Physics-guided Diffusion Neural Operators for Solving Forward and Inverse PDEs

Medha Sawhney[1], Abhilash Neog[1], Mridul Khurana[1], Amartya Dutta[1], Arka Daw[2], Anuj Karpatne[1]
[1]Virginia Tech, [2]Oak Ridge National Lab

Motivation

- PDEs are fundamental to modeling physical phenomena but often computationally expensive to solve
- ML approaches can provide fast approximations to PDE solutions after training

Why Diffusion Models?

Limitations of prior works:

- FNOs require uniform grids → can't easily handle irregular domains or arbitrary sensor locations
- Real-world measurements are sparse & noisy → purely deterministic operators struggle with uncertainty

Advantages of using Diffusion Models:

- Resolution-agnostic super-resolution
- Naturally handle multiple solutions for ill-posed inverse problems
- Provide flexible conditional generation capabilities
- Uncertainty quantification via sampling ensembles

Challenges

- Current diffusion-based PDE solvers require **hundreds to thousands** of denoising steps
- Existing solvers typically inject PDE residuals in pixel/pointwise form—**lacking multi-scale** or global spectral enforcement, which can miss **long-range physical dependencies.**
- Under very sparse or noisy measurements, current models can produce **unstable**, **non-physical artifacts** and lack principled uncertainty quantification.

