

A variationally mimetic operator network

Deep Ray

Department of Mathematics
University of Maryland, College Park

Email: deepray@umd.edu

Website: deepray.github.io

*Jointly with Dhruv Patel, Michael Abdelmalik, Assad A. Oberai
& Thomas J. R. Hughes*

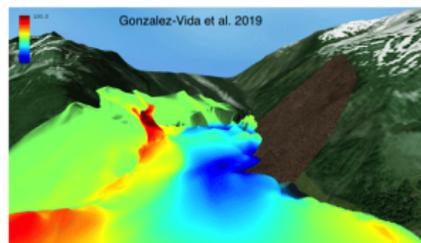
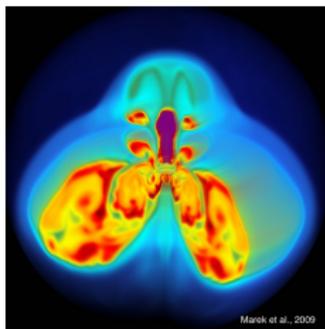
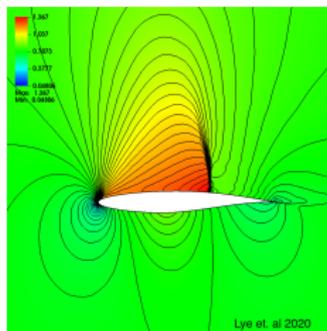
Supported by ARO grant W911NF2010050.

ICMTA-2024

SRM Institute of Science and Technology

January 2-3, 2024





- ▶ Problems of interest solved using finite difference/volume, finite element, spectral methods, ...
- ▶ **Many-query** applications need many calls to the solver: e.g.
 - Uncertainty quantification
 - PDE-constrained optimization
 - Inverse problems
- ▶ Single solve expensive for large-scale applications → **multiple solves $\mathcal{O}(10^5)$ prohibitively expensive!**
- ▶ Need to design **efficient, robust surrogates**.
- ▶ Recent interest in **deep learning-based surrogates** → focus of this talk!

- ▶ Deep learning with neural networks
- ▶ Solving PDEs using deep learning
- ▶ **Variationally Mimetic Operator Network (VarMiON)**
- ▶ Error estimates
- ▶ Numerical results
- ▶ Conclusion

A neural network is a parametrized mapping

$$NN_{\psi} : \Omega_x \rightarrow \Omega_y$$

typically formed by alternating composition

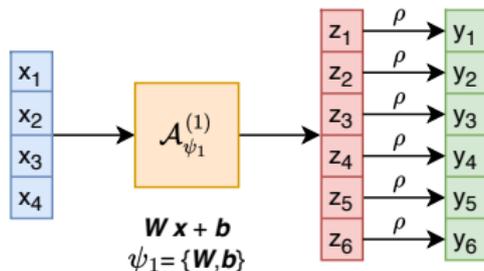
$$NN_{\psi} := \rho \circ \mathcal{A}_{\psi_{L+1}}^{(L+1)} \circ \rho \circ \mathcal{A}_{\psi_L}^{(L)} \circ \rho \circ \mathcal{A}_{\psi_{L-1}}^{(L-1)} \circ \dots \circ \rho \circ \mathcal{A}_{\psi_1}^{(1)}$$

where

$\psi = \{\psi_k\}_{k=1}^L \rightarrow$ trainable weights and biases of the network

$\mathcal{A}_{\psi_k}^{(k)} \rightarrow$ parametrized affine transformation

$\rho \rightarrow$ non-linear activation function



A neural network is a parametrized mapping

$$NN_{\psi} : \Omega_x \rightarrow \Omega_y$$

typically formed by alternating composition

$$NN_{\psi} := \rho \circ \mathcal{A}_{\psi_{L+1}}^{(L+1)} \circ \rho \circ \mathcal{A}_{\psi_L}^{(L)} \circ \rho \circ \mathcal{A}_{\psi_{L-1}}^{(L-1)} \circ \dots \circ \rho \circ \mathcal{A}_{\psi_1}^{(1)}$$

where

$\psi = \{\psi_k\}_{k=1}^L \rightarrow$ trainable weights and biases of the network

$\mathcal{A}_{\psi_k}^{(k)} \rightarrow$ parametrized affine transformation

$\rho \rightarrow$ non-linear activation function

Usage:

- ▶ Let \mathbf{x} and \mathbf{y} be related in some manner, say $\mathbf{y} = \mathcal{F}(\mathbf{x})$.
- ▶ We are only given $\mathcal{S} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$.
- ▶ NN_{ψ} can be used to learn \mathcal{F} .

Consider a suitable measure

$$\mu : \Omega_x \times \Omega_y \times \Omega_x \times \Omega_y \rightarrow \mathbb{R}$$

s.t. $\mu(\mathbf{x}, \mathbf{y}, \mathbf{x}, NN_{\psi}(\mathbf{x}))$ is the error/discrepancy between $(\mathbf{x}, \mathbf{y}) \in \mathcal{S}$ and $(\mathbf{x}, NN_{\psi}(\mathbf{x}))$.

