

COMBINING NEURAL NETWORKS WITH GAUSSIAN PROCESSES FOR ROBUST AND INTERPRETABLE PREDICTION

Amalia Macali, Georgios Soimouris, Nathan Eskue, Vahid Yaghoubi

1

Q-VAIbe Lab, Department of Aerospace Structures and Materials,
Faculty of Aerospace Engineering,
Delft University of Technology Kluyverweg 1, Delft, 2629 HS, the Netherlands
a.macali@tudelft.nl
g.soimouris@tudelft.nl

Keywords: Uncertainty Quantification, Artificial Neural Network, Gaussian Process

Abstract. *This paper presents a hybrid model that combines Artificial Neural Networks (ANN) and Gaussian Processes (GP). The goal is to achieve high prediction accuracy while quantifying uncertainty. The proposed structure is a simple ANN used as the trend of the GP, particularly emphasizing the joint training of the two models. The ANN+GP exploits the ability of the ANN to capture complex, non-linear relationships in the data. At the same time, the GP provides an approach to uncertainty estimation, thus improving the accuracy of the predictions. This paper emphasizes the importance of concurrent training, which can improve the accuracy of the prediction model. The algorithm is suitable for any application where both accurate, robust predictions and uncertainty estimates are critical to enhance the interpretability of the model. The proposed method has been successfully applied to the frequency response functions of a simple structure.*

1 Introduction

Machine learning has become widely used for predictive modeling, from engineering and finance to the natural sciences. Two of the most commonly used techniques are Artificial Neural Networks (ANN) and Gaussian Processes (GP). ANNs identify complex, non-linear relationships in large datasets, making them ideal for classification, regression, and feature extraction [1]. The weakness of ANNs is that even if they are highly effective at making predictions, they do not naturally provide uncertainty estimates, a critical component in different fields like healthcare, finance, and engineering, where decision-making depends on knowing how confident one can be in a prediction [2].

While the GP is a probabilistic model, rather than providing a single point prediction, GPs furnish the underlying function as a distribution, allowing them to generate confidence intervals around their predictions [3]. This probabilistic capability is beneficial in scenarios where understanding a prediction’s uncertainty is just as important as the prediction itself. However, GP scalability is challenging; in fact, the complexity increases with the size of the data, which limits their applicability to more minor data sets [4].

Given these models’ complementary strengths and weaknesses, combining ANN and GP holds significant potential. Hybrid models have been proposed before, but the ANN and GP are continuously trained in different steps, optimized independently [5, 6, 7, 8]. Inspired by [9], this paper presents a combination technique in which ANN is used as the mean function of a GP and applied to two real-world spatiotemporal data sets. Their model uses stochastic variational inference and sparse GP approximations to manage scalability and uncertainty estimation. This integration brings a more effective interaction between the two models, leveraging their strengths in a unified framework.

2 Related Work

While hybrid ANN-GP models have been previously explored, few methods enable a fully unified training procedure where both components influence each other throughout optimization. Most approaches optimize ANN and GP separately or rely on variational approximations that decouple the learning process.

In contrast, the method proposed in this work integrates a simple ANN as a trend of an Exact GP. This architecture promotes tighter coupling between feature extraction and uncertainty modeling, improving calibration and prediction performance.

Deep Gaussian Processes (DGPs) enhance model capacity by increasing the depth within the GP framework, stacking multiple GP layers to capture hierarchical uncertainty [11]. Other approaches such as Gaussian Process Regression Networks (GPRNs) [12] aim to model complex, multi-output relationships by combining GPs with latent neural network structures. However, these methods introduce considerable inference complexity and treat the GP as the central component. On the other hand, several methods focus on introducing uncertainty directly into neural networks. Bayesian Neural Networks (BNNs) [13] adopt a probabilistic treatment of the network weights, while Monte Carlo Dropout [14] estimates uncertainty by performing stochastic forward passes at test time. These techniques maintain the neural architecture as the core model but add uncertainty through approximate inference mechanisms.

