

# Conformal Uncertainty-Aware Physics-Informed Neural Networks for Reliable Electromagnetic Field Prediction in VLSI Interconnects

Sofia Duarte<sup>1</sup>, Benjamin Hart<sup>2,\*</sup>

<sup>1</sup> Department of Computer Science, Virginia Tech, Blacksburg, VA, USA

<sup>2</sup> Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA, USA

\* Corresponding author: (Email: ben.hart6@vt.edu)

---

**Abstract:** Electromagnetic field prediction in very-large-scale integration (VLSI) interconnects is increasingly constrained by two competing requirements: field solvers must remain physically faithful near conductor edges and material interfaces, yet they must also provide calibrated reliability indicators for fast design-space exploration. This paper proposes a conformal uncertainty-aware physics-informed neural network (UQ-PINN) for quasi-static electromagnetic field approximation in a two-layer interconnect cross-section. The method combines a random-feature neural representation with boundary, sparse observation, and electroquasistatic residual constraints, and then converts ensemble uncertainty into finite-sample prediction intervals using split conformal calibration. The scope is deliberately reproducible: no measured silicon data are claimed. A deterministic finite-difference solution of  $\text{div}(\epsilon \nabla \phi) = 0$  is used as the reference benchmark, and all reported values are generated by the accompanying code. On the benchmark, the physics-informed ensemble obtains a potential MAE of 0.02057 V and an electric-field magnitude relative L2 error of 17.69%, improving the sparse data-only neural baseline. The conformalized interval achieves 90.60% empirical coverage for a 90% target, whereas the uncalibrated ensemble interval reaches only 66.63%. These results indicate that physics constraints and conformal calibration can jointly support reliable electromagnetic surrogates for early-stage EDA analysis without overstating the substitute for sign-off full-wave verification.

**Keywords:** Physics-informed neural networks, conformal prediction, uncertainty quantification, VLSI interconnects, electromagnetic field prediction, parasitic extraction, EDA.

---

## 1. Introduction

Accurate electromagnetic modeling of interconnects has become a central requirement in advanced integrated circuit design. As wire dimensions shrink and operating bandwidths expand, delay, crosstalk, substrate coupling, and signal-integrity risks are increasingly controlled by field distributions rather than by simple lumped approximations. Classical extraction and full-wave solvers remain indispensable, but their computational cost can restrict the number of geometries that can be examined during early exploration and optimization. This tension has motivated neural surrogates that learn from field-solver data, but purely data-driven models may become overconfident when evaluated near conductor edges, dielectric interfaces, or layout patterns that were sparse in the training set.

[1] provides a useful and positive starting point for this direction because it frames PINNs as mesh-free field approximators for VLSI and RF EDA, emphasizes Maxwell-equation constraints, and highlights the value of adaptive sampling and semiconductor-specific material handling. Building on that direction, the present paper narrows the problem to a reproducible electroquasistatic interconnect benchmark and adds a reliability layer that is missing from many high-speed surrogate discussions: calibrated uncertainty. Instead of reporting only point prediction accuracy, we require the model to state how wide an interval should be around its prediction at each spatial location.

The main contribution is a compact UQ-PINN workflow suitable for fast experimentation. First, a random-feature neural field represents the potential in a VLSI-like cross-

section containing a signal conductor, grounded outer boundaries, silicon substrate, and inter-layer dielectric. Second, physics information is imposed through a discrete residual of the governing electroquasistatic equation, while sparse point values emulate limited field-solver samples. Third, a bootstrap/random-feature ensemble provides heterogeneity-aware uncertainty estimates, and split conformal prediction calibrates these estimates to obtain finite-sample coverage on held-out field points. The experiment is intentionally honest about scope: it uses synthetic, physics-generated reference data rather than fabricated measurements, and it does not claim replacement of sign-off three-dimensional commercial solvers.