Define the loss/objective function

$$\Pi(\psi) = \frac{1}{N} \sum_{i=1}^N \mu(\mathbf{x}_i, \mathbf{y}_i, \mathbf{x}_i, NN_{\psi}(\mathbf{x}_i))$$

Solve the optimization problem

$$\psi^* = \arg \min_{\psi} \Pi(\psi)$$

Then $NN_{\psi^*} \approx \mathcal{F}$

Also need to tune network **hyper-parameters**:

- Width
- Depth (L)
- Activation function ρ
- Optimizer
- Loss function
- Dataset

Consider the generic PDE:

$$\mathcal{L}(u(\mathbf{x})) = f(\mathbf{x}), \quad \forall \mathbf{x} \in \Omega$$

- ▶ Solution approximated by a network $\hat{u}(\mathbf{x}; \boldsymbol{\psi})$ with parameters $\boldsymbol{\psi}$.
- ▶ Minimize PDE residual at collocation points $\{\mathbf{x}_i\}$ (Lagaris et al., 2000): Solve

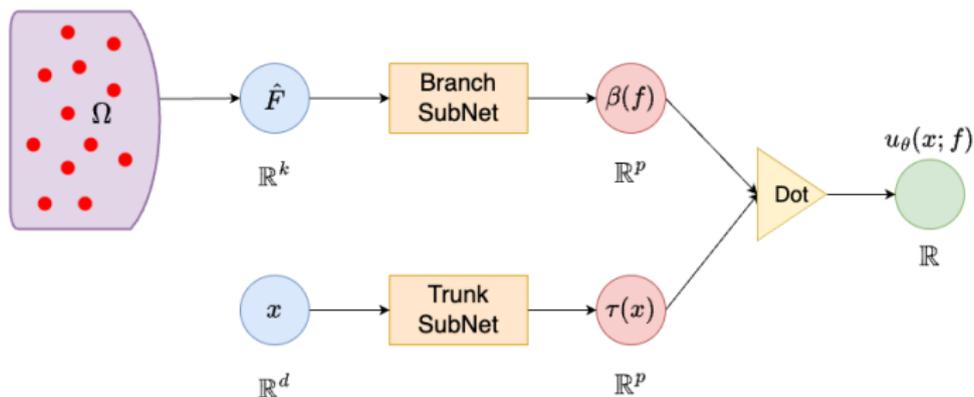
$$\boldsymbol{\psi}^* = \arg \min_{\boldsymbol{\psi}} \Pi(\boldsymbol{\psi}), \quad \Pi(\boldsymbol{\psi}) = \frac{1}{N} \sum_{i=1}^N \|\mathcal{L}(\hat{u}(\mathbf{x}_i; \boldsymbol{\psi})) - f(\mathbf{x}_i)\|^2$$

- ▶ Rediscovered as **Physics Informed Neural Nets (PINNs)** with deeper structures (Raissi et al., 2019).
- ▶ However, solves **one instance** of the PDE – must be retrained if f changes.

We are interested in approximating the **solution operator**

$$\mathcal{S} : \mathcal{F} \longrightarrow \mathcal{V}, \quad \mathcal{S}(f) = u(\cdot; f)$$

- ▶ Network-based operator learning with **shallow networks** [Chen & Chen, 1995].
- ▶ **DeepONets**: extension to deeper architectures [Lu et al., 2021]
- ▶ Comprise two subnetworks: the **trunk** (basis) and the **branch** (basis).



We are interested in approximating the **solution operator**

$$\mathcal{S} : \mathcal{F} \longrightarrow \mathcal{V}, \quad \mathcal{S}(f) = u(\cdot; f)$$

- ▶ Network-based operator learning with **shallow networks** [Chen & Chen, 1995].
- ▶ **DeepONets**: extension to deeper architectures [Lu et al., 2021]
- ▶ Comprise two subnetworks: the **trunk** (basis) and the **branch** (basis).
- ▶ **Neural operators** – an alternate strategy to approximate \mathcal{S}
- ▶ Come in many flavors: Fourier Neural Operators [Li et al., 2020], Graph Kernel Net [Li et al., 2020], PCA-NET [Bhattacharya et al., 2021], ...

We are interested in approximating the **solution operator**

$$\mathcal{S} : \mathcal{F} \longrightarrow \mathcal{V}, \quad \mathcal{S}(f) = u(\cdot; f)$$

- ▶ Network-based operator learning with **shallow networks** [Chen & Chen, 1995].
- ▶ **DeepONets**: extension to deeper architectures [Lu et al., 2021]
- ▶ Comprise two subnetworks: the **trunk** (basis) and the **branch** (basis).
- ▶ **Neural operators** – an alternate strategy to approximate \mathcal{S}
- ▶ Come in many flavors: Fourier Neural Operators [Li et al., 2020], Graph Kernel Net [Li et al., 2020], PCA-NET [Bhattacharya et al., 2021], ...
- ▶ This talk: Operator Networks that **mimic (approximate) variational form*** of the PDE.

[* Variationally Mimetic Operator Networks; Patel, Ray, Abdelmalik, Hughes, Oberai; CMAME, 2024]

- ▶ Consider a generic linear elliptic PDE:

$$\begin{aligned}\mathcal{L}(u(\mathbf{x}); \theta(\mathbf{x})) &= f(\mathbf{x}), & \forall \mathbf{x} \in \Omega, \\ \mathcal{B}(u(\mathbf{x}); \theta(\mathbf{x})) &= \eta(\mathbf{x}), & \forall \mathbf{x} \in \Gamma_\eta, \\ u(\mathbf{x}) &= 0, & \forall \mathbf{x} \in \Gamma_g,\end{aligned}$$

where $f \in \mathcal{F} \subset L^2(\Omega)$, $\eta \in \mathcal{N} \subset L^2(\Gamma_\eta)$, $\theta \in \mathcal{T} \subset L^\infty(\Omega)$.

