

Neural implicit representation for PDEs and hybrid numerical methods

J. Aghili¹², H. Barucq³, F. Foucher³, E. Franck¹², V. Michel-Dansac¹², L. Navoret¹², N. Victorion³, V. Vigon¹²

....

¹Inria Nancy Grand Est, France

²IRMA, Strasbourg university, France

³Inria Bordeaux, Pau center, France

Outline

Introduction

Physics-informed Neural Networks

New paradigm: Operator learning

Application to numerical methods

Conclusion


Numerical methods and implicit neural representation

Numerical methods

- We begin with a simple example:

$$\begin{cases} L_{t,x}u = \partial_t u - \Delta u = 0 \\ u(t=0, x) = u_0(x) \\ u(x) = g \text{ on } \partial\Omega \end{cases}$$

- Solving a PDE amounts to solving a infinite-dimensional problem.
- **Numerical method**: transform the PDE into a finite-dimensional problem of dimension N with convergence to the PDE solution when $N \rightarrow \infty$
- How to summarize most of numerical methods? (drawing from S. Mishra)



- Definitions:
 - \mathcal{E} , the **encoder**, transforms the data (initial conditions, RHS) into a finite dimensional vector. Transformed data are called **degree of freedoms** (DoF).
 - \mathcal{D} , the **decoder**, transforms degrees of freedom **into a function**.
 - \mathcal{A} , the **approximator**, transforms the DoF of the RHS into the DoF of the approximate solution.

Why numerical methods require a mesh?

Polynomial Lagrange interpolation

We consider a domain $[a, b]$. There exists a polynomial P of degree k such that, for any $f \in C^0([a, b])$,

$$|f(x) - P(x)| \leq |b - a|^k \max_{x \in [a, b]} |f^{k+1}(x)|.$$

- On small domains ($|b - a| \ll 1$) or for large k , this polynomial gives a very good approximation.
- Very high degrees k can generate oscillations.
- To enforce small domains: we **introduce a mesh and a cell-wise polynomial approximation**

First step: choose a parametric function

We define a mesh by splitting the geometry in small sub-intervals $[x_i, x_{i+1}]$, and we propose the following candidate to approximate the PDE solution u

$$u|_{[x_i, x_{i+1}]}(t, x) = \sum_{j=1}^k \alpha_j(t) \phi_j(x).$$

This is a **piecewise polynomial representation**.

Finite element, finite volume, discontinuous Galerkin

Finite element method

- **Encoder:** transforms the function f into $\alpha(t)$ the FE DoF (pointwise values, face/edge integral values, ...)
- **Decoder:** $D(\alpha)(t, x) = \sum_{i=1}^N \alpha_i(t) \phi_i(x)$ with $\phi_i(x)$ a compactly supported basis function defined on the whole mesh
- **Approximator:** we **plug the decoder** in the weak form of the equations to obtain an ODE or an algebraic system on α

Finite volume and discontinuous Galerkin method

- **Encoder:** transforms the function f into $\alpha(t)$ the FE DoF (average values, modal values, nodal values, ...)
- **Decoder:** $D(\alpha)(t, x)|_{\Omega_j} = \sum_{i=1}^N \alpha_i(t) \phi_i(x)$ with $\phi_i(x)$ a local cell-wise basis function.
- **Approximator:** we **plug the decoder** in the weak form of the equations to obtain an ODE or an algebraic system on α , in each cell
- For this method, the **decoder** generates a **finite-dimensional vector space**.
- The method projects a form of the equation on this finite-dimensional space. **Uniqueness** is ensured **by the Hilbert projection theorem**.
- **Convergence** is ensured: increasing the number of DoF (mesh, polynomial degree) makes the error decrease.

Spectral methods

Spectral theorem

The **spectral theorem** in Hilbert spaces proposes an approximation of any function in H by

$$u(x) = \sum_{i=1}^N \alpha_i \phi_i(x),$$

with $\phi_i(x)$ the orthonormal global Hilbert basis, and $\alpha_i = \langle f, \phi_i \rangle$.

Spectral method


- **Encoder:** Projection of the function f in the spectral basis. DoF: $\alpha_i = \langle f, \phi_i \rangle$
- **Decoder:** $D(\alpha)(t, x) = \sum_{i=1}^N \alpha_i(t) \phi_i(x)$ with $\phi_i(x)$ the first modes of the Hilbert basis.
- **Approximator:** we **plug the decoder** in the weak/strong form of the equations to obtain an ODE or an algebraic system on α .
- For this method, the **decoder** generates a **finite-dimensional vector space**.
- The method projects a form of the equation on this finite-dimensional space, using the **Unicity by Hilbert projection theorem**.
- **Convergence** is ensured: increasing the number of DoF (number of modes) makes the error decrease.

Idea

Represent the solution as a **sum of radial basis functions localized** at some points:

$$u(x) = \sum_{i=1}^N \alpha_i \phi_i(|x - x_i|)$$

with $\phi_i(r)$ a radial basis function such as $\phi(r) = e^{-(\varepsilon r)^2}$ or $\phi(r) = \frac{1}{1+(\varepsilon r)^2}$. Larger values of ε give more localized functions.



Radial basis method

- **Encoder:** Projection of the function f . DoF: **weights of the radial functions**
- **Decoder:** $D(\alpha)(t, x) = \sum_{i=1}^N \alpha_i(t) \phi(|x - x_i|)$ with $\phi(x)$ a radial basis function.
- **Approximator:** just like before, the decoder is plugged in the equation.
- Like before, we have a finite-dimensional function space.
- **Convergence:** increasing the number of points (DoF) makes the error decrease.

Reduced basis methods

- For many years, there has been research to propose reduced order models (including experts here in Bordeaux!).
- One of the classical approaches is the **Reduced basis method**.
- It represents a subset of solutions like


$$u(x) = \sum_{i=1}^N \alpha_i \phi_i(x),$$

where $\phi_i(x)$ is a **spectral basis computed to efficiently represent a subset of solutions** associated to a subset of parameters.

- Such methods can be viewed as data-driven Spectral methods.
- They have properties similar to the spectral method's.
- Example:

$$\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \frac{1}{Re} \partial_{xx} \rho$$

- Reynolds number $Re = 40$, 10 modes



Reduced basis methods

- For many years, there has been research to propose reduced order models (including experts here in Bordeaux!).
- One of the classical approaches is the **Reduced basis method**.
- It represents a subset of solutions like


$$u(x) = \sum_{i=1}^N \alpha_i \phi_i(x),$$

where $\phi_i(x)$ is a **spectral basis computed to efficiently represent a subset of solutions** associated to a subset of parameters.