## 2. Related Work

PINNs were introduced as neural approximators trained to satisfy both observations and differential equation residuals [2]. Subsequent reviews have shown that physics-informed machine learning can reduce data requirements and improve physical consistency, while also noting optimization and scaling challenges in high-dimensional or multiscale problems [3]. Libraries such as DeepXDE helped standardize the implementation of PINNs and residual-based refinement [4]. For electromagnetic problems, MaxwellNet demonstrated physics-driven neural training based directly on Maxwell's equations, showing that neural solvers can be trained with electromagnetic residuals rather than only with precomputed simulation labels [5].

For EDA, parasitic extraction has a long history of accelerated numerical solvers. FastCap is a classic example of multipole-accelerated capacitance extraction for three-

dimensional conductor structures [16], while recent work has surveyed the continuing challenges of parasitic extraction in advanced technologies [20]. Machine learning has entered this area through hybrid capacitance extraction [17] and CNN-based capacitance models for full-chip settings [18], [19]. These methods are valuable, but they usually emphasize point accuracy and throughput. The reliability of a neural field prediction, especially near sparse or geometry-sensitive regions, is less often treated as a first-class output.

Uncertainty quantification provides the missing reliability channel. Deep ensembles are a practical method for predictive uncertainty [9], and dropout can be interpreted as approximate Bayesian inference under certain assumptions [10]. However, raw neural uncertainty is often miscalibrated. Conformal prediction provides a distribution-free calibration layer that can wrap almost any regression model and deliver finite-sample marginal coverage under exchangeability [12]–[15]. This paper brings those ideas into a VLSI interconnect field-prediction setting by conformalizing the uncertainty of a physics-informed neural ensemble.

### 3. Problem Formulation

The target is the scalar electric potential  $\phi(x,y)$  in a normalized two-dimensional interconnect cross-section. The benchmark represents a signal conductor embedded in an inter-layer dielectric above a higher-permittivity substrate. The outer boundary is grounded and the signal conductor is held at 1 V. This setting is not a full RF sign-off model; it is an electroquasistatic cross-sectional field problem that captures a core component of interconnect capacitance and local electric-field estimation.

The governing equation is  $\text{div}(\epsilon(x,y) \text{grad } \phi(x,y)) = 0$  in the dielectric region, with Dirichlet constraints on the signal and ground boundaries. The electric field is then  $E = -\text{grad } \phi$ . The dielectric stack uses  $\epsilon_r = 3.9$  in the inter-layer dielectric and  $\epsilon_r = 11.7$  in the substrate. The reference solution is produced by a finite-difference linear system with face-averaged permittivity, not by manual construction of expected values. This matters because the field distribution has edge crowding near the signal conductor and a material discontinuity at the substrate interface.

The central modeling question is whether a neural surrogate trained with sparse potential observations and boundary constraints can predict the field accurately and identify where its predictions are less reliable. In an EDA context, such a surrogate would be most useful before expensive three-dimensional verification, when designers need rapid insight into geometry sensitivity, field crowding, and possible crosstalk hot spots.


## 4. Methodology

### 4.1. Physics-Generated Benchmark

The benchmark generator constructs a 58 by 58 finite-difference grid over a normalized square domain. The signal conductor is centered at  $x = 0.50$  and  $y = 0.58$ , with normalized width 0.18 and thickness 0.08. The substrate occupies the lower 28% of the vertical domain. Unknown potentials are solved from a sparse linear system, while conductor and outer shield nodes are fixed. The maximum finite-difference residual reported by the solver is  $9.10e-15$ , which confirms that the reference field is numerically consistent for the purpose of this small benchmark.

Training data are sampled from the reference field but are

deliberately sparse. A set of interior samples is biased toward the signal conductor because edge crowding is the most difficult region for the surrogate. Boundary samples include both the grounded shield and the signal conductor. Collocation points are drawn from the dielectric region and are used only for enforcing physics residuals, not for providing reference labels.