- ▶ The **variational formulation**: find $u \in \mathcal{V} \subset H_g^1$ such that $\forall w \in \mathcal{V}$,

$$a(w, u; \theta) = (w, f) + (w, \eta)_{\Gamma_\eta}.$$

- ▶ The solution operator is

$$\begin{aligned}\mathcal{S} : \mathcal{X} = \mathcal{F} \times \mathcal{T} \times \mathcal{N} &\longrightarrow \mathcal{V} \subset H_g^1 \\ (f, \theta, \eta) &\mapsto u(\cdot; f, \theta, \eta)\end{aligned}$$

This mapping can be **non-linear in θ** .

- ▶ Evaluate approximate solution in **finite-dimensional space**

$$\mathcal{V}^h = \text{span}\{\phi_i(\mathbf{x}) : 1 \leq i \leq q\}.$$

- ▶ Discrete weak formulation: find $u^h \in \mathcal{V}^h$ such that $\forall w^h \in \mathcal{V}^h$,

$$a(w^h, u^h; \theta^h) = (w^h, f^h) + (w, \eta^h)_{\Gamma_\eta}.$$

- ▶ Any function $v^h \in \mathcal{V}^h$ can be written as

$$v^h(\mathbf{x}) = \mathbf{V}^\top \boldsymbol{\Phi}(\mathbf{x}), \quad \mathbf{V} = (v_1, \dots, v_q)^\top, \quad \boldsymbol{\Phi}(\mathbf{x}) = (\phi_1(\mathbf{x}), \dots, \phi_q(\mathbf{x}))^\top.$$

Plugging this into discrete weak form gives...

- ▶ Linear system of equations,

$$\mathbf{K}(\theta^h)\mathbf{U} = \mathbf{M}\mathbf{F} + \widetilde{\mathbf{M}}\mathbf{N}$$

where the matrices are given by

$$K_{ij}(\theta^h) = a(\phi_i, \phi_j; \theta^h), \quad M_{ij} = (\phi_i, \phi_j), \quad \widetilde{M}_{ij} = (\phi_i, \phi_j)_{\Gamma_\eta} \quad 1 \leq i, j \leq q.$$

- ▶ Discrete solution operator is

$$\begin{aligned} \mathcal{S}^h : \mathcal{X}^h = \mathcal{F}^h \times \mathcal{T}^h \times \mathcal{N}^h &\longrightarrow \mathcal{V}^h \\ (f^h, \theta^h, \eta^h) &\mapsto u^h(\cdot; f^h, \theta^h, \eta^h) = \mathbf{B}(f^h, \eta^h, \theta^h)^\top \boldsymbol{\Phi} \end{aligned}$$

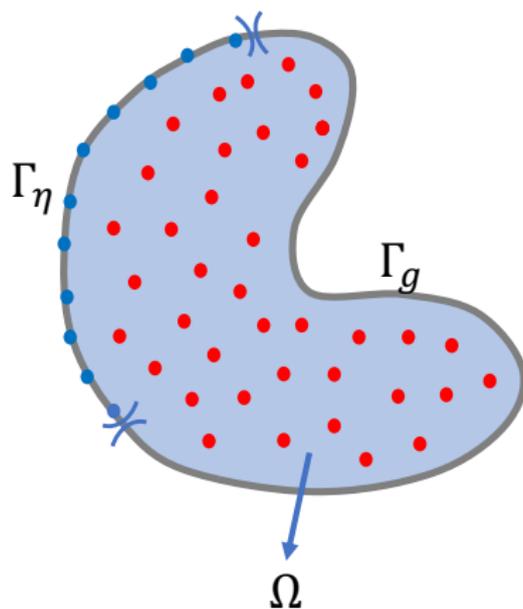
where

$$\mathbf{B}(f^h, \theta^h, \eta^h) = \mathbf{U} = \mathbf{K}^{-1}(\theta^h)(\mathbf{M}\mathbf{F} + \widetilde{\mathbf{M}}\mathbf{N}).$$

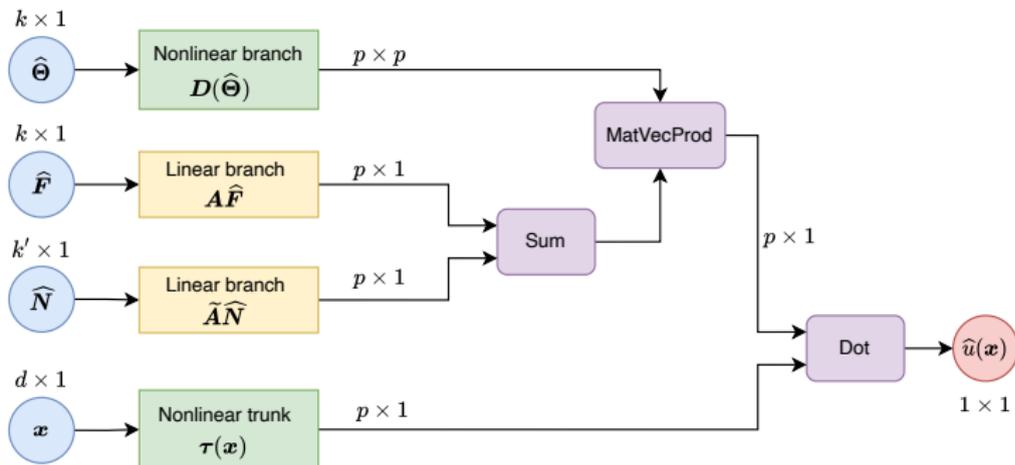
VarMiON will mimic this structure!

