbol{s}) = \boldsymbol{0}$$ **fixed point**

$$\boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{s}) + \nabla_x \boldsymbol{f}(x)|_{x=s}(\boldsymbol{x} - \boldsymbol{s}) + \ldots$$

$$\approx \boldsymbol{J}(\boldsymbol{x} - \boldsymbol{s})$$ **Jacobian matrix**

nonlinear stochastic dynamical system

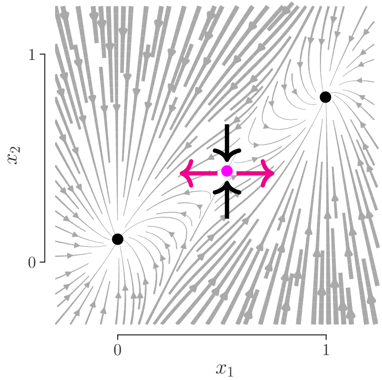

$$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma}d\boldsymbol{W}$$

$\boldsymbol{f}(\boldsymbol{s}) = \boldsymbol{0}$ **fixed point**

$$\boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{s}) + \nabla_x \boldsymbol{f}(x)|_{x=s}(\boldsymbol{x} - \boldsymbol{s}) + \dots$$

$$\approx \boldsymbol{J}(\boldsymbol{x} - \boldsymbol{s})$$ **Jacobian matrix**

nonlinear stochastic dynamical system

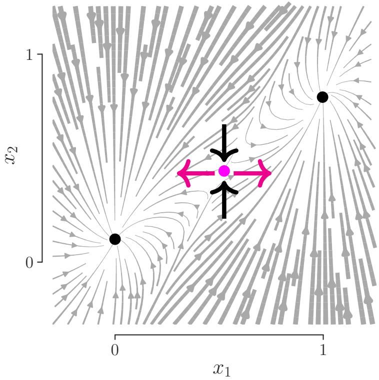

$$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma} d\boldsymbol{W}$$

$\boldsymbol{f}(\boldsymbol{s}) = \boldsymbol{0}$ **fixed point**

$$\boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{s}) + \nabla_x \boldsymbol{f}(x)|_{x=s}(\boldsymbol{x} - \boldsymbol{s}) + \ldots$$

$\approx J(\boldsymbol{x} - \boldsymbol{s})$ **Jacobian matrix**

interpretability:

- **stability analysis**
- **locally linearised dynamics**
- **. . .**

unevenly sampled high-d observations

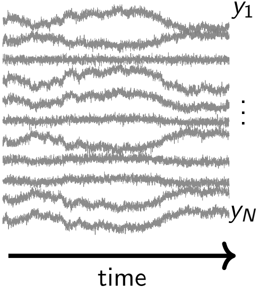

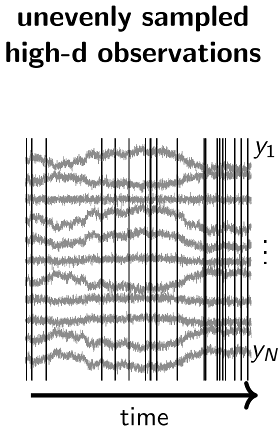

unevenly sampled
high-d observations
y_1
y_N
time

unevenly sampled high-d observations

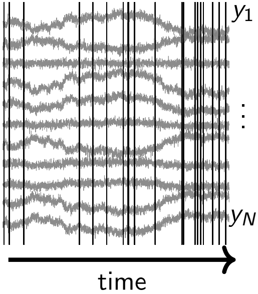

$$\langle \boldsymbol{y}(t_i) \rangle = g(C\boldsymbol{x}(t_i) + \boldsymbol{d})$$

latent low-d stochastic process | **unevenly sampled high-d observations**

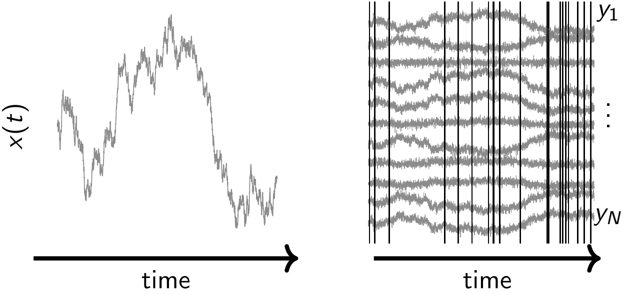

$$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma}d\boldsymbol{W} \qquad \langle \boldsymbol{y}(t_i) \rangle = g(C\boldsymbol{x}(t_i) + \boldsymbol{d})$$