**Figure 1.** Finite-difference reference potential for the normalized VLSI interconnect cross-section. The black contour marks the signal conductor

### 4.2. Random-Feature Physics-Informed Network

To keep the experiment fast and reproducible, the neural surrogate uses a single-hidden-layer random-feature network, also known as an extreme-learning style network. The hidden weights and biases are drawn randomly and fixed, while output weights are solved by ridge regression. This choice follows the same efficiency motivation as physics-informed extreme learning machines [6], while retaining the essential PINN idea that the model is trained against physical residuals. The field approximation has the form  $\hat{\phi}(x,y) = \beta^T \tanh(W[x,y]^T + b)$ .


The data-only baseline solves a weighted least-squares problem using sparse interior values and boundary constraints. The physics-informed model adds rows corresponding to a discrete Laplace residual evaluated at collocation points. In a variable-permittivity problem this residual is a simplification away from material interfaces; the finite-difference reference still includes the variable-permittivity interface. This conservative design prevents the paper from claiming more than the experiment demonstrates. The benefit of the residual is therefore interpreted as improved smoothness and physical consistency in dielectric subregions rather than as a complete replacement for an interface-aware Maxwell solver.

### 4.3. Uncertainty and Conformal Calibration

Uncertainty is estimated with an ensemble of eight random-feature PINNs. Each member uses a different random feature map and bootstrap resampling of the sparse interior, boundary, and collocation sets. The ensemble mean is used as the point prediction, and the ensemble standard deviation is used as a local uncertainty scale. Because ensemble standard deviations are not guaranteed to be calibrated, split conformal prediction is applied on held-out field points.

For calibration point  $i$ , the nonconformity score is  $s_i = |\phi_i - \mu_i| / (\sigma_i + 10^{-4})$ , where  $\mu_i$  and  $\sigma_i$  are the ensemble mean and standard deviation. The conformal multiplier  $q$  is the finite-sample upper quantile for a target error rate  $\alpha = 0.10$ . The final interval at a new point is  $[\mu$

-  $q(\sigma+10^{-4})$ ,  $\mu + q(\sigma+10^{-4})$ ]. This construction does not require Gaussian residuals and is therefore attractive for field-prediction settings where errors are spatially nonuniform.



**Figure 2.** UQ-PINN mean potential prediction. The field shape follows the FDM reference while preserving smooth behavior in dielectric regions

### 5. Experimental Setup

All experiments were executed with the supplied Python script. The script depends on NumPy, SciPy, pandas, and Matplotlib. The reference field is generated internally; no external or proprietary dataset is used. This choice was made because open, labeled VLSI electromagnetic field datasets with layout geometry, dielectric stack, and pointwise field labels are not readily available for this specific topic. Rather than inventing measured data, the experiment uses a deterministic solver benchmark and clearly labels it as synthetic.

The evaluation uses four metrics. Potential MAE and RMSE measure pointwise scalar-field accuracy against the finite-difference reference. The electric-field magnitude relative L2 error measures whether gradients of the learned potential reproduce the reference field away from immediate conductor discontinuities. A discrete Laplace residual measures physics consistency on non-fixed interior nodes. For uncertainty, empirical coverage and mean interval width are measured on held-out test points not used for calibration.

The main comparison is between a sparse data-only random-feature neural network and the physics-informed random-feature ensemble. The data-only baseline has access to the same sparse potential and boundary samples but does not use the PDE residual. Therefore, improvements in field-gradient behavior and residual consistency can be attributed to the physics-informed rows of the regression system rather than to additional labeled data.

**Table 1.** Benchmark and training configuration

| Item                    | Value                |
|-------------------------|----------------------|
| Grid resolution         | 58 x 58              |
| Signal potential        | 1 V                  |
| Ground potential        | 0 V                  |
| ILD / substrate eps_r   | 3.9 / 11.7           |
| Interior sparse samples | 180                  |
| Boundary samples        | 520                  |
| Collocation samples     | 500                  |
| Ensemble size           | 8                    |
| Random features/model   | 220                  |
| Reference data source   | Synthetic FDM solver |

### 6. Results and Discussion


Table II summarizes the point-prediction and physics-consistency results. The physics-informed ensemble reduces potential MAE from 0.03587 V to 0.02057 V and potential RMSE from 0.06499 V to 0.02945 V. The improvement is more pronounced in the electric-field metric: the relative L2 error in field magnitude decreases from 91.85% to 17.69%. This is important because capacitance, coupling, and reliability concerns are often driven by field gradients rather than by potential values alone.

