BARNN: A Bayesian Autoregressive and Recurrent Neural Network

<u>Dario Coscia</u>[1, 2], Max Welling[2], Nicola Demo[1, 3], Gianluigi Rozza[1]

1. mathLab, SISSA, Italy 2. University of Amsterdam, The Netherlands 3. FAST Computing, Italy

AI4Science Needs UQ

Autoregressive and recurrent models power breakthroughs from weather to molecular design, but lack principled uncertainty — a critical gap in scientific applications. The natural question then arises: c***an we TRUST AI surrogates?***

We introduce BARNN, a simple Bayesian approach that integrates seamlessly with SOTA architectures and provides calibrated uncertainty while matching or improving non-Bayesian accuracy.

BARNN in a Nutshell

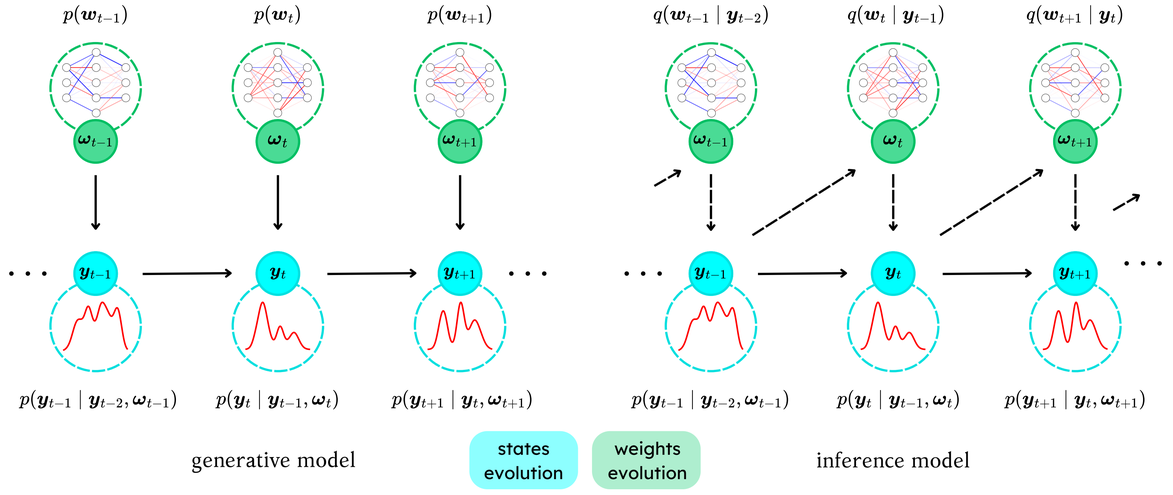

Key Insights. BARNN is a joint Bayesian model where network weights evolve with observable states. It enables structured uncertainty quantification for any recurrent or autoregressive model

Scaling to Large Networks

Key Insights. The posterior introduce variable in time dropout rates combined with Variational Dropout. We derive the tVAMP prior to make inference calibrated and training stable & efficient

$$q(\boldsymbol{\omega}_t^l \mid \boldsymbol{y}_{<t}) = \mathcal{N}(\alpha_t^l \boldsymbol{\Omega}^l, (\alpha_t^l \boldsymbol{\Omega}^l)^2) \quad \alpha_t^l = f(\boldsymbol{y}_{<t}; \boldsymbol{\delta})$$

$$p(\boldsymbol{\omega}_t) = \int q(\boldsymbol{\omega}_t \mid \boldsymbol{y}_{<t}) p(\boldsymbol{y}_{<t}) d\boldsymbol{y}_{<t}$$

Experiments

We validate BARNN on synthetic ***Time Series***, showing the tVAMP prior outperforms standard choices; ***Neural PDE Solvers***, BARNN enables both accurate solutions and UQ; and M***olecular Generation,*** BARNN-powered RNNs outperform existing models using SMILES.

$$\begin{cases} x_t &= x_{t-1} + \frac{3\pi}{100}, \quad x_0 = 0 \\ y_t &= \frac{1}{5}\sum_{j=1}^{5} \sin(\alpha_i x_t + \beta_i) \end{cases} \quad \text{with } \alpha_i \sim U[0.5, 1.5],\ \beta_i \sim U[0, 3\pi]$$