Evaluate (f, θ, η) at some fixed **sensor nodes** to get the discrete sample vectors

$$\hat{\mathbf{F}} = (f(\hat{\mathbf{x}}_1), \dots, f(\hat{\mathbf{x}}_k))^\top, \quad \hat{\Theta} = (\theta(\hat{\mathbf{x}}_1), \dots, \theta(\hat{\mathbf{x}}_k))^\top, \quad \hat{\mathbf{N}} = (\eta(\hat{\mathbf{x}}_1^b), \dots, \eta(\hat{\mathbf{x}}_{k'}^b))^\top$$

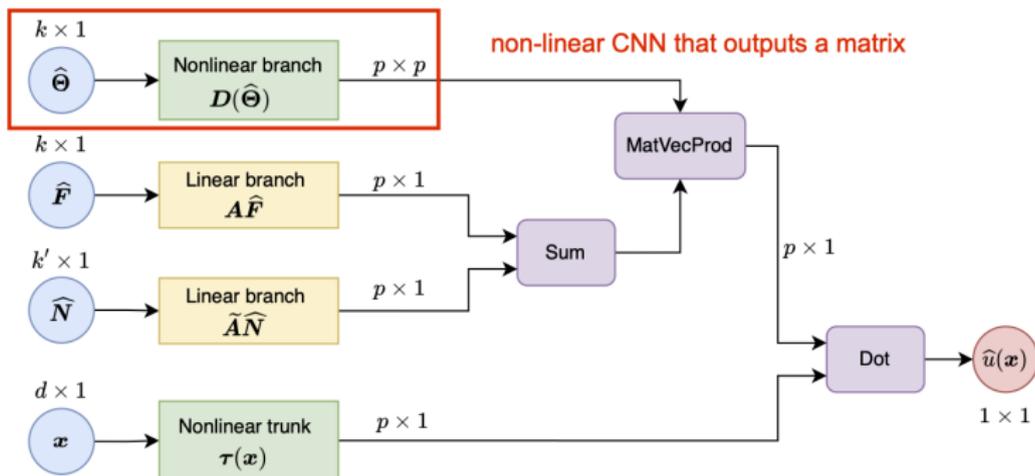


The network is comprises several sub-networks with **latent dimension p** :



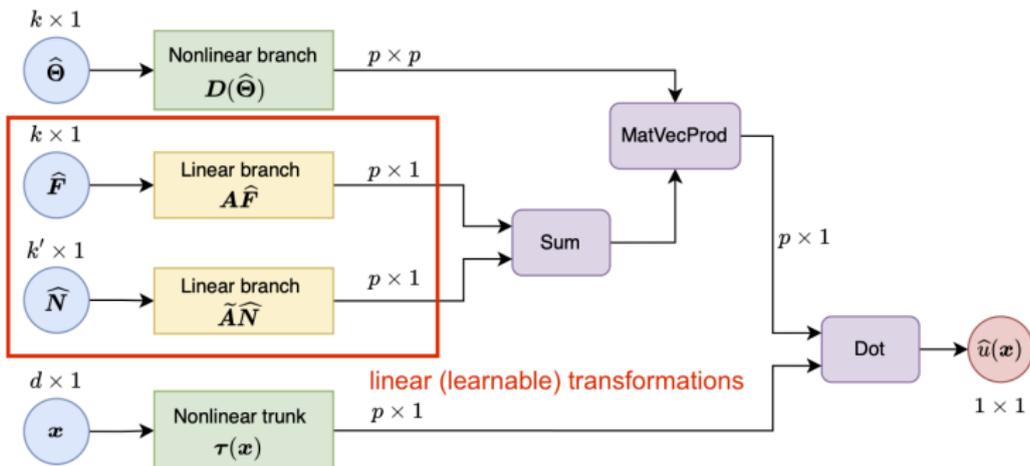
The network is comprised of several sub-networks with latent dimension p :

- **Non-linear** branch net: $\hat{\Theta} \mapsto D(\hat{\Theta}) \in \mathbb{R}^{p \times p}$



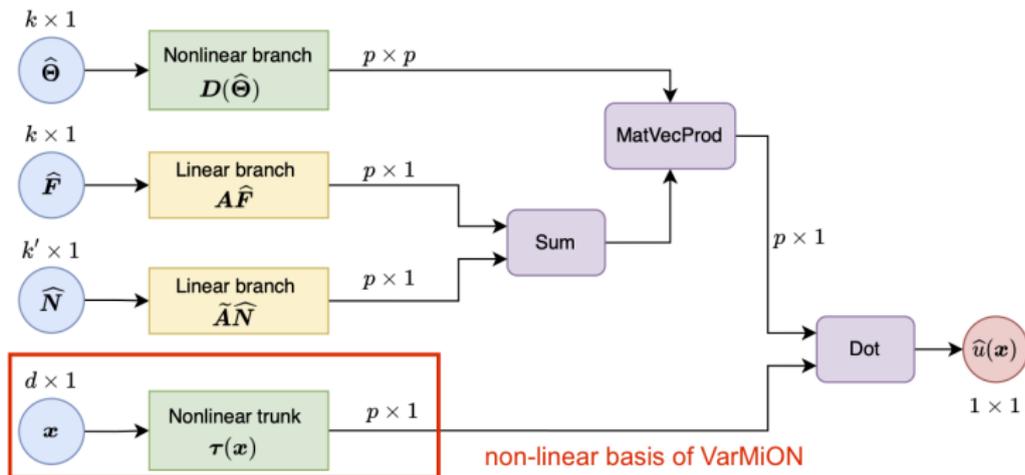
The network is comprised of several sub-networks with **latent dimension p** :