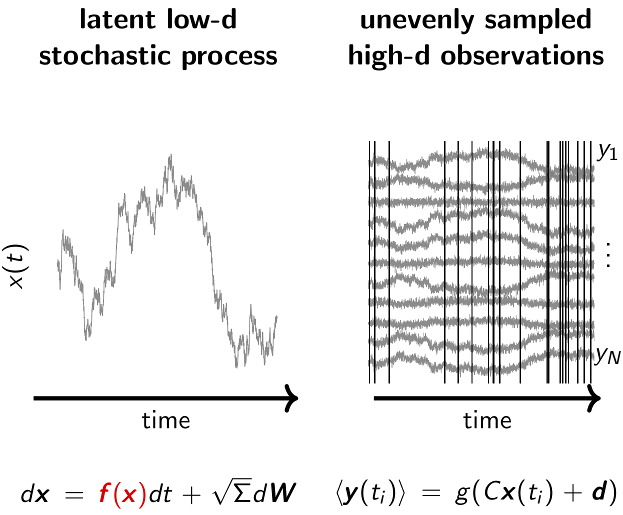

latent low-d
stochastic process
unevenly sampled
high-d observations
$x(t)$
y_1
y_N
time
time
$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma}d\boldsymbol{W}$
$\langle \boldsymbol{y}(t_i) \rangle = g(C\boldsymbol{x}(t_i) + \boldsymbol{d})$

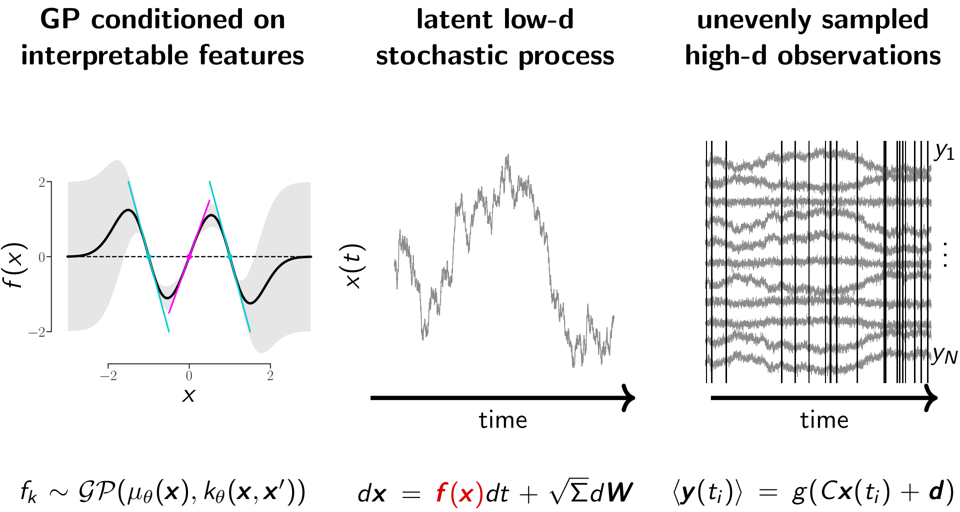
GP conditioned on interpretable features
latent low-d stochastic process
unevenly sampled high-d observations
f(x)
x
x(t)
time
time
y_1
y_N
$f_k \sim \mathcal{GP}(\mu_\theta(\boldsymbol{x}), k_\theta(\boldsymbol{x}, \boldsymbol{x}'))$
$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x})dt + \sqrt{\Sigma}d\boldsymbol{W}$
$\langle \boldsymbol{y}(t_i) \rangle = g(C\boldsymbol{x}(t_i) + \boldsymbol{d})$

$$q(\boldsymbol{x}, \boldsymbol{f}) = q_x(\boldsymbol{x}) \; q_f(\boldsymbol{f})$$

