Master Thesis Project Proposals:

Generative Models for Probabilistic Forecasting

Martin Andrae[1], Erik Larsson[1], Fredrik Lindsten[1], and Tomas Landelius[2]

[1]Division of Statistics and Machine Learning (STIMA),
[2]Swedish Meteorological and Hydrological Institute (SMHI)

October 8, 2025

1 Background

Much of modern deep learning research focuses on generative models. These models aim to sample from an unknown, highly complex distribution $p(x)$. *Sampling* can mean generating text (as in large language models), creating new images, or, in our context, producing plausible future states of a dynamical system.

For image generation the goal is to learn to draw new samples from $p(x)$ using only the examples we have observed. Several classes of generative models exist, including autoregressive models, variational autoencoders, diffusion models and flow-matching approaches. In recent years, continuous-time formulations such as *score-based diffusion* and *flow matching* have emerged as powerful tools for high-quality generation. The central idea behind these approaches is to construct a continuous transformation between a base distribution (e.g. a Gaussian) and the target distribution $p(x)$ by learning a vector field that defines an ordinary differential equation (ODE) or stochastic differential equation (SDE).

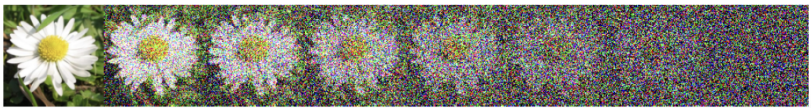

(a) Diffusion model transforming noise into an image.

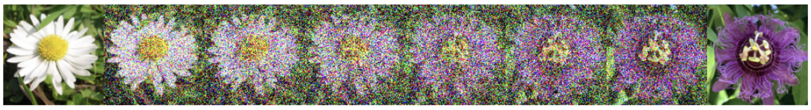

(b) Flow matching model transforming one image into another.

These continuous generative frameworks can be naturally transferred to spatiotemporal physical systems (see Figure 2a). Instead of sampling a single realistic image we now ask the model to produce a probabilistic forecast of the next state $p(x_t \mid x_{t-1})$, or a conditional forecast given observations $p(x_t \mid y_t, x_{t-1})$, the latter formally connects to data assimilation/Bayesian filtering. Generative models provide a flexible route to represent multimodal uncertainties, non-Gaussian error structures, and complex spatiotemporal correlations that are typical for chaotic PDE-driven systems.

One example of a spatiotemporal forecasting problem is weather forecasting. The Earth's atmosphere is a high-dimensional and inherently chaotic system, where accurate and efficient weather forecasting is essential for detecting extreme events and issuing timely warnings. Recent years have witnessed remarkable progress in applying machine learning to weather prediction (Ben-Bouallegue et al., 2023). These data-driven

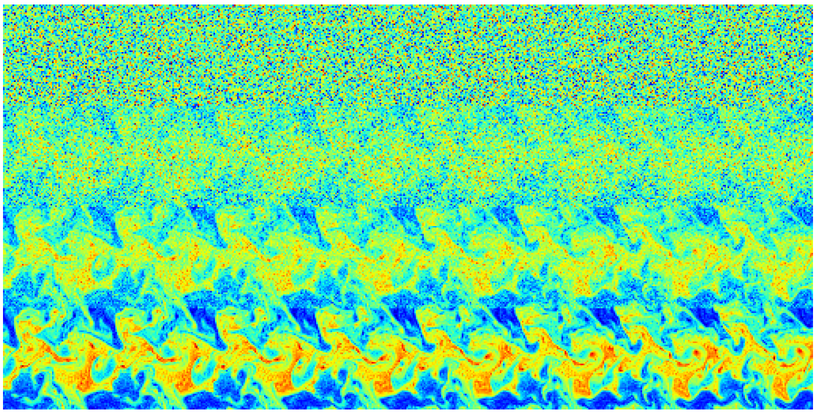

(a) Diffusion model transforming noise into a probabilistic forecast.

models can match or even outperform existing Numerical Weather Prediction (NWP) systems, while producing forecasts in a fraction of the time (Lam et al., 2022; Bi et al., 2023; Pathak et al., 2022). While early approaches focused on deterministic predictions, the field has increasingly shifted toward probabilistic ensemble forecasting (Price et al., 2023; Bonev et al., 2025; Alet et al., 2025). In our group, we have explored probabilistic weather prediction using both latent-variable and diffusion-based approaches. For an overview of our work so far see (Oskarsson et al., 2023, 2024; Andrae et al., 2025; Larsson et al., 2025).

