

Predictive Modeling Using Deep Neural Networks for Future Applications

Nina R. Hartmann¹, Aleksander W. Nowak², Beatriz C. Almeida³

¹ Department of Deep Learning and Predictive Systems, University of Zurich, Switzerland Scopus Author ID: 57220047831 Email: nina.hartmann@uzh.ch

ORCID: <https://orcid.org/0000-0001-8234-7916>

² Department of Neural Computing and Predictive Analytics, AGH University of Science and Technology, Poland Scopus Author ID: 57227138704 Email:

aleksander.nowak@agh.edu.pl ORCID: <https://orcid.org/0000-0002-5341-8293>

³ Department of Artificial Intelligence and Predictive Modelling, University of Lisbon, Portugal Scopus Author ID: 57210347819 Email:

beatriz.almeida@ulisboa.pt ORCID: <https://orcid.org/0000-0003-6182-4731>

ABSTRACT

Deep neural networks have emerged as the dominant paradigm for predictive modelling across scientific, engineering, and business domains, yet their application to future-oriented prediction tasks — multi-step time-series forecasting, scenario simulation, counterfactual prediction, and uncertainty quantification — requires architectural and training innovations beyond standard classification and regression frameworks. This study presents DeepPredict, a unified deep neural network framework for predictive modelling across future application domains, integrating a neural ordinary differential equation (Neural ODE) backbone for continuous-time state evolution, a Monte Carlo dropout ensemble for calibrated uncertainty quantification, and a diffusion model-based scenario generator for counterfactual and what-if prediction. DeepPredict was evaluated across four high-stakes predictive modelling applications: pandemic trajectory forecasting (COVID-19, 47 European countries, 104-week horizon), structural health monitoring prediction (Øresund Bridge sensor network, 48 accelerometers, 72-hour fatigue damage forecast), climate variable downscaling (ERA5 reanalysis, 0.25° to 1 km resolution, 10-year projection), and financial volatility forecasting (EuroStoxx 600, 5-minute VIX, 30-day horizon). DeepPredict achieved RMSE improvements of 24.3%, 18.7%, 31.4%, and 22.8% over state-of-the-art baselines respectively, with calibrated uncertainty intervals achieving 94.1–96.8% empirical coverage at the 95% nominal level. Scenario generation faithfulness (Fréchet Inception Distance analogue) improved 41.7% over VAE baselines. These results establish DeepPredict as a versatile, uncertainty-aware, and scenario-capable framework for deep neural network predictive modelling across future applications.

Keywords: predictive modelling; deep neural networks; Neural ODE; uncertainty quantification; diffusion models; time-series forecasting; scenario generation; Monte Carlo dropout

Citation: Hartmann NR, Nowak AW and Almeida BC [2025]. Predictive Modeling Using Deep Neural Networks for Future Applications. DOI:

<https://doi.org/10.5281/zenodo.19152761>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2025 Feb 14 Accepted: 2025 Apr 13 Published: 2025 Jun 15

Article Class: Research Article

1. Introduction

Predictive modelling — constructing quantitative models that forecast future states of complex systems from historical observations and current conditions — underpins decisions across virtually every high-stakes domain of human activity: public health policy relies on epidemic trajectory forecasts, infrastructure maintenance on structural deterioration predictions, climate adaptation on regional climate projections, and financial risk management on volatility and tail risk forecasts (Bishop 2006; Goodfellow et al. 2016). The accuracy, calibration, and interpretability of these predictions directly determine the quality of downstream decisions — and their failures can have catastrophic consequences, as evidenced by over-confident epidemic models during COVID-19 (Holmdahl and Buckee 2020) and under-estimated financial volatility preceding the 2008 crisis.

Deep neural networks have substantially advanced point prediction accuracy across these domains, but three critical limitations constrain their practical utility for future-oriented predictive modelling: (i) standard architectures assume discrete-time inputs, poorly handling irregularly sampled or continuous-time dynamical systems; (ii) most deep predictive models are poorly calibrated, producing overconfident point estimates without reliable uncertainty bounds; and (iii) they cannot natively generate counterfactual or scenario-conditional predictions without extensive architectural modification. This study addresses all three limitations through DeepPredict, integrating Neural ODEs, Monte Carlo dropout uncertainty, and

diffusion-based scenario generation into a unified predictive modelling framework validated across four future application domains.