Variational Bayes

Archambeau (2007), Titsias (2009)

Variational Bayes

$$q(\boldsymbol{x}, \boldsymbol{f}) = q_x(\boldsymbol{x})\ q_f(\boldsymbol{f})$$

Archambeau (2007), Titsias (2009)

$$q(\boldsymbol{x}, \boldsymbol{f}) = q_x(\boldsymbol{x})\ q_f(\boldsymbol{f})$$

Variational Bayes

Gaussian Process Dynamics

$$= \int P\left(\right.$$ 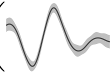 $$\left. \middle| \boldsymbol{u}, \boldsymbol{\theta} \right) q_u(\boldsymbol{u})\, d\boldsymbol{u}$$

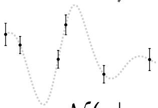

$\mathcal{N}(\boldsymbol{u} | \boldsymbol{m}_u, \boldsymbol{S}_u)$

sparse approx.
with inducing variables

Archambeau (2007), Titsias (2009)

$$q(\boldsymbol{x}, \boldsymbol{f}) = q_x(\boldsymbol{x})\; q_f(\boldsymbol{f})$$

Variational Bayes

Latent SDE path

$$d\boldsymbol{x} = (-A(t)\boldsymbol{x} + \boldsymbol{b}(t))dt + \sqrt{\Sigma} d\boldsymbol{W}$$

$$q(\boldsymbol{x}(t)) = \mathcal{N}\left(\boldsymbol{x}(t) | \boldsymbol{m}_x(t), \boldsymbol{S}_x(t)\right)$$

$$\dot{\boldsymbol{m}}_x = -A(t)\boldsymbol{m}_x + \boldsymbol{b}(t)$$

$$\dot{\boldsymbol{S}}_x = -A(t)\boldsymbol{S}_x - \boldsymbol{S}_x A(t)^{\mathsf{T}} + \Sigma$$

Gaussian approx.
with Markov structure

Gaussian Process Dynamics

$$= \int P\Big(\quad \Big| \boldsymbol{u}, \boldsymbol{\theta} \Big)\, q_u(\boldsymbol{u})\, d\boldsymbol{u}$$

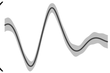

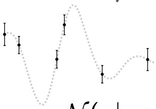

$\mathcal{N}(\boldsymbol{u} | \boldsymbol{m}_u, \boldsymbol{S}_u)$

sparse approx.
with inducing variables

Archambeau (2007), Titsias (2009)

Example: Van der Pol's Oscillator

dynamics: $f_1(\boldsymbol{x}) = 2\tau(x_1 - \frac{1}{3}x_1^3 - x_2)$

$f_2(\boldsymbol{x}) = \frac{\tau}{2}\, x_1$

Example: Van der Pol's Oscillator

dynamics: $$f_1(\boldsymbol{x}) = 2\tau(x_1 - \tfrac{1}{3}x_1^3 - x_2)$$
$$f_2(\boldsymbol{x}) = \tfrac{\tau}{2}\, x_1$$