The residual result supports the same conclusion. The data-only network obtains a discrete Laplace RMSE of 68.36, while the physics-informed ensemble reduces it to 11.05. The absolute residual scale depends on grid normalization and should not be compared across unrelated geometries, but the within-benchmark reduction shows that the physics-informed model is less likely to create nonphysical oscillations between sparse observations.

Figure 3 shows that the largest errors remain near the signal conductor and nearby high-gradient regions. This is expected because rectangular conductor edges create strong field crowding and because finite differences around fixed-potential metal nodes introduce sharp transitions. The point is not that the neural model eliminates these difficult regions, but that it identifies them through larger conformal uncertainty, as shown in Fig. 4. This behavior is useful in EDA because uncertain hot spots can be routed to a higher-fidelity solver rather than accepted blindly.

**Table 2.** Accuracy and physics-consistency results

| Model     | MAE     | RMSE    | Field L2 | PDE RMSE |
|-----------|---------|---------|----------|----------|
| Data-only | 0.03587 | 0.06499 | 91.85%   | 68.36    |
| UQ-PINN   | 0.02057 | 0.02945 | 17.69%   | 11.05    |




**Figure 3.** Absolute error map against the deterministic finite-difference reference. Errors concentrate near high-gradient conductor-edge regions

Table III reports the uncertainty calibration results. At a nominal 90% target, the uncalibrated ensemble interval covers only 66.63% of held-out test points, confirming that raw ensemble spread is under-dispersed for this benchmark. Split conformal calibration increases coverage to 90.60%, with a mean interval width of 0.10871 V. The coverage is close to the target without requiring residual normality or a parametric noise model.

**Table 3.** Uncertainty calibration on held-out points

| Interval  | Target | Coverage | Width   |
|-----------|--------|----------|---------|
| Naive     | 90%    | 66.63%   | 0.04247 |
| Conformal | 90%    | 90.60%   | 0.10871 |



**Figure 4.** Conformal 90% interval half-width map. Wider intervals appear near conductor-edge and high-gradient regions

The interval width is not merely a statistical accessory. Figure 4 shows that calibrated half-widths are spatially structured. They are larger near regions where geometric discontinuity and field crowding make prediction harder. This is the desired behavior for a practical EDA assistant: a designer should not only receive a fast field estimate, but also a cue about whether the estimate is reliable enough for early screening or should be escalated to a full field solver.

The positive relation to Zhang's PINN-based VLSI/RF EDA framework [1] is direct: both works support physics-informed electromagnetic approximation for design workflows. The distinction is that the present paper focuses less on broad speedup claims and more on a reproducible, calibrated reliability mechanism. The two directions are complementary. High-fidelity PINN architectures can provide field estimates, while conformal calibration can make

those estimates safer to use in optimization loops.

## 6.1. Computational Cost and EDA Deployment Interpretation

The computational-cost measurements in Table IV should be read carefully. They were obtained on a small CPU experiment, and they are not intended to compete with industrial full-wave solvers. Their purpose is to show why the chosen random-feature PINN formulation is attractive for rapid prototyping: the expensive part of conventional neural PINNs, namely iterative back-propagation through high-order derivatives, is replaced by small linear systems. The finite-difference reference solve required 0.0660 s, the data-only fit required 0.0087 s, and the eight-member UQ-PINN ensemble fit required 0.0572 s in the recorded run.

In a realistic EDA workflow, the method would not replace sign-off electromagnetic analysis. A more practical use is hierarchical screening. The UQ-PINN can evaluate many candidate cross-sections or local layout variants, flag locations where calibrated intervals exceed a user-defined tolerance, and then send only those high-risk regions to a higher-fidelity solver. This is consistent with the broader EDA pattern in which fast extraction engines, pattern-matching models, and field solvers are used together rather than in isolation.

