M2 Internship: Mechanistic-Statistical Modeling with Physics-Informed Neural Networks

Contacts:
Hugo Gangloff - hugo.gangloff@inrae.fr
Nicolas Jouvin - nicolas.jouvin@inrae.fr
Pierre Gloaguen - pierre.gloaguen@univ-ubs.fr

Keywords: PINNs, MCMC, mechanistic-statistical models, deep learning, ecology

To apply send an email with object "Application for PINN MCMC internship" with your CV, an academic transcript for M1, and a motivation letter.

1 Context

While many natural phenomena can be modeled as dynamical systems described through Partial Differential Equations (PDEs), we usually only access partial and noisy observations about them. For example, in spatial ecology, researchers are interested in modeling the spatio-temporal structure of populations in response to various and heterogeneous environmental covariates, given sparse and possibly indirect observations. In this context, mechanistic-statistical models (see Roques 2013, for an introduction) are a classical approach and have been successfully applied in various contexts (Roques et al. 2011; Soubeyrand et al. 2014; Louvrier et al. 2020). In a nutshell, they consist in the combination of a deterministic model described by a PDE, with a statistical model conditioned on the solution. We are interested in estimating the parameters describing the PDE, having no direct observations of its solution. This is a form of so-called inverse problems, which can be addressed in a Bayesian context thanks to the statistical model part, by sampling from the posterior distribution of the parameters given the observations.

Following Roques (2013), the internship will focus on the model where the physical process $u(t, x)$, with $t \in [0, T], x \in \mathbb{R}^2$, is modeled as a reaction-diffusion process of the form

$$\frac{\partial u(t,x)}{\partial t} = D(t,x)\Delta u(t,x) + u(t,x)(r(t,x) - \gamma u(t,x)). \tag{1}$$

Here, the PDE parameters are the diffusion coefficient D and the parameter r, which could be dependent on some additionally observed time and/or spatial covariates. Then, given some observations $\{y_l\}_{l=1}^n$ at space-time locations $\{(t_l, x_l)\}_{l=1}^n$, the solution u of the PDE is linked to the observations via a statistical model $Y \mid u, D, r$ which gives rise to the likelihood $p(Y|u, D, r)$. Finally, MCMC approaches are used to sample from the posterior $p(D, r|Y)$, which allows for biological and ecological interpretation regarding the species response to its environment.

Mechanistic-statistical models have gained interest during the past decade, as they propose innovative approaches for combining statistical modeling and expert knowledge on the underlying dynamical process. In addition, the Bayesian approach also allows to incorporate priors and measure uncertainty on the PDE parameters. However, the main bottleneck is the computational cost of the MCMC sampling. Indeed, each new proposal step requires a call to some PDE solver (*e.g.* finite elements or volumes). Reducing this computational cost using machine-learning approaches is the main motivation of this internship.

2 Goal of this internship: mechanistic-statistical modeling with PINNs

This internship aims at substituting or combining traditional PDE solvers with Physics-Informed Neural Networks (PINNs, Raissi et al. 2019) in order to approximate the PDE solution. They are a general approach to solve PDEs, defined as the identity $\mathcal{N}_\theta[u] = 0$, where $\mathcal{N}_\theta$ is an arbitrary differential operator with parameters θ. For example, in the case of the reaction-diffusion equation (1) we have $\theta = (D, r)$. One seeks to find the best neural network u_ν (ν being the set of weights and biases) approaching the solution by minimizing its PDE residuals computed at randomly sampled space-time locations. This is a mesh-less approach, as opposed to finite elements or volumes, and this has proven useful in a variety of contexts over the last years.

