Learning summary features of time series for likelihood-free inference

Pedro L. C. Rodrigues and Alexandre Gramfort, INRIA, France – pedro.rodrigues@inria.fr

1 CONTEXT

Likelihood-free inference (LFI) methods make **Bayesian inference** on modern physical simulators possible.

Most algorithms use a set of **handcrafted summary features** to describe data.

We propose a **data-driven** approach that learns the most appropriate summary features for LFI on time series data generated by linear and non-linear models

2 PROBLEM FORMULATION

We have a simulator parametrized by $\boldsymbol{\theta}$ that generates time series $\boldsymbol{x}(t)$

Our goal is to determine an approximation to the posterior distribution of the parameters for a given observed time series:

$$p(\boldsymbol{\theta}|\boldsymbol{x}_o) \approx q_{\boldsymbol{\phi}}(\boldsymbol{\theta}|f_{\boldsymbol{\lambda}}(\boldsymbol{x}_o))$$

$\boldsymbol{\phi}$ parametrizes the density estimator

$\boldsymbol{\lambda}$ parametrizes the feature extractor

3 OUR PROPOSAL

We present the **YuleNet**, a two-layer neural network based on temporal convolutions.

The architecture comes from [1] and has been applied for classifying EEG sleep signals.

We **jointly** estimate ϕ and $\boldsymbol{\lambda}$ by minimizing

$$\mathcal{L}(\boldsymbol{\phi},\boldsymbol{\lambda}) = \mathbb{E}_{(\boldsymbol{x},\boldsymbol{\theta})\sim p(\boldsymbol{x},\boldsymbol{\theta})}\Big[-\log(q_{\boldsymbol{\phi}}(\boldsymbol{\theta}|f_{\boldsymbol{\lambda}}(\boldsymbol{x})))\Big]$$

with a multiround procedure (SNPE-C) [2].