In contrast, the method proposed in this work avoids sequential or independent optimization by enabling joint training of both components, ANN and GP. This unified structure improves generalization, ensures stable convergence, and provides well-calibrated uncertainty estimates.

By optimizing both algorithms simultaneously, the generalization ability of the ANN is en-

hanced while providing robust uncertainty quantification through the GP. This unified optimization improves both predictive accuracy and uncertainty estimation. The results are obtained using frequency response prediction (FRF) data for mechanical systems [10].

The proposed method is compared with baseline models, including a standalone Gaussian Process and a standalone Artificial Neural Network. Although inspired by this architecture [9], the implementation adopted in this work is deliberately simplified, relying on an exact GP and avoiding variational inference or sparse approximations.

3 Methodology

3.1 GP with Neural Network Mean Function (ANN+GP)

This configuration uses the neural network as a parametric mean function within a Gaussian Process. The GP models the residual structure around the predictions made by the ANN, allowing for a probabilistic treatment of the output. An ANN captures the underlying data trend, while the GP accounts for local variations and provides uncertainty estimates. However, unlike their implementation, which relies on sparse approximations and variational inference for scalability, an exact GP model was chosen in the following experiments. A Gaussian Process is defined as:

$$f(x) \sim \mathcal{GP}(m(x), k(x, x')) \quad (1)$$

where $m(x)$ is the mean function and $k(x, x')$ is the covariance function.

In this configuration, the mean function $m(x)$ is modeled by a neural network:

$$m(x) = NN(x) = W_2 \sigma(W_1 x + b_1) + b_2 \quad (2)$$

where $\sigma(\cdot)$ denotes a non-linear activation function, such as ReLU.

The covariance function is defined using a Matérn kernel:

$$k_\nu(r) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}r}{\ell} \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}r}{\ell} \right) \quad (3)$$

where $r = \|x - x'\|$, ℓ is the lengthscale, and K_ν is the modified Bessel function of the second kind.

The resulting model combines the expressiveness of the neural network with the probabilistic nature of the GP:

$$f(x) \sim \mathcal{GP}(NN(x), k(x, x')) \quad (4)$$

This formulation allows the GP to model uncertainty around a learned trend. This work implements the model using an exact GP, avoiding variational approximations or sparse methods.

4 Results

4.1 Structural Dynamics Dataset

The structural dynamics application is based on a two-degree-of-freedom (2-DOF) mechanical system composed of two masses connected by linear springs and dampers. A harmonic excitation is applied to the first mass, and the system response is measured at both degrees of freedom, resulting in two complex-valued output channels. Uncertainty is introduced in the stiffness parameter, modeled as $k = k_0(1 + \delta_k \xi)$, where $\delta_k = 5\%$ and $\xi \sim \mathcal{N}(0, 1)$.

FRFs are computed over a frequency range of 10–35 Hz with a resolution of 0.01 Hz, resulting in 2501 frequency points per output for 100 samples. The resulting dataset has dimensions

$2 \times 2501 \times 100$. Each complex-valued output is decomposed into real and imaginary components, which are then normalized using a MinMax scaler.

Principal Component Analysis (PCA) is applied separately to the real and imaginary parts, retaining 99% of the variance. This reduces the dimensionality to 18 components for the real part and 19 for the imaginary part, maintaining computational efficiency while preserving relevant information.

4.2 Neural Architecture and Gaussian Process Configuration

To ensure a fair comparison between models, the same neural architecture is used for both the standalone ANN and the hybrid ANN+GP models. The architecture consists of a feedforward neural network with two hidden layers with 8 and 16 neurons, respectively, each followed by a ReLU activation function.

For the ANN+GP model, this neural network serves as the mean function of the Gaussian Process. The GP component uses a Multitask Matern kernel with $\nu = 2.5$, allowing it to jointly model all output components while capturing smooth variations in the signal. For the standalone GP, the same kernel and likelihood structure are used. In both cases, the model is trained using the exact GP formulation with a multitask Gaussian likelihood.