- Such methods can be viewed as data-driven Spectral methods.
- They have properties similar to the spectral method's.
- Example:

$$\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \frac{1}{Re} \partial_{xx} \rho$$

- Reynolds number $Re = 400000$, 40 modes



Space and space-time decoder

- Classical methods (FE/FV/DG/...) involve a decoder where only the space representation is fixed:

$$u(t, x) = \sum_{i=1}^N \alpha_i(t) \phi_i(x).$$

- Plugging this decoder in the equation, we obtain an **ODE** to solve.
- A more recent approach, **space-time methods**, proposes to fix both space and time representations:

$$u(t, x) = \sum_{i=1}^N \alpha_i \phi_i(t, x).$$

- Plugging this decoder in the equation we obtain an **algebraic system** to solve.

Explicit vs implicit representations

- Representations are called explicit if the degrees of freedom can be explicitly computed and understood from the function.
- FE/FV/DG/spectral methods use **explicit representations** (average value, ...).
- The **radial basis method**, however, uses a **partially explicit representation**. It is difficult to understand the DoF from the function, but they can easily be computed by inverting the mass matrix (projector).

Key idea

Summary

Every previously mentioned space and space-time methods consists in:

1. choosing a **linear representation** (linear combination of basis functions), either local (on a mesh) or global;
2. plugging this representation **into the equation** to obtain algebraic relations (linear for linear problems, nonlinear for nonlinear problems) or ODEs.
3. solving this **algebraic relation** with a linear solver or Newton's method, using a time scheme to solve the ODE.

In all these cases, **the decoder is linear with respect to the DoFs, and the representation is either explicit or partially explicit.**

Idea

Choose a **nonlinear representation** given by a neural network. We replace a sum of simple functions with a **composition of simple functions**.

Important points

Finite-dimensional spaces associated to a nonlinear decoder are **not vector spaces**. So:

- the projector is not unique, and **the representations will be implicit**.
- **Existence and uniqueness?** algebraic system replaced with **non-convex optimization**.
- Convergence is harder to study and understand.

Nonlinear models

- Nonlinear version of classical models: f is represented by the DoF α_i , μ_i , ω_i or Σ_i :

$$f(x; \alpha, \mu, \Sigma) = \sum_{i=1} \alpha_i e^{(x-\mu_i)\Sigma_i^{-1}(x-\mu_i)}, \quad f(x; \alpha, \omega) = \sum_{i=1} \alpha_i \sin(\omega_i x)$$

- **Neural networks** (NN).

Layer

A layer is a function $L_l(\mathbf{x}_l) : \mathbb{R}^{d_l} \rightarrow \mathbb{R}^{d_{l+1}}$ given by

$$L_l(\mathbf{x}_l) = \sigma(A_l \mathbf{x}_l + \mathbf{b}_l),$$

$A_l \in \mathbb{R}^{d_{l+1} \times d_l}$, $\mathbf{b} \in \mathbb{R}^{d_{l+1}}$ and $\sigma()$ a nonlinear function applied component by component.

Neural network

A neural network is **parametric function obtained by composition** of layers:

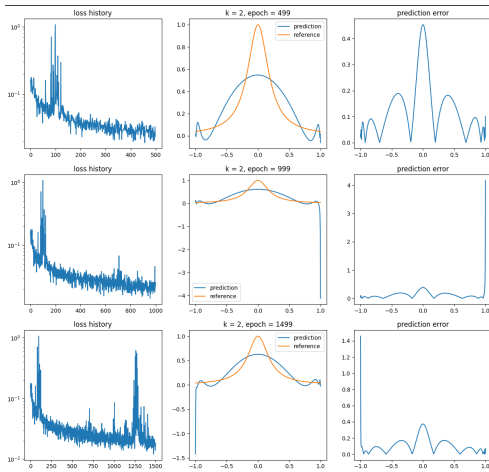
$$f_\theta(\mathbf{x}) = L_n \circ \dots \circ L_1(\mathbf{x})$$

with θ the trainable parameters composed of all the matrices $A_{l,l+1}$ and biases \mathbf{b}_l .

- **Goal:** using these models, we expect **to require fewer DoFs, not to require a mesh, and to deal with larger dimensions.**
- **Key point:** in the NN framework, **derivatives can be exactly computed through automatic differentiation tools.**


NN vs Polynomial

- We compare over-parametrized NN and polynomial regression on the Runge function.
- **Regression:** 120 data and approximately 800 parameters in each model.



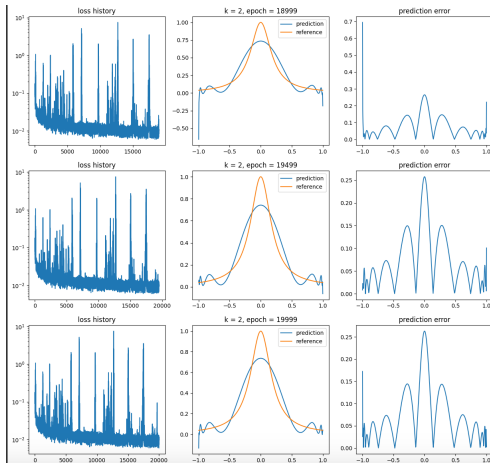
NN vs Polynomial

- We compare over-parametrized NN and polynomial regression on the Runge function.
- **Regression:** 120 data and approximately 800 parameters in each model.



NN vs Polynomial

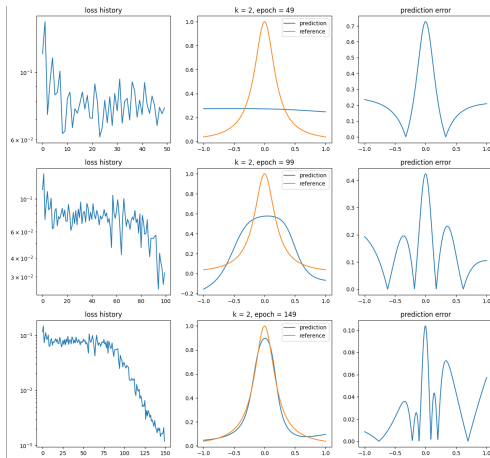
- We compare over-parametrized NN and polynomial regression on the Runge function.
- **Regression:** 120 data and approximately 800 parameters in each model.



