Neuro-Symbolic Partial Differential Equation Solver

Pouria A. Mistani[†*]
Nvidia
pmistani@nvidia.com

Samira Pakravan [†]
UCSB
spakravan@ucsb.edu

Rajesh Ilango
Nvidia
rilango@nvidia.com

Sanjay Choudhry
Nvidia
schoudhry@nvidia.com

Frederic Gibou
UCSB
fgibou@ucsb.edu

Figure 1: The proposed Neural Bootstrapping Method applied to electrostatics. Left: streamlines. Right: warped cross-section showing the physically correct jump in the field. The neural approach readily enables a 1024^3 effective resolution on a single NVIDIA A6000 GPU. Once trained, it takes sub-milliseconds for the network to evaluate such a simulation [30].

Abstract

We present a highly scalable strategy for developing mesh-free neuro-symbolic partial differential equation solvers from existing numerical discretizations found in scientific computing. This strategy is unique in that it can be used to efficiently train neural network surrogate models for the solution functions and the differential operators, while retaining the accuracy and convergence properties of state-of-the-art numerical solvers. This neural bootstrapping method is based on minimizing residuals of discretized differential systems on a set of random collocation points with respect to the trainable parameters of the neural network, achieving unprecedented resolution and optimal scaling for solving physical and biological systems.

1 Introduction

Most modern physical, biological and engineering systems are described by partial differential equations on irregular, often moving, boundaries. The difficulties in solving those problems stem from approximating the equations while respecting the physically correct discontinuous nature of the solution across the boundaries. Smoothing strategies are straightforward to design, but unfortunately introduce unphysical characteristics in the solution, which lead to systemic errors.

Since the 1990s, artificial neural networks have been used for solving differential equations using two general strategies: (i) mapping the algebraic operations of the discretized systems onto specialized

---

[*]corresponding author.

[†]These authors contributed equally to this work.

neural network architectures and minimizing the network energy, or (ii) treating the whole neural network as the basic approximation unit, with parameters trained to minimize a specialized error function that includes the differential equation itself as well as its boundary and initial conditions.

In the first category, neurons output the discretized solution values over a set of grid points, and minimizing the network energy drives the neuronal values towards the solution at the mesh points. The neural network energy is the residual of the finite discretization, summed over all neurons [22]. A strong feature is the preservation of the finite discretization convergence; however, the computational cost grows with increasing resolution and dimensionality. Early examples include [15, 10, 9].

The second strategy, proposed by Lagaris *et al.* [21], relies on the function approximation capabilities of the neural networks. Encoding the solution everywhere in the domain within a neural network offers a mesh-free, compact, and memory efficient surrogate model for the solution function, which can be used in subsequent inference tasks. This method has recently re-emerged as the physics-informed neural networks (PINNs) [35] and is widely used. Despite their advantages, these methods lack the controllable convergence properties of traditional numerical discretizations, and are biased towards the lower frequency features of the solutions [41, 34, 20].

Hybridization frameworks seek to combine the performance of neural network inference on modern accelerated hardware with the guaranteed accuracy of traditional discretizations developed by the scientific community. The hybridization efforts are algorithmic or architectural.

One important algorithmic method is the deep Galerkin method (DGM) [38], a neural network extension of the mesh-free Galerkin method where the solution is represented as a deep neural network rather than a linear combination of basis functions. Being mesh-free, it enables the solution of high-dimensional problems by training the neural network model to satisfy the differential operator and its initial and boundary conditions on a randomly sampled set of points, rather than on an exponentially large grid. Although the number of points can be very large in higher dimensions, the training is done sequentially on smaller batches of data points and second-order derivatives are calculated by a scalable Monte Carlo method. Another important algorithmic method is the deep Ritz [42]. It implements a deep neural network approximation of the trial function that is constrained by numerical quadrature rules for the variational functional, followed by stochastic gradient descent.

Architectural hybridization is based on differentiable numerical linear algebra. One emerging class involves implementing differentiable finite discretization solvers and embedding them in the neural architectures that enable application of end-to-end differentiable gradient based optimization. Differentiable solvers have been developed in JAX [7] for fluid dynamic problems, e.g. `Phi-Flow` [17], `JAX-CFD` [19], and `JAX-FLUIDS` [2]. These methods are suitable for inverse problems where an unknown field is modeled by the neural network, while the model influence is propagated by the differentiable solver into a measurable residual [33, 12, 26]. We also note the classic strategy for solving inverse problems, the adjoint method, to obtain the gradient of the loss without differentiation across the solver [1]; however, deriving analytic expression for the adjoint equations can be tedious or impractical. Other notable use of differentiable solvers is to model and correct for the solution errors of finite discretizations [40], and learning and controlling differentiable systems [13, 16].