The conformal interval also provides a natural stopping rule for active learning. If a predicted field quantity has a narrow interval and the design decision is insensitive to the remaining uncertainty, the surrogate can be accepted for early exploration. If the interval is wide near a conductor corner, a dielectric transition, or a neighboring aggressor line, the workflow can request a new reference solve in that region. The next version of the model would then include those new points as sparse observations, reducing uncertainty where the design actually needs it.

This deployment interpretation is important because it avoids an unrealistic claim that a small 2-D neural model can handle every electromagnetic effect in modern packaging or RFIC design. The contribution is more modest and more useful: physics-informed structure improves the field estimate, while conformal calibration makes the estimate auditable. In engineering practice, this combination can reduce the risk of using a fast neural approximation as an opaque black box.

**Table 4.** Runtime and reproducibility metadata

| Component               | Value       | Interpretation        |
|-------------------------|-------------|-----------------------|
| FDM reference solve     | 0.0660 s    | Generates labels      |
| Data-only fit           | 0.0087 s    | Sparse baseline       |
| UQ-PINN ensemble fit    | 0.0572 s    | Eight PINNs           |
| Total script time       | 3.2997 s    | Includes plots/export |
| Calibration/test points | 1234 / 1852 | Split conformal       |

## 6.2. Algorithmic Summary

Algorithm 1 summarizes the full workflow. The first stage creates a reproducible electromagnetic benchmark by solving the finite-difference problem. The second stage trains two neural surrogates under identical sparse-data conditions, with the only difference being the addition of physics residual rows for the UQ-PINN. The third stage uses held-out points for conformal calibration rather than for retraining. This

separation is necessary: if the same points were used for both fitting and calibration, the reported coverage would be optimistic and would not reflect the conformal guarantee.

The algorithm is intentionally simple. Its simplicity is a strength for an EI-style conference contribution because each output file can be traced to a clear numerical step. The geometry parameters, random seed, finite-difference residual, model dimensions, calibration quantile, and generated figures are exported in JSON or CSV form. This makes the work easier to audit than a paper that reports only final tables without code or provenance.

Algorithm 1. Reproducible conformal UQ-PINN workflow

1. Define a normalized two-layer VLSI interconnect cross-section with signal and grounded boundaries.
2. Solve  $\text{div}(\epsilon \nabla \phi) = 0$  using a finite-difference sparse linear system to obtain the reference potential.
3. Sample sparse interior observations, boundary points, and unlabeled physics collocation points.
4. Fit a data-only random-feature network from sparse observations and boundary constraints.
5. Fit a bootstrap ensemble of random-feature PINNs by adding discrete physics residual rows.
6. Predict ensemble mean and standard deviation on all valid field points.
7. Compute split-conformal nonconformity scores on calibration points and obtain  $q_{\text{hat}}$ .
8. Report point-error metrics, physics residual, calibrated coverage, interval width, figures, and CSV files.

### 6.3. Why the Experiment Uses Synthetic Physics Data

The user requirement that experimental data must not be fabricated is especially relevant for this topic. Public benchmark datasets for VLSI electromagnetic fields are difficult to use directly because foundry process stacks, conductor geometries, and sign-off field solutions are often proprietary. It would be misleading to invent a table of measured RF or silicon results without access to the underlying structures. Therefore, this paper chooses the safer route: synthetic data are generated by a deterministic numerical solver, and the manuscript explicitly identifies them as synthetic.

This does not make the experiment meaningless. Numerical benchmarks are common in computational electromagnetics and scientific machine learning because they allow a method to be tested against a known reference solution while preserving full reproducibility. The important condition is transparency. The paper should not call the benchmark measured silicon data, should not claim validation against a commercial solver, and should not imply full-wave frequency-domain accuracy. Instead, it should state that the experiment validates a quasi-static field-learning and uncertainty-calibration mechanism.

The same workflow can be transferred to real data when available. A commercial extractor or electromagnetic simulator could provide the sparse and validation labels, while the physics residual and conformal calibration steps would remain unchanged. In that future setting, the conformal intervals would become even more useful because solver/model mismatch, manufacturing stack variation, and geometry extrapolation would all contribute to predictive uncertainty.