Physics-guided Diffusion Neural Operators

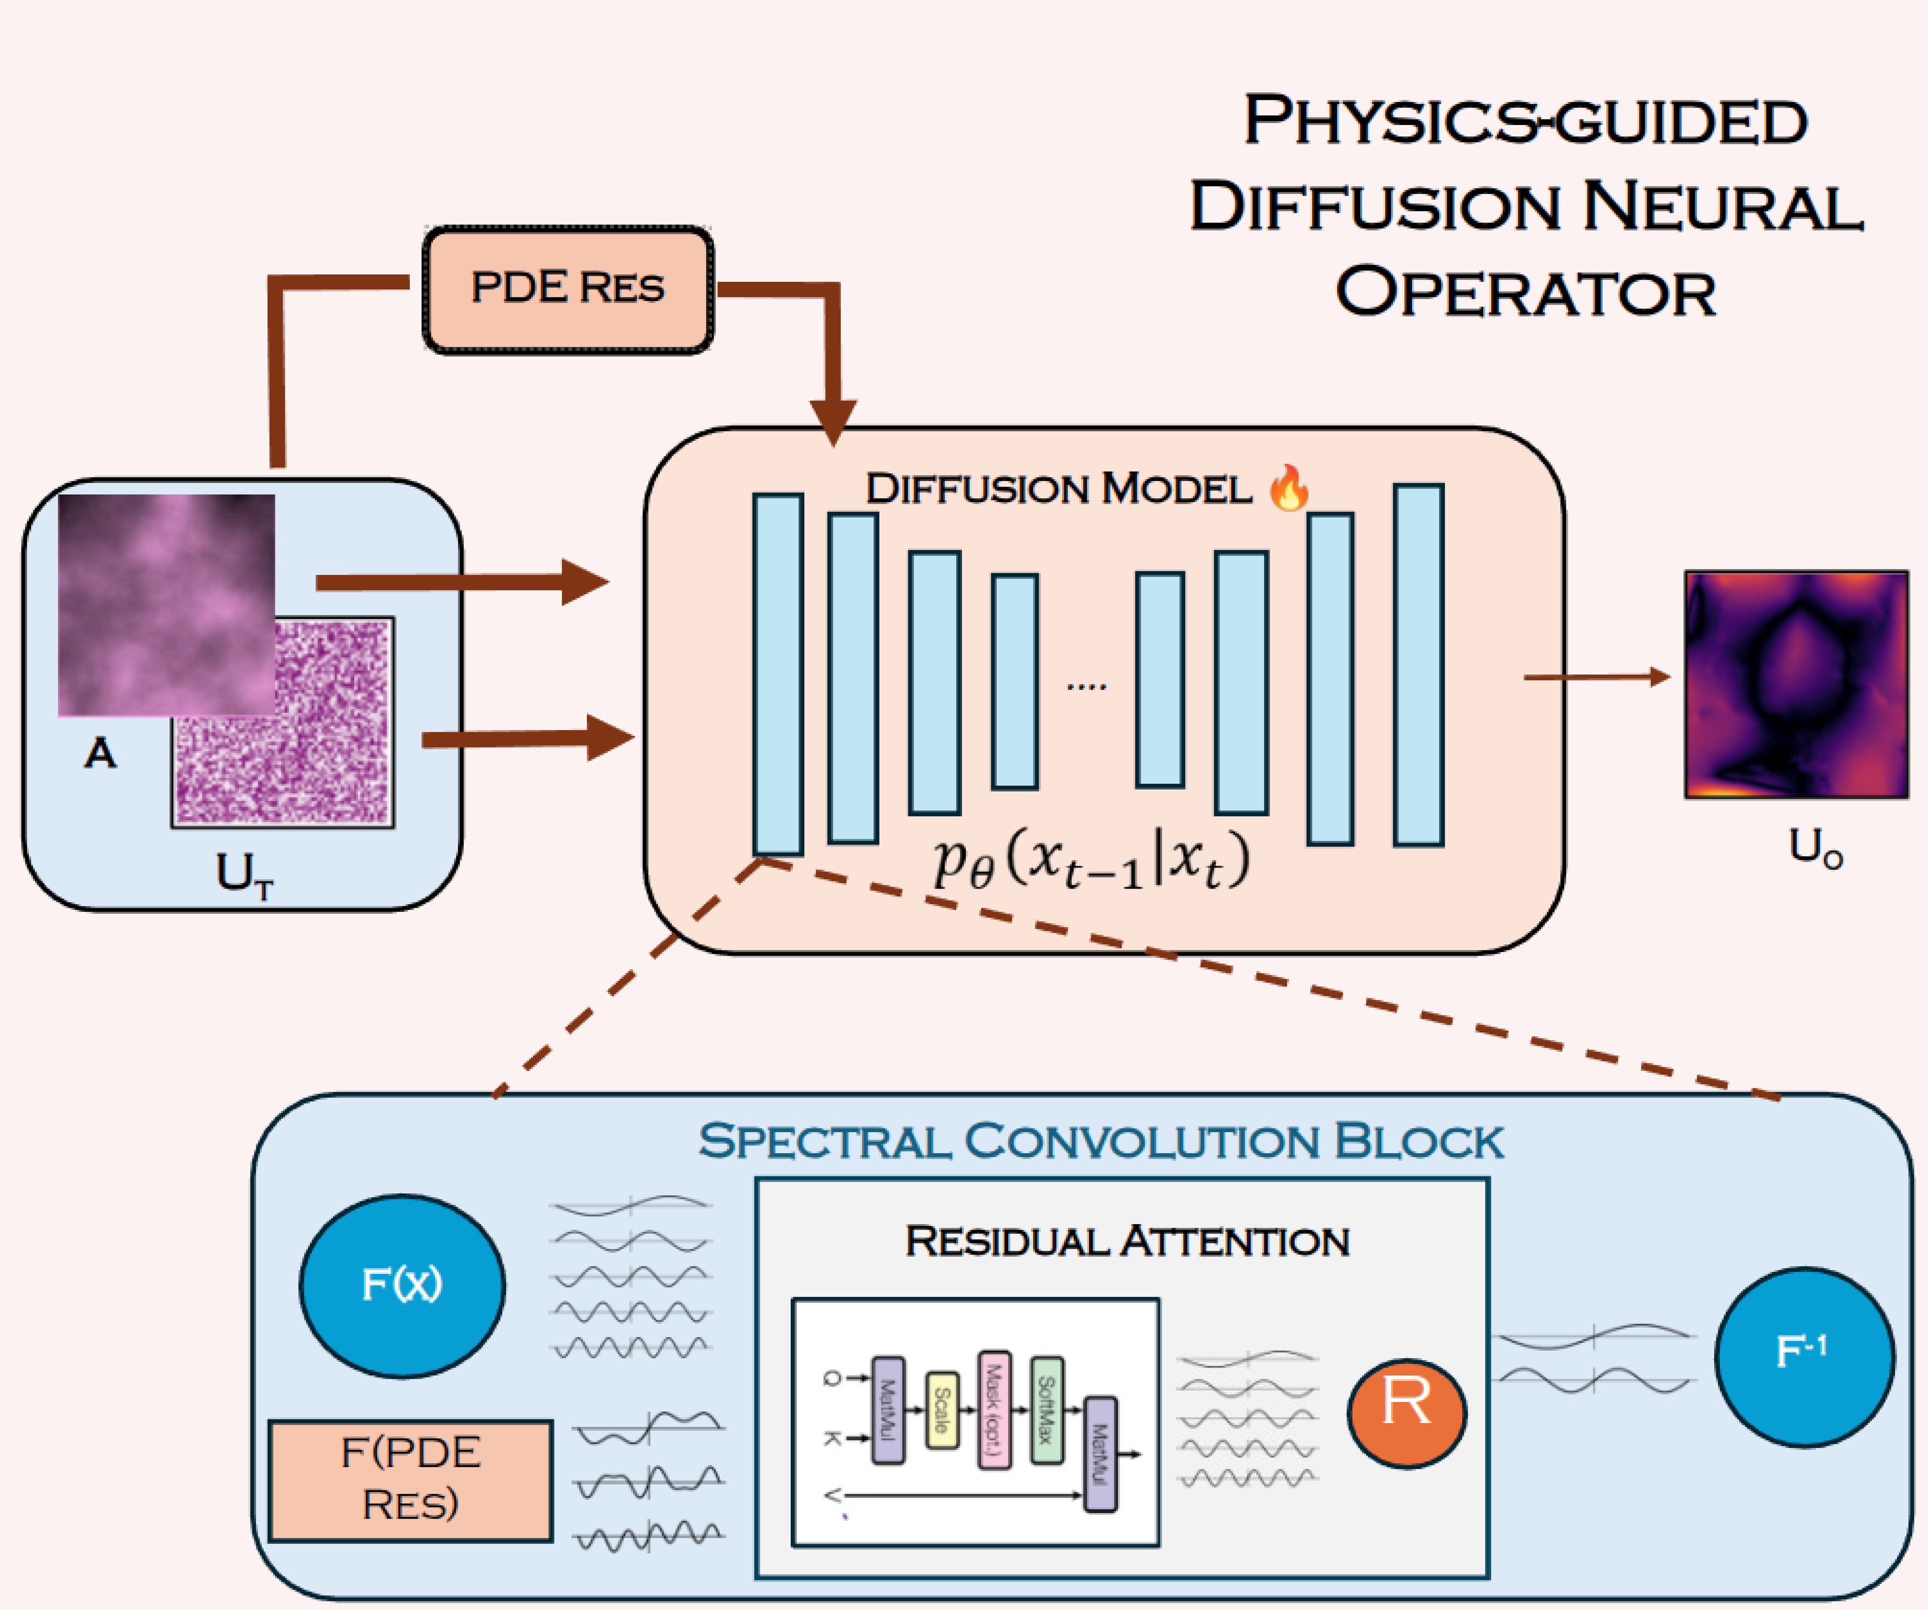

Key Highlights

- **Integrating spectral neural operators with diffusion model to** capture global physical dependencies
- **PDE-informed regularization directly in spectral space during training and sampling**
- Model learns to maintain physical consistency
- Noise-Residual Gating to fuse the current diffusion noise level with spectral residual
- PgDNO-RFA: Frequency based PDE residual attention in the spectral domain
- PgDNO Concat: RFA + PDE residual concatenated to Diffusion model input

Preliminary Findings: Results

Model	Steps	Darcy (Fwd)	Darcy (Inv)	Poisson (Fwd)	Poisson (Inv)	Helmholtz (Fwd)	Helmholtz (Inv)
FunDPS	200	2.88%	6.78%	2.04%	24.04%	2.20%	20.07%
FunDPS	500	2.49%	5.18%	1.99%	20.47%	2.13%	17.16%
DiffusionPDE	2000	6.07%	7.87%	4.88%	21.10%	12.64%	19.07%
FNO	—	28.2%	49.3%	100.9%	232.7%	98.2%	218.2%
PINO	—	35.2%	49.2%	107.1%	231.9%	106.5%	216.9%
DeepONet	—	38.3%	41.1%	155.5%	105.8%	123.1%	132.8%
PINN	—	48.8%	59.7%	128.1%	130.0%	142.3%	160.0%
$PgDNO_{Hybrid}$	18	2.50%	7.40%	4.90%	—	6.00%	—
$PgDNO_{RFA}$L	18	2.80%	20.0%	5.40%	3.90%	13.0%	80.0%
$PgDNO_{RFA}$S	18	6.35%	7.00%	3.99%	—	6.70%	36.0%

Table 1: Comparison of different models on PDE problems (in ℓ_2 relative error) on 3% function data. Green: least value; Yellow: second-least; Red: highest (worst). Lower is better. RFA = Residual Frequency Attention; L = Large, S = Small.

