Towards Fully Bayesian Hybrid Recurrent Neural Networks for Bioprocess Modelling

Author: Lucas Cosier
Affiliation: DataHow AG, R&D Team
Supervisor: Dr. Alessandro Butte

Scope: Thesis / Internship / Semester Project
Date: March 11, 2026

Contact: Harini Narayanan at h.narayanan@datahow.ch

Executive Summary

This project proposes extending the current hybrid recurrent neural network (RNN) framework to a fully Bayesian formulation, with the goal of providing principled, calibrated uncertainty quantification (UQ) for bioprocess modelling. By placing priors over model weights and learning posterior distributions, we aim to disentangle observation noise (aleatoric) from model uncertainty (epistemic), improve calibration across process conditions, and lay a foundation for uncertainty-aware process monitoring and control.

Background

Brief Introduction / Context

Bioprocesses are inherently stochastic dynamic systems governed by complex metabolic interactions. Hybrid models have emerged as a powerful paradigm, combining a mechanistic mass balance backbone with a data-driven component to learn residual dynamics:

$$\frac{dX_i}{dt} = \mathrm{NN}(X, Z) + \mathrm{In}_i - \mathrm{Out}_i \tag{1}$$

In practice, the continuous ODE is discretised via Euler integration, and a recurrent model learns $f(X, Z)$. Reliable uncertainty estimates over predicted trajectories are essential for process monitoring, fault detection, and optimal control.

Current Approaches

The existing framework estimates uncertainty through a heteroscedastic output head and Monte Carlo sampling, with variance calibrated post-hoc via Gaussian NLL. While this yields empirical prediction intervals, it does not place priors over model weights, leaving epistemic uncertainty uncaptured and aleatoric and epistemic contributions entangled. Bayesian deep learning offers principled solutions through variational inference [1], MC Dropout [2], or deep ensembles [3], each with different trade-offs in approximation quality and computational cost.

Challenges

Bioprocess datasets are typically small by machine learning standards, with limited experimental replicates and high annotation cost — making epistemic uncertainty the dominant concern in practice. Approximate posterior inference over recurrent weights is more involved than for feed-forward networks, as uncertainty must propagate consistently through the integration scheme. Discontinuities from feeding events further complicate posterior approximation, as they introduce non-smooth dynamics that stress gradient-based inference methods.

Project Rationale

The core proposal is to replace the current variance-head approach with a proper Bayesian treatment of model weights, enabling principled decomposition of predictive uncertainty. Aleatoric uncertainty will be retained via a heteroscedastic output head, while epistemic uncertainty will be captured through approximate posterior inference over network weights. Several inference strategies will be evaluated, including variational inference and MC Dropout, with training redesigned for joint optimisation of both uncertainty sources.

Objectives

1. **Baseline characterisation:** Evaluate calibration properties of the current framework and establish benchmark metrics.
2. **Bayesian RNN implementation:** Implement and compare approximate inference strategies for the recurrent model.
3. **Joint training:** Redesign the training procedure to jointly optimise predictive performance and uncertainty calibration.
4. **Uncertainty decomposition:** Develop routines to separate aleatoric and epistemic contributions to total predictive variance.
5. **Benchmarking:** Compare Bayesian variants against the current implementation across datasets of varying noise and data density.

Timeline

Phase	**Tasks**	**Target Dates**
Exploration	Literature review on Bayesian RNNs, variational inference, and calibration metrics. Baseline calibration analysis of the current framework.	Wk 1–3
Implementation	Implement MC Dropout and Bayes-by-Backprop variants. Adapt integrator to propagate weight uncertainty. Redesign joint training loop.	Wk 4–7
Analysis	Uncertainty decomposition; calibration benchmarking against current framework. Ablation over inference methods and prior choices.	Wk 8–10
Applications	Demonstrate uncertainty-aware predictions on held-out processes; stress test on out-of-distribution conditions.	Wk 11
Wrap-up	Documentation, evaluation summary, final report and presentation.	Wk 12

Expected Outcome

At the end of this project, we will have a fully Bayesian extension of the hybrid RNN framework with principled, decomposed uncertainty estimates over bioprocess trajectories. The project will deliver a clear empirical comparison of Bayesian inference strategies, quantifying gains in calibration and epistemic coverage, and providing a foundation for uncertainty-aware process control.

References

[1] Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. (2015). Weight uncertainty in neural networks. *Proceedings of the 32nd ICML*, PMLR 37, 1613–1622.

[2] Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. *Proceedings of the 33rd ICML*, PMLR 48, 1050–1059.

[3] Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems*, 30 (NeurIPS 2017).

[4] Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? *Advances in Neural Information Processing Systems*, 30 (NeurIPS 2017).