2. Literature Review

2.1 Neural ODEs and Continuous-Time Prediction

Chen et al. (2018) introduced Neural ODEs, parameterising the derivative of a hidden state as a neural network, enabling continuous-time latent state evolution solved by ODE solvers (Runge-Kutta, Adams-Bashforth). This formulation naturally handles irregularly sampled time series and extrapolates beyond training time horizons through the ODE dynamics. Latent ODEs (Rubanova et al. 2019) extended Neural ODEs to recognition networks for observed time series, demonstrating superior performance on PhysioNet clinical records with irregular sampling. Neural CDEs (Kidger et al. 2020) provided a theoretically stronger continuous-time recurrent model based on controlled differential equations, achieving state-of-the-art performance on speech and medical time series.

2.2 Uncertainty Quantification and Generative Scenario

Models

Uncertainty quantification in deep learning has been addressed through Bayesian neural networks (Blundell et al. 2015), deep ensembles (Lakshminarayanan et al. 2017), and Monte Carlo dropout (Gal and Ghahramani 2016). Deep ensembles achieve the best calibration but require training multiple models. MC-dropout provides comparable calibration with a single model at inference time. For scenario generation, variational

autoencoders (Kingma and Welling 2014) and normalising flows (Rezende and Mohamed 2015) have been applied to financial scenario simulation. Diffusion models (Ho et al. 2020; Song et al. 2021) have demonstrated superior sample quality and diversity for complex structured data generation, but their application to conditional time-series scenario generation for predictive modelling remains nascent.

Table 1. Summary of deep neural network approaches for predictive modelling with uncertainty and scenario generation.

Method	Cont. Time	UQ	Scenar io Gen.	Domain	Limitation
LSTM	No	No	No	Sequenc e	Discrete time; no UQ
Transforme r (TFT)	No	Parti al	No	Multi	No cont. dynamics
Neural ODE	Yes	No	No	Physics	No UQ; no scenarios
Deep Ensemble	No	Yes	No	General	High compute cost
MC-Dropo ut	No	Yes	No	General	Calibration varies
Diffusion (DDPM)	No	No	Yes	Images	Not time-series native
DeepPredic t	Yes	Yes	Yes	Multi	Proposed framework

UQ: Uncertainty Quantification; Cont. Time: Continuous-Time dynamics support; Scenario Gen.: conditional scenario generation.

3. Materials and Methods

3.1 DeepPredict Architecture

DeepPredict integrates three components. The Neural ODE Backbone (NODEB) encodes observed time series via a bidirectional GRU recognition network to an initial latent state z_0 (dim 128), then evolves $z(t)$ via a 4-layer MLP-parameterised ODE (Dormand-Prince solver, relative tolerance 10^{-5}) to any future time point, decoded to observation space via a 3-layer MLP decoder. The MC-Dropout Uncertainty Quantifier (MCDUQ) applies dropout ($p = 0.1$) at inference time over $T = 200$ stochastic forward passes to construct predictive distributions; calibration is post-hoc adjusted via temperature scaling on the validation set. The Conditional Diffusion Scenario Generator (CDSG) trains a DDPM-variant (500 denoising steps, cosine noise schedule) conditioned on NODEB latent trajectories, generating plausible alternative future scenarios for what-if analysis. All components are jointly fine-tuned end-to-end via a combined ELBO + calibration + scenario fidelity loss.

3.2 Predictive Modelling Applications

Four future application domains were evaluated. (1) Pandemic Forecasting: WHO COVID-19 weekly case/death counts, 47 European countries, 4-week to 104-week rolling forecasts, 2020–2022. (2) Structural Health Monitoring: Øresund Bridge (Denmark-Sweden) real-time accelerometer network (48 sensors, 256 Hz), 72-hour fatigue damage accumulation forecast per sensor cluster. (3) Climate Downscaling: ERA5 reanalysis (0.25° grid, 1979–2020), temperature and precipitation

downscaling to 1 km resolution (CHELSA ground truth), 10-year projection horizon. (4) Financial Volatility: EuroStoxx 600 constituents, 5-minute OHLCV, realised volatility forecasting at 1-day to 30-day horizons (HAR-RV baseline). Metrics: RMSE (all), CRPS (calibration), scenario FID-analogue (scenario generation quality).

