PhD Position in Engineering and Computer Science, Sorbonne Université, Paris, France

Learning Generative World Models of Physical Dynamics

Contact : Patrick Gallinari, patrick.gallinari@sorbonne-universite.fr

Location: Sorbonne Université, Pierre et Marie Curie Campus, 4 Place Jussieu, Paris, Fr.
Machine Learning and Information Access team.

Candidate profile: Master degree in computer science or applied mathematics, Engineering school. Background and experience in machine learning. Good technical skills in programming.

How to apply: please send a cv, motivation letter, grades obtained in master, recommendation letters when possible to patrick.gallinari@sorbonne-universite.fr

Start date: October/November 2026 for three years

Note: The research topic is open and depending on the candidate profile could be oriented more on the theory or on the application side

Keywords: Deep Learning, Scientific AI, Generative Models, Representation Learning

---

Context

AI4Science is an emerging scientific paradigm that leverages the potential of AI to advance scientific discovery and model complex natural phenomena. Within this context, generative modeling is opening new avenues for simulating and understanding complex physical systems. This PhD project aims to explore and advance generative deep learning architectures for modeling complex physical dynamical systems arising in domains such as fluid mechanics and climate science. These surrogate models hold strong potential for learning flexible, data-driven representations of physical laws. By developing generalizable, cross-physics generative models, this research contributes to the broader vision of AI4Science: accelerating scientific discovery through learned simulation and abstraction.

Research Objectives

Can we develop generative models that learn structured, physically grounded representations of dynamical systems, enabling synthesis, adaptation, and generalization across diverse physical regimes and heterogeneous physical domains? This research question unfolds into several complementary directions, including learning compact representations of physical dynamics, modeling uncertainty through generative approaches, and enforcing the scientific validity of generated solutions. Together, these directions aim to contribute to the development of generative world models capable of representing complex physical dynamics in a scalable and scientifically consistent manner.

Learning Latent representation of physical dynamics.

Learning representations of physical dynamics is a crucial step toward developing scalable neural surrogates. To handle the complexity of physical systems, this is generally performed in a reduced-dimensional latent space. Classical approaches pretrain models using forecasting objectives, for example with neural operators (Serrano et al., 2024; Serrano et al., 2025) or foundation models (McCabe et al., 2025). However, such objectives often entangle dynamics estimation and future-state prediction. An alternative approach is to decouple representation learning from downstream tasks in order to learn abstract representations that can support multiple tasks, as in the JEPA family of models (Balestriero et al., 2025; Maes et al., 2026). Downstream tasks can then leverage these learned representations in zero-shot or few-shot settings. This research direction aims to investigate latent representations that capture the structure of physical dynamics while remaining transferable across tasks and physical regimes.

Generative models and uncertainty modeling

A complementary research direction focuses on modeling uncertainty and the distribution of possible trajectories. In physical systems, one is often interested in predicting distributions of trajectories rather than single trajectories. Physical systems may be chaotic, as in climate and weather forecasting; observations are often incomplete; system dynamics may evolve due to external factors; and neural surrogates accumulate errors during forecasting. All these sources of uncertainty must therefore be modeled. Current approaches leverage recent generative paradigms such as diffusion models, flow models, and stochastic interpolants (Li et al., 2025; Koupai et al., 2025; Hao et al., 2024; Zhou et al., 2025). How to build foundation generative models operating on structured latent representations remains largely open. This research direction will investigate efficient generative models capable of capturing uncertainty and multimodality in learned physical solutions.

Physically consistent models

Another important challenge concerns the scientific validity of generated trajectories. Black-box neural surrogates may accurately predict system evolution according to standard metrics such as RMSE while remaining physically inconsistent. Different approaches aim to introduce constraints that enforce physical consistency. Generative models open new challenges and opportunities for generating physically consistent distributions. Recent approaches attempt to draw inspiration from self-consistency ideas developed in language models to enforce physical consistency (Wu et al., 2026).

More broadly, the PhD will explore generative and physically consistent models of physical dynamics, leveraging recent advances in self-supervised representation learning, generative modeling, and neural surrogate modeling. The objective is to develop models capable of producing solutions that are not only plausible and diverse, but also scientifically valid and transferable across physical regimes.

.Position and Working Environment

The PhD studentship is a three years position starting in October/November 2026. It does not include teaching obligation, but it is possible to engage if desired. The PhD candidate will work at Sorbonne Université (S.U.), in the center of Paris. He/She will integrate the MLIA team (Machine Learning and Deep Learning for Information Access) at ISIR (Institut des Systèmes Intelligents et de Robotique).

References

Balestriero, R., & LeCun, Y. (2025, November 14). *LeJEPA: Provable and Scalable Self-Supervised Learning Without the Heuristics*. http://arxiv.org/abs/2511.08544

Chen, Z., & Deng, S. (2026). *Latent Generative Solvers for Generalizable Long-Term Physics Simulation*. http://arxiv.org/abs/2602.11229

Hao, Z., Su, C., Liu, S., Berner, J., Ying, C., Su, H., Anandkumar, A., Song, J., & Zhu, J. (2024). DPOT: Auto-Regressive Denoising Operator Transformer for Large-Scale PDE Pre-Training. Icml. http://arxiv.org/abs/2403.03542

Koupaï, A. K., le Boudec, L., Serrano, L., & Gallinari, P. (2025). ENMA: Tokenwise Autoregression for Generative Neural PDE Operators. *NeurIPS*.

Li, Z., Zhou, A., & Farimani, A. B. (2025, March 28). *Generative Latent Neural PDE Solver using Flow Matching*. http://arxiv.org/abs/2503.22600

Maes, L., Lidec, Q. le, Scieur, D., LeCun, Y., & Balestriero, R. (2026, March 13). *LeWorldModel: Stable End-to-End Joint-Embedding Predictive Architecture from Pixels*. http://arxiv.org/abs/2603.19312

McCabe, M., Blancard, B. R.-S., Parker, L. H., Ohana, R., Cranmer, M., Bietti, A., Eickenberg, M., Golkar, S., Krawezik, G., Lanusse, F., Pettee, M., Tesileanu, T., Cho, K., & Ho, S. (2024). Multiple Physics Pretraining for Physical Surrogate Models. 1–25 http://arxiv.org/abs/2310.02994

Serrano, L., Kassai, A., Wang, T., Erbacher P., Gallinari, P., (2025) Zebra: In-Context Generative Pretraining for Solving Parametric PDEs.

Serrano, L., Wang, T., le Naour, E., Vittaut, J.-N., & Gallinari, P. (2024). AROMA : Preserving Spatial Structure for Latent PDE Modeling with Local Neural Fields. NeurIPS.

Wu, H., Xu, F., Gao, Y., Zhang, F., Wang, K., Huang, X., & Song, Q. (2026). *Direct Preference Optimization for Dynamical System Modeling*.

Zhou, A., Li, Z., Schneier, M., Buchanan Jr, J. R., & Farimani, A. B. (2025). TEXT2PDE: Latent Diffusion Models for Accessible Physics Simulation. ICLR.