Neural networks are not only universal approximators of continuous functions, but also of nonlinear operators [8]. Although this fact has been leveraged using data-driven strategies for learning differential operators by many authors [25, 3, 24, 23], researchers have demonstrated the ability of differentiable solvers to effectively train nonlinear operators without any data in a completely physics-driven fashion, see section on learning inverse transforms in [33].

We propose a novel framework for solving PDEs based on deep neural networks, that enables lifting any existing mesh-based finite discretization method off of its underlying grid and extend it into a mesh-free and embarrassingly parallel method that can be applied to high dimensional problems on unstructured random points. In addition, discontinuous solutions can be readily considered.

2 Problem statement

We illustrate our approach by considering a closed irregular interface (Γ) that partitions an interior (Ω^-) and an exterior (Ω^+) subdomain (see figure 2). The coupled solution $u^\pm \in \Omega^\pm$ satisfy the Helmholtz equation $k^\pm u^\pm - \nabla \cdot (\mu^\pm \nabla u^\pm) = f^\pm$ with jump conditions $[u] = \alpha$ and $[\mu \partial_{\mathbf{n}} u] = \beta$,
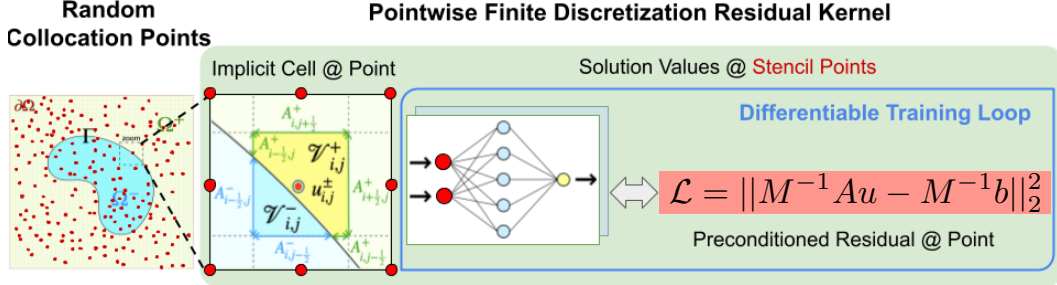
Figure 2: Neural Bootstrapping Method (NBM). NBM kernels compute the residual contribution at each collocation point per GPU thread. Kernel operations involve placing implicit cells at multi-resolutions and calculating interface-cell crossings. Point-wise residuals are locally preconditioned (M^{-1} matrix). A, b are the left/right-hand sides of the linear system obtained by finite discretization method. u is the vector of solution values predicted by the neural network model at the stencil points.

where $f^\pm$, $\mu^\pm$ and $k^\pm$ are spatially varying coefficients. For simplicity, we consider Dirichlet conditions at the boundary of a cubic computational domain.

This class of problems describes core components of diffusion dominated processes, where sharp irregular interfaces regulate transport across regions with different properties. Examples include Poisson-Boltzmann equation for describing electrostatic properties of biomolecules with jump in dielectric permittivities [37, 27], in electroporation of cell aggregates with nonlinear membrane jump conditions [29], or in epitaxial growth [28]. Other important applications are found in solidification of multicomponent alloys in additive manufacturing [39, 6], directed self-assembly of diblock copolymers for next generation lithography [14, 32, 5], and multiphase flows with phase change.

3   Scalable and Mesh-Freeing Neuro-Symbolic Differential Solver

We use neural networks as surrogates for the solution function that are iteratively adjusted to minimize discretization residuals at a set of randomly sampled points and at arbitrary resolutions. The key idea is that neural networks can be evaluated over vertices of any discretization stencils centered at any point in the domain, effectively emulating any structured mesh without actually materializing it. Therefore, we use neural networks to create mesh-free neural differential solvers from existing mesh-based discretization methods . We call this the Neural Bootstrapping Method (NBM) (figure 2).