2 Data description

You will work with trajectories from a spatiotemporal system governed by a chaotic PDE, called the surface quasi-geostrophic equations (SQG). This PDE is a simplified model[a] of atmospheric dynamics and is therefore challenging to forecast accurately. We have worked extensively with this system and have used it ourselves in our research. This means that it's easy to generate data and train models within our codebase.

[a] The main dataset used in these projects will be generated using the framework available at `https://github.com/jswhit/sqgturb`. The dataset can be flexibly produced to match the specific requirements of each project.

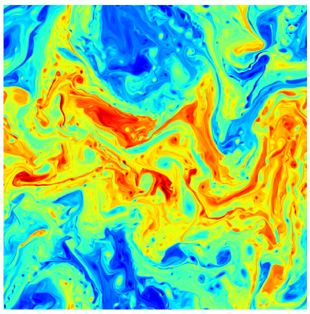

(a) A SQG datapoint

3 Projects

We offer several projects related to Probabilistic Spatiotemporal Forecasting, as detailed below, accompanied by some related work. Our ambition is to find suitable candidates for multiple projects. We envision to then set up a collaborative environment for the involved students, to encourage collaboration on generic questions related to working with the data and getting familiar with the application area.

Note that you can apply to multiple projects at the same time.

3.1 New types of probabilistic forecasting models

Probabilistic forecasting of entire trajectories.
Many current approaches model $p(x_{t+1} \mid x_t)$ and obtain multi-step trajectories by autoregressive sampling. While effective in practice, autoregressive schemes can struggle with long-range dependencies and error accumulation. This project targets direct modeling of the joint conditional distribution $p(x_{t+1:T} \mid x_t)$, allowing generation of entire future trajectories in a single (or few) pass(es). Ho et al. (2022); Liu et al. (2024)

Forecasting trajectories from marginals using correlated noise.
Building on a recent method that learns $p(x_{t+\Delta t} \mid x_t, \Delta t)$, one can generate full trajectories by injecting strongly correlated Gaussian noise during sampling. The student will implement and train such a diffusion-based model, then systematically study how the temporal correlation structure of the noise affects trajectory quality, ensemble spread, and calibration. Andrae et al. (2025)

Training and evaluating a CRPS-based model.
Proper scoring rules such as the Continuous Ranked Probability Score (CRPS) provide a principled alternative to SDE-based methods such as diffusion models. Training models using CRPS can yield very fast predictive sampling, but it is not obvious which architectures and design choices make this approach effective in high-dimensional spatiotemporal settings. This project will train CRPS-optimized models, evaluate their performance, and explore architectural or training modifications. Alet et al. (2025)

Probabilistic Forecasting with Stochastic Interpolants and Föllmer Processes
Flow matching and stochastic interpolants generalize diffusion models to arbitrary starting distributions. For forecasting, these methods can be used to construct diffusion-like transforms that map x_t to x_{t+1} (see Figure 1b). The project will implement such interpolant-based forecast models and compare their sample efficiency and predictive performance to standard diffusion approaches. Chen et al. (2024)

3.2 Adapting new generative models to probabilistic forecasting

Distributional Flow/Diffusion models compared to Flow/Diffusion models
In this project we aim to compare forcasting models based on standard flow/diffusion models with distributional diffusion models. The focus is on assessing whether this leads to improved uncertainty quantification and forecast skill. Bortoli et al. (2025)

Exploring Techniques for Faster Sampling in Flow Models
Sampling from flow-based models typically requires numerous forward passes through a neural network, which can be computationally expensive. This project investigates approaches such as consistency models and rectified flows to accelerate sampling while maintaining model performance. Boffi et al. (2025)

Pyramidal Flow Matching
Pyramidal flow matching has recently been applied successfully to video generation. In this project, we aim to adapt this framework for forecasting dynamical systems, where multi-scale temporal dynamics play a central role. The approach will be compared against traditional autoregressive flow matching models to assess potential improvements in forecast skill and efficiency. Jin et al. (2025)