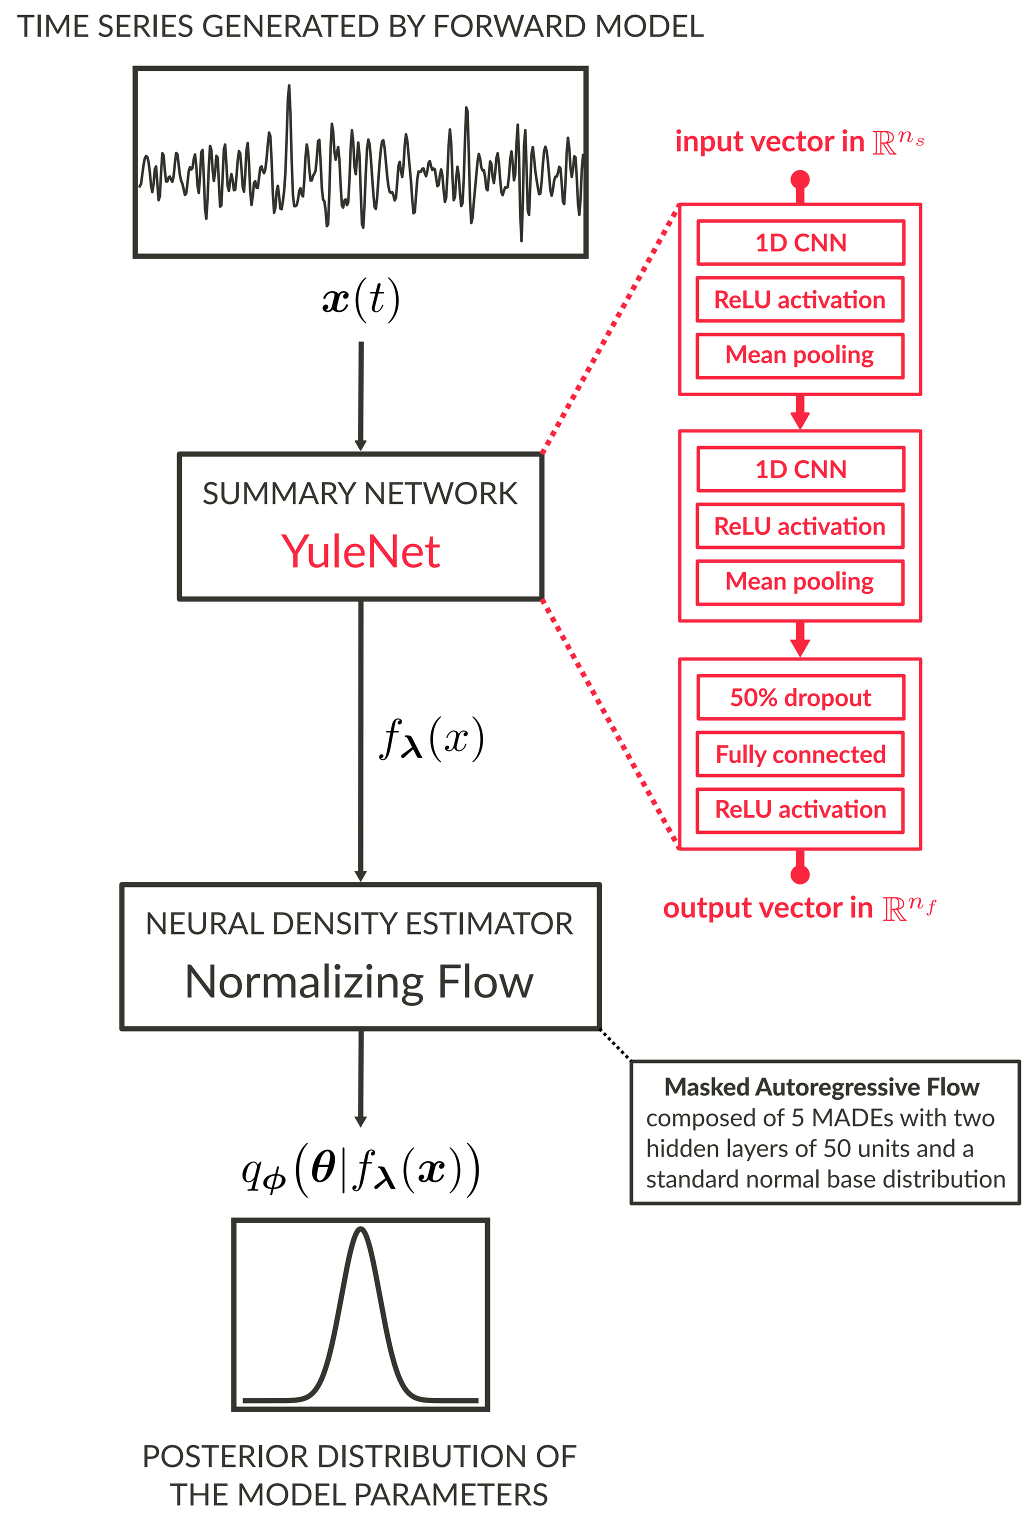

4 NUMERICAL EXPERIMENTS

(1) **Validation** on simulations with linear Gaussian time series models and comparison with inference based on **autocorrelations** as summary features.

↳ Uniformly better results with data driven features than handcrafted ones

(2) **Comparison** of YuleNet to *Partially Exchangeable Networks* (PEN) from [3].

↳ Equivalent results on the Gaussian example but YuleNet has less parameters and demands less computations

(3) Evaluation on a **non-linear** dynamical system and comparison of autocorrelations to YuleNet

↳ Autocorr gives poor results whereas YuleNet finds the ground truth parameters on every situation

The stochastic Van der Pol Oscillator

$$\boxed{\ddot{x} = \varepsilon(1 - x^2)\dot{x} - x + \sigma\dot{w}}$$

Observed time series: $x(t)$

($\dot{w}(t)$ is an input Gaussian white noise)

Parameters to infer:

- σ is the variance of the stochastic input
- ε is the degree of non-linearity of the VdP model