- The polynomial model tends to oscillate in the over parameterized regime. Problematic for overfitting.

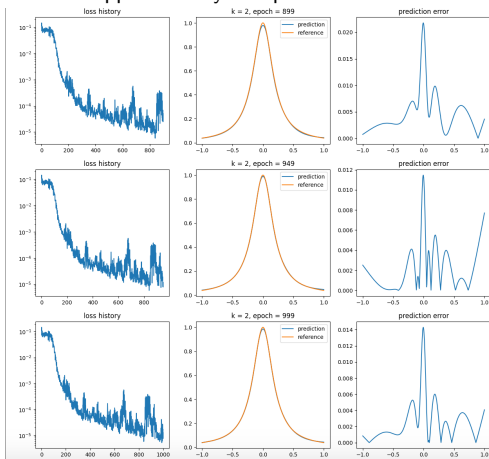
NN vs Polynomial

- We compare over-parametrized NN and polynomial regression on the Runge function.
- **Regression:** 120 data and approximately 800 parameters in each model.



NN vs Polynomial

- We compare over-parametrized NN and polynomial regression on the Runge function.
- **Regression:** 120 data and approximately 800 parameters in each model.



- The ANN generates very smooth/low frequency approximations.
- It is related to the **spectral bias**. The low frequencies are learned before the high frequencies.
- Seems very helpful to use it for global and high dimensional representation.

Physics-informed Neural Networks

Space-time approach: PINNs I

Idea of PINNs

- For u in some function space \mathcal{H} , we wish to solve the following PDE:

$$\partial_t u = \mathcal{F}(u, \nabla u, \Delta u) = F(u).$$

- Classical representation for space-time approach: $u(t, x) = \sum_{i=1}^N \theta_i \phi_i(x, t)$
- **Deep representation:** $u(t, x) = u_{nn}(x, t; \theta)$ with u_{nn} a NN with trainable parameters θ .

- Since ANNs are C^p functions, we can compute $\partial_t u_{nn}(x, t; \theta)$, $\partial_{x^p} u_{nn}(x, t; \theta)$ and

$$r(x, t) = \partial_t u_{nn}(x, t; \theta) - \mathcal{F}(u_{nn}(x, t; \theta), \nabla u_{nn}(x, t; \theta), \Delta u_{nn}(x, t; \theta))$$

- First idea: we solve the nonlinear problem

$$r(x_j, t_n) = 0, \quad \forall 1 \leq j \leq N_x, \quad \forall 1 \leq n \leq N_t$$

with $N_t * N_x$ equal to the number of parameters.

- **Problem:** The subspace of **NN functions is not a vector space**. The **existence** of the solution to the discrete problem cannot be ensured.

Conclusion

We move away from solving algebraic equations on the parameters, and go towards **non-convex optimization**.

Space-time approach: PINNs II

- We define the residual of the PDE:

$$R(t, x) = \partial_t u_{nn}(t, x; \theta) - \mathcal{F}(u_{nn}(t, x; \theta), \partial_x u_{nn}(t, x; \theta), \partial_{xx} u_{nn}(t, x; \theta))$$

- To learn the parameters θ in $u_{nn}(t, x; \theta)$, we minimize:

$$\theta = \arg \min_{\theta} \left(J_r(\theta) + J_b(\theta) + J_i(\theta) \right),$$

with

$$J_r(\theta) = \int_0^T \int_{\Omega} |R(t, x)|^2 dx dt$$

and

$$J_b(\theta) = \int_0^T \int_{\partial\Omega} \|u_{nn}(t, x; \theta) - g(x)\|_2^2 dx dt, \quad J_i(\theta) = \int_{\Omega} \|u_{nn}(0, x; \theta) - u_0(x)\|_2^2 dx.$$

- If these residuals are all equal to zero, then $u_{nn}(t, x; \theta)$ is a solution of the PDE.
- To complete the determination of the method, we need a way to compute the integrals.
- **Important point:** the derivatives are computed exactly using **automatic differentiation tools and back propagation**. Valid for any decoder proposed.

- How to compute the integrals? **With a Monte Carlo approach.**
 - The Monte-Carlo method stems from the **Law of large numbers**.
 - We consider a function $g : \mathbb{R}^d \rightarrow \mathbb{R}$. We define X a random variable with law μ .
 - The method comes from:

$$\frac{\text{Var}(\mu)}{\sqrt{N}} \left(\frac{1}{N} \sum_{i=1}^N f(X_i) - \mathbb{E}_{\mu}[f(X)] \right) \rightarrow \mathcal{N}(0, 1)$$

with X_i an random example sampled with the law μ

- It makes it possible to compute integrals. Indeed:

$$\int_{\Omega} f(x) dx = \int_{\mathbb{R}^d} f(x) \mathcal{U}_{\Omega} dx = \mathbb{E}[f(X)]$$

with \mathcal{U}_{Ω} the density of the uniform law Ω and X random variable following this law.

- The variance can be reduced through importance sampling:

$$\mathbb{E}[f(X)] = \int_{\Omega} f(x) dx = \int_{\Omega} \frac{f(x)}{g(x)} g(x) dx = \mathbb{E}_g \left[\frac{f(X)}{g(X)} \right]$$

- If $\text{Var}(\mathcal{U}_{\Omega}) > \text{Var}(g)$, the **error is reduced**.

Space-time approach: PINNs III

- We define the residual of the PDE:

$$R(t, x) = \partial_t u_{nn}(t, x; \theta) - \mathcal{F}(u_{nn}(t, x; \theta), \partial_x u_{nn}(t, x; \theta), \partial_{xx} u_{nn}(t, x; \theta))$$

To learn $u_{nn}(t, x; \theta)$, we minimize:

$$\min_{\theta} (J_r(\theta) + J_b(\theta) + J_i(\theta)),$$

with

$$J_r(\theta) = \sum_{n=1}^N \sum_{i=1}^N |R(t_n, x_i)|^2$$

with (t_n, x_i) **sampled uniformly or through importance sampling**, and

$$J_b(\theta) = \sum_{n=1}^{N_b} \sum_{i=1}^{N_b} |u_{nn}(t_n, x_i; \theta) - g(x_i)|^2, \quad J_i(\theta) = \sum_{i=1}^{N_i} |u_{nn}(0, x_i; \theta) - u_0(x_i)|^2.$$

- To avoid an extra loss for the BC and initial conditions, we use:

$$\bar{u}_{\theta}(t, x) = u_0(x) + t(\phi(x) * u_{\theta}(x)),$$




with $\phi(x) = g(x)$ on the boundary, and taking some other value within the domain.