- PgDNO maintain consistent performance advantage even with limited (3%) observation data, particularly on Darcy Flow problems

Sparse Observations

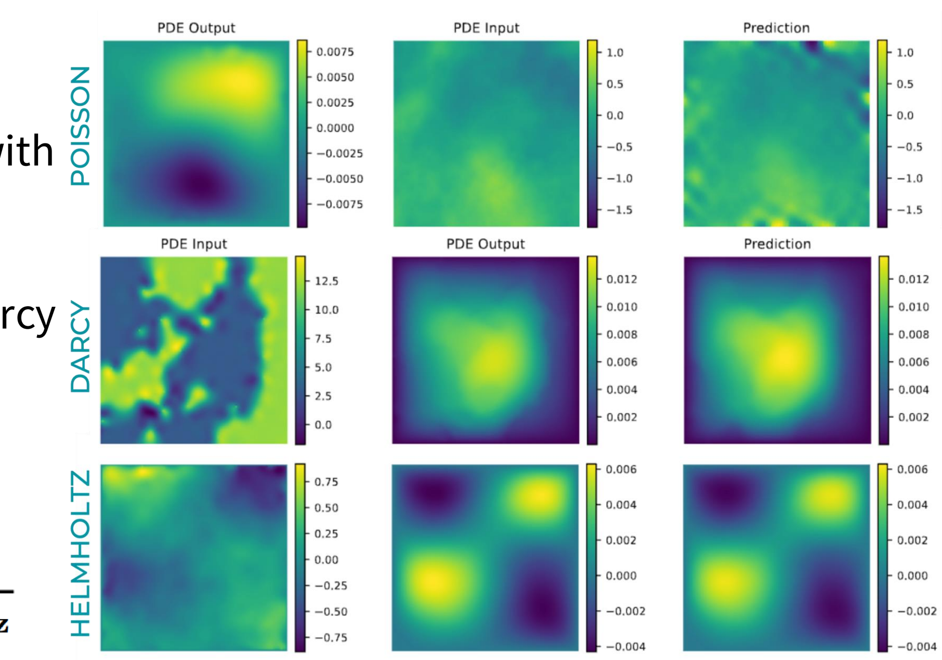

Model	Steps	Darcy (Fwd)	Darcy (Inv)	Poisson (Fwd)	Poisson (Inv)	Helmholtz (Fwd)	Helmholtz (Inv)
PINO	—	4.00%	2.1%	3.70%	10.2%	4.9%	4.9%
DeepONet	—	12.3%	8.4%	14.3%	29%	17.8%	28.1%
PINNs	—	15.4%	10.1%	16.1%	28.5%	18.1%	29.2%
FNO	—	5.3%	5.6%	8.2%	13.6%	11.1%	5.0%
DiffusionPDE	2000	5.98%	14.5%	15.27%	21.21%	10.9%	18.97%
FunDPS	200	1.1%	4.2%	-	-	-	-
FunDPS	500	1.4%	3.0%	-	-	-	-
FunDPS	2000	0.9%	2.1%	-	-	-	-
$PgDNO_{Hybrid}$	18	0.9%	3.67%	4.65%	10.9%	5.74%	9.8%
$PgDNO_{RFA}$L	18	1.73%	5.0%	5.9%	10.3%	2.1%	9.9%
$PgDNO_{RFA}$S	18	5.8%	5.0%	3.4%	13.2%	6.0%	14.7%

Table 2: Comparison of different models on PDE problems (in ℓ_2 relative error) on Full Observation Data. Green: least value; Yellow: second-least; Red: highest (worst). Lower is better. RFA = Residual Frequency Attention; L = Large, S = Small.