Despite the fact that less theoretical guarantees are available as compared to standard numerical methods, PINNs are an interesting approach for the following reasons:

- a trained PINN can be marginally adapted to a small change in the PDE parameters θ, using a small computational budget to re-train it. This is particularly useful in the context of sequential MCMC algorithms which usually needs to recompute the PDE solution for each proposal of the parameters θ.
- PINNs and hypernetworks have been combined in models coined HyperPINNs (Avila Belbute-Peres et al. 2021). The idea is to design a so-called *metamodel*, where the PINN is trained to approximate the solutions of the family of PDEs indexed by the differential operator parameters θ. At test time, for any set of parameters θ, computing the solution only requires a forward pass in the neural network which is several orders of magnitude more efficient. Thus, this would greatly improve the computational bottleneck of MCMC algorithms involving traditional solvers. Recently, a low-rank adaptations of HyperPINNs have been proposed (Cho et al. 2024; Majumdar et al. 2023), based on the highly popular low rank approaches for Large Language Models (Hu et al. 2022). Such extensions should also be explored during this internship.
- Moreover, as a machine learning approach, PINNs also provide a flexible framework for combining observations into the loss. Frequentist approaches are an alternative direction, trying to maximize the likelihood of observed data regularized by the physics prior on the PDE solution.

This internship will focus on PINNs in the context of mechanistic-statistical models. If successful, these new approaches could pave the way to new methodological developments of great interest for researchers within and outside INRAE.

Organization The internship will mainly revolve around the following tasks

- Bibliography on the recent literature on PINNs for metamodeling and MCMC approaches for mechanistic-statistical modeling.
- Conception a new mechanistic-statistical approach involving PINNs for parameter inference;
- Implementation of the new model using the `jinns` Python library, developed at MIA Paris-Saclay and and based on the JAX ecosystem[1].
- Comparison between the new model and classical approaches on synthetic and real-world data.

3 Profile & environment

The candidate should be a 2nd year master or last year engineer student, in Statistics/Machine Learning, with courses on Bayesian/latent variable modeling, computational statistics and deep learning. Scientific programming skills in Python are required, while familiarity with the JAX ecosystem is a bonus.

[1] https://jax.readthedocs.io/en/latest/

- Location: UMR MIA Paris-Saclay, Palaiseau Campus, 22 place de l'agronomie, 91120 Palaiseau, France
- Supervision: Hugo Gangloff & Nicolas Jouvin do their research on PINNs and are the developers of the `jinns` Python package. Pierre Gloaguen is a specialist on statistical modeling, with a specific focus on applications in ecology.
- Starting date: flexible, starting in February or after.
- Duration: 5-6 months
- Salary: as an intern, you'll receive a "gratification" which is unfortunately capped around 700 euros/month.

The candidate will have an office, and benefit from the work environment of the MIA Paris-Saclay laboratory, with many PhD students & postdocs working on statistical modeling and machine learning for the life sciences.

References

Roques, Lionel (2013). "Modèles de réaction-diffusion pour l'écologie spatiale: Avec exercices dirigés". In.

Roques, Lionel, Samuel Soubeyrand, and Jérôme Rousselet (2011). "A statistical-reaction–diffusion approach for analyzing expansion processes". In: *Journal of Theoretical Biology* 274.1, pp. 43–51.

Soubeyrand, Samuel and Lionel Roques (2014). "Parameter estimation for reaction-diffusion models of biological invasions". In: *Population ecology* 56, pp. 427–434.

Louvrier, Julie et al. (2020). "A mechanistic–statistical species distribution model to explain and forecast wolf (Canis lupus) colonization in South-Eastern France". In: *Spatial Statistics* 36, p. 100428.

Raissi, Maziar, Paris Perdikaris, and George E Karniadakis (2019). "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations". In: *Journal of Computational physics* 378, pp. 686–707.

Avila Belbute-Peres, Filipe de, Yi-Fan Chen, and Fei Sha (2021). "HyperPINN: Learning parameterized differential equations with physics-informed hypernetworks". In: *CoRR* abs/2111.01008.

Cho, Woojin et al. (2024). "Hypernetwork-based meta-learning for low-rank physics-informed neural networks". In: *Advances in Neural Information Processing Systems* 36.

Majumdar, Ritam et al. (2023). "HyperLoRA for PDEs". In: *arXiv preprint arXiv:2308.09290*.

Hu, Edward J. et al. (2022). "LoRA: Low-Rank Adaptation of Large Language Models". In: *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*.