- We solve the problem with usual gradient-type methods from Deep-Learning.




Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.1}{\pi}$. 10000 pts, medium-sized NN.
- beginning of training

```
iter = 200  
loss = 0.0774  
L2 error: 4.7142e-01
```






```
iter = 400  
loss = 0.0318  
L2 error: 4.0850e-01
```






Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.1}{\pi}$. 10000 pts, medium-sized NN.
- middle of training

```
iter = 2000  
loss = 0.0000  
L2 error: 9.0829e-03
```






```
iter = 2200  
loss = 0.0000  
L2 error: 8.2614e-03
```






Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.1}{\pi}$. 10000 pts, medium-sized NN.
- end of training

```
iter = 4800  
loss = 0.0000  
L2 error: 4.6718e-03
```






```
iter = 5000  
loss = 0.0000  
L2 error: 4.7307e-03
```






Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.01}{\pi}$. 10000 pts, medium NN.
- beginning of training

```
iter = 600  
loss = 0.0885  
L2 error: 4.3601e-01
```






```
iter = 800  
loss = 0.0233  
L2 error: 2.0901e-01
```






Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.01}{\pi}$. 10000 pts, medium NN.
- middle of training

```
iter = 2000  
loss = 0.0003  
L2 error: 1.8053e-02
```






```
iter = 2200  
loss = 0.0002  
L2 error: 1.7773e-02
```






Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.01}{\pi}$. 10000 pts, medium NN.
- end of training

```
iter = 4800  
loss = 0.0001  
L2 error: 5.9728e-03
```






```
iter = 5000  
loss = 0.0001  
L2 error: 5.2593e-03
```






Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.002}{\pi}$. 10000 pts, medium NN.
- beginning of training

```
iter = 200  
loss = 0.0774  
L2 error: 4.7142e-01
```






```
iter = 400  
loss = 0.0318  
L2 error: 4.0850e-01
```






Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.002}{\pi}$. 10000 pts, medium NN.
- middle of training

```
iter = 2000  
loss = 0.2076  
L2 error: 6.2666e-01
```






```
iter = 2200  
loss = 0.1361  
L2 error: 6.0138e-01
```






Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.002}{\pi}$. 10000 pts, medium NN.
- end of training

```
iter = 4800  
loss = 0.0272  
L2 error: 4.0909e-01
```






```
iter = 5000  
loss = 0.0212  
L2 error: 4.0300e-01
```






Example: viscous Burgers equation

- Application: viscous Burgers equation $\partial_t \rho + \partial_x \left(\frac{\rho^2}{2} \right) = \nu \partial_{xx} \rho$.
- Solving for different values of the μ parameters:
- $\nu = \frac{0.002}{\pi}$. 40000 pts, larger NN.
- end of training

```
iter = 4800  
loss = 0.0006  
L2 error: 2.3011e-01
```



```
iter = 5000  
loss = 0.0004  
L2 error: 2.2456e-01
```



PINNs for parametric PDEs

- **Advantages of PINNs:** mesh-less approach, not too sensitive to the dimension.
- **Drawbacks of PINNs:** they are often not competitive with classical methods.
- Interesting possibility: use the strengths of PINNs to solve PDEs parameterized by some μ .
- The neural network becomes $u_{nn}(t, x, \mu; \theta)$.


New Optimization problem for PINNs

$\min_{\theta} J_r(\theta) + \dots$, with

$$J_r(\theta) = \int_{V_{\mu}} \int_0^T \int_{\Omega} \|\partial_t u_{nn} - \mathcal{L}(u_{nn}(t, x, \mu), \partial_x u_{nn}(t, x, \mu), \partial_{xx} u_{nn}(t, x, \mu))\|_2^2 dx dt$$

with V_{μ} a subspace of the parameters μ .

- Application to the Burgers equations with many viscosities $[10^{-2}, 10^{-4}]$:



- Training for $\mu = 10^{-4}$: 2h. Training for the full viscosity subset: 2h.

Spatial approach: Neural Galerkin I

- We solve the following PDE:

$$\partial_t u = \mathcal{F}(u, \nabla u, \Delta u) = F(u).$$

- Classical representation: $u(t, x) = \sum_{i=1}^N \theta_i(t) \phi_i(x)$
- **Deep representation:** $u(t, x) = u_{nn}(x; \theta(t))$ with u_{nn} a neural network, with parameters $\theta(t)$, taking x as input.
- We want that:

$$F(u_{nn}(x; \theta(t))) = \partial_t u_{nn}(x; \theta(t)) = \left\langle \nabla_{\theta} u_{nn}(x; \theta), \frac{d\theta(t)}{dt} \right\rangle$$

- How to find an equation for $\frac{d\theta(t)}{dt}$?
- We solve the minimization problem:

$$\frac{d\theta(t)}{dt} = \arg \min_{\eta} J(\eta) = \arg \min_{\eta} \int_{\Omega} |\langle \nabla_{\theta} u_{nn}(x; \theta), \eta \rangle - F(u_{nn}(x; \theta(t)))|^2 dx.$$

- The solution is given by

$$M(\theta(t)) \frac{d\theta(t)}{dt} = F(x, \theta(t))$$

with

$$M(\theta(t)) = \int_{\Omega} \nabla_{\theta} u_{nn}(x; \theta) \otimes \nabla_{\theta} u_{nn}(x; \theta) dx, \quad F(x, \theta(t)) = \int_{\Omega} \nabla_{\theta} u_{nn}(x; \theta) F(u_{nn}(x; \theta)) dx.$$

Spatial approach: Neural Galerkin II

- How to estimate $M(\theta(t))$ and $F(x, \theta(t))$?
- **Firstly**: we need to differentiate the network with respect to θ and to x (in the function F). This can easily be done with automatic differentiation.
- **Secondly**: How to compute the integrals? **Monte Carlo approach**.
- So, we use:

$$M(\theta(t)) \approx \sum_{i=1}^N \nabla_{\theta} u_{nn}(x_i; \theta) \otimes \nabla_{\theta} u_{nn}(x_i; \theta)$$

and the same for $F(x, \theta(t))$.