Full Observations FWD

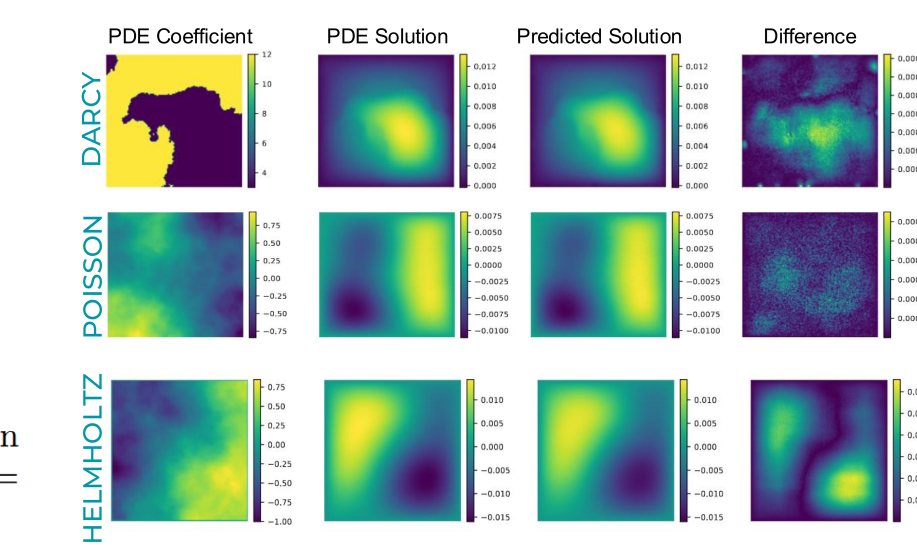

- PgDNO models achieve competitive results with only 18 steps vs 2000 steps for comparable diffusion models (100x speedup)
- On full observations, PgDNOHybrid matches or exceeds FunDPS (2000) while PgDNORFAL shows specialized strength on Helmholtz problems
- All PgDNO variants outperform traditional neural operators (FNO, PINO, DeepONet, PINNs) on forward problems

Future Work

- Extend the physics-guided diffusion neural operator to handle spatio-temporal domain
- Improve model performance under extremely sparse observations by incorporating advanced uncertainty quantification and adaptive sampling techniques
- Addressing spectral instabilities during training by developing regularization techniques specifically tailored for Fourier-based operators
- Investigating the complex loss landscapes that emerge from the interaction between physical constraints and frequency domain operations

Full Observations INV

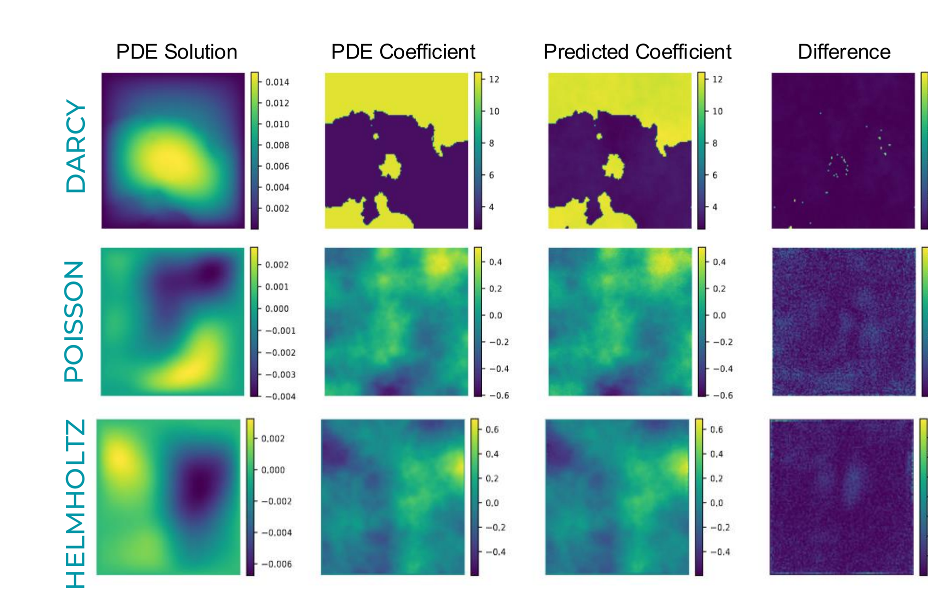

Acknowledgements: This work was supported in part by NSF awards IIS-2239328 and IIS-2107332. We are grateful to the Advanced Research Computing (ARC) Center at Virginia Tech for providing access to GPU compute resources for this project. This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). This work has also used Bridges-2 at Pittsburgh Supercomputing Center through allocation CIS250315 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.