Numerical methods offer guaranteed accuracy and controllable convergence properties for the training of neural network surrogate models. NBM offers a straightforward path for applying mesh-based numerical methods on random points. This is an important capability for augmenting observational data in the training pipelines. NBM is also a highly parallelizable strategy and the point-wise nature of its kernels is ideally suited for GPU-accelerated computing. The difficult problem of multi-GPU parallel solutions of differential systems is thus reduced to the much simpler task of data-parallel training using existing machine learning frameworks. Data parallelism involves distributing collocation points across multiple processors to compute gradient updates and then aggregating these locally computed updates [36].

Here we bootstrap the numerical scheme proposed by [4] to solve the Helmholtz equation, hence the loss function is with the coefficients described in figure 2. The operations in differentiable NBM kernels are strictly local. NBM starts by placing implicit compute cells of a specified resolutions at the collocation point. At discontinuities, two neural networks are used to represent the solution on each side of Γ and a separate coarse mesh encapsulates a level-set function [31] that provides the necessary geometric information for the numerical kernel and the preconditioner (denoted M^{-1}). The numerical kernel is applied on the compute cell where the solution values are evaluated by the neural network. Each kernel contributes a local L^2-norm residual $r_p = ||M^{-1}Au - M^{-1}b||$ at one point p. Preconditioning balances the relative magnitude of contributions from all points before aggregating the residuals to form a global loss value. Finally, gradient based optimization methods used in machine learning, *e.g.* `Adam` optimizer [18], are applied to adjust neural network parameters. The automatic differentiation loop passes across the NBM kernels.

3

4 Numerical results

4.1 Convergence and accuracy

We consider a sphere centered at the origin with radius 0.5 in a domain $[-1, 1]^3$, a discontinuous exact solution of $u^- = e^z$ and $u^+ = \cos(x)\sin(y)$. In addition to the jump in solution, we consider a jump in the variable diffusion coefficient to be $\mu^- = y^2 \ln(x + 2) + 4$ and $\mu^+ = e^{-z}$. Table 1 reports convergence in the L^∞-norm. Two neural networks were used to represent solutions inside and outside the sphere with 5 hidden layers and 10 sine-activated neurons each, for a total of 982 trainable parameters.

Table 1: Convergence and overall time to solution for our JAX implementation, with $10,000$ epochs in each case, and initial compilation time. Measurements are on a single NVIDIA A6000 GPU.

	RMSE		L^∞		GPU Statistics	
$N_{x,y,z}$	Solution	Order	Solution	Order	t (sec/epoch)	VRAM (GB)
2^3	3.7×10^{-2}	-	3.25×10^{-1}	-	0.0306	1.05
2^4	7.1×10^{-3}	2.38	1.10×10^{-1}	1.56	0.056	1.72
2^5	5.9×10^{-3}	0.27	8.36×10^{-2}	0.4	0.053	2.15
2^6	4.1×10^{-3}	0.53	6.44×10^{-2}	0.38	0.287	5.57
2^7	2.64×10^{-3}	0.64	3.53×10^{-2}	0.87	2.125	32.1

4.2 Time complexity and parallel scaling on GPU clusters

We simulate a Helmholtz problem with discontinuities on the Dragon geometry presented in [11]. We consider a solution with jumps across the dragon's surface and a spatially varying diffusion coefficient. The results are shown in figure 1, with a L^∞-error of 0.5 and RMSE of 0.06 after 1000 epochs on a base resolution of 64^3 and implicitly refined onto multi-resolutions $128^3, 256^3, 512^3$. The neural network pair have only 1 hidden layer with 100 sine-activated neurons, although investigating more complex networks (transformers, symmetry preserving, *etc*.) would likely improve accuracy.

In table 2 we report scaling results on NVIDIA A100 GPUs at four base resolutions with three levels of implicit refinement. We used a batchsize of $32 \times 32 \times 16$ in all cases. At fixed number of GPUs, training time scales linearly (*i.e.*, optimal scaling) with the number of grid points. At a fixed resolution, increasing the number of GPUs accelerates training roughly with epoch time $\sim 1/\sqrt{\# \text{ GPUs}}$, although the advantage is more effective at higher resolutions. Compile time increases with resolution and decreases with number of GPUs. A maximum grid size of 1024^3 at multi-resolutions $1024^3, 2048^3, 4096^3, 8192^3$ was simulated on one NVIDIA DGX with 8 A100 GPUs.