- **Summary**: we obtain an **ODE in time (as usual) and a mesh-less method in space**.
- Like in the case of PINNs, we can apply this framework to parametric PDEs and larger dimensions.
- We solve the following PDE:

$$\partial_t u = \mathcal{F}(u, \nabla u, \Delta u, \alpha) = F(u; \mu).$$

- **Deep representation**: $u(t, x, \mu) = u_{nn}(x, \mu; \theta(t))$
- The solution is given by

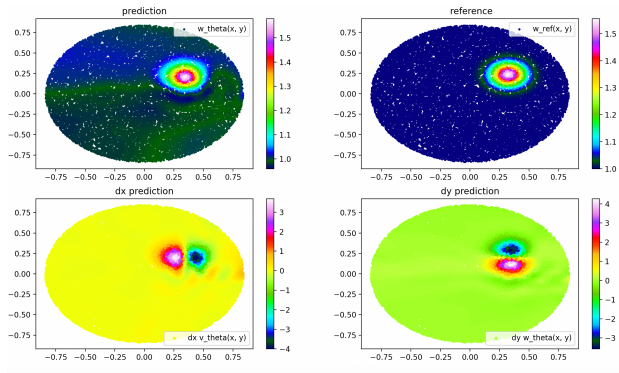
$$M(\theta(t)) \frac{d\theta(t)}{dt} = F(x, \theta(t), \mu)$$

with

$$M(\theta(t)) = \int_{V_{\mu}} \int_{\Omega} \nabla_{\theta} u_{nn}(x, \mu; \theta) \otimes \nabla_{\theta} u_{nn}(x, \mu; \theta) dx d\mu.$$

Spatial approach: Neural Galerkin III

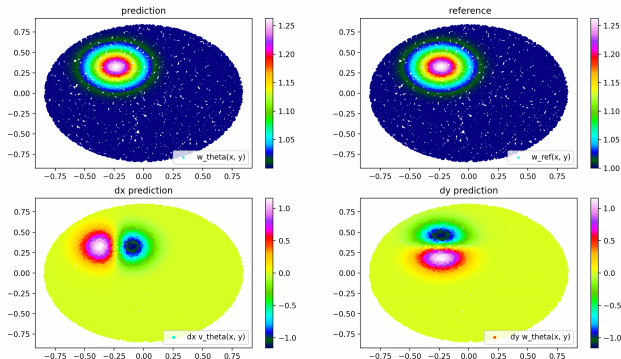
- We solve the advection-diffusion equation $\partial_t \rho + \mathbf{a} \cdot \nabla \rho = D \Delta \rho$ with a Gaussian function as initial condition.
- Case 1: with a neural network



- 3 minutes on CPU, L^2 error around 0.4. Bad initialization method, naive solve, many tricks to decrease the cost. Naive and small network.

Spatial approach: Neural Galerkin III

- We solve the advection-diffusion equation $\partial_t \rho + \mathbf{a} \cdot \nabla \rho = D \Delta \rho$ with a Gaussian function as initial condition.
- Case 2: with a Gaussian mixture (one Gaussian):



- 5 sec on CPU. Error around 5.0^{-4} . Decoder perfect to represent this test case.

Summary

New numerical methods

New numerical methods are derived using with neural networks.

- Same spirit as classical methods: plug an Ansatz into the equation to obtain equations on DoFs
- However, NN-based methods introduce a **nonlinear Ansatz with implicit representation**.

Many questions remain

- Accuracy? How to limit the GPU cost? Preconditioning? Domain decomposition?
- Positivity? capture of shocks? structure preservation (symplecticity)?
- **Convergence and stability?** (this is really trickier)

Drawbacks

- less accurate than classical approaches,
- convergence and theoretical study difficult,
- expensive in low dimensions and on CPU

Advantages

- mesh free
- more efficient in large dimension and for parametric PDEs, perfect for GPUs
- more general decoder coupled the exact computation of the gradient allows to impose some properties easier.

Operator Learning

Operator learning

- New paradigm for **reduced modeling**. We consider the following problem:

$$\begin{cases} G_{\alpha(x,t)}(u(t,x)) = \partial_t u(t,x) + \mathcal{L}_{\alpha(x)}(u(t,x)) = 0 & \text{on } \Omega, \\ u(t,x) = g(x) & \text{on } \partial\Omega, \\ u(t=0,x) = u_0(x). \end{cases}$$

- We denote by $\mu(t,x) = (\alpha(x,t), g(x), u_0(x))$ the parameters.
- Formally, there exists a pseudo-inverse operator G^+ , such that $G^+(\mu) = u(t,x)$.

Objective

Approximate G^+ by a neural network on a subspace of the data where the **results do not depend on the mesh resolution of the input/output functions**.

Problem

We construct a neural network $G_\theta^+(\mu(t,x))$, which minimizes $\mathcal{J}(\theta) = \mathcal{J}_1(\theta) + \mathcal{J}_2(\theta)$, with

$$\mathcal{J}_1(\theta) = \int_{\mathbf{V}_\mu} \int_{\Omega} \int_0^T \|G_\theta^+(t,x,\mu_h(x,t)) - \mathbf{u}(x,t)\|_2^2 dt dx d\mu$$

and

$$\mathcal{J}_2(\theta) = \int_{\mathbf{V}_\mu} \int_{\Omega} \int_0^T \|G_\mu(G_\theta^+(t,x,\mu_h(x,t)))\|_2^2 dt dx d\mu,$$

where the integrals are approximated by MC.

- **DeepOnet**: Encode the function of parameters, and decode with parametric PINNS.

Green theory and Neural operator

■ Beginning with a simple case:

- Linear elliptic equation: $\Delta u = f$. On unbounded domains, the solution is given by

$$u(x) = \int_{\mathbb{R}} G(x, y) f(y) dy$$

- Linear heat equation: $\partial_t u - \Delta u = f(t, x)$. On unbounded domains, the solution is given by

$$u(t, x) = \int_{\mathbb{R}} G_i(t, x, y) u_0(y) dy + \int_{\mathbb{R}} \int_{\mathbb{R}} G_e(t, \tau, x, y) f(\tau, y) dy d\tau$$

Idea for linear problem

For bounded (and potentially complex) domains, **learn the Green functions**.

- No Green theory for nonlinear PDEs. How can this be adapted to nonlinear PDEs?

Idea of Neural Operators

- To approximate nonlinear functions, NNs compose parametric linear maps and nonlinear local functions.
- To approximate nonlinear operators, **NOs compose parametric linear operators and nonlinear local functions**.