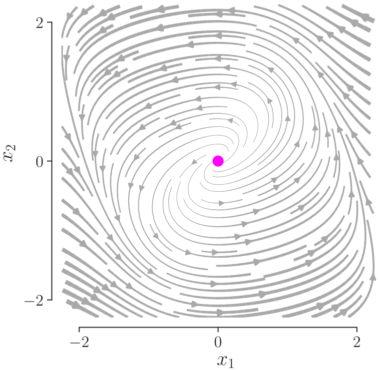

true dynamics

Example: Van der Pol's Oscillator

dynamics: $f_1(\boldsymbol{x}) = 2\tau(x_1 - \frac{1}{3}x_1^3 - x_2)$

$f_2(\boldsymbol{x}) = \frac{\tau}{2}\, x_1$

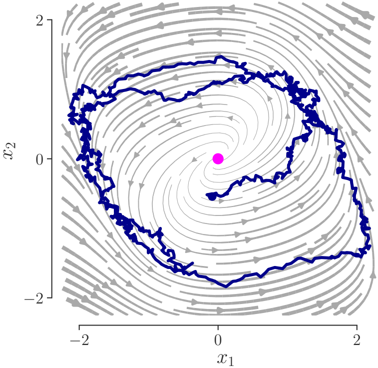

true dynamics

Example: Van der Pol's Oscillator

dynamics: $f_1(\boldsymbol{x}) = 2\tau(x_1 - \frac{1}{3}x_1^3 - x_2)$

$f_2(\boldsymbol{x}) = \frac{\tau}{2}\, x_1$

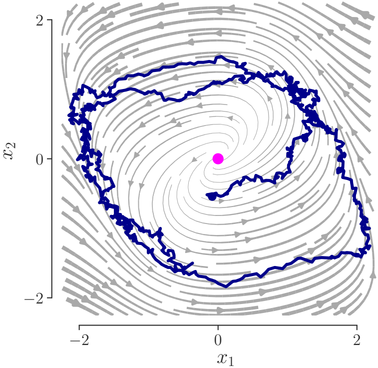

true dynamics

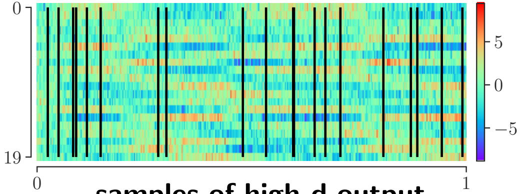

samples of high-d output on 20 trials

Example: Van der Pol's Oscillator

dynamics: $f_1(\boldsymbol{x}) = 2\tau(x_1 - \frac{1}{3}x_1^3 - x_2)$

$f_2(\boldsymbol{x}) = \frac{\tau}{2}\, x_1$

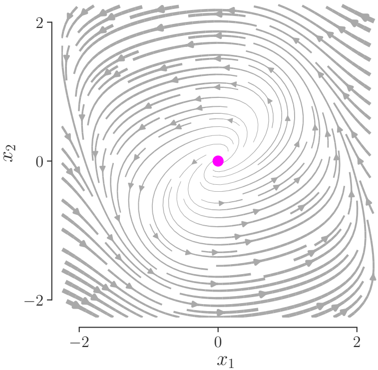

true dynamics

Example: Van der Pol's Oscillator

dynamics: $f_1(\boldsymbol{x}) = 2\tau(x_1 - \frac{1}{3}x_1^3 - x_2)$

$f_2(\boldsymbol{x}) = \frac{\tau}{2}\, x_1$