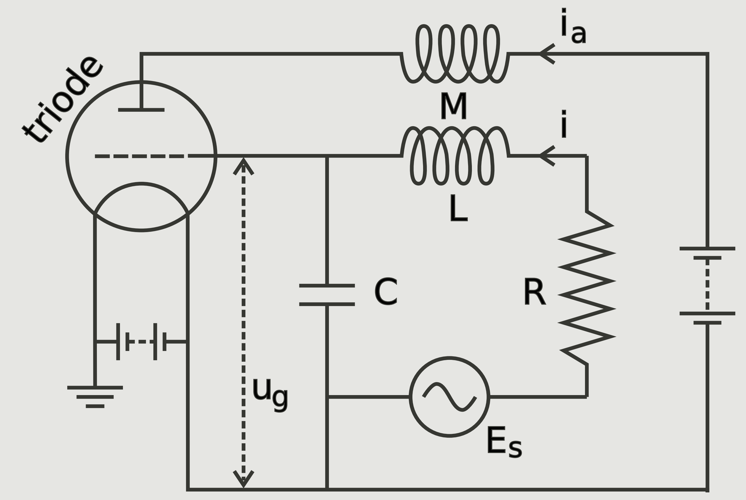

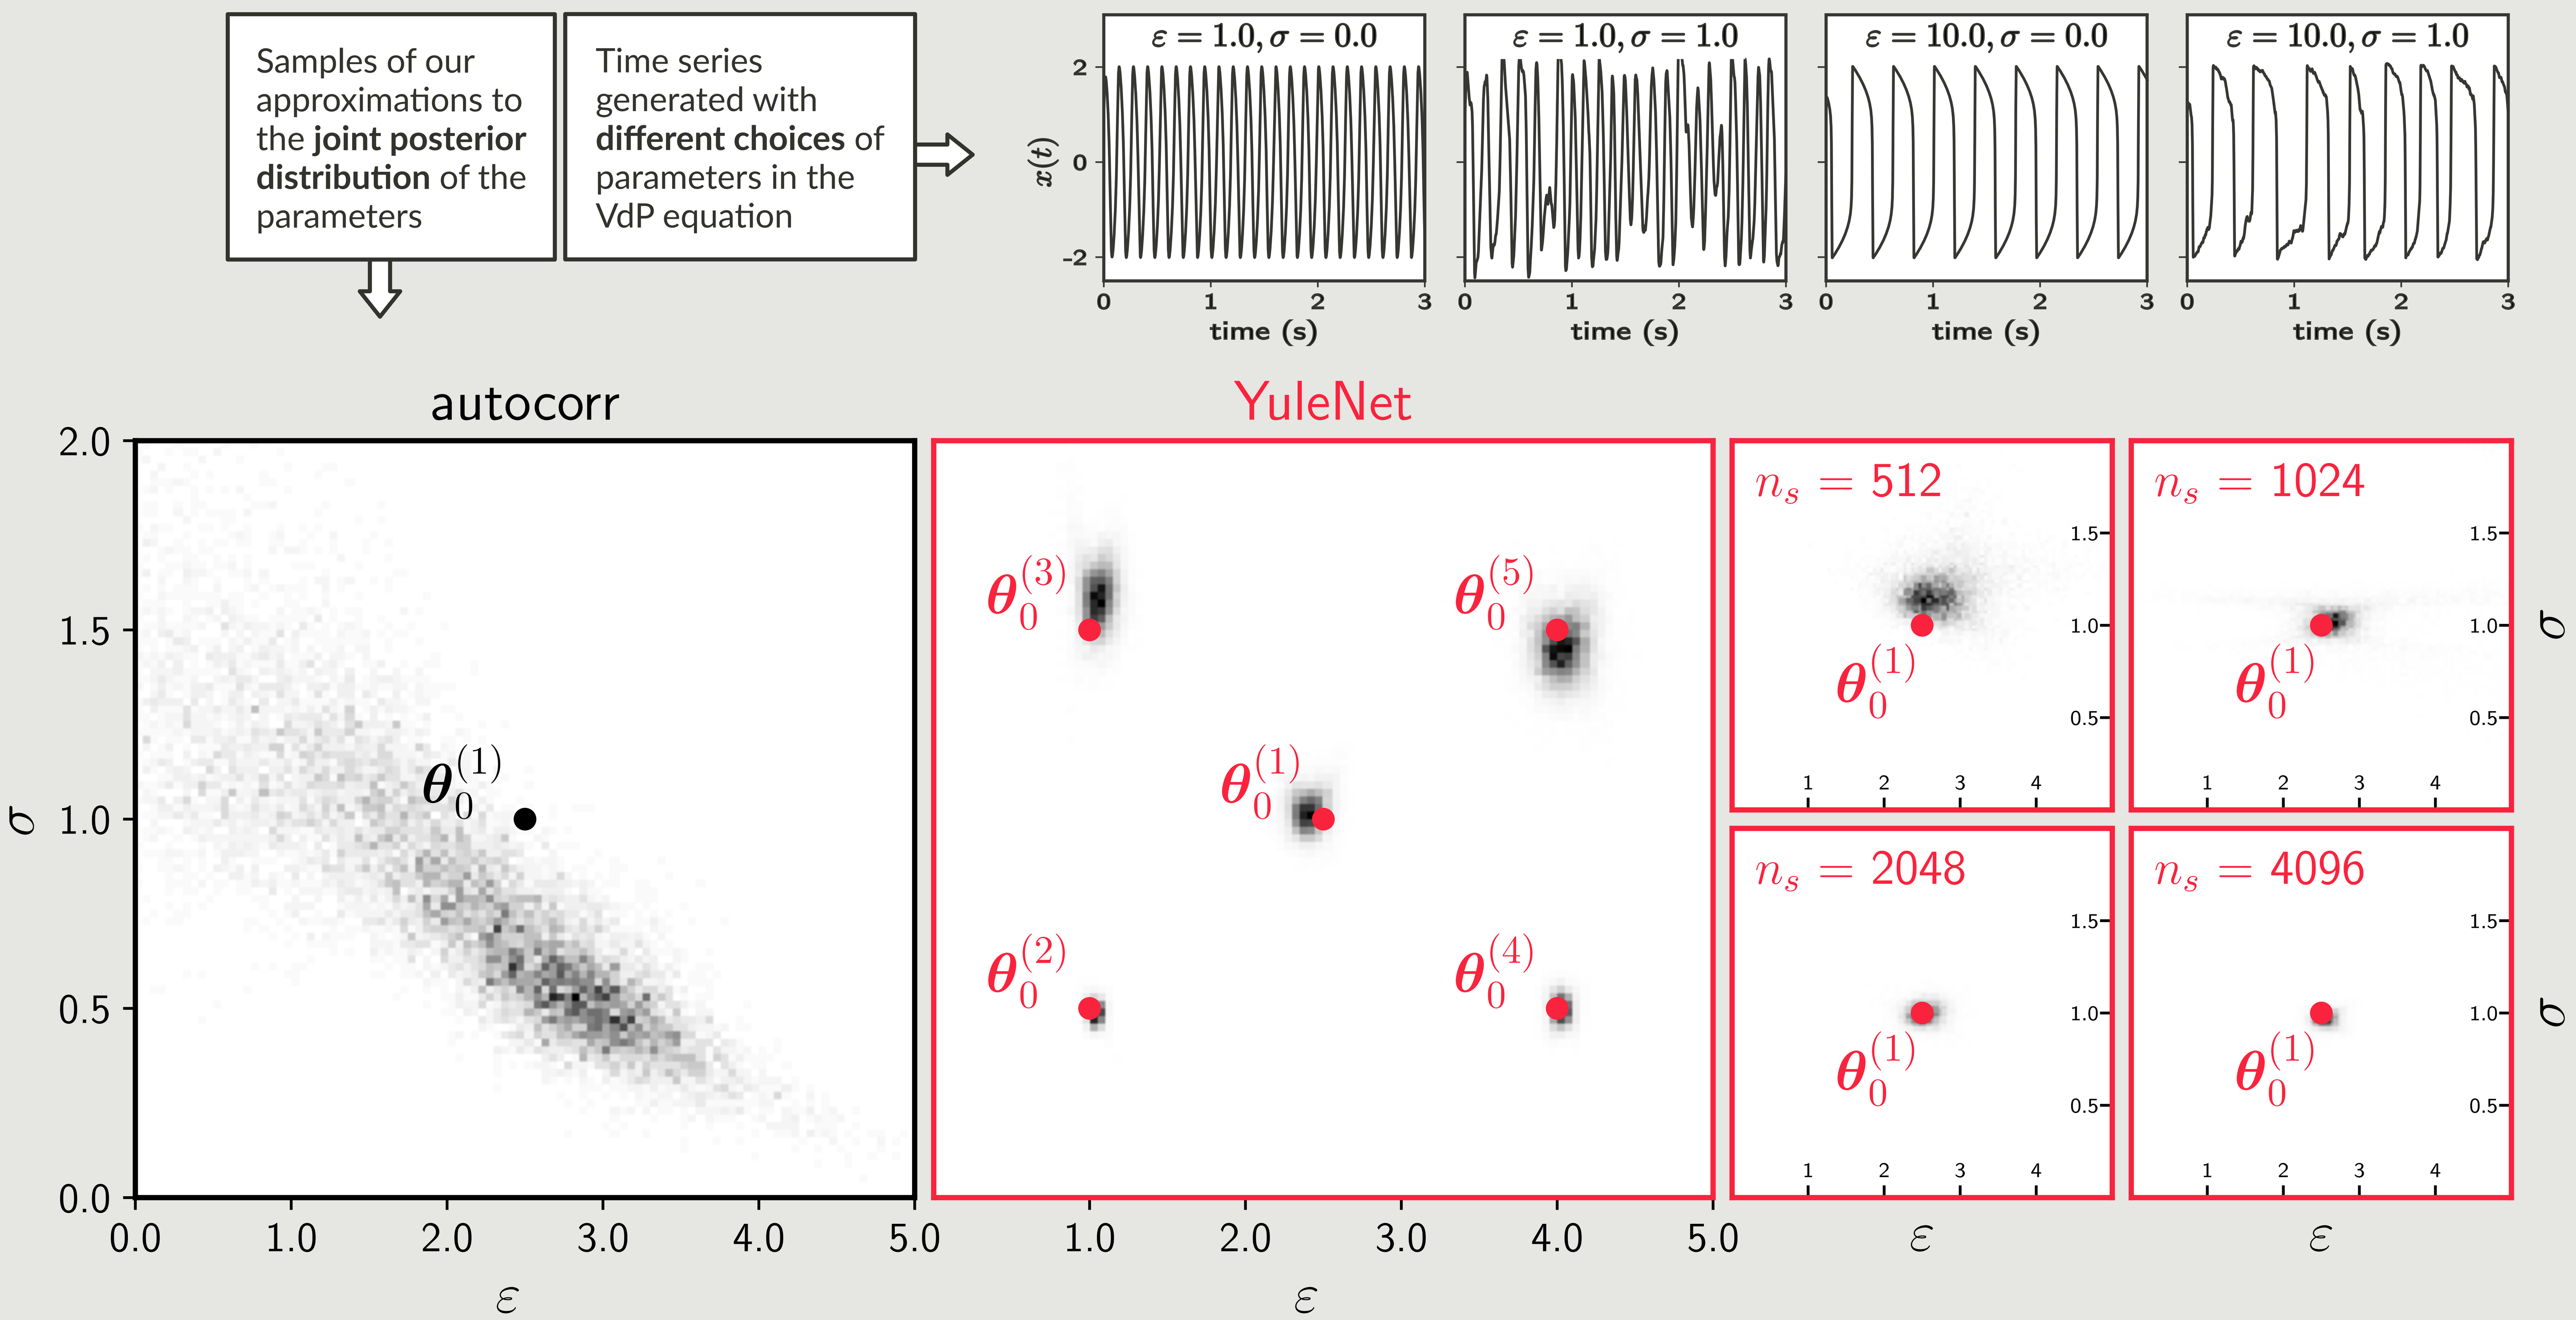

5 DISCUSSION AND CONCLUSION

- YuleNet has good results on **linear** and **non-linear** stochastic models. Better think twice before using **autocorrelations**.
- YuleNet requires **less parameters** to train than PEN and is computationally **efficient**.
- YuleNet is a promising **tool** for studying non-linear dynamical systems such as those used in computational **neuroscience**.

REFERENCES

[1] Chambon et al. (2017) - arxiv.org/abs/1707.03321

[2] Greenberg et al. (2019) - arxiv.org/abs/1905.07488

[3] Wiqvist et al. (2019) - arxiv.org/abs/1901.10230