Learning from Observations
In real-world applications, the full state of a dynamical system is rarely observed. Instead, we only have access to sparse and partial observations. This project aims to investigate methods for learning data assimilation models, which reconstruct the com-

plete current state from limited observations, or forecasting models, which predict future states while only observing sparse partial observations of states. Xiang et al. (2024); Alexe et al. (2024); Vaughan et al. (2024)

3.3 Latent spaces

Learning latent dynamics for probabilistic forecasting with flow-based generative models.
Flow-based generative models are invertible and map data to a latent space in which samples are approximately Gaussian. This latent representation enables smooth interpolation and compact parametrization of complex fields. The project aims to learn spatiotemporal dynamics in that latent space: build forecasting models, perform temporal interpolation, and generate ensemble perturbations from single observations. Bodin et al. (2025)

Latent Diffusion or Flow Matching for Forecasting Physical Systems
Training and sampling flow-based models directly in data space is computationally expensive. Recent advances demonstrate that learning a compressed latent representation of the data and performing diffusion or flow matching in this latent space can offer a favorable trade-off between efficiency and accuracy. This project will focus on systematically comparing diffusion or flow matching models with their latent-space counterparts. Rombach et al. (2022)

Latent Space Forecasting
Most forecasting models operate in data space, predicting the next state $t+1$ from the current state t. Longer trajectories are then obtained through autoregressive rollouts, where each forecasted output serves as the next input. This project explores whether learning a latent representation of the data enables forecasting directly in latent space, allowing for arbitrarily many forecasting steps in the latent space before decoding back to data space. Du et al. (2025)

4 How to apply

Apply by sending

- A **CV** and a **transcript of grades**.
- A brief motivation on why you are interested in the project(s) and which ones you find the most interesting.
- Pointers to any past experience that makes you a particularly suitable candidate.

to both `martin.andrae@liu.se, erik.larsson@liu.se`. Write "*Thesis Project Application*" in the subject line. We will review applications in two rounds, with deadlines on **November 1st** and **December 1st**. Projects will be allocated based on the applications received by each date. If you want to maximize your chances of getting your preferred project, we recommend applying early.

If you have any questions about the projects or application feel free to reach out.

References

F. Alet, I. Price, A. El-Kadi, D. Masters, S. Markou, T. R. Andersson, J. Stott, R. Lam, M. Willson, A. Sanchez-Gonzalez, and P. Battaglia. Skillful joint probabilistic weather forecasting from marginals, 2025. URL `https://arxiv.org/abs/2506.10772`.

M. Alexe, E. Boucher, P. Lean, E. Pinnington, P. Laloyaux, A. McNally, S. Lang, M. Chantry, C. Burrows, M. Chrust, F. Pinault, E. Villeneuve, N. Bormann, and S. Healy. Graphdop: Towards skilful data-driven medium-range weather forecasts learnt and initialised directly from observations, 2024. URL `https://arxiv.org/abs/2412.15687`.

M. Andrae, T. Landelius, J. Oskarsson, and F. Lindsten. Continuous ensemble weather forecasting with diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL `https://openreview.net/forum?id=ePEZvQNFDW`.

Z. Ben-Bouallegue, M. C. A. Clare, L. Magnusson, E. Gascon, M. Maier-Gerber, M. Janousek, M. Rodwell, F. Pinault, J. S. Dramsch, S. T. K. Lang, B. Raoult, F. Rabier, M. Chevallier, I. Sandu, P. Dueben, M. Chantry, and F. Pappenberger. The rise of data-driven weather forecasting. *arXiv preprint arXiv:2307.10128*, 2023.

K. Bi, L. Xie, H. Zhang, X. Chen, X. Gu, and Q. Tian. Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, pages 1–6, 2023.

E. Bodin, A. I. Stere, D. D. Margineantu, C. H. Ek, and H. Moss. Linear combinations of latents in generative models: subspaces and beyond. In *The Thirteenth International Conference on Learning Representations*, 2025. URL `https://openreview.net/forum?id=n5PrId7pk5`.

N. M. Boffi, M. S. Albergo, and E. Vanden-Eijnden. How to build a consistency model: Learning flow maps via self-distillation, 2025. URL `https://arxiv.org/abs/2505.18825`.