Neural operator II

Neural Operator layer

Consider the vector function given by the previous layer $\mathbf{v}_I(x)$. A kernel layer is given by

$$\mathbf{v}_{I+1}(x) = \sigma \left(W_I(t) \mathbf{v}_I(x) + \int_{\Omega} G_{\theta_I}(t, x, y) \mathbf{v}_I(y) dy \right)$$

with $W(t)$ a learnable weight matrix and G_{θ} a learnable kernel.

- We can also introduce a kernel like $G_{\theta}(t, \tau, x, y)$ for source terms, etc.
- In general, we add a first layer Q which transforms $u_0(x)$ into $\mathbf{v}_0(x)$ to increase the dimension of the function, and a layer performing the inverse transform at the end.

Question

How to learn this object independently of the discretization?

First possibility: MoDnet

- **Trainable parameters:** (W_I, θ_I) . G_{θ_I} is a classical network. In the space-time case, W_I is a network; when the equation does not depend on time, W_I is a matrix.
- **Evaluation of the integral:** **Monte Carlo** for large dimensions (also low rank or multigrid versions), or Gauss quadrature (for low dimensions).
- **Learning:** data + physics-informed loss function.
- General geometry. Large complexity, but computations are perfectly suited to GPUs.

Second possibility: Fourier approach

- We can use

$$\int_{\Omega} G(x, y) f(y) dy \approx \mathcal{F}^{-1}(\mathcal{F}(G(x, \cdot)) \mathcal{F}(f(\cdot)))$$

with \mathcal{F} the Fourier transform. It is exact for translation-invariant kernels (e.g. convolution).

- In practice, we use

$$\mathcal{F}^{-1}(K_{\theta} \mathcal{F}(f(\cdot)))$$

and we learn K_{θ} , limiting frequency to k_{max} .

- Faster than previous approach, but limited to Cartesian grids to use FFTs. There exist variants for general geometries.

More general Spectral approach

- We can use

$$\int_{\Omega} G(x, y) f(y) dy \approx \sum_{m=1}^M \langle K_{\theta_m} \phi_k(x) \rangle_{L^2} \psi_k(x)$$

- Variants include: Chebyshev or Legendre polynomials, Laplace transform, Fourier transform on manifolds, etc.

Application to numerical methods

Hybrid predictor-corrector methods

Hybrid methods

In this context, **hybrid methods** combine classical numerical methods and numerical methods based on **Implicit Neural representation** (IRM).

Objectives

Taking the best of both worlds: the accuracy of classical numerical methods, and the mesh-free large-dimensional capabilities of IRM-based numerical methods.

General Idea

- **Offline process:** train a Neural Network (PINNs, NGs or NOs) to **obtain a large family of approximate solutions**.
- **Online process:** **predict** the solution associated to our test case using the NN.
- **Online process:** **correct** the solution with a numerical method.

Predictor-Corrector Newton method for elliptic problems

- We consider a nonlinear elliptic problem:

$$u(x) - \alpha_0 \partial_{xx}(\alpha(x)|u(x)|^p \partial_x u) = f(x)$$

- To solve this PDE we use FE or FD solver + **Jacobian-Free Newton-Krylov method**.

Idea

Train a neural operator (Fourier neural operator on a large data set) and use its **prediction as an initial guess for Newton's method**.

- We only compare the average results:

mesh	$\alpha_0 = 2$ (40 sim)	$\alpha_0 = 5$ (25 sim)	$\alpha_0 = 8$ (25 sim)
100 cells	+500%	+1800%	+5000%
200 cells	+88%	+230%	+620%
400 cells	+82%	+150%	+220%
600 cells	+92%	+220%	+250%

Table: Comparison of the mean “gain” in CPU time for different values of α_0 .

- **Failures:** on all the tests, **we have 0% of fail** (our method being less efficient than the classical one) in terms of number of iterations, and **around 2% of fail** in terms of CPU time.
- On more refined meshes, the gain is smaller (the network acts only at the beginning of the convergence).
- In 2D, gains are also lower, and the classical method converges more quickly (why?).

Predictor-Corrector: using PINNs in a FE method

- We consider the following elliptic problem:

$$\begin{cases} Lu = -\partial_{xx}u + v\partial_xu + ru = f, & \forall x \in \Omega \\ u = g, & \forall x \in \partial\Omega \end{cases}$$

- We assume that we have a **continuous** prior of the solution given by a **parametric PINN** $u_\theta(x)$
- We propose the following corrections of the finite element basis functions:

$$u(x) = u_\theta(x) + p_h(x), \quad u(x) = u_\theta(x)p_h(x),$$

with $p_h(x)$ a perturbation discretized using **P_k Lagrange finite element**.

- For the **first approach (additive prior)**, we solve in practice:

$$\begin{cases} Lp_h(x) = f - Lu_\theta(x), & \forall x \in \Omega \\ p_h(x) = g - u_\theta(x), & \forall x \in \partial\Omega \end{cases}$$

- For the **second approach (multiplicative prior)**, we need $u_\theta(x) \neq 0$, so we take $M > 0$ and we solve:

$$\begin{cases} L(u_\theta(x)p_h(x)) = f, & \forall x \in \Omega \\ p_h(x) = \frac{g}{u_\theta(x)} + M, & \forall x \in \partial\Omega \end{cases}$$

Theory for hybrid EF

- Approach one: we rewrite the Cea lemma for $u_h(x) = u_\theta(x) + p_h(x)$. We obtain

$$\|u - u_h\| \leq \frac{M}{\alpha} \|u - u_\theta - I_h(u - u_\theta)\|$$

with I_h the interpolator. Using the classical result of P_k Lagrange interpolator we obtain

$$\|u - u_h\|_{H^m} \leq \frac{M}{\alpha} Ch^{k+1-m} \underbrace{\left(\frac{|u - u_\theta|_{H^m}}{|u|_{H^m}} \right)}_{\text{gain}} |u|_{H^m}$$

- Approach two: $u_h(x) = u_\theta(x)p_h(x)$. We use a modified interpolator:

$$I_{mod,h}(f) = \sum_{i=1}^N \frac{f(x_i)}{u_\theta(x_i)} \phi_i(x) u_\theta(x)$$

using $I_{mod,f}(f) = I_h\left(\frac{f}{u_\theta}\right)u_\theta(x)$, the Cea lemma and interpolation estimate we have:

$$\|u - u_h\|_{H^m} \leq \frac{M}{\alpha} Ch^{k+1-m} \underbrace{\left(\frac{|\frac{u}{u_\theta}|_{H^m} \|u_\theta(x)\|_{L^\infty}}{|u|_{H^m}} \right)}_{\text{gain}} |u|_{H^m}$$

- The prior must give a good approximation of the m^{th} derivative.