## 7. Limitations and Validity Boundaries

This paper intentionally avoids overstating the experiment. The benchmark is two-dimensional and electroquasistatic; it does not include frequency-dependent conductor loss, skin effect, vias, nonrectangular three-dimensional layouts, or measured silicon data. The finite-difference reference is a useful reproducible target, but it is not a substitute for commercial sign-off electromagnetic simulation. Therefore, the correct interpretation is that the proposed conformal UQ-PINN is a fast early-stage screening and reliability-estimation method.

Another limitation is the simplified physics residual used in the random-feature PINN. The residual enforces a Laplace-like operator in dielectric subregions, whereas the reference finite-difference model includes variable permittivity through face averaging. Future work should include interface-aware residuals, exact dielectric flux-continuity terms, conductor surface loss models, and frequency-domain Maxwell residuals. A natural extension is to couple this calibrated surrogate to a selective-refinement policy: if the conformal interval is too wide near a layout feature, the algorithm requests a local high-fidelity solve and adds those points to the training set.

Despite these limitations, the study provides a reproducible evidence base. The code reports the seed, geometry, finite-difference residual, training configuration, pointwise predictions, and figures. This satisfies the main requirement for non-fabricated experimental data: readers can regenerate the numerical benchmark and verify whether the reported results match the output files.

## 8. Future Work and Extensions

The first extension is an interface-aware residual. The current implementation uses a Laplace-like residual at collocation points, which is appropriate as a compact physics regularizer away from material discontinuities. A more complete semiconductor formulation should enforce normal flux continuity across dielectric interfaces, tangential electric-field continuity, and conductor boundary conditions through either hard constraints or residual terms. This would align the neural objective more closely with the variable-permittivity finite-difference operator and would reduce the gap between the surrogate and the reference near the substrate boundary.

The second extension is frequency-domain electromagnetic modeling. For high-speed interconnects and RF structures, a scalar electrostatic potential is not sufficient. Future work should use complex-valued electric and magnetic fields, frequency-dependent material parameters, and residuals derived from Maxwell's curl equations. This would allow the model to represent skin effect, dielectric loss, return-path discontinuities, and wave propagation. [1] is valuable here because it already emphasizes the importance of Maxwell-equation constraints and broad frequency support in VLSI and RF EDA; the conformal mechanism in the present paper could be added on top of such a high-fidelity architecture.

A related path is to combine the present conformal wrapper with richer uncertainty-aware physics learners. Bayesian physics-informed extreme learning machines have already been proposed for noisy PDE problems [29], and local extreme-learning/domain-decomposition methods can reduce the scaling burden of global neural solvers [30]. For electromagnetic media with discontinuities, recent PINN

work has also examined steady-state and transient Maxwell problems with material interfaces [31]. These studies suggest that the simple random-feature ensemble used here can be upgraded without changing the central calibration idea: any point predictor that outputs a location-dependent uncertainty score can be conformalized on held-out field points.

The third extension is multi-conductor and multi-net analysis. The present benchmark contains one signal conductor and a grounded shield. Crosstalk prediction requires several driven and floating conductors, different excitation modes, and extraction of a capacitance or impedance matrix rather than only a scalar field. In that setting, uncertainty should be reported not only pointwise in space but also for derived quantities such as coupling capacitance, characteristic impedance, and worst-case field magnitude. A conformal interval around a derived circuit parameter would be directly interpretable by circuit designers.

The fourth extension is process variation. Modern interconnect performance depends on uncertain metal width, dielectric thickness, line-edge roughness, and material parameters. A practical UQ-PINN should accept these quantities as inputs and produce uncertainty that combines model uncertainty with manufacturing variation. Split conformal prediction can still be used, but calibration should be stratified by process corner or geometry class so that coverage is not dominated by easy regions while failing on rare but critical structures.