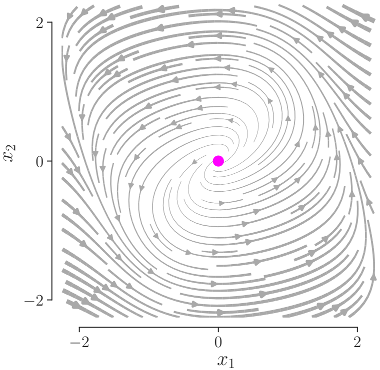

true dynamics

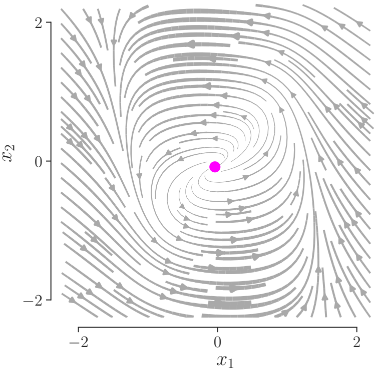

learnt dynamics

Example: Van der Pol's Oscillator

dynamics: $f_1(\boldsymbol{x}) = 2\tau(x_1 - \frac{1}{3}x_1^3 - x_2)$

$f_2(\boldsymbol{x}) = \frac{\tau}{2}\, x_1$

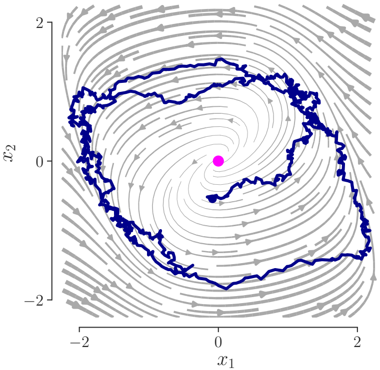

true dynamics

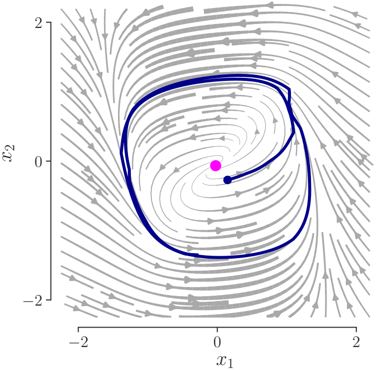

learnt dynamics

limit cycles

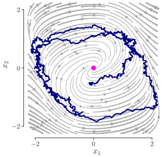

true dynamics

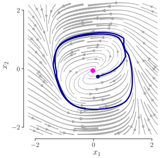

learnt dynamics

limit cycles

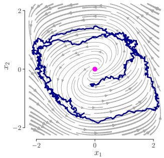

true dynamics

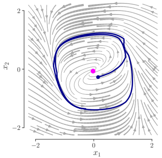

learnt dynamics

double-well dynamics

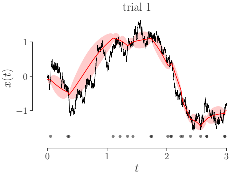

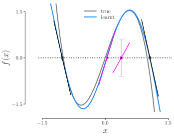

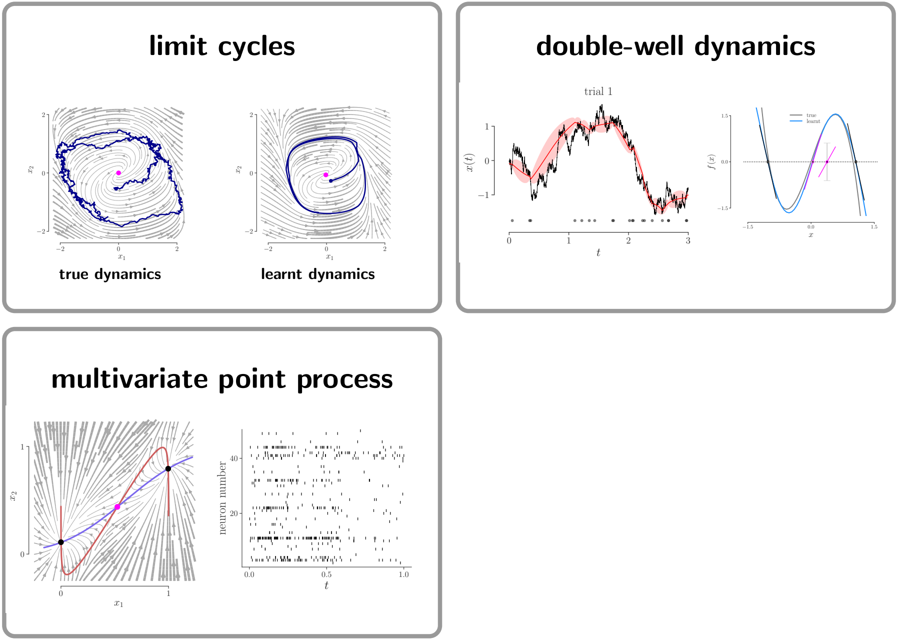
limit cycles
true dynamics
learnt dynamics
double-well dynamics
trial 1
$x(t)$
t
true
learnt
$f(x)$
x
multivariate point process
neuron number
x_1
x_2

