Graph-Informed Neural Networks

Søren Taverniers,[1] Eric J. Hall,[2] Markos A. Katsoulakis,[3] Daniel M. Tartakovsky[4]

[1] Palo Alto Research Center (PARC), 3333 Coyote Hill Road, Palo Alto, CA 94304, USA
[2] Division of Mathematics, University of Dundee, Dundee, DD1 4HN, UK
[1] Department of Mathematics and Statistics, University of Massachusetts Amherst, Amherst, MA 01003, USA
[4] Department of Energy Resources Engineering, Stanford University, Stanford, CA 94305, USA
ehall001@dundee.ac.uk (Eric J. Hall), tartakovsky@stanford.edu (Daniel M. Tartakovsky)

Abstract

Graph-Informed Neural Networks (GINNs) present a strategy for incorporating domain knowledge into scientific machine learning for complex physical systems. The construction utilizes probabilistic graphical models (PGMs) to incorporate expert knowledge, available data, constraints, etc. with physics-based models such as systems of ordinary and partial differential equations (ODEs and PDEs). Computationally intensive nodes in this hybrid model are replaced by the hidden nodes of a neural network (i.e., learned features). Once trained, the resulting GINN surrogate can cheaply generate physically-relevant predictions *at scale* thereby enabling robust sensitivity analysis and uncertainty quantification (UQ). As proof of concept, we build a GINN for a multiscale model of electrical double-layer capacitor dynamics embedded into a Bayesian network (BN) PDE hybrid model.

In recent years, several approaches have been proposed to inform deep neural networks (DNNs) of physical laws and constraints to ensure they produce physically sound predictions. Two main classes of DNNs for building surrogate representations of physics-based models described by PDEs have emerged: physics-informed NNs (PINNs) (Raissi, Perdikaris, and Karniadakis 2019) and “data-free” physics-constrained NNs (Zhu et al. 2019). Our approach uses the well-known concept of PGMs to embed domain knowledge, including correlations between control variables (CVs), into standard DNNs by only modifying their input layer structure and enabling the use of a standard penalty in the loss function, e.g., ℓ_1 (lasso regression) or ℓ_2 (ridge regression) regularization. This non-intrusive approach permits the use of off-the-shelf software like `TensorFlow` or `PyTorch` with minimal effort from the user, while remaining compatible with PINNs and other customized NN architectures which can be used to replace individual computational bottlenecks in the physics-based representation.

GINNs are particularly suited to enhance the computational workflow for complex systems featuring intrinsic computational bottlenecks and intricate physical relations among input CVs. Hence, to showcase the potential of this approach, we apply a GINN to simulation-based decision-making in electrical double-layer (EDL) supercapacitors, where it is deployed to build highly accurate kernel density estimators (KDEs) for the probability density functions (PDFs) of relevant output quantities of interest (QoIs).

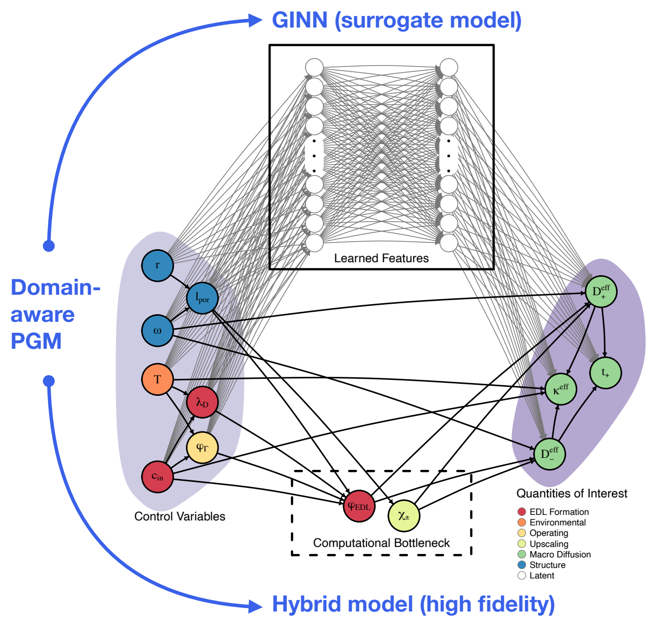

Figure 1: A domain-aware PGM encoding structured priors on CVs serves as input to both the BN PDE (lower route) and trained GINN (upper route) for a homogenized model of ion diffusion in supercapacitors (Taverniers et al. 2020).

Constructing and training a GINN

Simulation-based decision-making for design tasks involving complex multiscale/multiphysics systems requires predicting the impact of tunable CVs on the system’s QoIs. Typically, this is modeled by recasting the problem in a probabilistic framework where CVs and QoIs are represented as random quantities that can be sampled from their corresponding probability distributions. For most real-world applications, these are continuous, non-Gaussian variables that need to be characterized by their full PDF rather than through a finite set of moments.

Figure 1 visualizes the construction of a GINN surrogate for a multiscale model of EDL supercapacitor dynamics. A BN, a type of directed acyclic PGM, systematically incorporates domain knowledge into the physics-based model through structured priors on CVs, resulting in a hybrid BN PDE model for macroscopic diffusion QoIs. The GINN retains the structured priors as inputs but replaces the hybrid model's computationally intensive nodes, related to upscaling via homogenization, with learned features to speed up the generation of QoIs while maintaining physical relevance.