- ▶ **Non-linear** branch net: $\widehat{\Theta} \mapsto D(\widehat{\Theta}) \in \mathbb{R}^{p \times p}$
- ▶ Two **linear** branches with learnable matrices \mathbf{A} , $\widetilde{\mathbf{A}}$



The network is comprised of several sub-networks with **latent dimension p** :

- ▶ **Non-linear** branch net: $\widehat{\Theta} \mapsto D(\widehat{\Theta}) \in \mathbb{R}^{p \times p}$
- ▶ Two **linear** branches with learnable matrices \mathbf{A} , $\widetilde{\mathbf{A}}$
- ▶ **Non-linear** trunk (basis of VarMiON): $\mathbf{x} \mapsto \boldsymbol{\tau}(\mathbf{x}) = (\tau_1(\mathbf{x}), \dots, \tau_p(\mathbf{x}))^\top$



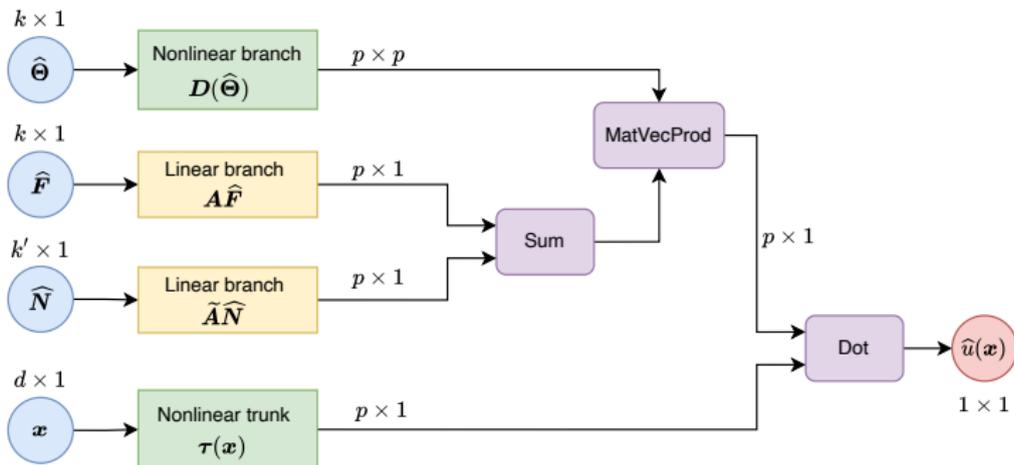
VarMiON operator is

$$\widehat{\mathcal{S}} : \mathbb{R}^k \times \mathbb{R}^k \times \mathbb{R}^{k'} \longrightarrow \mathcal{V}^\tau = \text{span}\{\tau_i(\mathbf{x}) : 1 \leq i \leq p\}$$

$$(\widehat{\mathbf{F}}, \widehat{\Theta}, \widehat{\mathbf{N}}) \mapsto \widehat{u}(:, \widehat{\mathbf{F}}, \widehat{\Theta}, \widehat{\mathbf{N}}) = \beta(f^h, \eta^h, \theta^h)^\top \boldsymbol{\tau}$$

where

$$\beta(\widehat{\mathbf{F}}, \widehat{\Theta}, \widehat{\mathbf{H}}) = D(\widehat{\Theta})(\mathbf{A}\widehat{\mathbf{F}} + \widetilde{\mathbf{A}}\widehat{\mathbf{N}}).$$



VarMiON operator is

$$\widehat{S} : \mathbb{R}^k \times \mathbb{R}^k \times \mathbb{R}^{k'} \longrightarrow \mathcal{V}^\tau = \text{span}\{\tau_i(\mathbf{x}) : 1 \leq i \leq p\}$$

$$(\widehat{F}, \widehat{\Theta}, \widehat{N}) \mapsto \widehat{u}(\cdot; \widehat{F}, \widehat{\Theta}, \widehat{N}) = \beta(f^h, \eta^h, \theta^h)^\top \tau$$

where

$$\beta(\widehat{F}, \widehat{\Theta}, \widehat{H}) = D(\widehat{\Theta})(A\widehat{F} + \widetilde{A}\widehat{N}).$$

Compare this to the discrete solution operator:

$$S^h : \mathcal{F}^h \times \mathcal{T}^h \times \mathcal{N}^h \longrightarrow \mathcal{V}^h$$

$$(f^h, \theta^h, \eta^h) \mapsto u^h(\cdot; f^h, \theta^h, \eta^h) = B(f^h, \eta^h, \theta^h)^\top \Phi$$

where

$$B(f^h, \theta^h, \eta^h) = K^{-1}(\theta^h)(M\mathbf{F} + \widetilde{M}\mathbf{N}).$$

In comparison with the variational formulation

- ▶ Can prove D is the **reduced order** counterpart of K^{-1} ($p \ll q$)
- ▶ While the basis Φ are **fixed**, τ are **learned** from training data.

In comparison with a vanilla DeepONet

- ▶ DeepONet typically has a **single nonlinear branch** for inputs.
- ▶ VarMiON explicitly **constructs matrix** operators. DeepONet does not.

