State-consistent physics guidance for Diffusion Models: a joint parameter covariance-based approach to inverse problems

Type of offer :
Ph.D.

Location	
Mathematical & Numerical Modelling Laboratory Applied Mathematics & Statistics Department Conservatoire Nationale Arts et Métiers 2 rue Conté 75003, Paris	ONERA Département d'Aérodynamique Aéroélasticité Acoustique Meudon

Context :
The integration of generative priors into scientific machine learning has enabled the solution of ill-posed inverse problems with unprecedented fidelity. Diffusion models, in particular, have emerged as a powerful tool for reconstructing physical fields from sparse observations. Despite their success, current state-of-the-art methods often rely on heuristic modifications to the sampling process rather than rigorous probabilistic derivations. DiffusionPDE [3] introduces a PDE-residual term into the diffusion sampler but treats it as an auxiliary loss function weighted by a scalar hyperparameter. This formulation lacks a clear statistical interpretation of the physical uncertainty and assumes isotropic confidence in the physics across the entire domain.
More recently, PIDDM [5] critiques standard guidance methods for suffering from "Jensen's Gap", the discrepancy between the physical residual of the expected state versus the expected residual of the state. While they propose a distillation-based solution, their approach requires training task-specific student networks, sacrificing the flexibility of the generative prior. We propose a principled Bayesian formulation that addresses these limitations without requiring retraining or distillation.
The key contributions of this project will be:

1. **Rigorous Joint Posterior:** we define the physics term not as a loss, but as a likelihood function derived from a virtual observation of the PDE residual.
2. **Structured Physics Covariance:** we move beyond scalar weighting by introducing a spatially adaptive covariance matrix. This allows the model to "trust" the physics less in regions of high numerical error (e.g., shocks).
3. **Implicit Parameter Estimation:** we derive a gradient-based update for physical parameters that backpropagates through the denoising network's Jacobian. This "Implicit Gradient" ensures that parameter updates are consistent with the learned data manifold.
4. **Dynamic Physics Gating:** we introduce an adaptive computation strategy that monitors the signal-to-noise ratio of the physical residual. By disabling expensive gradient corrections during high-noise regimes or when the prior is sufficient, we significantly reduce inference cost without sacrificing accuracy.

This thesis is part of the PostGenAI@Paris cluster program hosted by SCAI (Sorbonne University), and funded by ANR: https://www.sorbonne-universite.fr/actualites/le-cluster-postgenaiparis-projet-porte-par-sorbonne-universite-obtient-35-millions

Scientific partners: CNAM & ONERA

Profile :
The candidate should have a MSc degree or equivalent (engineering diploma) in mechanics, computer science or applied mathematics, with experience in scientific machine learning.

Skills :
Programming experience and expertise in data-driven techniques will be considered very positively.

Duration and start date :
The position is offered for the duration of 36 months, from October 1st, 2026 to September 30th, 2029.
Deadline for applications: 01/05/2026

Required documents :
The applicant should include a CV, and a motivation letter.

Contacts :
Taraneh Sayadi - taraneh.sayadi@lecnam.net
Samir Beneddine - samir.beneddine@onera.fr
Denis Sipp - denis.sipp@onera.fr

References :
[1] Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In International Conference on Learning Representations (ICLR), 2023.
[2] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. Advances in neural information processing systems, 34:8780–8794, 2021.
[3] Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. Diffusionpde: Generative PDE-solving under partial observation. In Advances in Neural Information Processing Systems (NeurIPS), 2024. NeurIPS 2024.
[4] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 10684–10695, 2022.
[5] Yi Zhang and Difan Zou. Physics-informed distillation of diffusion models for pde-constrained generation. arXiv preprint arXiv:2505.22391, 2025.