EF for elliptic problems

- First test:

$$-\partial_{xx}u = \alpha \sin(2\pi x) + \beta \sin(4\pi x) + \gamma \sin(8\pi x)$$

We train with $(a, b, c) \in [0, 1]^3$ and test with $(a, b, c) \in [0, 1.2]^3$.

method:	average gain	variance gain
additive prior with PINNs	273	13000
Multiplicative prior $M = 3$ with PINNs	92	4000
Multiplicative prior $M = 100$ with PINNs	272	13000
additive prior with NN	15	18
Multiplicative prior $M = 3$ with NN	11	17.5
Multiplicative prior $M = 100$ with NN	15	18

- The PINN is trained with the physical loss, the NN with only data, no physics.
- The NN is able to better learn the solution itself, but the approximation of derivatives is less accurate than with the PINN.

EF for elliptic problems

- Second test:

$$v \partial_x u - \frac{1}{P_e} \partial_{xx} u = r$$

We train with $r \in [1, 2]$, $Pe \in [10, 100]$. We test with $(r, Pe) = (1.2, 40)$ and $(r, Pe) = (1.5, 90)$

Case 1	Classical FE		Additive prior			Multiplicative prior		
	error	order	error	order	gain	error	order	gain
10	$1.07e^{-1}$	–	$2.70e^{-3}$	–	40	$2.29e^{-4}$	–	467
20	$3.36e^{-2}$	1.97	$8.00e^{-4}$	1.76	42	$9.06e^{-5}$	1.93	371
40	$9.09e^{-3}$	1.89	$2.01e^{-4}$	2.00	45	$2.63e^{-5}$	1.97	345
80	$2.32e^{-3}$	1.97	$5.01e^{-5}$	1.99	46	$6.37e^{-6}$	1.99	365
160	$5.82e^{-4}$	1.99	$1.30e^{-6}$	1.97	45	$1.77e^{-6}$	2.0	289



Case 2	Classic		additive prior			Multiplicative prior		
	error	order	error	order	gain	error	order	gain
10	$2.65e^{-1}$	–	$1.51e^{-1}$	–	1.7	$9.33e^{-4}$	–	284
20	$1.06e^{-1}$	1.32	$6.04e^{-2}$	1.33	1.7	$3.84e^{-4}$	1.28	276
40	$3.46e^{-2}$	1.62	$1.96e^{-2}$	1.62	1.8	$1.13e^{-4}$	1.76	305
80	$9.50e^{-3}$	1.86	$5.32e^{-3}$	1.87	1.8	$3.26e^{-5}$	1.80	291
160	$2.43e^{-3}$	1.86	$2.43e^{-3}$	1.86	1.8	$8.67e^{-6}$	1.91	280

Hyperbolic systems with source terms

- In the team, most of us are interested in hyperbolic systems:

$$\partial_t \mathbf{U} + \nabla \cdot \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U})$$

- It is important to have a good preservation of the steady state $\nabla \cdot \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U})$.
- **Example:** Lake at rest for shallow water:
- **Exactly Well-Balanced schemes:** exact preservation of the steady state.
Approximately Well-Balanced schemes: preserve with a high-accuracy than the scheme the steady state.
- Building exact WB schemes is difficult for some equilibria, or for 2D flows.



Idea

Compute offline a family of equilibria with parametric PINNs (or NOs) and **plug the equilibrium in the DG basis to obtain a more accurate scheme around steady states.**

Theory for hybrid DG

- Theory for the scalar case.
- The classical modal DG scheme uses the local representation:

$$u|_{\Omega_k}(x) = \sum_{l=0}^q \alpha_l \phi_l(x)^k, \text{ with } [\phi_1^k, \dots, \phi_q^k] = [1, (x - x_k), \dots, (x - x_k)^q]$$

- If $u_\theta(x)$ is an approximation of the equilibrium, we propose to take as basis:

$$V_1 = [u_\theta(x), (x - x_k), \dots, (x - x_k)^q], \text{ or } V_2 = u_\theta(x)[1, (x - x_k), \dots, (x - x_k)^q]$$

Lemma [Yuan Shu 2006]

Consider an nonlocal basis $(v_{k,0}, \dots, v_{k,q})$. If there exists constant real numbers $a_{j\ell}$ and b_j independent of the size of the cell Δx_k such that, in each cell Ω_k ,

$$\forall j \in \{0, \dots, q\}, \quad \left| v_{k,j}(x) - \sum_{\ell=0}^q a_{j\ell} (x - x_k)^\ell \right| \leq b_j (\Delta x_k)^{q+1},$$

then for any function $u \in H^{q+1}(\Omega_k)$, there exist a constant real number C independent of Δx_k , such that: $\|P_h(u) - u\|_{L^\infty(\Omega_k)} \leq C \|u\|_{H^{q+1}(\Omega_k)} (\Delta x_k)^{q+\frac{1}{2}}$.

- This lemma is sufficient to prove the convergence. Both bases satisfy the assumption.

Theory for hybrid DG

- Theory for the scalar case.
- The classical modal DG scheme uses the local representation:

$$u|_{\Omega_k}(x) = \sum_{l=0}^q \alpha_l \phi_l(x)^k, \text{ with } [\phi_1^k, \dots, \phi_q^k] = [1, (x - x_k), \dots, (x - x_k)^q]$$

- If $u_\theta(x)$ is an approximation of the equilibrium, we propose to take as basis:

$$V_1 = [u_\theta(x), (x - x_k), \dots, (x - x_k)^q], \text{ or } V_2 = u_\theta(x)[1, (x - x_k), \dots, (x - x_k)^q]$$

More accurate estimate

Assume that the prior u_θ satisfies

$$u_\theta(x; \mu)^2 > m^2 > 0, \quad \forall x \in \Omega, \quad \forall \mu \in \mathbb{P}.$$

and still consider the vector space V_2 . For any function $u \in H^{q+1}(\Omega)$,

$$\|u - P_h(u)\|_{L^2(\Omega)} \lesssim \left| \frac{u}{u_\theta} \right|_{H^{q+1}(\Omega)} (\Delta x_k)^{q+1} \|u_\theta\|_{L^\infty(\Omega)}.$$

- Adding a stability estimate, we can also prove the convergence.