1. For $1 \leq j \leq J$, consider distinct samples $(f_j, \theta_j, \eta_j) \in \mathcal{X}$.
2. Obtain the discrete approximations $(f_j^h, \theta_j^h, \eta_j^h) \in \mathcal{X}^h$.
3. Find the discrete numerical solution $u_j^h = \mathcal{S}^h(f_j^h, \theta_j^h, \eta_j^h)$.
4. Choose output nodes $\{\mathbf{x}_l\}_{l=1}^L$ to sample the numerical solution $u_{jl}^h = u_j^h(\mathbf{x}_l)$.
5. Generate the input vectors $(\widehat{\mathbf{F}}_j, \widehat{\Theta}_j, \widehat{\mathbf{N}}_j)$.
6. Collect input & output to form training set with $J \times L$ samples

$$\mathbb{S} = \{(\widehat{\mathbf{F}}_j, \widehat{\Theta}_j, \widehat{\mathbf{N}}_j, \mathbf{x}_l, u_{jl}^h) : 1 \leq j \leq J, 1 \leq l \leq L\},$$

Find the network weights that **minimize the loss** function

$$\Pi(\boldsymbol{\psi}) = \frac{1}{J} \sum_{j=1}^J \Pi_j(\boldsymbol{\psi}), \quad \Pi_j(\boldsymbol{\psi}) = \sum_{l=1}^L w_l \left(u_{jl}^h - \widehat{\mathcal{S}}_{\boldsymbol{\psi}}(\widehat{\mathbf{F}}_j, \widehat{\Theta}_j, \widehat{\mathbf{N}}_j)[\mathbf{x}_l] \right)^2.$$

where $\boldsymbol{\psi}$ are all the trainable parameters of the VarMiON.

The **generalization error** for any $(f, \theta, \eta) \in \mathcal{X}$

$$\mathcal{E}(f, \theta, \eta) := \|\mathcal{S}(f, \theta, \eta) - \widehat{\mathcal{S}}(\widehat{\mathbf{F}}, \widehat{\Theta}, \widehat{\mathbf{N}})\|_{L^2}.$$

Split into four errors:

$$\begin{aligned} \mathcal{E}(f, \theta, \eta) \leq & \|\mathcal{S}(f, \theta, \eta) - \mathcal{S}(f_j, \theta_j, \eta_j)\|_{L^2} \longrightarrow \text{Stability of } \mathcal{S} \\ & + \|\mathcal{S}(f_j, \theta_j, \eta_j) - \mathcal{S}^h(f_j^h, \theta_j^h, \eta_j^h)\|_{L^2} \longrightarrow \text{Numerical error in generating data} \\ & + \|\mathcal{S}^h(f_j^h, \theta_j^h, \eta_j^h) - \widehat{\mathcal{S}}(\widehat{\mathbf{F}}_j, \widehat{\Theta}_j, \widehat{\mathbf{N}}_j)\|_{L^2} \longrightarrow \text{Training error of VarMiON} \\ & + \|\widehat{\mathcal{S}}(\widehat{\mathbf{F}}_j, \widehat{\Theta}_j, \widehat{\mathbf{N}}_j) - \widehat{\mathcal{S}}(\widehat{\mathbf{F}}, \widehat{\Theta}, \widehat{\mathbf{N}})\|_{L^2} \longrightarrow \text{Stability of VarMiON } \widehat{\mathcal{S}} \end{aligned}$$

Generalization error estimate (Patel et al., 2022)

If $\mathcal{X} := \mathcal{F} \times \mathcal{T} \times \mathcal{N}$ is compact, the non-linear branch is Lipschitz, i.e.,

$$\|D(\widehat{\Theta}) - D(\widehat{\Theta}')\|_2 \leq L_D \|\widehat{\Theta} - \widehat{\Theta}'\|_2.$$

Then, the generalization error can be bounded as

$$\mathcal{E}(f, \theta, \eta) \leq C \left(\epsilon_h + \epsilon_s + \sqrt{\epsilon_t} + \frac{1}{k^{\alpha/2}} + \frac{1}{(k')^{\alpha'/2}} + \frac{1}{L^{\gamma/2}} \right)$$

where

$\epsilon_h \rightarrow$ numerical error in training data

$\epsilon_s \rightarrow$ covering estimate

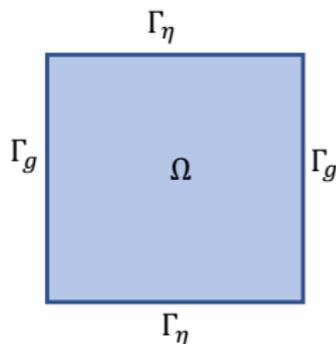
$\epsilon_t \rightarrow$ training error

$\alpha, \alpha', \gamma \rightarrow$ quadrature convergence rates

The constant C depends on the stability constants of \mathcal{S} and $\widehat{\mathcal{S}}$.