Model	Prior	MSE (↓)	NLL (↓)	ECE (↓)
Static	-	$0.490^{\pm 0.000}$	-	-
MLP	-	$0.081^{\pm 0.011}$	-	-
Dropout (p=0.5)	-	$0.072^{\pm 0.004}$	$0.593^{\pm 0.461}$	$0.084^{\pm 0.010}$
Dropout (p=0.2)	-	$0.048^{\pm 0.004}$	$-0.075^{\pm 0.004}$	$0.068^{\pm 0.009}$
BARNN	log-uniform	$0.045^{\pm 0.003}$	$-0.092^{\pm 0.064}$	$0.050^{\pm 0.016}$
BARNN	tVAMP	$\mathbf{0.043}^{\pm 0.001}$	$\mathbf{-0.166}^{\pm 0.019}$	$\mathbf{0.049}^{\pm 0.008}$

Ablation Study on tVAMP

On synthetic time series, BARNN + tVAMP prior outperforms log-uniform, dropout, and deterministic baselines in prediction (lower RMSE), and uncertainty quantification (lower ECE and NLL)

Accurately Solving PDEs with Uncertainty

BARNN outperforms existing neural PDE solvers on NLL and ECE, offering better uncertainty quantification even when RMSE is similar. It dynamically adjusts confidence intervals based on local uncertainty, demonstrating robust, interpretable, and adaptive UQ across PDEs.

Model	NLL (↓)			ECE (↓)		
	Burgers	*KS*	*KdV*	*Burgers*	*KS*	*KdV*
BARNN	$\mathbf{-2.51}^{\pm 0.29}$	$\mathbf{-0.87}^{\pm 0.52}$	$\mathbf{-0.94}^{\pm 1.23}$	$\mathbf{0.05}^{\pm 0.03}$	$\mathbf{0.04}^{\pm 0.03}$	$\mathbf{0.10}^{\pm 0.01}$
ARD	$5.00^{\pm 7.11}$	$10.7^{\pm 11.6}$	break	$0.26^{\pm 0.10}$	$0.17^{\pm 0.01}$	$0.19^{\pm 0.05}$
Dropout	$0.06^{\pm 4.59}$	$-0.48^{\pm 0.65}$	$0.40^{\pm 4.57}$	$0.17^{\pm 0.15}$	$0.10^{\pm 0.01}$	$0.14^{\pm 0.07}$
Perturb	$-2.01^{\pm 0.08}$	$-0.65^{\pm 1.07}$	$-0.37^{\pm 1.32}$	$0.16^{\pm 0.01}$	$0.05^{\pm 0.02}$	$0.11^{\pm 0.06}$
Refiner	break	break	break	$0.30^{\pm 0.08}$	$0.23^{\pm 0.02}$	$0.34^{\pm 0.09}$

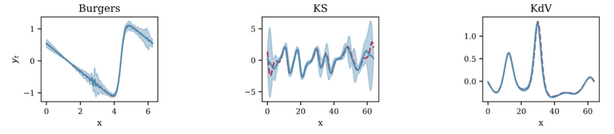

Model	Validity (↑)	Diversity (↑)	Novelty (↑)	Uniqueness (↑)
BARNN	$\mathbf{95.09}^{\pm 0.34}$	$\mathbf{88.63}^{\pm 0.01}$	$\mathbf{95.41}^{\pm 0.24}$	$\mathbf{95.06}^{\pm 0.34}$
SMILES LSTM	$94.60^{\pm 0.27}$	$88.57^{\pm 0.09}$	$94.25^{\pm 0.15}$	$94.58^{\pm 0.28}$
LSTM	$93.02^{\pm 0.33}$	$88.57^{\pm 0.02}$	$92.81^{\pm 0.34}$	$92.14^{\pm 0.15}$

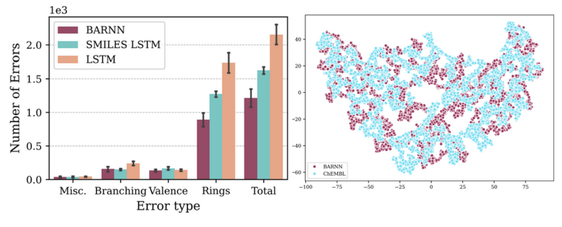

Enhancing Molecular Generation

BARNN generates molecules with higher validity, novelty, and uniqueness than baselines. It effectively captures long-range dependencies, improving ring closure accuracy by up to 50%. It better preserves molecular property distributions, achieving broader and more accurate chemical space coverage.

Perspectives and References

BARNN is a scalable, calibrated Bayesian extension for autoregressive and recurrent models. BARNN is flexible and widely applicable, paving the way for Bayesian deep learning in language, audio, and beyond.

MORE ABOUT BARNN *MORE ABOUT PINA*