B. Bonev, T. Kurth, A. Mahesh, M. Bisson, J. Kossaifi, K. Kashinath, A. Anandkumar, W. D. Collins, M. S. Pritchard, and A. Keller. FourCastNet 3: A geometric approach to probabilistic machine-learning weather forecasting at scale, 2025. URL `http://arxiv.org/abs/2507.12144`.

V. D. Bortoli, A. Galashov, J. S. Guntupalli, G. Zhou, K. P. Murphy, A. Gretton, and A. Doucet. Distributional diffusion models with scoring rules. In *Forty-second International Conference on Machine Learning*, 2025. URL `https://openreview.net/forum?id=N82967FcVK`.

Y. Chen, M. Goldstein, M. Hua, M. S. Albergo, N. M. Boffi, and E. Vanden-Eijnden. Probabilistic forecasting with stochastic interpolants and föllmer processes. In *Forty-first International Conference on Machine Learning*, 2024. URL `https://openreview.net/forum?id=UQYXZdca92`.

H. Du, L. Kim, J. Creus-Costa, J. Michaels, A. Shetty, T. Hutchinson, C. Riedel, and J. Dean. Weathermesh-3: Fast and accurate operational global weather forecasting, 2025. URL `https://arxiv.org/abs/2503.22235`.

J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet. Video diffusion models, 2022. URL `https://arxiv.org/abs/2204.03458`.

Y. Jin, Z. Sun, N. Li, K. Xu, K. Xu, H. Jiang, N. Zhuang, Q. Huang, Y. Song, Y. MU, and Z. Lin. Pyramidal flow matching for efficient video generative modeling. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=66NzcRQuOq.

R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirnsberger, M. Fortunato, A. Pritzel, S. Ravuri, T. Ewalds, F. Alet, Z. Eaton-Rosen, W. Hu, A. Merose, S. Hoyer, G. Holland, J. Stott, O. Vinyals, S. Mohamed, and P. Battaglia. GraphCast: Learning skillful medium-range global weather forecasting. *arXiv preprint arXiv:2212.12794*, 2022.

E. Larsson, J. Oskarsson, T. Landelius, and F. Lindsten. Diffusion-lam: Probabilistic limited area weather forecasting with diffusion. In *ICLR 2025 Workshop on Tackling Climate Change with Machine Learning*, 2025. URL https://www.climatechange.ai/papers/iclr2025/36.

Y. Liu, K. Zhang, Y. Li, Z. Yan, C. Gao, R. Chen, Z. Yuan, Y. Huang, H. Sun, J. Gao, L. He, and L. Sun. Sora: A review on background, technology, limitations, and opportunities of large vision models, 2024. URL https://arxiv.org/abs/2402.17177.

J. Oskarsson, T. Landelius, and F. Lindsten. Graph-based neural weather prediction for limited area modeling, 2023. URL http://arxiv.org/abs/2309.17370.

J. Oskarsson, T. Landelius, M. P. Deisenroth, and F. Lindsten. Probabilistic weather forecasting with hierarchical graph neural networks, 2024. URL https://arxiv.org/abs/2406.04759.

J. Pathak, S. Subramanian, P. Harrington, S. Raja, A. Chattopadhyay, M. Mardani, T. Kurth, D. Hall, Z. Li, K. Azizzadenesheli, P. Hassanzadeh, K. Kashinath, and A. Anandkumar. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*, 2022.

I. Price, A. Sanchez-Gonzalez, F. Alet, T. R. Andersson, A. El-Kadi, D. Masters, T. Ewalds, J. Stott, S. Mohamed, P. Battaglia, R. Lam, and M. Willson. Gencast: Diffusion-based ensemble forecasting for medium-range weather, 2023. URL https://arxiv.org/abs/2312.15796.

R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, June 2022.

A. Vaughan, S. Markou, W. Tebbutt, J. Requeima, W. P. Bruinsma, T. R. Andersson, M. Herzog, N. D. Lane, M. Chantry, J. S. Hosking, and R. E. Turner. Aardvark weather: end-to-end data-driven weather forecasting, 2024. URL https://arxiv.org/abs/2404.00411.

Y. Xiang, W. Jin, H. Dong, M. Bai, Z. Fang, P. Zhao, H. Sun, K. Thambiratnam, Q. Zhang, and X. Huang. Adaf: An artificial intelligence data assimilation framework for weather forecasting, 2024. URL https://arxiv.org/abs/2411.16807.