Steady-state heat conduction

$$\begin{aligned} -\nabla \cdot (\theta(\mathbf{x}) \nabla u(\mathbf{x})) &= f(\mathbf{x}), & \forall \mathbf{x} \in \Omega, \\ \theta(\mathbf{x}) \nabla u(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x}) &= \eta(\mathbf{x}), & \forall \mathbf{x} \in \Gamma_\eta, \\ u(\mathbf{x}) &= 0, & \forall \mathbf{x} \in \Gamma_g. \end{aligned}$$



- ▶ Input: thermal conductivity θ , heat sources f , and heat flux η . Output: temperature u .
- ▶ Inputs: Gaussian Random Fields.
- ▶ Networks with two inputs (θ, f) and three inputs (θ, f, η) .
- ▶ Compare VarMiON ($p = 100$) and vanilla DeepONet.
- ▶ Similar number of network parameters. Identical trunk architecture.
- ▶ Robustness: sampling (spatially uniform or random) and trunk functions (ReLU or RBF).

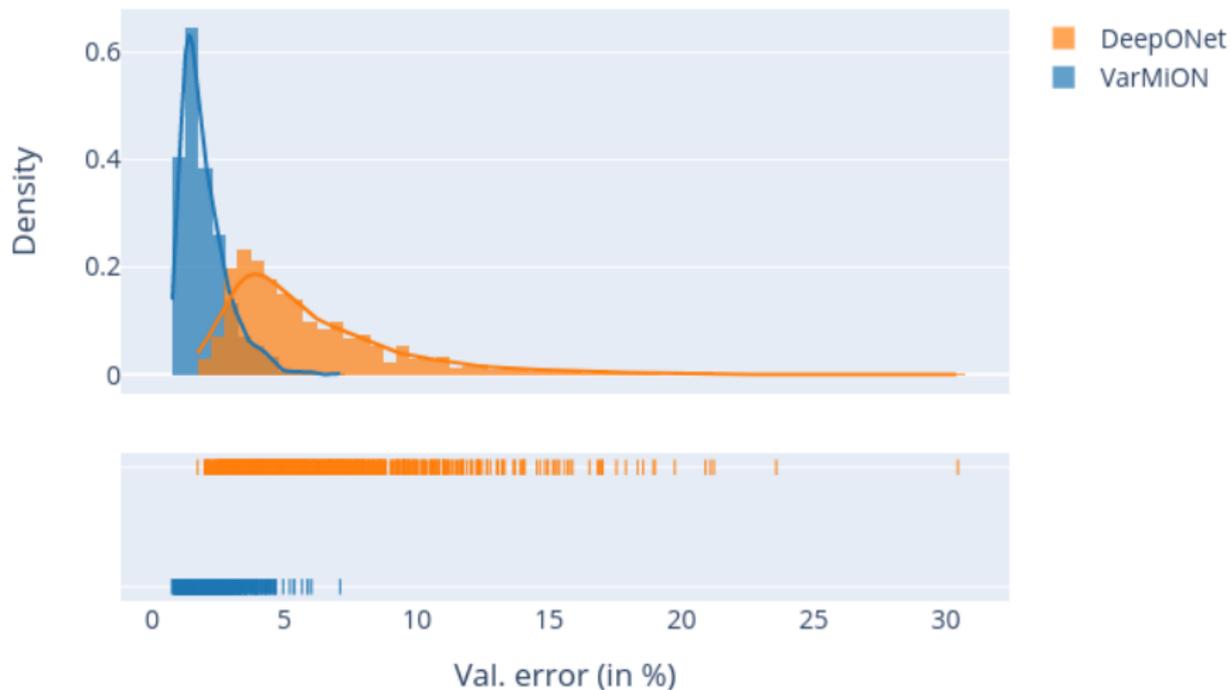
Numerical Example

- ▶ 10,000 samples generated using Fenics.
- ▶ 9,000 for training/validation and 1,000 for testing.
- ▶ Average value of relative L_2 error reported below.

Case	Model	Number of parameters	Relative L_2 error
Two input functions with randomly sampled input	DeepONet (w/ ReLU trunk)	111,248	1.07 ± 0.39 %
	VarMiON(w/ ReLU trunk)	109,013	0.93 ± 0.28 %
Two input functions with uniformly sampled input	DeepONet (w/ ReLU trunk)	49,928	1.98 ± 0.79 %
	VarMiON(w/ ReLU trunk)	46,281	1.05 ± 0.42 %
Two input functions with uniformly sampled input	DeepONet (w/ RBF trunk)	17,911	1.39 ± 0.60 %
	VarMiON(w/ RBF trunk)	17,345	0.64 ± 0.35 %
Three input functions with uniformly sampled input	DeepONet (w/ RBF trunk)	24,543	5.99 ± 4.24 %
	VarMiON(w/ RBF trunk)	23,065	2.06 ± 0.90 %