Table 2: Scaling test. Time per epoch (sec) and JAX compile time for different configurations.

base resolution:	64^3		128^3		256^3		512^3	
A100 GPUs	epoch	compile	epoch	compile	epoch	compile	epoch	compile
1	0.908	9.027	6.960	9.288	55.287	12.164	438.45	49.020
2	0.657	7.575	5.893	7.823	47.360	10.045	378.98	39.815
4	0.405	7.480	3.629	7.863	28.261	9.129	226.73	27.405
8	0.384	7.983	3.340	7.901	26.799	9.154	204.88	20.632

5 Conclusion

We presented a neural bootstrapping method and applied it to the problem of solving elliptic PDEs with discontinuities. NBM is a differentiable computing method that creates scalable and mesh-free numerical methods from mesh-based finite discretizations. It represents the solution by training neural networks using automatic differentiation of the discretized residual at collocation points. We implemented the method using JAX and showed accuracy and parallel scaling in three dimensions.

Acknowledgments and Disclosure of Funding

Broader Impact

The work presented provides a systematic framework for creating neural-based solvers for partial differential equations with potential jump conditions across sharp interfaces, which describe a plethora of important systems in the physical and life sciences. While the driving motivation for our work is the development of algorithms for multi-scale simulations to accelerate drug discovery and formulation pipelines in the pharmaceutical industry, we believe that the methodology will positively impact other industries that face similar computational challenges. We do not believe that the methodology we have introduced will have an adverse effect to society or will have negative ethical implications.

References

[1] J. Berg and K. Nyström. Neural network augmented inverse problems for pdes. *arXiv preprint arXiv:1712.09685*, 2017.

[2] D. A. Bezgin, A. B. Buhendwa, and N. A. Adams. Jax-fluids: A fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. *arXiv preprint arXiv:2203.13760*, 2022.

[3] K. Bhattacharya, B. Hosseini, N. B. Kovachki, and A. M. Stuart. Model reduction and neural networks for parametric pdes. *arXiv preprint arXiv:2005.03180*, 2020.

[4] D. Bochkov and F. Gibou. Solving elliptic interface problems with jump conditions on cartesian grids. *Journal of Computational Physics*, 407:109269, 2020.

[5] D. Bochkov and F. Gibou. A non-parametric shape optimization approach for solving inverse problems in directed self-assembly of block copolymers. *arXiv preprint arXiv:2112.09615*, 2021.

[6] D. Bochkov, T. Pollock, and F. Gibou. Sharp-interface simulations of multicomponent alloy solidification. *arXiv preprint arXiv:2112.08650*, 2021.

[7] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy programs, 2018.

[8] T. Chen and H. Chen. Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems. *IEEE Transactions on Neural Networks*, 6(4):911–917, 1995.

[9] L. O. Chua and L. Yang. Cellular neural networks: Applications. *IEEE Transactions on circuits and systems*, 35(10):1273–1290, 1988.

[10] L. O. Chua and L. Yang. Cellular neural networks: Theory. *IEEE Transactions on circuits and systems*, 35(10):1257–1272, 1988.

[11] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996.

[12] N. Dal Santo, S. Deparis, and L. Pegolotti. Data driven approximation of parametrized pdes by reduced basis and neural networks. *Journal of Computational Physics*, 416:109550, 2020.

[13] F. de Avila Belbute-Peres, K. Smith, K. Allen, J. Tenenbaum, and J. Z. Kolter. End-to-end differentiable physics for learning and control. *Advances in neural information processing systems*, 31, 2018.

[14] K. Galatsis, K. L. Wang, M. Ozkan, C. S. Ozkan, Y. Huang, J. P. Chang, H. G. Monbouquette, Y. Chen, P. Nealey, and Y. Botros. Patterning and templating for nanoelectronics. *Advanced Materials*, 22(6):769–778, 2010.

[15] D. Gobovic and M. Zaghloul. Design of locally connected cmos neural cells to solve the steady-state heat flow problem. In *Proceedings of 36th Midwest Symposium on Circuits and Systems*, pages 755–757. IEEE, 1993.