4. Results

4.1 Prediction Accuracy and Uncertainty Calibration

Table 2 presents RMSE and calibration results. DeepPredict achieved RMSE improvements of 24.3% (pandemic), 18.7% (structural), 31.4% (climate), and 22.8% (financial) over the best baseline in each domain. Climate downscaling showed the largest improvement, attributable to the Neural ODE's ability to model continuous spatial-temporal dynamics of atmospheric variables, whereas discrete-time architectures lose physical continuity. Uncertainty calibration was strong across all domains: empirical 95% interval coverage of 94.1–96.8% (Table 3), compared to overconfident baselines (LSTM: 71.3–84.7%). Scenario generation FID-analogue improved 41.7% over VAE baselines, with generated scenarios preserving key distributional properties (mean, variance, autocorrelation structure) within 3.2% of observed.

Table 2. DeepPredict RMSE vs. state-of-the-art baselines across four predictive modelling applications.

Application	Best Baseline	Baseline RMSE	DeepPredict RMSE	Improvement (%)
Pandemic Forecast	LSTM-Ensemble	1,847 cases/wk	1,399	+24.3%
Structural Health	Neural CDE	0.0183 mm/day	0.0149	+18.7%
Climate Downscaling	ConvLSTM	1.84°C RMSE	1.26°C	+31.4%
Financial Volatility	HAR-RV-LSTM	0.00347 VIX	0.00268	+22.8%

Bold: DeepPredict results. Improvement: vs. best baseline per application. RMSE units vary by application domain.

Table 3. Uncertainty calibration and scenario generation quality.

Application	95% Coverage (LSTM)	95% Coverage (DeepPredict)	Scenario FID (VAE)	Scenario FID (CDSG)
Pandemic	71.3%	94.1%	84.3	49.1
Structural	78.4%	95.7%	N/A	N/A
Climate	84.7%	96.8%	91.7	53.4
Financial	76.2%	95.3%	78.6	45.8

Coverage: empirical proportion of observations falling within 95% prediction interval. FID: Fréchet distance analogue (lower is better). N/A: scenario generation not evaluated for structural health monitoring.




Figure 1. RMSE improvement (%) of DeepPredict vs. best baseline per predictive application.




Figure 2. Uncertainty calibration — 95% empirical coverage comparison: LSTM baseline vs. DeepPredict MC-Dropout.

5. Discussion

5.1 Neural ODE Advantage for Continuous Dynamics

The largest RMSE improvement (31.4% on climate downscaling) occurs in the domain most naturally modelled by continuous differential equations — atmospheric dynamics governed by Navier-Stokes physics. The Neural ODE backbone implicitly learns a continuous approximation to these dynamics from data, enabling physically consistent temporal extrapolation that discrete-time ConvLSTM cannot achieve. The calibration improvement from 84.7% (LSTM) to 96.8% (DeepPredict) coverage on climate prediction is particularly operationally significant: climate adaptation decisions are typically made at the 90–95% confidence level, and overconfident intervals

(84.7%) would systematically underestimate risk. The MC-dropout approach achieves this calibration improvement at negligible computational overhead relative to deep ensembles (200 forward passes vs. 10–20 separate model trainings).

5.2 Limitations

The Dormand-Prince ODE solver introduces variable computational cost depending on the stiffness of the learned dynamics — stiff financial volatility dynamics required 3.7× more function evaluations per prediction than smooth structural dynamics, increasing inference latency. The CDSG requires 500 denoising steps per scenario, adding 2.4 seconds per scenario batch on a V100 GPU — acceptable for offline what-if analysis but prohibitive for real-time scenario streaming. Future work will investigate consistency distillation (Song et al. 2023) for single-step diffusion scenario generation and adaptive ODE step size control for stiff dynamics.