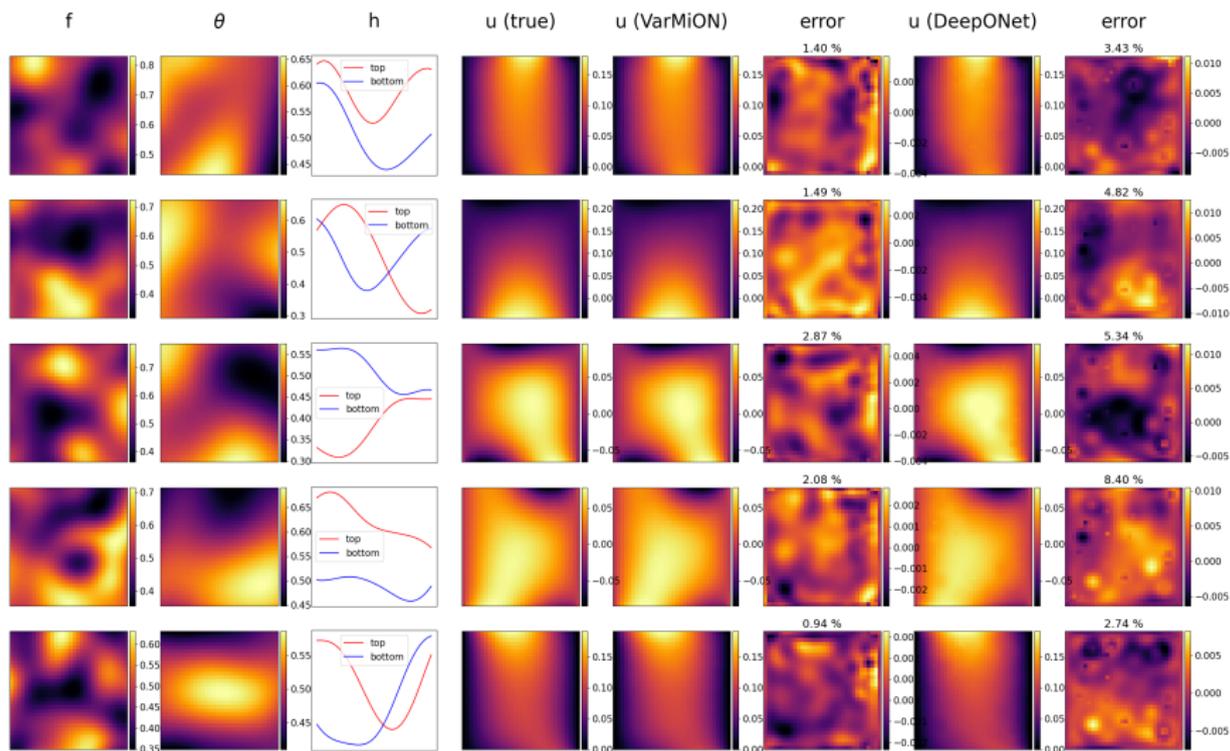
Numerical Example: Three inputs

Density of scaled L_2 error (RBF trunk and uniform spatial sampling).



Numerical Example: Three inputs

Predictions by VarMiON and DeepONet.

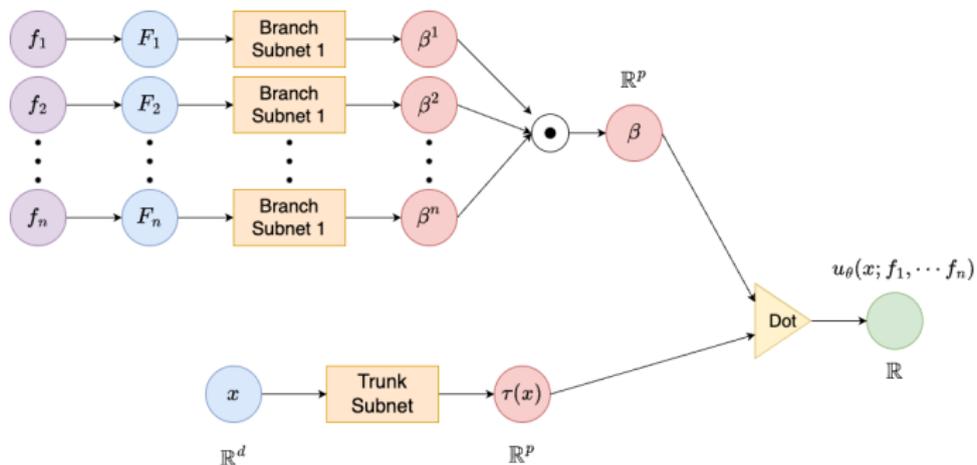


Numerical Example: Comparison with MIONet

MIONet has been proposed [Jin et al., 2022] to handle multiple input functions.

Separate branch for each f_i , followed by a **Hadamard** product

$$\beta = (\beta^1 \odot \beta^2 \odot \dots \odot \beta^n), \quad u_\theta(x; f_1, \dots, f_n) = \beta^\top \tau(x)$$



Comparison with MIONet as the number of training samples n is varied

Training dataset size	Model	Relative L_2 error
$n = 1000$	DeepONet	10.23 ± 5.20
	MIONet	88.07 ± 69.68
	VarMiON	4.27 ± 2.23
$n = 2000$	DeepONet	9.00 ± 5.63
	MIONet	85.03 ± 123.77
	VarMiON	4.04 ± 1.86
$n = 4000$	DeepONet	7.19 ± 3.75
	MIONet	88.39 ± 62.10
	VarMiON	2.90 ± 1.50
$n = 6000$	DeepONet	6.28 ± 3.55
	MIONet	82.89 ± 34.19
	VarMiON	2.74 ± 1.31

- ▶ VarMiON needs to be carefully constructed due to the **dependence** on the structure of the **variational form**.
- ▶ Have constructed a VarMiON for the **regularized Eikonal equation**.
- ▶ We believe the VarMiON architecture needs to be **customized** depending on the **type of PDE** (ongoing work)

- ▶ VarMiON: an operator network that **mimics variational formulation**.
- ▶ Takes the form of a **reduced order model**.
- ▶ Precise specification of the branch network – **depends on weak form** of PDE.
- ▶ **Error analysis** reveals important components – currently investigating training with appropriate **Sobolev loss**.
- ▶ Numerical results point to better and more **robust** performance.
- ▶ Several extensions: other nonlinear operators, physics-informed residuals, time-dependent problems, hyperbolic systems, and specification of geometry.