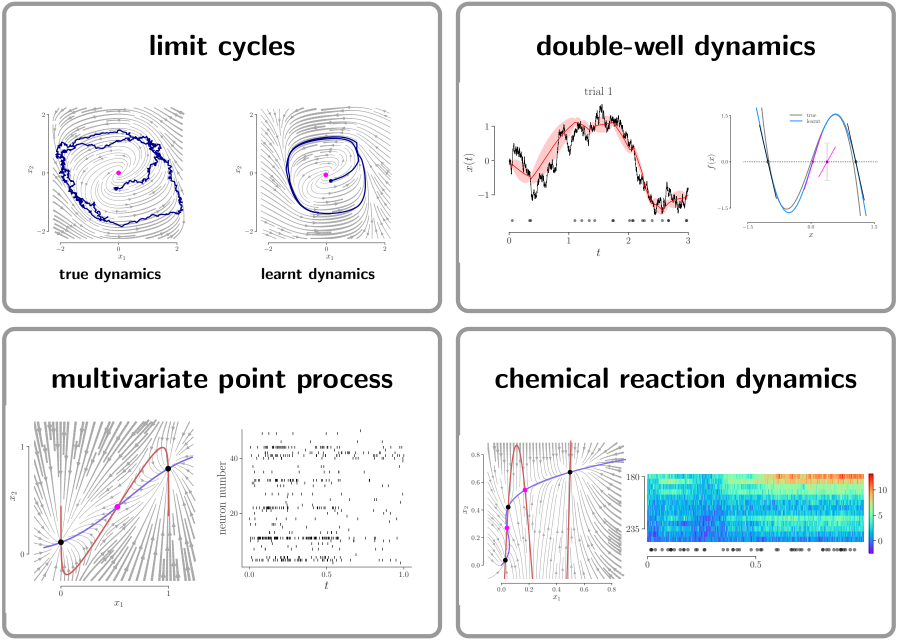

limit cycles
true dynamics
learnt dynamics
double-well dynamics
trial 1
true
learnt
multivariate point process
neuron number
chemical reaction dynamics

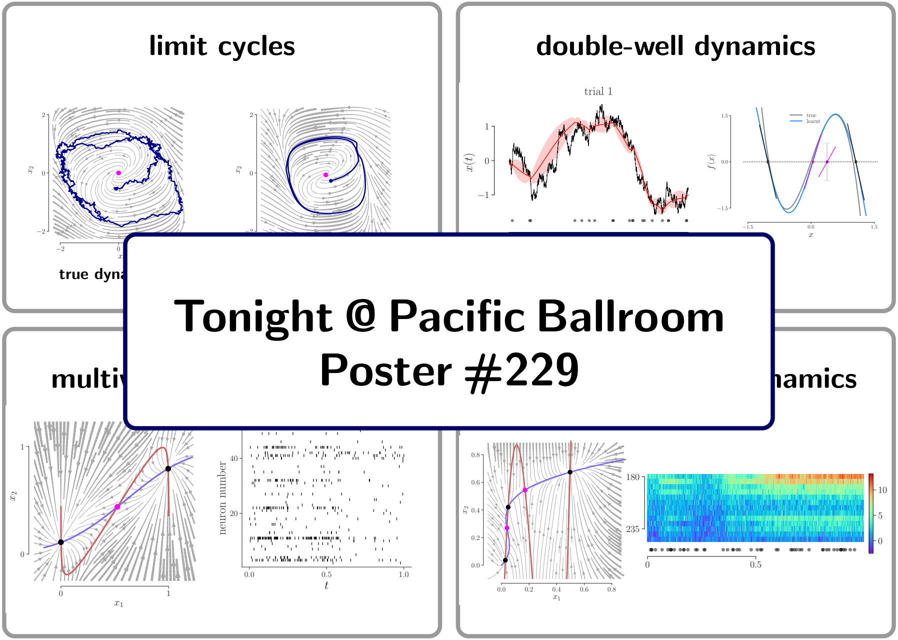

limit cycles
double-well dynamics
trial 1
true
learnt
Tonight @ Pacific Ballroom
Poster #229
neuron number

Gergő Bohner

Julien Boussard

Maneesh Sahani

GATSBY

SIMONS FOUNDATION