This architectural alignment ensures that any performance differences can be attributed to the modeling strategy (pure GP, pure ANN, or hybrid) rather than differences in network capacity or training procedure.

4.3 Results

The performance of the proposed ANN+GP architecture is evaluated against two baseline models: a standalone Artificial Neural Network (ANN) and a standalone Gaussian Process (GP).

The evaluation includes the following metrics: Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), coefficient of determination (R^2), and execution time.

It is important to highlight that the R^2 values reported in this work are directly obtained from the training pipeline, as computed within the PCA-transformed space (i.e., the reduced latent representation). In contrast, MSE, RMSE, and MAE are computed after reconstructing the full FRF signal from PCA, and thus reflect prediction accuracy on the original, full-resolution dataset.

Figure 1 shows a comparison of MSE, RMSE, and MAE values across the two datasets. The ANN+GP model achieves a test MSE of approximately 11.2 (real) and 11.6 (imag), clearly outperforming the ANN baseline (23.6 real, 24.1 imag) and closely matching the performance of the standalone GP model (11.5 real, 11.0 imag).

These results show that the GP model achieves the best overall performance regarding both error metrics. However, the ANN+GP model closely matches the GP regarding MSE, RMSE, and MAE on the reconstructed signal while significantly outperforming the standalone ANN across all metrics. Although not as expressive as GP in the latent space, ANN+GP still provides consistent predictions with lower error and better generalization than ANN.

Figure 2 presents the R^2 scores obtained on the test sets. The GP model achieves the highest values with $R^2 = 0.61$ (real) and $R^2 = 0.59$ (imag), followed by ANN+GP with $R^2 = 0.48$ (real) and $R^2 = 0.35$ (imag). The ANN model performs worst, with $R^2 = 0.20$ (real) and $R^2 = 0.11$ (imag).




Figure 1: Comparison of MSE, RMSE, and MAE values for each model on real and imaginary components of the FRF.

This suggests that GP best captures the latent dynamics in the reduced space, while ANN+GP maintains a strong balance between accuracy and robustness—significantly outperforming the ANN.




Figure 2: R^2 scores on the test sets for the real and imaginary components. GP achieves the highest latent-space performance, while ANN+GP offers a solid compromise between expressiveness and generalization.

Figures 3 and 4 provide a qualitative visualization of the reconstructed signals for the imaginary and real components, respectively. The top panel in each figure shows the true signal, followed by the predicted signals for each model.

Both ANN+GP and GP accurately capture the key dynamics of the FRF, especially around resonance peaks. The ANN+GP results appear smoother and more consistent with the ground truth, with well-calibrated uncertainty. In contrast, the ANN model struggles with signal complexity and produces unstable predictions, particularly around sharp transitions.

5 Conclusion

This work presents a hybrid modeling framework that integrates Artificial Neural Networks (ANNs) with Gaussian Processes (GPs) by using the ANN as the mean function of an exact GP. The proposed ANN+GP architecture is jointly trained and aims to combine the expressiveness of neural networks with the uncertainty quantification of GPs.




Figure 3: Predicted vs true values for the imaginary part of the FRF. ANN+GP and GP closely match the ground truth; ANN introduces artifacts and underestimates uncertainty.

The methodology was evaluated using a structural dynamics application involving predicting frequency response functions (FRFs) under parametric uncertainty. Results show that the




Figure 4: Predicted vs true values for the real part of the FRF. ANN+GP achieves the best visual match to the ground truth, with stable predictions and well-modeled uncertainty.

standalone GP consistently achieves the highest R^2 scores in the latent PCA space and slightly outperforms the hybrid model regarding MSE and other error metrics. However, the ANN+GP

model performs comparably to the GP while significantly improving over the standalone ANN across all evaluation metrics.

Although the ANN+GP framework has shown promising results in spatio-temporal datasets, its performance in structural dynamics remains limited. This suggests that alternative modeling strategies may be required to fully capture the underlying physics of such systems. As a potential remedy, future work will explore architectures where the Gaussian Process is used together with the neural network, allowing for more flexible uncertainty-aware predictions tailored to the dynamic behavior of structural systems.