[16] P. Holl, V. Koltun, and N. Thuerey. Learning to control pdes with differentiable physics. *arXiv preprint arXiv:2001.07457*, 2020.

[17] P. Holl, V. Koltun, K. Um, and N. Thuerey. phiflow: A differentiable pde solving framework for deep learning via physical simulations. In *NeurIPS Workshop*, volume 2, 2020.

[18] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[19] D. Kochkov, J. A. Smith, A. Alieva, Q. Wang, M. P. Brenner, and S. Hoyer. Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21), 2021.

[20] A. Krishnapriyan, A. Gholami, S. Zhe, R. Kirby, and M. W. Mahoney. Characterizing possible failure modes in physics-informed neural networks. *Advances in Neural Information Processing Systems*, 34:26548–26560, 2021.

[21] I. E. Lagaris, A. Likas, and D. I. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE transactions on neural networks*, 9(5):987–1000, 1998.

[22] H. Lee and I. S. Kang. Neural algorithm for solving differential equations. *Journal of Computational Physics*, 91(1):110–131, 1990.

[23] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

[24] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. Neural operator: Graph kernel network for partial differential equations. *arXiv preprint arXiv:2003.03485*, 2020.

[25] L. Lu, P. Jin, and G. E. Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019.

[26] P. Y. Lu, S. Kim, and M. Soljačić. Extracting interpretable physical parameters from spatiotemporal systems using unsupervised learning. *Physical Review X*, 10(3):031056, 2020.

[27] M. Mirzadeh, M. Theillard, and F. Gibou. A second-order discretization of the nonlinear Poisson-Boltzmann equation over irregular geometries using non-graded adaptive Cartesian grids. *Journal of Computational Physics*, 230(5):2125–2140, Mar. 2011.

[28] P. Mistani, A. Guittet, D. Bochkov, J. Schneider, D. Margetis, C. Ratsch, and F. Gibou. The island dynamics model on parallel quadtree grids. *Journal of Computational Physics*, 361:150–166, 2018.

[29] P. Mistani, A. Guittet, C. Poignard, and F. Gibou. A parallel voronoi-based approach for mesoscale simulations of cell aggregate electropermeabilization. *Journal of Computational Physics*, 380:48–64, 2019.

[30] T. Müller. Tiny CUDA neural network framework, 2021. https://github.com/nvlabs/tiny-cuda-nn.

[31] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of computational physics*, 79(1):12–49, 1988.

[32] G. Y. Ouaknin, N. Laachi, K. Delaney, G. H. Fredrickson, and F. Gibou. Level-set strategy for inverse dsa-lithography. *Journal of Computational Physics*, 375:1159–1178, 2018.

[33] S. Pakravan, P. A. Mistani, M. A. Aragon-Calvo, and F. Gibou. Solving inverse-pde problems with physics-aware neural networks. *Journal of Computational Physics*, 440:110414, 2021.

[34] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310. PMLR, 2019.

[35] M. Raissi, P. Perdikaris, and G. Karniadakis. 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, 2019.

[36] C. J. Shallue, J. Lee, J. Antognini, J. Sohl-Dickstein, R. Frostig, and G. E. Dahl. Measuring the effects of data parallelism on neural network training. *arXiv preprint arXiv:1811.03600*, 2018.

[37] K. A. Sharp and B. Honig. Calculating total electrostatic energies with the nonlinear poisson-boltzmann equation. *Journal of Physical Chemistry*, 94(19):7684–7692, 1990.

[38] J. Sirignano and K. Spiliopoulos. Dgm: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018.

[39] M. Theillard, F. Gibou, and T. Pollock. A sharp computational method for the simulation of the solidification of binary alloys. *Journal of scientific computing*, 63(2):330–354, 2015.

[40] K. Um, R. Brand, Y. R. Fei, P. Holl, and N. Thuerey. Solver-in-the-loop: Learning from differentiable physics to interact with iterative pde-solvers. *Advances in Neural Information Processing Systems*, 33:6111–6122, 2020.

[41] S. Wang, X. Yu, and P. Perdikaris. When and why pinns fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022.

[42] B. Yu et al. The deep ritz method: a deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] see section 4.2 on exploring more complex neural network architectures for improving accuracy.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See section on Broder Impact.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]

3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] The code is proprietary.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] It is a completely physics driven work without any data.

5. If you used crowdsourcing or conducted research with human subjects...

   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]