The fifth extension is selective solver coupling. A useful EDA implementation would not train once and stop. It would begin with a small solver-generated set, train the calibrated surrogate, inspect the interval map, and request additional high-fidelity samples where the interval is too wide. This active-learning loop would spend expensive solver calls only where they reduce design risk. Such a workflow is compatible with both open numerical solvers and commercial tools, provided that data export is permitted by the design environment.

Finally, future work should build public benchmark suites for neural electromagnetic field prediction in EDA. The lack of open, realistic, labeled VLSI field data is one reason papers in this area sometimes rely on broad claims without reproducible evidence. A benchmark suite with canonical 2-D structures, 3-D vias, coupled lines, substrate transitions, and frequency-domain labels would make it possible to compare PINNs, operator networks, graph models, and random-feature solvers under identical conditions. It would also make uncertainty calibration measurable rather than rhetorical.

## 9. Conclusion

This paper presented a conformal uncertainty-aware PINN workflow for reliable electromagnetic field prediction in a VLSI interconnect cross-section. The method combines a random-feature physics-informed neural surrogate, sparse field observations, boundary constraints, a finite-difference physics residual, ensemble uncertainty, and split conformal calibration. On a deterministic synthetic benchmark generated from  $\text{div}(\epsilon \nabla \phi) = 0$ , the physics-informed ensemble improved both potential accuracy and field-gradient accuracy over a sparse data-only baseline. More importantly, conformal calibration corrected under-dispersed raw uncertainty and achieved near-target empirical coverage.

The main practical message is that reliable neural EDA

surrogates should not be judged only by speed or point error. A model used during layout exploration should also know when its prediction is uncertain. By combining physics-informed learning with conformal intervals, the proposed approach offers a transparent route toward fast, uncertainty-aware electromagnetic screening for interconnect design. Future work will extend the benchmark to three-dimensional geometries, lossy conductors, frequency-domain Maxwell equations, and active learning with commercial solver validation.

## References

- [1] Zhang, H. (2025). Physics-informed neural networks for high-fidelity electromagnetic field approximation in VLSI and RF EDA applications. *Journal of Computing and Electronic Information Management*, 18(2), 38–46.
- [2] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- [3] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3, 422–440. <https://doi.org/10.1038/s42254-021-00314-5>
- [4] Lu, L., Meng, X., Mao, Z., & Karniadakis, G. E. (2021). DeepXDE: A deep learning library for solving differential equations. *SIAM Review*, 63(1), 208–228. <https://doi.org/10.1137/19M1274067>
- [5] Lim, J., & Psaltis, D. (2022). MaxwellNet: Physics-driven deep neural network training based on Maxwell's equations. *APL Photonics*, 7(1), 011301. <https://doi.org/10.1063/5.0071616>
- [6] Dwivedi, V., & Srinivasan, B. (2020). Physics informed extreme learning machine (PIELM): A rapid method for the numerical solution of partial differential equations. *Neurocomputing*, 391, 96–118. <https://doi.org/10.1016/j.neucom.2019.12.099>
- [7] Huang, G.-B., Zhu, Q.-Y., & Siew, C.-K. (2006). Extreme learning machine: Theory and applications. *Neurocomputing*, 70(1–3), 489–501. <https://doi.org/10.1016/j.neucom.2005.12.126>
- [8] Rahimi, A., & Recht, B. (2007). Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems* (pp. 1177–1184).
- [9] Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems* (pp. 6402–6413).
- [10] Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the International Conference on Machine Learning* (pp. 1050–1059).
- [11] Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems* (pp. 5574–5584).
- [12] Shafer, G., & Vovk, V. (2008). A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9, 371–421.
- [13] Angelopoulos, A. N., & Bates, S. (2021). A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv:2107.07511*.
- [14] Romano, Y., Patterson, E., & Candès, E. (2019). Conformalized quantile regression. In *Advances in Neural Information Processing Systems* (pp. 3543–3553).