Acknowledgements

This project is made possible in part by a contribution from the National Growth Fund program NXTGEN HIGHTECH, which is about to be launched. Its program will invest as much as 1 billion euros until 2030 with over 330 partners, in more than 60 projects and in six essential domains. I doing so, NXTGEN HIGHTECH will make a significant contribution to the structural and sustainable economic growth in the Netherlands and offer solutions for the major societal challenges in the areas of energy transition, health, safety and food. For more information, please visit NXTGEN hightech.

REFERENCES

- [1] Giulia Vilone and Luca Longo, *Explainable Artificial Intelligence: a Systematic Review*, arXiv, 2020. <https://arxiv.org/abs/2006.00093>
- [2] Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U Rajendra Acharya, and others, *A review of uncertainty quantification in deep learning: Techniques, applications and challenges*, **Information Fusion**, 76:243–297, 2021. Elsevier.
- [3] Matthias Seeger, *Gaussian Processes for Machine Learning*, **International Journal of Neural Systems**, 14(02):69–106, 2004. <https://doi.org/10.1142/S0129065704001899>. PMID: 15112367.
- [4] Alexander C. Ogren, Berthy T. Feng, Katherine L. Bouman, and Chiara Daraio, *Gaussian process regression as a surrogate model for the computation of dispersion relations*, **Computer Methods in Applied Mechanics and Engineering**, 420:116661, 2024. Elsevier.
- [5] Ruzicka, J., Koza, J., Tumpach, J., Pitra, Z., & Holena, M. (2021). Combining Gaussian Processes with Neural Networks for Active Learning in Optimization. I *IAL@PKDD/ECML* (pp. 105–120).
- [6] Gao, J., Wang, J., Xu, Z., Wang, C., & Yan, S. (2023). Multiaxial fatigue prediction and uncertainty quantification based on back propagation neural network and Gaussian process regression. *International Journal of Fatigue*, 168, 107361.
- [7] Farid, M. (2022). Data-driven method for real-time prediction and uncertainty quantification of fatigue failure under stochastic loading using artificial neural networks and Gaussian process regression. *International Journal of Fatigue*, 155, 106415.

- [8] Y. Ide, S. Ozaki, S. Izutsu, T. Kotoura, M. Yamashiro, and M. Kodama, "Hybrid real-time wave forecasting model combining Gaussian process regression and neural networks," *Ocean Engineering*, vol. 312, p. 119028, 2024.
- [9] Iwata, T., & Ghahramani, Z. (2017). Improving output uncertainty estimation and generalization in deep learning via neural network Gaussian processes. *arXiv preprint arXiv:1707.05922*.
- [10] V. Yaghoubi, S. Marelli, B. Sudret, and T. Abrahamsson, "Sparse polynomial chaos expansions of frequency response functions using stochastic frequency transformation," *Probabilistic Engineering Mechanics*, vol. 48, pp. 39–58, 2017, doi: <https://doi.org/10.1016/j.probengmech.2017.04.003>.
- [11] Andreas Damianou and Neil D. Lawrence, *Deep Gaussian Processes*, **Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS)**, PMLR 31:207–215, Scottsdale, Arizona, USA, 2013.
- [12] Andrew Gordon Wilson, David A. Knowles, and Zoubin Ghahramani, *Gaussian Process Regression Networks*, arXiv preprint arXiv:1110.4411, 2011. Available at: <https://arxiv.org/abs/1110.4411>.
- [13] Igor Kononenko, *Bayesian neural networks*, **Biological Cybernetics**, 61(5):361–370, 1989. Springer.
- [14] Daily Milanés-Hermosilla, Rafael Trujillo Codorníu, René López-Baracaldo, Roberto Sagaró-Zamora, Denis Delisle-Rodriguez, John Jairo Villarejo-Mayor, and José Ricardo Núñez-Álvarez, *Monte Carlo Dropout for Uncertainty Estimation and Motor Imagery Classification*, **Sensors**, 21(21):7241, 2021. MDPI.