The GINN workflow, summarized in Fig. 2, consists of:

1. **Data generation:** Generate N_{sam} input-output (io) samples, divided into N_{train} training and N_{test} test samples.
2. **Training:** Train the GINN with N_{train} training samples.
3. **Testing:** Test the trained GINN's ability to handle unseen data using the N_{test} test samples.
4. Repeat steps 1 through 3 (modifying N_{train}) until both the training and test error tolerance are satisfied.
5. **Prediction:** Draw $N_{\text{sam}}^{\text{pred}}$ inputs from the structured priors on the CVs and predict corresponding QoIs with the trained GINN surrogate.

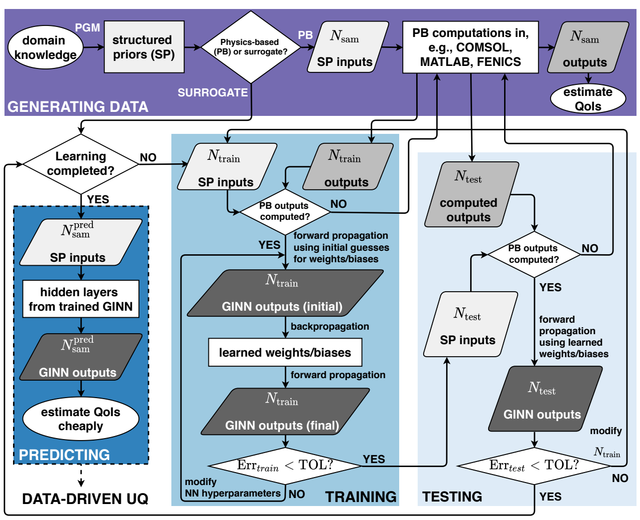

Figure 2: Overview of the global algorithm for GINN-based training, testing, and predicting (Hall et al. 2021).

GINN-based decision-making

A GINN's ability to cheaply generate io sample pairs can be leveraged to construct KDEs for the marginal and joint PDFs of QoIs with appropriate confidence intervals. Such nonparametric estimators form the building blocks for UQ tasks such as sensitivity analysis.

In Fig. 3, we plot KDEs for QoIs based on 8×10^3 samples simulated using the BN PDE (the minimum amount of io data needed to train the GINN) and on 10^7 samples predicted with the GINN. We find that the GINN-predicted KDEs do not include spurious features observed with the smaller, expensive-to-compute data set generated with the physics-based model, and achieve much tighter confidence intervals for an equivalent computational cost (since learning the GINN's parameters and predicting new data with the GINN carries a negligible computational expense).

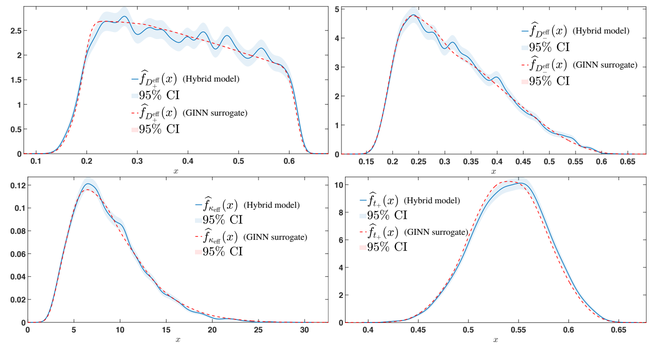

Figure 3: Estimated marginal densities for the QoIs in the supercapacitor testbed based on 8×10^3 samples computed with the hybrid BN PDE (solid/blue) or 10^7 samples computed with the GINN (dashed/red) (Hall et al. 2021).

Conclusions

Our full analysis, in (Hall et al. 2021; Taverniers et al. 2020), suggests that GINNs, which take structured PGMs as inputs, produce physically relevant QoIs that can be used to generate KDEs for robust and reliable sensitivity analysis and further UQ. Trained on a small set of high-fidelity input-output data from a domain-aware hybrid model, GINNs can quickly generate large amounts of output predictions, yielding an approach that is orders of magnitude faster than counterparts that rely on physics-based models alone.

Acknowledgments

This work was performed while S. T. was employed by Stanford University.

References

Hall, E. J.; Taverniers, S.; Katsoulakis, M. A.; and Tartakovsky, D. M. 2021. GINNs: Graph-Informed Neural Networks for Multiscale Physics. *J. Comput. Phys.* 433: 110192. doi:10.1016/j.jcp.2021.110192. Share link `authors.elsevier.com/a/1ccIO508Hokch` valid until 2021-04-10.

Raissi, M.; Perdikaris, P.; and Karniadakis, G. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378: 686–707.

Taverniers, S.; Hall, E. J.; Katsoulakis, M. A.; and Tartakovsky, D. M. 2020. Mutual Information for Explainable Deep Learning of Multiscale Systems. ArXiv:2009.04570.

Zhu, Y.; Zabaras, N.; Koutsourelakis, P.-S.; and Perdikaris, P. 2019. Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *J. Comput. Phys.* 394: 56–81.