Euler-Poisson system in spherical geometry

- We consider the Euler-Poisson system in spherical geometry

$$\begin{cases} \partial_t \rho + \partial_r q = -\frac{2}{r} q, \\ \partial_t q + \partial_r \left(\frac{q^2}{\rho} + p \right) = -\frac{2}{r} \frac{q^2}{\rho} - \rho \partial_r \phi, \\ \partial_t E + \partial_r \left(\frac{q}{\rho} (E + p) \right) = -\frac{2}{r} \frac{q}{\rho} (E + p) - q \partial_r \phi, \\ \frac{1}{r^2} \partial_{rr} (r^2 \phi) = 4\pi G \rho, \end{cases}$$

- **First application:** we consider the barotropic pressure law $p(\rho; \kappa, \gamma) = \kappa \rho^\gamma$ such that the steady solutions satisfy

$$\frac{d}{dr} \left(r^2 \kappa \gamma \rho^{\gamma-2} \frac{d\rho}{dr} \right) = 4\pi r^2 G \rho.$$

- The PINN yields an approximation of $\rho_\theta(x, \kappa, \gamma)$
- **Second application:** we consider the ideal gas pressure law $p(\rho; \kappa, \gamma) = \kappa \rho T(r)$, with $T(r) = e^{(-\alpha r)}$, such that the steady solutions satisfy

$$\frac{d}{dr} \left(r^2 \kappa \frac{T}{\rho} \frac{d\rho}{dr} \right) + \frac{d}{dr} \left(r^2 \kappa \frac{dT}{dr} \right) = 4\pi r^2 G \rho,$$

- The PINN yields an approximation of $\rho_\theta(x, \kappa, \alpha)$
- To simulate a flow around a steady solution, we need a scheme that is very accurate on the steady solution.

Results


- Training takes about 10 minutes on an old GPU, with **no data**, only the PINN loss.
- We take a quadrature of degree $n_Q = n_G + 1$ (sometimes, more accurate quadrature formulas are needed).
- Barotropic case:

q	minimum gain			average gain			maximum gain		
	ρ	Q	E	ρ	Q	E	ρ	Q	E
0	19.14	2.33	17.04	233.48	3.73	197.28	510.42	4.48	371.87
1	7.61	8.28	6.98	158.25	188.92	130.57	1095.68	1291.90	1024.59
2	0.14	0.22	2.99	12.11	16.55	23.73	89.47	109.93	169.28

- ideal gas case:

q	minimum gain			average gain			maximum gain		
	ρ	Q	E	ρ	Q	E	ρ	Q	E
0	13.30	1.05	16.24	151.96	1.88	150.63	600.13	2.91	473.83
1	6.30	7.53	5.40	72.63	77.20	51.09	321.20	302.58	257.19
2	3.35	3.45	2.20	18.96	22.58	13.56	55.47	63.45	47.83

- 2D shallow water equations: equilibrium with $\mathbf{u} \neq 0$ + small perturbation. Plot the deviation to equilibrium:



Conclusion

Conclusion

Short conclusion

Using **nonlinear implicit representations**, we proposed **new numerical/reduced modeling methods** whose advantages/drawbacks are very different to those of classical approaches. We will continue to investigate **hybrid approaches**.

Current work: Neural operators

We investigate the modification/extension of Neural Operator methods on general grids, to multiscale problems, and to preserve some structures (PEPR NUMPEX).

Current work: Continuous ROMs

- Using PINNs or Neural Galerkin approaches, we wish to construct discretization-independent **continuous ROMs**.

- Encoder:

$$E_{\theta}(f(x_1), \dots, f(x_n)) \rightarrow \beta \in \mathbb{R}^d$$

where the (x_1, \dots, x_n) is a random point cloud.

- Decoder:

$$D_{\theta}(\beta) = \sum_{i=1} \beta_i(t) \phi_{\theta_i}(x), \text{ or } D_{\theta}(\beta) = u_{\theta}(x; \beta)$$

- Coupling with Neural Galerkin, hyper-reduction and structure/property preserving approaches.

Scimba

- For the PEPR Numpex, we are currently writing the **Scimba** code. It contains for PINNs, Neural Galerkin, Neural operator methods, ...; the goal is for this code to be shared by different teams.
- If you are interested to try these methods, play with Scimba, or participate contact us!

Macaron

- Our Inria team TONUS/MACARON will specialize in the hybridation between ML and numerical methods for PDEs.
- We regularly have PhD, post-doc and even permanent positions open on these subjects. If you are interested, contact us :)

Main references

■ PINNs:

- *Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations*, M. Raissi, P. Perdikaris, G.E. Karniadakis
- *An Expert's Guide to Training Physics-informed Neural Networks*, S. Wang, S. Sankaran, H. Wang, P. Perdikaris
- *Estimates on the generalization error of Physics Informed Neural Networks (PINNs) for approximating PDEs*, S. Mishra, R. Molinaro

■ Neural Galerkin:

- *Neural Galerkin Scheme with Active Learning for High-Dimensional Evolution Equations*, J. Bruna, B. Peherstorfer, E. Vanden-Eijnden
- *A Stable and Scalable Method for Solving Initial Value PDEs with Neural Networks*, M. Finzi, A. Potapczynski, M. Choptuik, A. Gordon Wilson

■ Neural Operator:

- *Fourier Neural Operator for Parametric Partial Differential Equations*, Z.i Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, A. Anandkumar
- *Neural Operator: Learning Maps Between Function Spaces*, N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, A. Anandkumar
- *MOD-Net: A Machine Learning Approach via Model-Operator-Data Network for Solving PDE*, L. Zhang, T. Luo, Y. Zhang, Weinan E, Z. Xu, Z. Ma

■ Deep Predictor for Newton:

- *Accelerating hypersonic reentry simulations using deep learning-based hybridization (with guarantees)*, P. Novello, G. Poëtte, D. Lugato, S. Peluchon, P. Marco Congedo
- *DeepPhysics: a physics aware deep learning framework for real-time simulation*, A. Odot, R. Haferssas, S. Cotin
- *Accelerating Newton convergence for nonlinear elliptic PDE using neural operator approach*, E. Franck, R. Hild, V. Vigon, V. Michel-Dansac, J. Aghili. En cours de rédaction.

■ Hybrid methods:

- *Enhanced Finite element by neural networks for elliptic problems*, H. Barucq, E. Franck, F. Faucher, N. Victorion. En cours de rédaction
- *Approximately well-balanced Discontinuous Galerkin methods using bases enriched with Physics-Informed Neural Networks*, E. Franck, V. Michel-Dansac, L. Navoret. Arxiv preprint.