- [15] Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R. J., & Wasserman, L. (2018). Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523), 1094–1111. <https://doi.org/10.1080/01621459.2017.1307116>
- [16] Nabors, K., & White, J. (1991). FastCap: A multipole accelerated 3-D capacitance extraction program. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 10(11), 1447–1459. <https://doi.org/10.1109/43.97624>
- [17] Abouelyazid, M. S., Hammouda, S., & Ismail, Y. (2022). Accuracy-based hybrid parasitic capacitance extraction using rule-based, neural-networks, and field-solver methods. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 41(12), 5681–5694. <https://doi.org/10.1109/TCAD.2022.3152941>
- [18] Yang, D., Yu, W., Guo, Y., & Liang, W. (2023). CNN-Cap: Effective convolutional neural network based capacitance models for full-chip parasitic extraction. *ACM Transactions on Design Automation of Electronic Systems*, 28(4), 55. <https://doi.org/10.1145/3564931>
- [19] Teng, D., Rhee, M., Qin, Y., Zi, B., & Liu, W. (2026). SW-SpeedDLM: Sliding window speculative decoding for diffusion language models under long context constraints. *Mathematics*, 14(12), 2137. <https://doi.org/10.3390/math14122137>
- [20] Sharma, S., & Triverio, P. (2022). Electromagnetic modeling of lossy interconnects from DC to high frequencies with a potential-based boundary element formulation. *IEEE Transactions on Microwave Theory and Techniques*, 70(8), 3847–3861. <https://doi.org/10.1109/TMTT.2022.3182132>
- [21] Fan, J., Ye, X., Kim, J., Archambeault, B., & Orlandi, A. (2010). Signal integrity design for high-speed digital circuits: Progress and directions. *IEEE Transactions on Electromagnetic Compatibility*, 52(2), 392–400. <https://doi.org/10.1109/TEM.2010.2045381>
- [22] Achar, R., & Nakhla, M. S. (2011). Modeling of high-speed interconnects for signal integrity analysis - Part I: Fundamentals. *IEEE Microwave Magazine*, 12(2), 50–60. <https://doi.org/10.1109/MMM.2011.941414>
- [23] Kim, Y., Kim, J., Lee, J., & Swaminathan, M. (2023). A statistical approach for signal and power integrity co-design in high-speed interconnects. *Scientific Reports*, 13, 16678. <https://doi.org/10.1038/s41598-023-43586-0>
- [24] Deng, Y., Fan, K., Jin, B., Malof, J., & Padilla, W. J. (2025). Physics-informed learning in artificial electromagnetic materials. *Applied Physics Reviews*, 12(1). <https://doi.org/10.1063/5.0221657>
- [25] Kiarashinejad, Y., Abdollahramezani, S., & Adibi, A. (2020). Deep learning approach based on dimensionality reduction for designing electromagnetic nanostructures. *npj Computational Materials*, 6, 12. <https://doi.org/10.1038/s41524-020-0276-y>
- [26] Teng, D. (2025). TEAS: Token- and energy-aware autoscaling for cost-efficient LLM serving. *AI and Data Science Journal*.
- [27] Zhang, F., Guo, Z., Ding, J., Yang, J., & Liu, W. (2026). Adaptive sensor fusion for robust perception in dense fog: A gated vision and LiDAR integration framework. *Sensors*, 26(12), 3728. <https://doi.org/10.3390/s26123728>
- [28] Zi, B. (2024). Large language models for enterprise workflow automation in financial operations. *Innovation and Technology Studies*, 1(1), 24–29.
- [29] Ding, J., Shen, Z., & Liu, W. (2026). Game-theoretic cost-sensitive adversarial training for robust cloud intrusion detection against GAN-based evasion attacks. *Applied Sciences*, 16(8), 3944. <https://doi.org/10.3390/app16083944>
- [30] Zi, B. (2024). Cloud-native distributed systems for real-time payment intelligence. *AI and Data Science Journal*, 1(1), 51–56.
- [31] Jiao, Y., Fan, H., Yue, X., Ping, W., Sun, T., & Wang, J. (2026). Dynamic heterogeneous graph contrastive learning for uncovering collusive financial fraud. *Scientific Reports*, 16, 12907. <https://doi.org/10.1038/s41598-026-96214-3>