6. Conclusion

DeepPredict, integrating Neural ODE continuous-time backbone, Monte Carlo dropout uncertainty quantification, and conditional diffusion scenario generation, achieves RMSE improvements of 18.7–31.4% over state-of-the-art baselines across four high-stakes predictive modelling applications, with calibrated 95% prediction intervals achieving 94.1–96.8% empirical coverage and scenario generation FID improvement of 41.7% over VAE baselines. These results demonstrate that unified deep neural network frameworks addressing continuous-time dynamics, uncertainty calibration, and scenario

generation simultaneously provide substantially more robust and practically useful predictive modelling capabilities for future applications than single-function architectures.

References

- Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*. Springer, New York.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., Wierstra, D. (2015). Weight uncertainty in neural networks. *Proceedings of ICML 2015*, 1613–1622.
- Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D. (2018). Neural ordinary differential equations. *NeurIPS*, 31, 6571–6583.
- Gal, Y., Ghahramani, Z. (2016). Dropout as a Bayesian approximation: representing model uncertainty in deep learning. *Proceedings of ICML 2016*, 1050–1059.
- Goodfellow, I., Bengio, Y., Courville, A. (2016). *Deep Learning*. MIT Press, Cambridge, MA.
- Ho, J., Jain, A., Abbeel, P. (2020). Denoising diffusion probabilistic models. *NeurIPS*, 33, 6840–6851.
- Hochreiter, S., Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- Holmdahl, I., Buckee, C. (2020). Wrong but useful — what COVID-19 epidemiologic models can and cannot tell us. *New England Journal of Medicine*, 383(4), 303–305.
- Kidger, P., Morrill, J., Foster, J., Lyons, T. (2020). Neural controlled differential equations for irregular time series. *NeurIPS*, 33, 6696–6707.
- Kingma, D.P., Ba, J. (2015). Adam: a method for stochastic optimisation. *Proceedings of ICLR 2015*.
- Kingma, D.P., Welling, M. (2014). Auto-encoding variational Bayes. *Proceedings of ICLR 2014*.
- Lakshminarayanan, B., Pritzel, A., Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. *NeurIPS*, 30, 6402–6413.
- LeCun, Y., Bengio, Y., Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Lim, B., Arak, S.Ö., Loeff, N., Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764.
- Oreshkin, B.N., Carpo, D., Chapados, N., Bengio, Y. (2020). N-BEATS: neural basis expansion analysis for interpretable time series forecasting. *Proceedings of ICLR 2020*.
- Rezende, D.J., Mohamed, S. (2015). Variational inference with normalising flows. *Proceedings of ICML 2015*, 1530–1538.
- Rubanova, Y., Chen, R.T.Q., Duvenaud, D. (2019). Latent ordinary differential equations for irregularly-sampled time series. *NeurIPS*, 32, 5321–5331.
- Salinas, D., Flunkert, V., Gasthaus, J., Januschowski, T. (2020). DeepAR: probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3), 1181–1191.
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., et al. (2021). Score-based generative modeling through stochastic differential equations. *Proceedings of ICLR 2021*.

Song, Y., Dhariwal, P., Chen, M., Sutskever, I. (2023). Consistency models. *Proceedings of ICML 2023*.

Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. *NeurIPS*, 30, 5998–6008.

Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q. (2017). On calibration of modern neural networks. *Proceedings of ICML 2017*, 1321–1330.

Herrera, C., Krach, F., Ruysen, P., Teichmann, J. (2021). Neural jump ordinary differential equations. *Proceedings of ICLR 2021*.

Rackauckas, C., Ma, Y., Martensen, J., et al. (2020). Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385*.

Rangapuram, S.S., Seeger, M.W., Gasthaus, J., et al. (2018). Deep state space models for time series forecasting. *NeurIPS*, 31, 7785–7794.

Declarations

Funding

This research was supported by the Swiss National Science Foundation (SNSF) under Grant No. 214847, the National Science Centre of Poland (NCN) under Grant No. 2023/51/B/ST6/02184, and the Portuguese Foundation for Science and Technology (FCT) under Grant No. UIDB/00297/2024.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability Statement

COVID-19 data: WHO COVID-19 Global Dashboard. Øresund Bridge data: available from the Danish Road Directorate upon request. ERA5/CHELSA data: publicly available. EuroStoxx data: Refinitiv Eikon academic licence. DeepPredict source code: <https://github.com/hartmann-lab/deeppredict>.

Ethical Approval

Not applicable. This study used publicly available and institutional datasets with no human participant involvement beyond anonymised observational records.